

Thinking about the body as subject

Canadian Journal of Philosophy, forthcoming.

The notion of immunity to error through misidentification (IEM) has played a central role in discussions of first-person thought. It seems to be a way of making precise the idea of thinking about oneself 'as subject'. Asking whether bodily first-person judgments (e.g. 'My legs are crossed') can be IEM is a way of asking whether one can think about oneself simultaneously as a subject and as a bodily thing. The majority view is that one cannot. I rebut that view, arguing that on all the notions of IEM that have so far been successfully defined, bodily first-person judgments can be IEM.

Key words: *de se*, first-person, immunity to error through misidentification, memory, quasi-memory, proprioception, privileged access, introspection

The notion of *immunity to error through misidentification* (IEM) has played a central role in discussions of first-person thought. Its claim to centrality rests, I think, on two main planks.

First, asking whether certain first-person thoughts are IEM has appeared to be a way of getting at the same kind of maximally deep philosophical issue as is raised by asking whether certain first-person judgments are about oneself *as subject*.¹ So, the *interest* of this latter question is inherited by claims about which first-person judgments are IEM. But the notion of IEM is reasonably precisely defined, in a way that the notion of thought about oneself as subject is not. So, the *obscurity* of this question is *not* inherited by claims about which first-person judgments are IEM. IEM is an essential bit of philosophical apparatus because it allows us to make tractable a question that is of maximal philosophical interest but that, formulated as it previously had been, was impossible to settle.

Second, the distribution of IEM – which judgments have the property – has seemed to pattern in a way that connects strikingly with the distinction between mental and bodily properties of persons. The mental-bodily distinction has obviously played a huge role in the debate about the *metaphysics* of persons. But there is a separate question about the place the distinction has in the *epistemology* of first-person thought. Most philosophers who have discussed the distinction have thought that the only first-person judgments that are IEM in the fullest sense are *mental* first-person judgments. ‘I am in pain’, based on introspection, or ‘I am thinking that *p*’, based on introspection, will be IEM. ‘My knees are crossed’, based on proprioception, or ‘I am moving through a wooded grove’, based on vision, will not be. If they are right about that, then there is an epistemological asymmetry between our relationship to ourselves as bearers of

¹ For the distinction between uses of “I” as subject and as object, see Wittgenstein (1958, 66-7).

bodily properties and our relationship to ourselves as bearers of mental properties (one which one might then use to try to establish metaphysical conclusions, or to diagnose metaphysical illusions, according to taste). If they are wrong about that, because both bodily and mental judgments can be IEM, then our bodily nature is deeply ingrained in the way we think about ourselves. One can think about oneself simultaneously as a subject, and as a bodily thing.²

My discussion of IEM will focus on the following thesis:

Bodily first-person judgments, just like mental first-person judgments, can be IEM.
(*Parity Thesis*)

The majority view, as just outlined, is that this thesis is false, and that it is false because bodily first-person judgments can't be IEM in the fullest sense, whereas mental first-person judgments can be. I'll argue that the two most influential notion attempts to craft a notion of IEM that supports this verdict fail and that, as things stand, we have been given no reason to reject the *Parity Thesis*.

Before getting into the main discussion, I take care of three preliminaries. First, I offer a quick recap of the debate about IEM thus far, highlighting the points relevant to the *Parity Thesis*. Second, I note that there are *different* notions of IEM, and I say why the *Parity Thesis* is best understood as a claim about the most demanding notion of IEM. Third, I say why *no* notion of IEM should be *so* demanding that it collapses into the notion of infallibility.

(i) Recap of the debate

The rejection of the *Parity Thesis* is clearly embedded in Wittgenstein's original presentation of the distinction between uses of 'I' *as subject* and uses of 'I' *as object*.

²This issue is also discussed outside the analytic tradition. For example, the idea that our bodily nature is deeply ingrained in the way we think about ourselves is explicit in the work of Heidegger and Merleau-Ponty. See O' Brien (2007, 201) for discussion.

All of his examples of uses of 'I' as *subject* involve mental predications — 'I see so and so', 'I hear so and so', 'I try to lift my arm', 'I think it will rain', 'I have a toothache'.

All his examples of uses of 'I' as *object* involve bodily predications — 'My arm is broken', 'I have grown six inches', 'I have a bump on my forehead', 'The wind blows my hair about' (Wittgenstein 1958, 66-7). Wittgenstein did not say *why* his examples exhibited the striking mental/physical pattern he took them too.

Shoemaker (1968) introduced the idea of a quasi-memory to the discussion. He claimed that memory judgments are vulnerable to error through misidentification because they are vulnerable to quasi-memories, and quasi-memories, Shoemaker thought, involve a kind of misidentification. It just isn't a kind of misidentification that *de facto* we have to worry about. Shoemaker's notion of a quasi-memory generalizes to the notion of a 'quasi-case', which will be important in what follows. The generalization of Shoemaker's suggestion about the significance of quasi-memories is that any judgment that is vulnerable to a quasi-case — and that, importantly, will include every bodily first-person judgment — will *ipso facto* be vulnerable to misidentification.

Gareth Evans (1982, 241-8) defended the *Parity Thesis* against Shoemaker's attack, arguing that, on a fairly flatfooted way of thinking about the notion of a misidentification that Shoemaker himself seemed to be operating with, quasi-cases do not involve misidentification.

