

Forthcoming in *Episteme*.

Speaking Sense: A Hybrid Source of Justification for Self-Knowledge

Daniel Gregory

University of Salzburg

Abstract

Nico Silins (2012, 2013, 2020) argues that conscious judgments justify self-attribution of belief in the content judged. In defending his view, he makes use of Moore's Paradox, seeking to show how his theory can explain what seems irrational or absurd about sentences of the form, ' p and I do not believe that p '. I show why his argument strategy is not available to defend the view that conscious judgments can justify the self-attribution of belief in the content judged. I then propose an amended version of his theory, which holds that sincerely asserting a proposition—whether aloud or silently—justifies self-ascribing belief in the proposition expressed. In doing so, I draw on an argument which I made in Gregory (2018) that there is something it is like to make a sincere assertion which is different from what it is like to make an insincere assertion. The phenomenology of sincere assertion provides *immediate* justification for self-ascription of belief in a proposition which has been sincerely asserted; nonetheless, it may be that we need to *interpret* our own assertions in order to determine which propositions they express. This paves the way for showing how two competing schools of thought about self-knowledge—one which holds that self-knowledge is immediate and one which holds that self-knowledge is inferential—might be combined.

1. Introduction

Moore's Paradox is that there are conjunctions of particular kinds which it seems irrational to assert or believe even though they are not contradictory, e.g., ' p and I do not believe that p '. It is relevant to the epistemological question of how we can be justified in self-ascribing beliefs for the following reason. If it is not the logical form of a sentence such as ' p and I do not believe that p ' which makes it seem irrational to assert or believe it, it must be something else. But it is also not anything to do with the content of the sentence or the context in which it is asserted or believed; one can substitute any proposition for p in any context.¹ Rather, the reason that it strikes us as somehow irrational to assert or believe such a sentence is that it involves denying—or at least being disposed to deny—that one is in a particular mental state when one is plainly not justified in making such a denial. Many existing theories of self-knowledge offer explanations of Moore's Paradox (e.g., Shoemaker (1995), Williams (2004),

¹ Or almost any proposition in almost any context. See footnote 16.

and Fernández (2005)). Accordingly, a good check of the basic plausibility of any theory of self-knowledge is whether it can offer a plausible treatment of Moore’s Paradox. If not, it is probably out of the running.²

Nico Silins (2012, 2013, 2020) applies roughly this strategy.³ Silins holds that consciously judging that a proposition is true provides justification for self-ascribing belief in that proposition. To see the significance of this suggestion, one has to notice how it runs against conventional wisdom. Silins observes that:

[a]lthough there is much disagreement about self-knowledge, philosophers in the literature currently tend to agree that it should not be understood on the model of perception. They focus on the metaphysics of perceptual states, and on the absence of appropriately similar states in the case of introspection. The idea is that there is no good sense in which we perceive our mental states. (2012: 323)

Within this context, Silins wants to emphasize that we should not lose sight of the possibility that experiences can justify self-attributions of beliefs, just as experiences can justify first-order beliefs. Perceptual experiences can justify beliefs about the external world. Conscious judgments are also experiences. Silins is urging that these experiences can justify beliefs that we have certain beliefs.

Silins believes his view can explain the problem with a sentence such as ‘*p* and I do not believe that *p*’—which he calls ‘MP’—as follows. Consciously judging MP entails judging that *p*. For Silins, judging that *p* justifies believing ‘I believe that *p*’. ‘The overall upshot,’ he writes, ‘will be that, when one judges the MP conjunction, one judges the second conjunct, while having justification to believe the negation of the second conjunct’ (2012: 303). This does seem irrational.⁴ It is an elegant and eye-opening piece of philosophy.

I am going to suggest that there is nonetheless a problem with it, though I will then offer an amended version of the account which avoids the problem. In developing this amended position, I will draw on an observation which I offered in Gregory (2018) that there is a phenomenological difference between producing a sincere assertion and producing an insincere assertion. This phenomenological difference may be a difference in cognitive phenomenology; at any rate, it is a difference in non-sensory phenomenology. The amended version of Silins’s theory holds that making a sincere assertion (rather than a judgment), combined with experiencing the distinctive phenomenology of making a sincere assertion, justify the self-ascription of belief.

There are, of course, two versions of Moore’s Paradox. The version I have set out, involving sentences of the form, ‘*p* and I do not believe that *p*’, is called the ‘omissive’ version. The other version, involving sentences of the form, ‘*p* and I believe that not-*p*’, is called the ‘commissive’ version. Silins

² Thanks to both Luke Roelofs and an anonymous referee for this.

³ I say ‘roughly’ because Silins claims that we can ‘support’ (2012: 302) or ‘defend’ (2020: 333) his theory by demonstrating its application to Moore’s Paradox. Perhaps these locutions—especially ‘support’—suggest something stronger than the strategy just described in the text. But this makes no difference to anything that follows.

⁴ I deal with some complications after describing the argument again in Section 2.

only discusses the omissive version. I will focus on the omissive version as well, but I will also show how the theory which I offer explains the commissive version as well.

Traditionally, the theory that justification for self-knowledge is immediate,⁵ based only on something like the experience of cognitive phenomenology, and the theory that self-knowledge depends on interpreting one's speech and behaviour, are rivals. Within this context, the amended version of Silins's account which I will suggest is significant, as it shows how aspects of the two theories can be combined. This compels a rethink of current approaches to self-knowledge.⁶

2. Evans, Moore, and Silins

Gareth Evans (1982) famously wrote the following:

[I]n making a self-ascription of belief, one's eyes are, so to speak, or occasionally literally, directed outward—upon the world. If someone asks me 'Do you think there is going to be a third world war?', I must attend, in answering him, to precisely the same outward phenomena as I would attend to if I were answering the question, 'Will there be a third world war?' I get myself in a position to answer the question whether I believe that *p* by putting into operation whatever procedure I have for answering the question whether *p*. (225).

This is the classic statement of the 'transparency' view of self-knowledge.⁷ What the classic statement offers, though, is an account as to how one comes to know that they have a particular belief, not an account as to how one might be justified in self-ascribing that belief. These are different questions.

Nico Silins (2012, 2013) notes that the transparency view as articulated by Evans has certain vulnerabilities precisely because it is an account of how we come to know our own mental states.⁸ One issue Silins highlights is that it seems inapt to begin an inquiry into some matter if one wants to know what they already believe about it, as the inquiry might result in the formation of a different belief. '[I]f you wondered whether you had a prior belief that God exists', Silins writes, 'you shouldn't answer that question by considering considerations for and against the existence of God' (2012: 300).

Silins frames his own topic of interest as follows:

I will largely bracket questions about your access to your antecedent beliefs. My main focus will instead be on the situation once you do answer the question whether *p*, and how you stand with respect to your current beliefs once you do. Since I am bracketing questions about one's prior beliefs, I am also interested in the introspective upshot of answering the question whether *p*, whether or not you started out wondering about your beliefs before you asked the question whether *p*. (2012: 300; see also 2013: 295)

⁵ By immediate, I mean 'non-interpretative'. Silins makes use of this notion in making a refinement to his own view, which I will come to in Section 5.1. Citing Pryor (2005), he writes: 'you have immediate justification to believe that *p* just in case you have justification to believe that *p*, and you do so in a way that does not rely on your justification to hold any other belief' (Silins 2012: 298).

⁶ I am extremely grateful to Melissa Fusco and the Australian National University Work-in-Progress Group (at a meeting in November 2016) for suggesting this way of looking at the ideas in the paper.

⁷ Silins (2020) writes that, '[w]hile commonly assigned to Evans (1982), the idea of Transparency arguably traces back to Edg[e]y' (333, footnote 11; see also Silins (2012: 297, 2013: 293).

⁸ He refers to Peacocke (1999), Shah & Velleman (2005), and Byrne (2005); see Silins (2012: 300, footnote 6); and see also Silins (2013: 295, footnote 6), citing Shah & Velleman (2005) and Boyle (2009).

He provides the following provisional statement of his position: ‘If you judge that p , then you have justification to believe that you believe that p ’ (2012: 302).⁹ (Silins makes some refinements to this provisional statement. I will return to these, but it will be convenient initially to work with this provisional statement.) So, his concern is with justification for the self-ascription of beliefs, rather than with any procedure for self-ascribing beliefs. In particular, his concern is with propositional justification, rather than doxastic justification, for the self-ascription of beliefs. That is, he is concerned with the kind of justification which is available for the self-ascription of a belief, even if no such self-ascription is made. Still, his theory is explicitly derived from Evans, insofar as he is concerned with the position you are in ‘with respect to your current beliefs’ once you ‘answer the question whether p ’, i.e., once you consciously judge whether p (2012: 300).

What are conscious judgments? For Silins, they are ‘those which modify what it is like for you at the time you make them’ (2012: 298). He elaborates:

To make a judgment of the kind I am interested in, you might sincerely assert to someone that p . But you can consciously judge that p without performing the linguistic act of assertion. In many cases conscious judgment will be the ‘inner analogue’ of assertion, although it may well be that one can judge that p without in any way vocalizing or imagining a sentence with the content that p . (2012: 298).

Consider the words, ‘[i]n many cases conscious judgment will be the “inner analogue” of assertion’. I take it Silins means the ‘inner analogue’ of *sincere* assertion, given that he explicitly speaks of sincere assertion immediately prior in the case of external speech. I also take it that, by “‘inner analogue’ of assertion’, he means (sincere) assertions made in inner speech, given that, directly afterwards, he contrasts judgments made in this way with judgments made ‘without in any way vocalizing or imagining a sentence’, i.e., judgments made without inner speech. This is more explicit in Silins (2020: 333).

Another feature of conscious judgments is that they involve consciously taking something to be true. If you consciously judge that p , you take p to be true. Reasonably enough, Silins leaves this mundane point unsaid. I make it explicit here, however, as it will be important in what follows.

I described in the Introduction how Silins makes use of Moore’s Paradox. Assuming that his theory is at least *prima facie* plausible, he seeks to show that it can explain the irrationality involved in making a judgment with the form of MP. If his theory can explain this, there would at least be reason to examine it further.¹⁰

There is actually something a little unusual about the way Silins sets up Moore’s Paradox. He seeks to show that his account can explain the irrationality involved in making a conscious *judgment* of the form MP. Canonically, Moore’s Paradox is set up not in terms of judgments but in terms of assertions, whether sincere or insincere, or beliefs. Silins holds that judgments can take the form of sincere assertions but he also allows that they might take other forms: ‘it may well be that one can judge

⁹ In the following, I will primarily refer to Silins’s (2012) presentation of his theory.

¹⁰ Again, perhaps he says something a little stronger. See footnote 3.

that p without in any way vocalizing or imagining a sentence with the content that p ' (298). In any event, to recapitulate, he reasons as follows.

Suppose you judge that a sentence of the form MP is true. That is, you make the conjunctive judgment: ' p and I do not believe that p '. For Silins, if you judge that p (the first conjunct), then you have justification to believe that you believe that p . So, when you come to judge 'I do not believe that p ' (the second conjunct), you already have justification to believe its negation. This means that judging that a sentence of the form MP is true entails judging that you do not believe a proposition when you in fact have justification to believe that you believe that proposition, which does seem irrational. As Silins says, if you judge that a sentence of the form MP is true, 'you flout the evidence provided by your judgment that p ' (2013: 297). So the account yields the correct result.

Notably, on Silins's account, if you judge that p , then, even if you also have some other source of justification to believe the second conjunct, it would nonetheless be irrational to judge that the second conjunct is true. You now have some justification to believe that p , and you have some justification to believe that you do not believe that p . In such a situation, one should presumably refrain from making any judgment as to whether the second conjunct is true, rather than judging that it is. Accordingly, judging that a whole sentence of the form MP is true will be irrational, even if one has some independent justification to believe that the second conjunct is true.

A small caveat Silins places on his theory is that the justification he is describing is only *prima facie* justification. 'When one judges that p ,' he writes, 'one's justification to believe that one believes that p might well be defeated, say by opposing testimony from one's therapist' (2012: 302; see also 2013: 297–298). The difference between this case and the one described in the previous paragraph is that, in this case, one receives not only some evidence against their belief that they believe that p , but defeating evidence. So, they will have lost their justification to believe that they believe that p entirely. But such circumstances will be 'highly unusual' and 'exceptional' (2013: 297), according to Silins. So judging that p may not always justify the belief that one believes that p , but it will almost always do so. After making this caveat explicit, Silins omits further reference to it, asking the reader to take it as assumed.¹¹

¹¹ Notwithstanding Silins's caveat, one might ask: Is it too strong to say that judging that you do not believe a proposition when you in fact have (even *prima facie*) justification to believe that you believe that proposition seems irrational? A referee points out that we are sometimes rationally mistaken, so perhaps we can be rationally mistaken about the justification which is available to us.