The success of Evans's defense has dictated the strategy of subsequent opponents of the *Parity Thesis*. Their strategy is to attempt to craft rather less flatfooted definitions of misidentification than Shoemaker's on which quasi-cases *do* count as involving misidentification. Two main versions can be distinguished. The first invokes the notion of an *undermining defeater*. It is defended by James Pryor (1999). The second centers on the notion of a *background presupposition*. It is developed first and in

most detail by Annalisa Coliva (2006) (2012), and has also been explored by Crispin Wright (2012). This paper argues that neither strategy is successful.³

(ii) Different notions of IEM

Many theorists think, and I agree, that there are *different kinds* of IEM. Most theorists of IEM would say that it only the most demanding kind of IEM which bodily first-person judgments lack. For example, James Pryor writes:

‘There are several different sorts of error through misidentification, and they are importantly different. This paper aims to clarify and distinguish these different kinds of identification and misidentification and the correspondingly different notions of immunity to error through misidentification. I will argue that immunity to error comes in two main sorts...I will show how we can use these distinctions in assessing a debate between Shoemaker and Evans about whether first-person memory based judgments are immune to error through misidentification. Evans says they are immune. Shoemaker says they are not. I will argue that, with respect to the most interesting sort of immunity, Shoemaker is right. Our first-person memory judgments do not have that sort of immunity.’ (1999, 272).

If there are several different notions of IEM, then there are different ways of reading the claim that bodily first-person judgments, just like mental first-person judgments, can be IEM. I intend the *Parity Thesis* to be understood as concerning the most demanding grade of IEM, whatever that turns out to be. It would be highly unimpressive to show that there is *a* grade of IEM that bodily as well as mental first-person judgment can have, while allowing that there is another that only mental first-person judgments can have. Against the background of the presumption that there is *some* interesting epistemic asymmetry between our relationship to our mental and our physical

³ Other prominent theorists who reject the *Parity Thesis* include Parfit (1984) and Recanati (2012). These are not discussed here since they engage in far less detail with the question of how the notion of IEM is to be defined.

properties, it would be easy to insist in this situation that it is the more demanding grade of IEM that should be used to cash out the notion of thought about oneself *as subject*.

(iii) IEM and infallibility

On the other hand, I think that no adequate notion of IEM will be so demanding as to collapse into the notion of infallibility. This assumption does work in my argument. For example, I argue that Pryor's attempt to define a notion of IEM cannot be right because, if it were, his notion would collapse into infallibility. What justifies this assumption that IEM and infallibility are different?

The first justification is that the collapse of one into the other would disqualify paradigm examples of IEM from being examples of the phenomenon. For example, 'I am in pain', based on introspection, is a paradigm example of an IEM judgment.

However, many will accept that it is an exaggeration to say that 'I am in pain', based on introspection, is infallible. It certainly seems as though that one might introspectively mistake an itch for a pain (See Williamson [2000]).

Not everyone will be convinced by this. Perhaps the mental is a domain in which the appearance/reality distinction doesn't really arise. If a property is properly mental, someone might say, then one will be introspectively inclined to self-ascribe it only if one really instantiates it. In that case, 'I am in pain', made on the basis on introspection, will be infallible (and we will need some other way of thinking about what the pain-itch case shows). The notion of IEM's collapsing into that of infallibility won't cause paradigm examples of the phenomenon to be ruled out, so long as all paradigm examples of the phenomenon are infallible.

There is however a different justification for not allowing the collapse. If being IEM and being infallible turn out to be the very same epistemic status, latched onto in different ways, what is the point of introducing the notion of IEM in the first place? Any

question that could be formulated using the notion of IEM could be formulated using the notion of infallibility instead. Since the notion of infallibility is older, and far simpler, introducing IEM would be a pointless complication. To that extent, accepting that IEM is a different epistemic status from infallibility is the price of admission to the debate about IEM. If they are the same status, then, on the face of it, the debate about IEM (*qua* IEM) shouldn't exist.⁴

Here is an outline of the argument that follows. Section (§1) sets out the notion of *immunity to error through false-belief misidentification*, and points out that it is an epistemic status that bodily first-person judgments can have. Section (§2) introduces the notion of a 'quasi-case', a kind of error-possibility to which every bodily first-person judgment seems vulnerable in principle. The remaining two sections focus on the two main attempts to craft a notion of IEM on which vulnerability to quasi-cases entails vulnerability to error through misidentification and on which *Parity Thesis* is therefore false. The first of these, discussed in section (§3), is Pryor's. Pryor's attempt focuses on *immunity to error through wh-misidentification*. I agree with Pryor that immunity to error through wh-misidentification is different from immunity to error through false belief misidentification. But I argue that his definition of immunity to error through wh-misidentification in terms of *undermining defeaters* is incorrect, and I argue that, on the correct definition, immunity to error through wh-misidentification is a status that bodily first-person judgments can have. The other attempt to undermine the *Parity Thesis*, discussed in section (§4), is offered by Coliva, and also Wright. They both focus on the notion of *immunity to error through false-belief-or-false-background-presupposition misidentification*. I argue that their definition picks out an epistemic status that no

⁴There are other arguments against the collapse, starting from premises about how the notion of IEM is supposed to be structured. E.g. IEM is supposed to relate particularly to the singular, as opposed to the predicative, aspect of a judgment – the 'a' part of 'a is F' not the 'F' part. But the notion of infallibility doesn't relate particularly to the singular, as opposed to the predicative, aspect of a judgment.

judgments have, so it definitely isn't an epistemic status that some mental first-person judgments, but no bodily first-person judgments, have. Section (§5) concludes.