It is true that rational mistakes are sometimes possible, but they are not always possible. If we have strong, clear evidence that we believe a proposition, then it likely will be irrational to judge that its negation is true. And Silins does seem to think that judging that p provides strong evidence that we believe that p , because he thinks there is a strong connection between judgment and belief: 'judgment is a guide to belief', as he puts it, appealing to the 'correlation between judging that p and having justification to believe you believe that p ' (2012: 306). This is perhaps why he says that judging the second conjunct to be true after judging the first conjunct to be true involves '*flout[ing]* evidence' provided when one judges the first conjunct (2013: 297, my emphasis), not merely neglecting or misunderstanding it. Flouting evidence is, presumably, irrational. In any case, I believe that this would be a Silinsian response to the challenge.

I am very grateful to Brett Topey for assistance in addressing this issue.

In any case, I think there is a problem with Silins's argument. Again, for Silins, conscious judgments might 'modify what it is like for you at the time you make them'—but conscious judgments also involve taking something to be true.¹² I suggest that there is no way in which consciously judging 'p and I do not believe that p' can modify what it is like for you at the time of judging, because it is simply not possible to consciously judge 'p and I do not believe that p'. You *can* believe that p and also believe that you do not believe that p, but with an important caveat: it is only possible if you are not aware that you believe that p and also believe that you do not believe that p (see Shoemaker (1995: 213–214) for related remarks). You might, for example, believe that a particular acquaintance is an unpleasant person. Your behaviour in generally avoiding them would be evidence of this belief. As well, you might not believe that you believe this person is unpleasant. Perhaps you had previously been fond of them and do not realise that you have changed your mind.¹³ Still, there is something wrong, if not quite absurd, about holding such a pair of beliefs. As soon as you become aware that you have been holding such a pair of beliefs, you cannot continue doing so. If you are not immediately aware which one is false, or if neither of them immediately seems false, then you suspend belief in both. If this were not so—if it were not unbearably irrational to persist in holding such a pair of beliefs once you become aware of it—then the version of Moore's Paradox which applies to beliefs would not have the force which it does. This shows us why it is not possible to consciously judge, e.g., that a person is unpleasant and that you do not believe that they are unpleasant. Consciously realising that you have been holding such a pair of beliefs compels you to adjust at least one of your beliefs; you cannot just continue holding both of them. But consciously judging any proposition of the form MP would amount to doing virtually the same thing: consciously taking both p and 'I do not believe that p' to be true at the same time. It would require the same degree of irrationality—a degree of irrationality which it seems we simply cannot tolerate, at least when we are conscious of it.¹⁴

In a word, then, Moore's Paradox simply does not arise within the context of conscious judgments, because the relevant conscious judgments are not possible.¹⁵ Silins fails to show that his theory gives an explanation of the irrationality involved in a version of Moore's Paradox, because he focusses on a version that does not really exist. The theory does not pass this plausibility test for theories of self-knowledge.

¹² Silins (2012, 2020) emphasises that he is not thinking of conscious judgments as a kind of conscious belief (in which case, consciously judging that p would entail believing that p). Nor am I. Like Silins, I take a belief to be a standing mental state which requires having dispositions to act in certain ways, and I take a judgment to be an episodic mental state which, in itself, does not require having any behavioural dispositions. This is why it is possible to make judgments which are inconsistent with (unconscious) beliefs. Nonetheless, I take it that making a judgment does involve taking something to be true, just as holding a belief involves taking something to be true.

¹³ Thanks to Frank Jackson for this point and suggesting an example from which this one derives.

¹⁴ I address an objection in Section 4.1.

¹⁵ On these last few points, cf. Shoemaker (1995).

3. A Different Route

It is easy to see how Moore's Paradox arises in the case of assertions as well as in the case of beliefs. One *can* assert ' p and I do not believe that p '. Of course, one cannot do so sincerely, precisely because one cannot consciously take ' p ' and 'I do not believe that p ' to be true at the same time. Such an assertion will be insincere—or, if one prefers, saying ' p and I do not believe that p ' will involve making one sincere assertion and one insincere assertion. Indeed, Moore's original example involves an assertion. He writes that 'to say such a thing as "I went to the pictures last Tuesday, but I don't believe that I did" is a perfectly absurd thing to say' (Moore 1942: 543).¹⁶

Silins writes that '[t]o make a judgment of the kind I am interested in, you might sincerely assert to someone that p ' (2012: 298). Even supposing this is true, Moore's Paradox would still not arise in the case of judgments, for the reason that one cannot sincerely assert ' p and I do not believe that p ' (notwithstanding that you can do so insincerely). Whereas judging that p necessarily involves taking p to be true, asserting that p does not; the assertion can be insincere. If this were not so—if assertions were necessarily sincere—then the paradox could not arise in the case of assertions either. (This point that assertions are not necessarily sincere bears emphasizing, as it is essential to what follows. So: One can assert something which one does not believe!)

Shoemaker (1995) approaches the issue differently from Moore, emphasizing beliefs. He points out that 'there is something paradoxical or logically peculiar about the idea of someone's believing the propositional content of a Moore-paradoxical sentence, whether or not the person gives linguistic expression to this belief' (1995: 213). He also proposes, more contentiously, that a solution to the paradox about belief would indirectly provide a solution to the paradox about assertion.

I will set aside the question of whether the paradox about asserting a sentence such as MP and the paradox about believing it are distinct problems or whether one reduces to the other. I am going to suggest a revised version of Silins's theory which generates a solution to the version of Moore's Paradox which relates to assertions, no matter how that version relates to the other.

A short detour is now necessary. In Gregory (2018), I argued that what it is like to make an assertion varies depending on whether or not the assertion is sincere, i.e., on whether or not you believe what you assert. The argument involves the strategy of phenomenal contrast (on which, see Bayne & Montague (2011), Siegel (2010)). I described two cases:

1. Candice, a meteorologist, is explaining the La Niña weather system to her friend, Cameron. At one point, Cameron says, 'This is interesting'. Cameron is enjoying the impromptu lesson; he does believe that the topic is interesting.

¹⁶ Crimmins (1992), Turri (2010), Pruss (2012), Coliva (2015) and Fileva & Brakel (2019) offer examples of cases where, arguably, one can sincerely assert a sentence which has the form MP or a very similar form. But the examples are very specific cases. The point that one cannot sincerely assert a sentence with the form MP is a very general one.

2. Candice, a meteorologist, is explaining the La Niña weather system to her friend, Cameron. At one point, Cameron says, ‘This is interesting’. Cameron is *not* enjoying the impromptu lesson; he *does not* believe that the topic is interesting. (2018: 227).