(§I) Immunity to error through false belief misidentification

To say that a judgment is *immune* to error through misidentification is to say that it is not possible for it to be in error as a result of a misidentification. The natural way of specifying a notion of IEM is therefore to say what the corresponding notion of misidentification is. In this, it is natural to be guided, at least provisionally, by the ordinary use of the term 'identify'. That word has two main uses. On one of them, to identify an object is, roughly, to distinguish it from its surroundings — to single it out. For example, I might identify the tiger in the grass when my attention seizes on a striking orangey form moving amid the surrounding green. On a different use, to identify an object is, roughly, to see that it is identical to something else. For example, I might identify the tiger I have just spotted in the grass as Tony Tiger, a tiger I know from way back.

It is the second use of the term 'identity' that is most relevant for the first kind of IEM I want to introduce. Plausibly, an identification, in the second sense, is some kind of attitude with a content involving identity. A bit more tendentiously, but still plausibly, an identification is a *belief* with a content involving an identity. Suppose I judge 'Tony Tiger is swishing his tail'. I may not have performed a conscious chain of reasoning involving the premises 'That tiger is swishing his tail' and 'That tiger is Tony Tiger'. Recognition of Tony Tiger may have been instantaneous, so that I came by the suite of beliefs 'That tiger is swishing his tail', 'That tiger is Tony Tiger' and 'Tony Tiger is swishing his tail' in one fell swoop. Still, the first two beliefs are the *basis* for my third belief, the belief that Tony Tiger is swishing his tail. The first two beliefs

causally sustain that third belief. And they are what I would appeal to if asked to give my reasons for coming to that belief.

One notion of IEM is the following:

A judgment ‘*a* is *F*’ is IEM, relative to basis *b*, if and only if *b* does not include an identity belief. (Immunity to error through false-belief misidentification)⁵

An important thing to note for what follows is that, on this notion of IEM, a belief’s not being based on *any* other beliefs — i.e. being non-inferential — guarantees that it will be IEM. This is so, even if the belief in question is vulnerable to quasi-cases.

(§II) Quasi-Cases

The most famous kind of quasi-case are cases of quasi-memory.

A subject quasi-remembers an event *e* iff *e* took place, the subject has an apparent memory of *e*, and this apparent memory causally derives from someone else’s past perceptions and experiences.

For example, imagine that two hemispheres of someone’s brain are separated and each is put in a new body. This produces two new people each of whom is psychologically continuous with the original person but neither of whom is identical with her. Because each of the new people has inherited half the original person’s brain they quasi-remember many events in the original person’s life. Say the original person was once on a ship. As a result, both of the new people may each now have a quasi-memory as of having been on a ship.

⁵This is somewhat similar to Pryor’s notion of ‘immunity to *de re* misidentification’ (see 1999: 279). One way in which it is different is that my notion is not restricted to *de re* identifications. This reflects my suspicion that the *de re*/ non-*de re* distinction does not have any useful role to play in the debate about IEM. Nothing in this paper, however, depends on this opinion. I do not discuss ‘immunity to *de re* misidentification’ in detail in this paper since Pryor regards it as the less interesting of his two notions of immunity and, relatedly, he thinks that bodily first-person judgments can satisfy it (see 1999, 293). My notion is even closer to Evans’s notion of ‘identification freedom’ (see 1982, 243-44).

Quasi-memory seems to be a species of a broader genus. For example, we can say that someone *quasi-proprioceives* the instantiation of a bodily property (e.g. having crossed knees) iff that bodily property actually is instantiated, the subject has apparent proprioceptive awareness as of *their* instantiating that property, and this apparent awareness derives from the state of somebody else's body. I might have apparent proprioceptive awareness as of my knees being crossed because I am appropriately hooked up to somebody else's body, and that person's knees are crossed. Or again, we can imagine a case of *quasi-vision* in which it visually appears to me that I am standing in front of a tree because I am hooked up to, and receiving visual information, through the eyes of somebody else who really is standing in front of a tree.

A 'quasi-case', we can say, is a case in which all these three elements figure. A subject has an impression that a property is instantiated. It seems to the subject as though *they* are the one who instantiates the property. In fact, it is because someone else instantiates the property that the subject has the impression they have.

It is very plausible that every bodily first-person judgment will be vulnerable to quasi-cases. The contentious question is whether quasi-cases involve misidentification, and whether a judgment's being vulnerable to quasi-case therefore makes it vulnerable to error through misidentification. It cannot mean that, if our notion of a misidentification is just that of a false identity belief. 'My knees are crossed', based on proprioception, is vulnerable to quasi-proprioception. But it may still be based directly on experience and, in that case, it won't be based on an identity belief. So the opponent of the *Parity Thesis* needs a different notion of IEM.

(§III) Immunity to error through *wh*-misidentification

There is a notion of misidentification that is different from the notion of a false identity belief. I will use Pryor's term for this different kind of misidentification: *wh*-

misidentification. The phenomenon of *wh*-misidentification can be explained in the abstract, and examples of it can be given, without mentioning quasi-cases. The question then arises of whether quasi-cases involve error through *wh*-misidentification. Pryor argues that the answer to this question is ‘Yes’.⁶ I first rebut his argument, which depends on defining immunity to error through *wh*-misidentification in terms of undermining defeaters. I then give an argument that the answer is ‘No’. So, the conclusion of the discussion of Pryor will be that immunity to error through *wh*-misidentification is indeed a different, and more demanding, notion from immunity to error through false belief misidentification. But it is still a notion on which the *Parity Thesis* holds.

To get a grip on what *wh*-misidentification is notice that if a judgment of the form ‘*a* is *F*’ is in error it seems a reasonable and potentially instructive question to ask whether the error is due to the singular part of the judgment, or to the predicative part of the judgment. A symptom of the error being due to the predicative part of the judgment would be if there was another belief about the same object that the subject retains knowledge of. E.g., if I judge ‘That vase is red’, because I see a white vase bathed in red light, I may nevertheless retain knowledge of ‘That vase *appears* red’. A symptom of the error being due to the singular part of the judgment would be if there was another belief about the same property that the subject retains knowledge of. E.g. if I judge ‘Tony Tiger is swishing his tail’, because I see Tommy Tiger swishing his tail and I take Tommy for Tony, I may retain knowledge of ‘Something is swishing its tail’.

A perfectly good notion of ‘misidentification’ is just the notion of an error that in some way involves the singular component of the judgment. One species of this genus is where the subject falsely identifies one thing as a different thing, e.g. by

⁶Pryor focuses on the case of memory judgments. See (1999, 304) for his account of how things are likely to extend to, e.g., bodily first-person judgments that are based on proprioception or perception. I will stick close to Pryor’s discussion by also using memory judgments as examples.

forming an identity belief. But, it turns out, there is a different species of the same genus — *wh*-misidentification.

We encounter *wh*-misidentification in the following example of Pryor's (1999, 297). In *The Little Girl Case* a subject is looking at a little girl who is standing 10 feet in front of them. The subject's visual experience justifies them in judging 'There is a little girl standing about 10 feet in front of me'. So the subject makes that judgment on a non-inferential basis. So, in particular, it is not based on an identity belief.

Suppose though that the subject is in a hall of mirrors, that there is a full-length mirror two feet in front of the subject, and that this mirror has interfered with the subject's capacity to accurately locate the little girl. The little girl is really some distance behind the subject. The subject may still know, on the basis of vision, 'There is a little girl somewhere or other in my vicinity', which is the existential generalization of 'There is a little girl standing about 10 feet in front of me'. We might describe the situation as follows: the subject has correctly detected a feature in their environment — the presence of a little girl — but they have gone wrong in identifying — or singling out — the location that instantiates that feature. But, to bring out the difference with the other species of misidentification, that isn't because they have singled out a different location, and formed a belief that *it* is the location of the feature.⁷

On Pryor's way of thinking about memory, memory judgments are similar to the judgment 'There is a little girl standing about 10 feet in front of me' in *The Little Girl Case*. If I take at face value a memory impression as of having been on a ship that is in fact a quasi-memory, then I will have succeeded in detecting a feature of the world's history — that someone was once on a ship — but I will just have gone wrong in

⁷ Campbell (1997, 69-70) gave the first case of *wh*-misidentification. A subject judges, on the basis of vision, 'That chair is yellow'. In fact though, the chair is transparent, but set against a yellow background. The subject has correctly detected the presence of yellowness but has gone wrong in identifying the chair as the thing that is yellow.

identifying myself as the person who has that feature. So, memory judgments will be vulnerable to error through *wh*-misidentification.

To work out whether this is true we need to motivate a definition of immunity to error through *wh*-misidentification, and then apply the definition to memory. I first argue that Pryor's definition, which is intended to yield the verdict that memory judgments are vulnerable to *wh*-misidentification, and which brings in the notion of an undermining defeater, is unacceptable. Then I motivate my own definition of immunity to error through *wh*-misidentification, which gives the opposite verdict.⁸

(i) Pryor's attempt to define immunity to error through wh-misidentification using undermining defeaters

Pryor's first attempt at a definition says that that a judgment will be immune to error through *wh*-misidentification, relative to certain grounds, if and only if:

'...there is no [proper] "part" of your justification for believing that *a* is F which could offer you knowledge that *something* is F, while leaving it an open question for you whether *a* is F.' (1999, 283)

Pryor is — admirably — not content to leave the notion of 'part' of your justification as an undefined primitive. He connects the notion with the possibility of one's justification, or grounds, being defeated. The crudest articulation of the connection would be to suggest that a part of your grounds for *p* offers you knowledge that *q* if and only if your grounds for *p* could be defeated in such a way that they still offer you knowledge that *q*. But, as Pryor himself points out, this theory is too crude. For any *p* and *q*, and for any grounds, you can imagine God (or some other known-to-be-reliable

⁸ Some have denied that there is any kind of misidentification that does not involve false identity belief, such as *wh*-misidentification is intended to be. This denial is crucial to the defense Smith (2006) gives of the claim that bodily first-person judgments can be IEM. Smith argues that non-inferential bodily first-person judgments must be IEM since they are IEM on the false belief based notion of IEM, and there is no alternative notion which (i) is an acceptable notion of IEM and (ii) differs in its extension from the false belief based notion of IEM. The examples Campbell and Pryor offer suggest that *wh*-misidentification is just such a notion.

giver of testimony) coming along and saying ‘ p is false, but q isn’t’. Once you have the expanded set of grounds that include God’s testimony, you will be offered knowledge of q . This does nothing to show that part of your original grounds offer you knowledge of q .

The moral Pryor takes from this is that we need to invoke a distinction between two kinds of defeater. An *undercutting* defeater for p , Pryor says, undermines one’s justification for p without adding to one’s case for not- p . For example, if it seems to me that the wall is blue and you tell me that my eyes aren’t working well today, that is an *undercutting* defeater. An *additive* defeater for p , Pryor says, undermines my justification for p and adds to my case for not- p . If you tell me straight out that the wall is not blue, that is an *additive* defeater.⁹ Pryor’s suggestion is this:

‘What we should say is this: “part” of your grounds justifies you in believing that p iff it is possible to take away some of those grounds, by undercutting them, in such a way that what is left will (still) justify you in believing that p .’ (1999, 284)

Feeding Pryor’s understanding of ‘part of your grounds’ into his first attempt at a definition, we get Pryor’s finished definition.