Cameron makes the same assertion in Cases 1 and 2; he believes it in Case 1 but not in Case 2. I take it to be clear that Cameron’s subjective experience will differ in the two cases. What it is like to make a sincere assertion is different from what it is like to make an insincere assertion, even when the assertion is exactly the same.¹⁷

As I also argued in that paper, what it is like to make a sincere assertion in inner speech is relevantly similar to what it is like to make a sincere assertion in external speech. There are obviously differences, insofar as one involves performing the physical motions necessary to speak aloud and hearing the sounds thereby generated, and the other involves producing and experiencing auditory imagery. However, if Cameron had produced the utterance, ‘This is interesting’, silently in Case 1, he would still have had the phenomenology of producing a sincere assertion.¹⁸ All of this applies regardless of whether the belief was held prior to one’s making the assertion or whether the belief forms as one makes the assertion. It also applies, *mutatis mutandis*, to insincere assertions in inner speech and external speech, which means that the phenomenological difference between sincere and insincere assertions in inner speech parallels the phenomenological difference between sincere and insincere assertions in external speech.

My claim in the (2018) paper is supported only by introspection. One will be convinced only if their own phenomenology varies depending on whether they make a sincere or an insincere assertion. So it is with all claims about phenomenology. The claim still seems clearly true to me, but slightly extending another point from the (2018) paper might help to make the phenomenal contrast more apparent for those in doubt. The phenomenology of sincerely asserting is constant across different sincere assertions, even though the sensory phenomenology will change; *mutatis mutandis* for insincere assertions. If you sincerely assert ‘snow is white’ and sincerely assert ‘grass is green’, there will be an aspect of your phenomenology which remains the same, even though the sounds you produce and hear will differ (or, in the context of inner speech, even though the auditory-imagistic sensations which you produce and experience differ). This non-sensory phenomenology is the phenomenology of sincerely asserting. The non-sensory phenomenology of sincerely asserting (or of insincerely asserting) does not vary with the content of the assertion made. This may make the phenomenological contrast between sincerely and insincerely asserting easier to identify for some. If one struggles to notice a

¹⁷ There are some related suggestions about the phenomenology of speech acts in Prinz (2007) and Cassam (2014). See Gregory (2018: 233, footnotes 10 and 14) for discussion of the difference between those suggestions and the suggestion being reviewed here. Cf. also Jorba (2016), who holds that cognitive attitudes (as opposed to the expression of them in speech) have a phenomenology which allows us to distinguish between them (referring also to Kim (1996)), and Munroe (2023) on the feeling of rightness which one can experiencing inner speech utterances.

¹⁸ In a later paper, I question whether utterances in inner speech are speech acts at all (Gregory 2020). I will not take up this issue here, and will continue to refer to ‘assertions’ in inner speech, setting aside whether or not they are not speech acts in a technical sense.

phenomenological difference between making a particular assertion depending on whether or not one believes it, one can look for a phenomenological feature which is present across sincere assertions and look for a phenomenological feature which is present across insincere assertions, and contrast these. Whether this helps, of course, will be an individual matter.

The strategy of phenomenal contrast is commonly used to argue for the existence of a distinctively cognitive phenomenology, i.e., phenomenology ‘that is associated with cognitive activities, such as thinking, reasoning, and understanding’ (Smithies 2013: 744). A standard example involves a scenario in which two people hear an utterance in a language which only one of them understands. Their sensory experiences will be the same. Nonetheless, their experiences will differ. One will have the experience of understanding and the other will not. So, there will be a difference in the phenomenology of their cognitive experiences—or so believers in cognitive phenomenology claim. This example, or versions of it, have been used regularly; Bayne & Montague (2011) cite Moore (1953), Strawson (1994), and Siewert (1998) as examples.

The existence of cognitive phenomenology is contentious. Some who deny its existence have claimed that the phenomenological difference in pairs of cases like the above are in fact differences in sensory phenomenology. Tye & Wright (2011), for example, emphasise that the ‘phonological processing of the sound stream is different from the processing that goes on when we hear someone speaking in a language we do not comprehend’ (2011: 337; see also Siegel (2010)). Assuming this is so, the auditory experience of the individual who understands the language might be different from the auditory experience of the individual who does not.

On the face of it, the difference between what it is like to make a sincere assertion and what it is like to make an insincere assertion is a difference between what it is like to perform one kind of action and what it is like to perform another kind of action. What distinguishes the action of sincerely asserting from the action of insincerely asserting is the mental state one is in at the time. Still, the difference in phenomenology could not be called a difference in *cognitive* phenomenology (assuming there is such a thing), at least in the case of assertions made in external speech, because asserting externally is not a purely mental activity. What it is like to make sincere (or insincere) assertions externally would be a matter of what it is like to perform a particular kind of action which has both physical and mental components.

An alternative analysis would be that the difference between what it is like when one asserts sincerely and what it is like when one asserts insincerely, whether internally or externally, is just a matter of what it is like to be in the respective mental states. If so, then the difference would be a difference in purely cognitive phenomenology. Especially in the case of inner speech, it seems very unlikely that there could be a difference in one’s sensory—or, rather, imagistic—experience, depending on whether a particular assertion is made sincerely or insincerely, so the Tye & Wright-style analysis on which differences in apparently cognitive phenomenology can ultimately be explained as differences in sensory phenomenology is not available.

In the present context, it is not essential to determine conclusively the nature of the phenomenological difference, though the answer might well be interesting. What really matters is that there is a difference in some form of non-sensory phenomenology when one produces sincere and insincere assertions, whether in inner speech or in external speech. This much does seem clear.

We can now state the revised version of Silins's theory. It holds that, if one asserts some proposition, externally or internally, and experiences the phenomenology which is distinctive of asserting sincerely—which I will call 'S-phenomenology' from here on—they will then have justification for self-ascribing belief in that proposition. Let us apply this revised theory to Moore's Paradox. Suppose you make an assertion of the form MP, externally or internally. If you believe p , you will experience S-phenomenology when you assert the first conjunct, and you will have justification for believing that you believe that p . If you then assert 'I do not believe that p ', you will be contradicting a proposition that you have justification to believe, viz., 'I believe that p '.¹⁹ This is why asserting MP seems irrational. Whereas we cannot judge MP, we can assert it, precisely because we can assert things which we do not believe, but this does not mean that asserting it will seem rational; it will not. The revised version of Silins's theory succeeds in delivering a solution to Moore's Paradox as it applies to assertions.

4. Matters Arising

Before making some refinements to this revised version of Silins's theory, here is a convenient point to address two issues.

4.1 The Impossibility of Judging MP

One might object to my claim that it is not possible to judge MP. I said that we cannot consciously take both p and 'I do not believe that p ' to be true at the same time; it would require the same degree of irrationality as consciously taking both p and 'not- p ' to be true at the same time, and this degree of irrationality is not possible for us. But, if this is wrong, then my argument against Silins's account—that it wrongly assumes that we can judge MP—fails. So what if someone were to say that they are capable of judging MP? Why should we not believe them?