A judgment ‘ a is F’, believed on grounds g , is vulnerable to error through wh misidentification if and only if:

‘It is possible for g to be defeated by undercutting evidence in such a way that:

- (i) the combination of g and that evidence *no longer* justifies you in believing that a is F, but
- (ii) the combination of g and that evidence *could*, by itself, offer you knowledge that something is F. (1999, 284)¹⁰

⁹These glosses appear on (1999, 284). See Pryor (2013) for a good outline of some of the issues, and potential difficulties, with the distinction between undercutting and additive defeaters.

¹⁰See Pryor (1999, 302 fn 28) for an explanation of the (rather technical) wrinkle that leads him to have ‘could’ here rather than ‘does’.

Obviously, a judgment is immune to error through *wh*-misidentification if and only if it is not vulnerable to *wh*-misidentification.

(ii) The argument against Pryor's definition

Suppose '*a* is F' is a judgment whose IEM status, relative to grounds *g*, is being adjudicated. Pryor's definition tells us that we should ask whether there is a defeater of an appropriate kind — the undercutting kind — that can be added to *g* so that the subject is no longer justified in judging '*a* is F' but is offered knowledge of 'Something is F'. How exactly should we understand the restriction to undercutting defeaters? Is the restriction merely that the defeater not add to the subject's case for '*a* is *not* F' (this is all that Pryor's own formulations, interpreted literally, require)? Or is it, more demanding, that the defeater not add to the subject's case for *any* other proposition, including, for example, 'Something is F'?

Suppose that we interpret the restriction in the less demanding way. Then there is a quick argument that Pryor's definition implies that any judgment of the form '*a* is F' will be vulnerable to error through *wh*-misidentification. The relevant defeat scenario isn't one in which God comes along and says to me: 'Well, actually, no – *a* isn't F. But I'll tell you this much. *Something* is F'. Rather, God says: 'Your judgment that *a* is F isn't based on its being the case that *a* is F in the right way for that judgment to be knowledge. But I'll tell you this much. *Something* is F'. In this kind of case, God's testimony meets Pryor's restriction, interpreted in the less demanding way. It does not add to my case that *a* is not F. The fact that my judgment isn't based on its being the case that *a* is F in the right way for that judgment to be knowledge is entirely compatible with its being the case that *a* is F. But I do end up in a position in which I am offered knowledge of 'Something is F'.

Suppose instead that we interpret the restriction in the more demanding way. Doing so will make it harder to craft a defeat scenario for a memory judgment that meets the restriction to undercutting defeaters. Consider as a first candidate a scenario in which a subject is told straight out that their memory impression as of having been on a ship is a quasi-memory. That is — once we unpack the notion of a quasi-memory — they are told that their memory impression (i) is not a memory of their own past but (ii) does causally derive from somebody else's having been on a ship. This testimony entails, and so clearly adds to the subject's case for, 'Someone was on a ship'. So it is not an undercutting defeater, on the more demanding interpretation.

However, the defeat scenario Pryor focuses on is more subtle. It involves a victim of quasi-memory who is told:

'... that some of his memories are quasi-memories, and that none of his memories as of being F [e.g. as of being on a ship — author] derive from actual events in his own life. We leave it open whether his memories as of being F are among the quasi-memories of someone else's past. Call this information D*. D* undercuts the justification that S's quasi-memories give him for believing that *he* was F.' (1999, 295)

It's a delicate matter whether D* really is undercutting in the more demanding sense. One might well insist that D* does add to the subject's case for 'Someone was on a ship'. The obvious difference from the scenario in which the subject is actually told outright that their memory is a quasi-memory is that D* (even if its accuracy is assumed) doesn't provide a conclusive basis for 'Someone was on a ship'. But it surely *adds something* to the subject's case for 'Someone was on a ship', which makes it an additive defeater.

It's also a delicate question whether the subject's memory impression, supplemented with D*, offers them knowledge of 'Someone was on a ship'. One might

argue that, since the testimony in D^* (even if accurate) is not a conclusive basis for ‘Someone was on a ship’, the subject does not end up with grounds that offer them knowledge of ‘Someone was on a ship’. Rather, they ought to regard the matter as left somewhat open.

These two worries should make us doubt whether Pryor has managed to describe a defeat scenario for a memory impression as of being on a ship that does all of following. First, it is undercutting (in the more demanding sense). Second, it undermines the subject’s justification for ‘I was on a ship’. Third, it leaves the subject with grounds that offer them knowledge of ‘Someone was on a ship’. But I do not want to insist on either worry.

What I would insist on rather is a conditional. If the scenario Pryor describes shows that memory judgments are vulnerable to error through *wh*-misidentification, then there will be analogous scenarios that reveal the same about any fallible judgment. So, granted that we don’t want the notion of IEM to collapse into that of infallibility, we must reject Pryor’s definition.

For any fallible judgment ‘*a* if *F*’, relative to any grounds *g*, we can distinguish two kinds of case in which the judgment is in error. One is a case in which something other than *a* really is *F* and that other thing’s being *F* is implicated in one’s judging ‘*a* is *F*’ (this is analogous to quasi-memory).¹¹ The other is a case in which there isn’t anything else whose being *F* is implicated in one’s judging ‘*a* is *F*’ (this is analogous to more ordinary kinds of memory failure). We can always imagine someone with some justification for ‘*a* is *F*’ being told that their judgment is not based on its being the case that *a* is *F*. It has gone wrong in either of the two distinguished ways, but we’re not saying which. The subject ends up with at least some justification for ‘Something is *F*’

¹¹ For example, suppose that ‘*a* is *F*’ is ‘I am in pain’. Then an example of the first kind of case, which Smith (2006, 279) discusses, would be if someone else is in pain and this causes you to feel a pain-like itch.

but his justification for ‘*a* is F’ is completely undermined. If, in general, the possibility of this kind of defeat scenario is enough to show that ‘*a* is F’ is vulnerable to error through *wh*-misidentification, then it will follow that every fallible judgment is vulnerable to error through *wh*-misidentification. If it isn’t, then the kind of defeat scenario Pryor describes will not show that memory judgments are vulnerable to error through *wh*-misidentification.

There are quite a few delicate issues involved in working out which judgments satisfy Pryor’s finished definition of immunity to error through *wh*-misidentification. There is the initial exegetical issue about the more and less demanding interpretation of ‘undercutting’. Once that has been resolved (on the more demanding interpretation, most plausibly), there is a substantive question about which defeat scenarios count as undercutting, and another substantive question about which defeat scenarios leave the subject with grounds that offer them knowledge of the relevant existentially general proposition. So I have do not have settled view on what the extension of Pryor’s definition is. But I have argued that the only way its extension is going to turn out to exclude memory judgments is if it excludes every fallible judgment.

(iii) A definition in terms of independent knowledge

Rebutting Pryor’s argument that bodily first-person judgments are vulnerable to error through *wh*-misidentification isn’t the same as establishing that they are immune to it. To establish that they are we need to argue for, and then apply, an alternative definition.

The definition I propose is as follows:

A judgment of the form ‘*a* is F’ is, relative to certain grounds *g*, immune to error through *wh*-misidentification if and only if *g* justifies ‘*a* is F’ without offering independent knowledge of ‘Something is F’. (*Immunity to error through wh-misidentification*).

Pryor's first attempt at a definition invoked the idea of one's grounds dividing into different parts and some part of one's grounds supporting 'Something is F' without supporting 'a is F'. We have seen that this 'part'-based approach goes astray when the notion of part is cashed out in terms of that of an undercutting defeater. I think we should focus on what seems to be the root idea behind the part-based approach: the idea of independence. If some part of my grounds supports one proposition without supporting a different, logically stronger, proposition that can only be because my grounds give some support to the weaker proposition that is *independent* of the support they give to the stronger proposition.

Invoking independence will only help if we have a firm basis for judgments of independence. I propose the following test:

Grounds g offer knowledge of q that is independent of the justification they provide for p if and only if one can *know* that q on the basis of g , even if p is false (*Independence Test*).

What makes this test appropriate? A certain fact about knowledge: the fact that it is not possible to gain knowledge by inferring from false premises.¹² Suppose that 'a is F' is false but my grounds for 'a is F', g , nevertheless offer me knowledge of 'Something is F'? Then it must be that g offer me knowledge of 'Something is F' that is *independent* of the justification they provide for 'a is F'. If g supported 'Something is F' only by supporting 'a is F', from which 'Something is F' can be inferred, the fact the 'a is F' is false would deprive me of knowledge of 'Something is F'.¹³

The definition I have argued for returns the intuitively correct verdict in *The Little Girl Case*. The subject's 'There is a little girl 10 feet in front of me' judgment is

¹²Some apparent counterexamples to this principle have been presented. For a convincing argument that the counterexamples are merely apparent see Ball and Blome-Tillman (2014).

¹³This is one reason it is important that the definition focus on knowledge of 'Something is F' rather than justification for 'Something is F'. Justification, unlike knowledge, can be gained by inference from false beliefs. In the Gettier cases, justification but not knowledge is gained by inferring from a false belief.

false. But the subject can still know ‘There is a little girl standing somewhere or other in my vicinity’. It follows that the subject’s grounds must offer *independent* knowledge of ‘There is a little girl standing somewhere or other in my vicinity’. So, on my definition, the judgment will count as vulnerable to error through *wh*-misidentification, as it intuitively ought to.

The verdict my definition returns about the ‘I was on a ship’ judgment, based on a memory impression as of having been on a ship, is that it will be immune to error through *wh*-misidentification. For the memory impression does not offer knowledge of ‘Someone was on a ship’ that is independent of the justification it provides for ‘I was on a ship’. If it did offer such independent knowledge, then the subject could know ‘Someone was on a ship’, *just on the basis of the memory impression*, even if ‘I was on a ship’ were false. But, as Evans persuasively argues, that isn’t the case.

Evans first points out that a victim of quasi-memory is unlikely to spontaneously come up with the hypothesis that they are a victim of quasi-memory. He then adds that, even if they did, their hypothesis would not be one that their memory impression itself justifies. So, the subject might judge ‘Someone was on a ship’, just on the unjustified suspicion that their memory impression is a quasi-memory. But in that case their judgment would merely be a (successful) ‘shot in the dark’ — true but not knowledge.

Evans’s assessment here has enough intuitive plausibility that Pryor concedes it.^{14 15} Having conceded it, Pryor is forced to argue for a definition of immunity to error

¹⁴ See Evans (1982, 246) and Pryor (1999, 295).