I am inclined to be slightly concessive here. I think it would be extraordinary if someone truthfully claimed that they can simultaneously, consciously judge both conjuncts of MP to be true, but suppose someone does. If they really can simultaneously, consciously judge both conjuncts of MP to be true, then I suggest that it will also not seem irrational to them to do so. Rationality, to the extent that one has it, imposes restrictions on what one can take to be true. This is the least that rationality is. So, someone who can consciously take MP to be true will be someone who finds that it does not violate

¹⁹ Of course, you will *not* experience S-phenomenology when you assert the second conjunct.

any principles of rationality. But this means that Moore's Paradox will not arise for such a person in the context of judgments (and maybe not even in the context of beliefs), but not because it is not possible for them to judge MP. It will not arise for them because it does not seem irrational to them to do so. Still, Moore's Paradox will not arise for them. So, Moore's Paradox will not arise in the context of judgments whether one thinks that it is not possible to judge MP (as I do), or whether one thinks that it is possible. And my complaint about Silins's account is just that: MP does not arise in the context of judgments. This is why the issue needs to be approached differently.

4.2 The Commissive Form

As I mentioned in the Introduction, there are two versions of Moore's Paradox. There is the omissive version which Silins focusses on and which I have been discussing: sentences of the form, ' p and I do not believe that p '. And there is the commissive version: sentences of the form, ' p and I believe that not- p '. The account I am proposing can also explain the irrationality involved in asserting sentences of the form, ' p and I believe that not- p '.

To see this, bear in mind that belief reverse-distributes through conjunction (or 'factors out' or 'agglomerates', in alternative terminologies)—if the number of beliefs involved is small. To say that belief distributes through conjunction is to say that, from 'I believe that a and b ', one can derive 'I believe that a and I believe that b '. To say that belief reverse-distributes through conjunction is to say that, from 'I believe that a and I believe that b ', one can derive 'I believe that a and b '. Belief is usually thought not to reverse-distribute through conjunction where the number of beliefs involved is large. Makinson's (1965) 'Preface Paradox' demonstrates this: An author can believe every statement they make in a book they have written taken independently but, recognising their human fallibility, acknowledge that they have likely made a mistake somewhere in the book. This being so, they can believe every statement in the book but not believe their conjunction. This seems rational. Critically, though, what gives the Preface Paradox force is that it involves large numbers of belief. It is *very* doubtful that a case can be constructed in which just a pair of beliefs do not reverse-distribute through conjunction (though more on this in a moment).

Now, if you assert p and experience S-phenomenology, you will be justified in self-ascribing belief in p . That is, you will be justified in believing 'I believe that p '. If, after attaining justification to believe 'I believe that p ', you then assert 'I believe that not- p ', you are not denying something which contradicts 'I believe that p '. If you were, then the commissive version of Moore's Paradox could be addressed in the same way as the omissive version. However, you are asserting a proposition which, together with the proposition 'I believe that p ' (which you have justification to believe), and on account of the reverse-distributivity of small numbers of belief, entails 'I believe that p and not- p '. This is why

asserting ‘ p and I believe that $\text{not-}p$ ’ seems irrational: it amounts to self-ascribing belief in both a proposition and its negation, i.e., a contradiction.²⁰

I said that it is very doubtful that a case can be constructed in which just a pair of beliefs do not reverse-distribute through conjunction. I cannot prove that it is impossible. But recall why it is rational for the author to believe all of their claims independently but not believe that their conjunction is true: it is because of the awareness of human fallibility, i.e., the awareness that we are vulnerable to error, and that the probability that we make an error when we make a significant number of claims increases. It really is essential to the paradox that it involves a large number of beliefs. It does not impugn the principle that small numbers of beliefs reverse-distribute through conjunction—or come anywhere near doing so.²¹

So, the revised version of Silins’s theory succeeds in delivering a solution to the commissive version of Moore’s Paradox as it applies to assertions, as well as the omissive version.²² This is good news. The commissive form also involves self-ascribing a belief ($\text{not-}p$) when one is plainly not entitled to do so, so it is also something we can expect an account of self-knowledge to explain.

5. Some Refinements—And Why All of This Matters

Some refinements are necessary. Silins’s position, broadly, is that consciously judging that p provides justification for believing that you believe that p . As mentioned earlier, however, Silins does offer a more precise statement of his position. He actually works through several statements of his view, making a series of amendments, to deal with various complications, before coming to this version: ‘If you judge that p , then your judgment that p gives you immediate fallible justification to believe that you believe that p ’ (2012: 309).

How many of Silins’s refinements transfer to the amended version of his account which I am proposing? Exploring this will reveal the particular significance of my proposal. I will first consider the question of *immediacy*, then of *fallibility*, then of *giving* justification.

²⁰ One might say that there would be irrationality even if one did not have justification to self-ascribe belief in one of the propositions. Insofar as small numbers of beliefs reverse-distribute through conjunction, asserting ‘ p and I believe that $\text{not-}p$ ’ will amount to self-ascribing belief in a contradiction, whether or not one has justification to self-ascribe belief in one of the propositions. Perhaps, but there is clearly something especially strange about asserting ‘ p and I believe that $\text{not-}p$ ’ which goes beyond what the sentence logically entails, and the theory I am proposing captures this. Asserting ‘ p and I believe that $\text{not-}p$ ’ amounts to self-ascribing belief in a contradiction when one has justification to believe one of the contradictory conjuncts in particular, and my theory explains where that justification comes from. See Silins’s closely related remarks about what is ‘distinctively defective about Moore Paradoxical judgments’ (2012: 303).

²¹ The relationship between Moore’s Paradox and the Preface Paradox has been observed previously. See, e.g., Hájek (2021), taking his inspiration from Edgington (1992, 2011). There are obviously some complex issues relating to the Preface Paradox, which I cannot go into here.

²² Silins’s account, if it were otherwise successful, could also be applied straightforwardly to the commissive version of Moore’s Paradox. This is to be expected, given that the theory I am suggesting is a revised version of his, replacing judgments with sincere assertions. However, as I hope I have shown, his theory is not otherwise successful.

5.1 Immediacy

For Silins (following Pryor (2005)), ‘you have immediate justification to believe that p just in case you have justification to believe that p , and you do so in a way that does not rely on your justification to hold any other belief’ (2012: 298). It is important for Silins that the justification which judging that p might give for self-ascribing belief in p is immediate in this sense, because his purpose is to show that experiences themselves can give justification for the self-ascription of beliefs, in much the same way that perceptual experiences can give justification for beliefs about the observable world.

On this definition of ‘immediate’, we can certainly say that someone who makes an assertion and experiences S-phenomenology does have immediate justification for self-ascribing the belief expressed. All that matters is the phenomenology one experiences, not any other beliefs that they hold. There is nothing surprising here. Silins holds that the experience of consciously judging can justify the self-ascription of a belief. I have altered his account only by replacing one kind of experience, judging, with another kind of experience, the experience of asserting sincerely.

But there is something further to say. An important difference between my account and Silins’s becomes salient at this point. Suppose that making an assertion and experiencing S-phenomenology does give one justification to self-ascribe belief in the proposition asserted. How does one know which proposition they have asserted? There is more than one possibility, but there is one in particular which should be discussed, because it suggests that an unorthodox position in debates about self-knowledge may deserve further attention.

Traditionally, those who hold that our speech (including our inner speech) is important to self-knowledge hold that we need to interpret our own speech (including our inner speech). This position was famously articulated by Ryle (1949); contemporary proponents include Carruthers (2009, 2011) and Cassam (2014). What does interpretation involve, in this context? Carruthers (2009) tells us that

an interpretative process ... is one that accesses information about the subject’s current circumstances, or the subject’s current or recent behaviour, as well as any other information about the subject’s current or recent mental life. For this is the sort of information that we must rely on when attributing mental states to other people. (2009: 123).

So, if our speech, including our inner speech, plays a role in the self-ascription of beliefs, we must interpret that speech, just as we interpret the speech of others. Whatever justification we then have for self-ascribing beliefs on the basis of our speech will be of the same kind as the justification we have for ascribing beliefs to others on the basis of their speech.

Plausibly, then, while making an assertion and experiencing S-phenomenology gives one justification to self-ascribe belief in the proposition asserted, one only knows which proposition they have asserted by interpreting the assertion. If this is right, then justification for self-ascribing a particular belief on the basis of a particular assertion will also require that one is justified in believing that they have interpreted the assertion correctly. Can one have immediate justification for this?

How one answers this will depend on what they think is involved in the interpretation of speech generally. Again, I will consider just one possibility, for the reason that it illuminates a new path in thinking about self-knowledge. On one view, the interpretation of speech is an inferential process, involving the application of a theory of mind, i.e., a set of beliefs about the relationship between behaviour, including verbal behaviour, and mental states (Cassam 2014). If this is so, then any justification one has for believing that they have interpreted an assertion correctly will not be immediate in the Pryor/Silins sense (Cassam (2014), though note that Cassam is concerned with knowledge of one's episodic mental states, not standing mental states, like beliefs).

This is why there is something significant about what I am proposing. Theories on which justification for self-ascribing beliefs is immediate, and theories on which justification for self-ascribing beliefs is not immediate, are rivals in a strong sense. They offer competing explanations and, on the face of it, they cannot both be right. What I am suggesting is that, in at least one important context, viz., self-ascriptions of belief involving one's own (external and internal) speech, the justification might be partly immediate and partly non-immediate. One's justification for believing that they have interpreted an assertion correctly is not immediate: it depends on the truth of the propositions in one's theory of mind. But then self-ascribing belief in the relevant proposition is immediate, if one has the phenomenology of asserting sincerely.²³

This hybrid theory has some advantages vis-à-vis the theories from which it derives. First, it succeeds in the objective of explaining the apparent irrationality of asserting MP. In contrast, Silins's account is unsuccessful in its treatment of Moore's Paradox because it focusses on conscious judgments, and conscious judgments of the relevant kind are simply not possible (subject to the qualification added in Section 4.1). Second, the theory that we acquire knowledge of our own mental states by interpreting our own behaviour faces the problem that it cannot account for what seems to be a 'profound asymmetry between self-knowledge and other-knowledge', as Gertler (2021, Section 3.3) puts it. It is highly counter-intuitive to think that our knowledge of our mental states has the same basis as our knowledge of the mental states of others, which is what follows if one thinks that the basis of self-knowledge is interpretation of one's behaviour. While still allowing a significant role for interpretation, the hybrid theory explains this asymmetry. An individual experiences S-phenomenology when they assert sincerely, but not when others do. Finally, those who believe in cognitive phenomenology typically claim that cognitive phenomenology by itself can be the basis for knowledge of one's judgments and, thus, of the beliefs one forms. However, even if there is a cognitive phenomenology of judging in general, it is doubtful that there is a distinct cognitive phenomenology for judging each possible proposition to be true (though see Pitt (2004) for the contrary view, which Silins seems to sympathize with (2012: 24)). That is, it is doubtful that what it is like to judge that p is true is different from what

²³ Again, my immense gratitude to Melissa Fusco and the Australian National University Work-in-Progress Group (at a meeting in November 2016) for suggesting this way of looking at the theory.

it is like to judge that q is true (where $p \neq q$). The hybrid account avoids this issue. One experiences distinct sensory (or imagistic) phenomenology when one makes distinct assertions; this is what allows one to interpret an assertion as expressing a particular proposition. It does not matter that S-phenomenology (whether it is cognitive phenomenology or some other non-sensory phenomenology) does not vary depending on the proposition asserted; whenever one experiences S-phenomenology, they will be justified in self-ascribing belief in the particular proposition which they have interpreted themselves as asserting.

5.2 Fallibility

Next, there is no problem in saying that one has only *fallible* justification for the self-ascription of a belief when they make an assertion expressing it and experience S-phenomenology. Silins actually makes his point about fallibility partly in terms of assertion:

[C]onsider the following: you judge that your flight leaves at noon, and then realize that you do not and did not believe that your flight leaves at noon. In such cases, your judgment that p is a kind of performance error which fails to reflect an underlying belief—‘what was I thinking?’, you might go on to say. You ‘blurted out’ that p , either in speech or merely in thought, consciously endorsing the proposition that p , yet failing to have a standing belief that p . (2012: 308)

It is a recognisable experience, I take it, to make an assertion and to experience S-phenomenology, only to realise moments later that you do not and actually never did believe the proposition which you interpret yourself as having asserted. So, someone who makes an assertion and experiences S-phenomenology has fallible justification for self-ascribing belief in the proposition asserted.

This may be slightly too quick. A referee suggests an alternative analysis of what happens when we make an assertion which we do not believe and yet experience S-phenomenology, which can be brought out by examining a revised version of Silins’s example. Suppose someone asks you, at 2 pm, what time your flight leaves. You answer, ‘9 am’, and you experience S-phenomenology. Your interlocutor, alarmed, replies ‘But it’s already 2 pm!’. You realize that what you should have said is that the flight leaves at 9 pm. On my (Silins-inspired) analysis, you do not and never did believe that the flight leaves at 9 am.