¹⁵ Even if Pryor made this concession, couldn’t a different defender of Pryor’s view about memory resist it, saying that it only holds with respect to worlds in which quasi-memories are rare? In what we might call a Shoemakerian world – a world in which quasi-memories are common -- a subject who spontaneously comes up with the hypothesis that their memory impression is a quasi-memory would be likely to be correct. *So, if they are correct, their hypothesis will be knowledge*. And, in that case, their judgment ‘Someone was on a ship’ will also be knowledge. So they are offered independent knowledge of ‘Someone was on a ship’. So their judgment ‘I was on a ship’ will be vulnerable to *wh*-misidentification.

I think that the problem with this reasoning is the crude reliabilism implicit in the italicized sentence. If we accept that kind of crude reliabilism, the notion of immunity to *wh*-misidentification collapses into

through *wh*-misidentification on which the crucial issue is what knowledge the subject is offered *when their memory impression is appropriately supplemented (by undermining defeaters)*. The core of my argument that memory judgments are immune to *wh*-misidentification has therefore consisted in an argument that the correct definition of immunity to *wh*-misidentification does not mention defeaters. Once the alternative, defeater-free definition is accepted, it is a fairly immediate consequence that the memory judgment ‘I was on a ship’ — despite being vulnerable to quasi-memory — will be immune to error through *wh*-misidentification. Similarly, other non-inferential first-person bodily judgments — such as those based on proprioception or perception — turn out to be immune to error through *wh*-misidentification, despite being vulnerable to quasi-cases.

I suggested at the beginning of this paper that it was the conjunct of the *Parity Thesis* that says that bodily first-person judgments can be IEM that is most controversial. I took for granted the other conjunct that says there are some mental first-person judgments — in particular, those based on introspection — that are IEM.

Now that we have in place a positive definition of immunity to error through *wh*-misidentification, it is worth testing it against the minority strand in the literature on IEM that claims that even introspective judgments fail to be IEM. Campbell (1999) initiated this strand, with his appeal to the phenomenon of schizophrenic patients who introspect their own thoughts, but who think those thoughts are someone else’s — so called cases of thought insertion. A typical example is one in which a subject introspects the thought ‘I am thinking *p*’ but is willing only to endorse ‘Someone is thinking *p*’ or

infallibility, and even introspective judgments fail to satisfy it. For there are analogues to Shoemakerian worlds for introspective judgments. Consider a subject who almost always experiences itches-that-are-nearly-pains rather than true pains, and whose itches are almost always caused by someone else’s actually being in pain. If such a subject is inclined to judge that they are in pain, and spontaneously comes up with the hypothesis that their inclination to judge that they are in pain is caused by someone else’s actually being in pain, then their hypothesis is likely to be correct. So, with the crude reliabilism in place, one can argue that their hypothesis will be knowledge. In that case, their judgment ‘Someone is in pain’ will be knowledge. So, ‘I am in pain’, based on introspection, will be vulnerable to *wh*-misidentification

‘That other person is thinking p ’. This kind of example fails to support the verdict that introspective judgments are vulnerable to error through *wh*-misidentification, on the definition I have argued for. On that definition, what is required would be a case in which the subject is offered knowledge of ‘Someone is thinking p ’ even though ‘I am thinking p ’ is false. But the thought-insertion cases are cases in which ‘I am thinking p ’ is true, not false.¹⁶

In summary, I’ve argued that immunity to error through *wh*-misidentification really is different from the notion of immunity to error through false belief misidentification. The former is the more demanding notion so, as between the two, we should prefer to read the *Parity Thesis* as concerning the former. However, even on that reading, the *Parity Thesis* is true.

(§IV) Immunity to error through false belief-or-background-presupposition misidentification

My notion of immunity to error through false belief misidentification mentioned the *basis* for one’s belief. My notion of immunity to error through *wh*-misidentification, as well as Pryor’s definition of immunity to error through *wh*-misidentification, mentioned the *grounds* for one’s belief. I didn’t highlight this distinction, since it was possible to give my argument against Pryor’s definition, and for my own, while ignoring it. In the definition of IEM under discussion in this section, however, the distinction will take center-stage.

Earlier, I provided a gloss on the notion of a ‘basis’ for belief. To bring out the difference with the notion of ‘grounds’ for belief, consider the following case, discussed by Pryor at (1999, 276). I have justification for ‘The gas meter reads “E”’ and, because

¹⁶I think similar remarks apply to Lane and Liang (2011) and Cappelen and Dever’s (2013) more recent attempts to show that not even introspective judgments are IEM, but will not attempt to defend that suggestion here.

of that, I also have justification for ‘I am out of gas’. Because of fears about an evil demon, I form neither belief. So, it is not the case that I have an ‘I am out of gas’ belief that is based on a ‘The gas meter reads “E”’ belief. Even so, there is a kind of epistemic dependence here. The justification I have for ‘I am out of gas’ depends on the justification I have for ‘The gas meter reads “E”’. One can use ‘grounds’ as a label for this relation of epistemic dependence. In this case, we can say that ‘The gas meter reads “E”’ is included in my grounds for ‘I am out of gas’.

One might expect that in cases in which a subject actually does form a belief, the two notions — bases and grounds — will march fairly closely in step. One might think that, if a subject is justified in believing p , and the grounds for p include q , then their belief in p will have to be based on belief in q . One might also think the converse of this: if a subject is justified in believing p , and their belief in p is not based on belief in q , then it cannot be that their grounds for p include q . Applying this to the case of memory, one might argue that since a subject may be justified in judging ‘I was on a ship’ just on the basis of a memory impression, their grounds for that judgment cannot include the identity ‘I am the person from whose past my memory impression derives’.