But perhaps what happened—the referee suggests—is simply that you misspoke. You said ‘9 am’, but you meant ‘9 pm’. The content of your assertion was that the flight leaves at 9 pm, which you sincerely believed. Your error was simply in saying the wrong word. If this analysis is right, then all that has happened is that you have asserted something you believe and experienced S-phenomenology. We will not have seen a reason to believe that one can experience S-phenomenology when asserting something they do not believe.

Probably, this does happen sometimes, and it could be the most likely analysis of this particular example. However, Silins (2020: 333), borrowing from Smithies (2019: 184), describes another case which, adapted slightly to exemplify the theory I am suggesting, is naturally analyzed as most likely

involving an assertion of the wrong proposition: ‘[s]uppose you have been gazing at photos of the Sydney Opera House, and are asked to identify the capital of Australia in a game of trivia. You might answer that the capital of Australia is Sydney [and experience S-phenomenology], but then soon correct your mistake, realizing that this is not what you believe’ (for what you actually believe, correctly, is that the capital of Australia is Canberra). It seems unlikely that you merely misspoke: you were thinking of Sydney because your recent perusal of photos of the Sydney Opera House made it salient to you. You made an assertion and experienced S-phenomenology, only to realize that you do not believe the proposition you asserted. So, an assertion and the relevant phenomenology only provide fallible justification for the self-attribution of belief.

5.3 Giving Justification

Finally, can we say that making an assertion and experiencing S-phenomenology *give* you justification for self-ascribing the belief expressed? That is, that they are the source of such justification? The reason Silins thinks that conscious judgments give justification for the self-ascription of belief is that ‘judgment is itself a guide to belief’ (2012: 306). ‘This claim’, he writes, ‘is highly plausible, since it explains the correlation between judging that p and having justification to believe you believe that p ’ (2012: 306). ‘[I]t is doubtful that there just happens to be a correlation between judging that p and having justification to believe that you believe that p ’ (2020: 334).

Silins’ theory as provisionally stated—‘If you judge that p , then you have justification to believe that you believe that p ’—does not pass the plausibility test of explaining Moore’s Paradox, so we do not need to go further in investigating whether judgment can actually be a *source* of justification. But the theory that assertion plus S-phenomenology justifies self-ascription of belief does pass the plausibility test. And now we can simply make Silins’s move, because there is also a correlation between asserting p and experiencing S-phenomenology and having justification to believe that you believe that p . It is highly plausible that assertion plus S-phenomenology gives one justification for the self-ascription, because it explains this correlation. So, the final modification to Silins’s theory also applies to the theory I am outlining.

For a final statement of the view which I am proposing as an amended version of Silins’s theory, then: If you assert that p and experience S-phenomenology, that gives you immediate fallible justification to believe that you believe that p .

5.4 Scope

Some of the suggestions in this section are made tentatively. As already acknowledged, it may be that we can know which propositions we assert, whether externally or internally, without having to interpret those assertions, or that we interpret them without applying a theory of mind. I have not argued for

these positions. But they are reasonable, familiar positions. This is enough to show that the amended version of Silins's theory which I have offered is one that should be taken seriously.

In setting out the amended version of Silins's theory, I have focussed on assertions and self-ascriptions of belief. Before concluding, I want to suggest that the theory might actually have considerably broader application. Consider other speech acts. How it feels to express a wish varies depending on whether or not it expresses a genuine desire. How it feels to make an exclamation varies depending on whether or not one is genuinely surprised. Plausibly, there is something it feels like to perform any kind of speech act authentically (I will set aside the question of whether this phenomenology varies between different kinds of speech acts). This applies in inner speech as much as in external speech. But it may be that we have to interpret our wishes and exclamations etc. just as we have to interpret our assertions (supposing we do). Plausibly, then, the performance of any speech act, plus the phenomenology of performing it authentically, justifies self-attributing the attitude which one interprets the speech act as expressing.

The application of the theory might be broader still. In particular, it may go beyond speech acts. It may be open to hold a general theory of self-knowledge on which 1) we learn the content of our propositional attitudes by interpreting not only our (external and internal) speech, but also other kinds of conscious, episodic mental states, and 2) we can know our attitude towards such contents because of some phenomenology which accompanies the episodic mental states, akin to the S-phenomenology which accompanies sincere assertions.²⁴ Suppose, for example, that one has visual imagery resembling a graph and interprets it as representing the rate of inflation for the first three months of the year. Perhaps what it is like to produce such a mental image varies depending on whether or not one believes that inflation did increase at that rate in the relevant time period. If one produces the mental image and also experiences a particular kind of non-sensory phenomenology, they would be justified in self-ascribing belief in this content.

But this is very speculative. I venture only to say that it is a possibility which should be taken seriously.

6. Conclusion

How much motivation have I provided for the amended version of Silins's position which I am suggesting? Well, as much motivation as Silins's argument would provide for his original position, if that argument were successful. I have shown that the theory I am suggesting is at least plausible. And I hope that, in doing so, I have illuminated the possibility of hybrid theories of self-knowledge more generally.

²⁴ Thanks to a referee for suggesting this.

One final question—which might have been lingering for some readers—should be answered. If one holds the internalist position that self-knowledge can only be justified if we are aware (or, minimally, can become aware of) the basis of the justification, then they may want to know how we can be aware of our assertions and of S-phenomenology (and, if the suggestions made in the last subsection are added to the account, the internalist may want to know how we can be aware of our conscious, episodic mental states generally). If no answer can be given, then they might suggest that I have merely pushed the problem of self-knowledge of one’s beliefs (and perhaps other propositional attitudes) back a step.

Like Silins, I have been developing a theory on which a particular kind of conscious experience can justify the self-ascription of a belief. We are inevitably aware of our conscious experiences (or, minimally, we can become aware of them). *How* this is so is a further question. But even the internalist should accept that we have not so much pushed the problem back a step as taken a step forward. If we find that the self-ascription of beliefs (or other propositional mental states) can be justified by a particular kind of conscious experience, then we should certainly ask how we are aware of conscious experiences of that kind, but we will nonetheless have learnt something we did not know.

Acknowledgements

Much of this paper originally appeared in a PhD dissertation at the Australian National University (Gregory 2017).

I am grateful to Miloud Belkoniene, Melissa Fusco, Philip Gerrans, Frank Jackson, Luke Roelofs, Daniel Stoljar, Tim Smartt, Brett Toppey, and Agustín Vicente for comments and feedback. I also received valuable feedback at a meeting of the Philosophy of Mind Work-in-Progress Group at the Australian National University in November 2016. This paper has been in progress for a long period of time and it is possible that I have forgotten some people who have assisted in one way or another. I regret any omissions.

This research has been supported by an Australian Postgraduate Award, a Humboldt Postdoctoral Fellowship, and a María Zambrano Postdoctoral Fellowship. It has also been supported by the Barcelona Institute of Analytic Philosophy (María de Maeztu Unit of Excellence grant CEX2021-001169-M, funded by MICIU/AEI/10.13039/501100011033); by the *Inner Speech in Action: New Perspectives* research project, which receives funding from the Spanish State Research Agency and the Spanish Ministry of Science and Innovation (grant number PID2020-115052GA-Ioo); and by the *How Can We Think in a Language?* research project, supported by a grant from the Austrian Science Fund (FWF) (Project I-6219).

Conflict of Interest Statement

None.

References

- Bayne, T. & M. Montague, (2011), ‘Cognitive phenomenology: An introduction’, in T Bayne & M Montague (eds.), *Cognitive Phenomenology*, Oxford: Oxford University Press.
- Boyle, M. (2009), ‘Two kinds of self-knowledge’, *Philosophy and Phenomenological Research*, 78: 133–164.
- Byrne, A. (2005), ‘Introspection’, *Philosophical Topics*, 33(1): 79–104.
- Carruthers, P. (2009), ‘How we know our minds: The relationship between mindreading and metacognition’, *Behavioral and Brain Sciences*, 32: 121–138.
- Carruthers, P. (2011), *The Opacity of Mind: An Integrative Theory of Self-Knowledge*, Oxford: Oxford University Press.
- Cassam, Q. (2014), *Self-Knowledge for Humans*, Oxford: Oxford University Press.
- Coliva, A. (2015), ‘How to commit Moore’s Paradox’, *Journal of Philosophy*, 112: 169–192.
- Crimmins, M. (1992), ‘I falsely believe that *p*’, *Analysis*, 52: 191.
- Edgington, D. (1992), ‘Validity, uncertainty, and vagueness’, *Analysis*, 52: 193–204.
- Edgington, D. (2011), handout for “Three Grades of Uncertainty”.
- Edgley, R. (1969), *Reason in Theory and Practice*, London: Hutchinson.
- Evans, G. (1982), *The Varieties of Reference*, New York: Oxford University Press.
- Fernández, J. (2005), ‘Self-knowledge, rationality and Moore’s Paradox’, *Philosophy and Phenomenological Research*, 71: 533–556.
- Fileva, I. & L.A.W. Brakel (2019), ‘Just another article on Moore’s Paradox, but we don’t believe that’, *Synthese*, 196: 5153–5167.
- Gertler, B. (2021), ‘Self-knowledge’, in E. Zalta (ed.), *Stanford Encyclopedia of Philosophy* (winter 2021 edition), available at <https://plato.stanford.edu/entries/self-knowledge/>.

- Gregory, D. (2017), *Inner Speech: A Philosophical Analysis*, PhD dissertation at the Australian National University.
- Gregory, D. (2018), 'The feeling of sincerity: Inner speech and the phenomenology of assertion', *Thought*, 7: 225–236.
- Gregory, D. (2020), 'Are inner speech utterances actions?', *Teorema*, 39(3): 55–78.
- Hájek, A. (2021), 'Hysteresis hypotheses', in L. Walters & J. Hawthorne (eds.), *Conditionals, Paradox, and Probability: Themes from the Philosophy of Dorothy Edgington*, Oxford: Oxford University Press.
- Jorba, M. (2016), 'Attitudinal cognitive phenomenology and the horizon of possibilities', in T. Breyer & C. Gutland (eds.), *The Phenomenology of Thinking: Philosophical Investigations into the Character of Cognitive Experiences*, London: Routledge.
- Kim, J. (1996), *Philosophy of Mind*, Boulder, CO: Westview Press.
- Makinson, D.C. (1965), 'The Paradox of the Preface', *Analysis*, 25: 205–207.
- Moore G.E. (1942), 'A reply to my critics', in P. A. Schilpp (ed.), *The Philosophy of G. E. Moore*, Menasha, WI: George Banta Publishing Company.
- Moore, G.E. (1953), 'Propositions', in G E Moore, *Some Main Problems of Philosophy*, London: George Allen & Unwin.
- Munroe, W. (2023), 'Thinking through talking to yourself: Inner speech as a vehicle of conscious reasoning', *Philosophical Psychology*, 36: 292–318.
- Peacocke, C. (1999), *Being Known*, Oxford: Oxford University Press.
- Pitt, D. (2004), 'The phenomenology of cognition or *What is it like to think that p*', *Philosophy and Phenomenological Research*, 69: 1–36.
- Prinz, J. (2007), 'All consciousness is perceptual', in B. McLaughlin & Jonathan Cohen (eds.), *Contemporary Debates in Philosophy of Mind*, Malden: Blackwell.
- Pruss, A.R. (2012), 'Sincerely asserting what you do not believe', *Australasian Journal of Philosophy*, 90: 541–546.
- Pryor, J. (2005), 'There is immediate justification', in M. Steup & E. Sosa (eds.), *Contemporary Debates on Epistemology*, Oxford: Blackwell.
- Ryle, G. (1949), *The Concept of Mind*, London: Hutchinson.
- Shah, N. & D. Velleman (2005), 'Doxastic deliberation', *Philosophical Review*, 114: 497–534.

- Shoemaker, S. (1995), 'Self-knowledge and Moore's Paradox', *Philosophical Studies*, 77: 211–228.
- Siegel, S. (2010), *The Contents of Visual Experience*, New York: Oxford University Press.
- Siewert, C. (1998), *The Significance of Consciousness*, Princeton, NJ: Princeton University Press.
- Silins, N. (2012), 'Judgment as a guide to belief', in D. Smithies & D. Stoljar (eds.), *Introspection and Consciousness*, Oxford: Oxford University Press.
- Silins, N. (2013), 'Introspection and inference', *Philosophical Studies*, 163: 291–315.
- Silins, N. (2020), 'The evil demon inside', *Philosophy and Phenomenological Research*, 100: 325–343.
- Smithies, D. (2013), 'The nature of cognitive phenomenology', *Philosophy Compass*, 8: 744–754.
- Smithies, D. (2019), *The Epistemic Role of Consciousness*, New York: Oxford University Press.
- Strawson, G. (1994), *Mental Reality*, Cambridge, MA: MIT Press.
- Turri, J. (2010), 'Refutation by elimination', *Analysis*, 70: 35–39.
- Tye, M. & B. Wright (2011), 'Is there a phenomenology of thought?', in T Bayne & M Montague (eds.), *Cognitive Phenomenology*, Oxford: Oxford University Press.
- Williams, J. N. (2004), 'Moore's Paradoxes, Evans's principle, and self-knowledge', *Analysis*, 64: 348–53.