The advocate of the notion of IEM that I discuss in this section says that this is too hasty. They think that ‘I was on a ship’, even if just based on a memory impression, *will* include that identity in its grounds. A subject can form a judgment that is justified, where the grounds on which the justification depends include a claim that is *not* part of the basis on which the subject relied. They can if that claim is a *background presupposition*.

The definition of IEM that this leads them to canvass is as follows:

A judgment ‘ a is F ’ is, relative to certain grounds g , IEM if and only if g (i) do not include an identity-belief and (ii) do not include a *background presupposition* of

identity. (*Immunity to error through false-belief-or-background-presupposition misidentification*)

This kind of definition was first set out in Coliva (2006) and followed up in Coliva (2012). It has also been explored in Wright (2012). There are some differences in formulation between Coliva and Wright, but none that make a difference to my argument. I shall target the proposal generically, while also noting a few small differences in the two presentations.

How are we to understand the notion of a background presupposition?

Background presuppositions are supposed to be a kind of epistemic ground, so they must be epistemically *relevant* to the thing being judged. The relevance, however, isn't foregrounded in the subject's own psychology: it is *in the background*. I'll explain these two elements in reverse order.

There are in fact two *different* ways that a ground's relevance could be foregrounded in the subject's own psychology and that *background* presuppositions therefore contrast with. I'll illustrate them using the case discussed at Coliva (2012, 417). Suppose I see a woman, my aunt Miriam, and form two separate judgments, 'That woman is wearing a red hat' and '*That woman is Aunt Miriam*', from which I *consciously infer* 'Aunt Miriam is wearing a red hat'. There is dependence on an identity here, of an extremely foregrounded variety.

Here is a slightly different possibility. I see Aunt Miriam. I don't form two separate judgments, one of which is an identity, and consciously infer 'Aunt Miriam is wearing a red hat'. I come by my 'That woman is wearing a hat' and 'Aunt Miriam is wearing a hat' beliefs in one fell swoop. Still, if I were asked to justify my 'Aunt Miriam is wearing a red hat', *I would appeal to the identity*. I might say, for example, 'It's perfectly clear that the woman I'm looking at is wearing a red hat. And I can see

that *that woman is my aunt Miriam*'. There is dependence on an identity here, of an at least somewhat foregrounded variety.

If *either* of these kinds of dependence is present, then 'That woman is aunt Miriam' is part of the *basis* on which 'Aunt Miriam is wearing a hat' is judged and is not a mere background presupposition. However, notice that in the case of the memory judgment 'I was on a ship' *neither* kind of dependence on an identity seems present. I definitely don't *consciously infer* 'I was on a ship' from 'I am having a memory impression as of being on a ship' and 'I am the one from whose past this memory impression derives'. In addition, it's very unlikely that that if someone challenged me to defend my 'I was on a ship' judgment, I would offer those premises. So, 'I am the one from whose past this memory impression derives' is not part of the *basis* I rely on.

That is why it is important for Coliva and Wright that there is a third kind of ground – background presuppositions. The test for a content's being a background presupposition is, first, that a belief in it isn't part of the *basis* the subject relies on (if it were, the dependence would be foregrounded in the subject's own psychology) and second, that a positive condition is met, whose obtaining captures why the content is nevertheless relevant. Coliva and Wright state the positive condition slightly differently.

For Coliva (see 2012, 417):

p is a *background presupposition* of *q* only if: were *p* to fall into question, a (rational and fully conceptually equipped subject) would be prepared to withdraw from *q*.

For Wright (see 2012, 269) (at least for judgments based directly on experience, which are the kind of judgment we are interested in):

p is a *background presupposition* of *q* only if: if the subject doubted *p*, they would doubt that their experience justified *q*.

One can have more than one set of grounds for a single judgment. So I take it that Coliva's 'withdraw from q ' is implicitly relativized to a set of grounds. E.g. if I thought that my memory was utterly unreliable, I might withdraw from judging 'I was on a ship' on the grounds of memory but still be reluctant to withdraw from 'I was on a ship' *tout court*, because I think that most people have been on ships at one point or another, and don't see why I should be any exception. Wright achieves the same kind of relativization by talking not about withdrawing from q but about doubting that one's experience justifies q . I will take the relativization as read.

Coliva specifies that our subject is rational and fully conceptual equipped, which Wright omits to do. I take it though that Wright should be read as assuming that his subject satisfies something like Coliva's restriction. It's utterly unpredictable what a sufficiently irrational subject would do, if they doubted p .

Suppose I judge, on the basis of a memory impression, 'I was on a ship'. Someone comes along and queries my 'I was on a ship' judgment. I very likely *won't* immediately start thinking about quasi-memory involving possibilities. The person who is putting me under pressure didn't specifically mention that kind of possibility, and it is far from the first thing anyone would think of unprompted.

Something different could happen though. Someone could come along and zero in on that possibility *themselves*. Perhaps they don't even mention my 'I was on a ship' judgment. They simply query whether, as they would put it, 'You are the person from whose past your memory impression derives'. I certainly wasn't thinking about that possibility before. But I can hardly claim that the question is *irrelevant* to my 'I was on a ship' judgment, which is based on memory. So if their query leads the identity to seriously fall into question for me in the sense that I actually start to doubt it, I surely had better retreat from 'I was on a ship'. So, the identity does meet Coliva and Wright's

test for being a background presupposition of ‘I was on a ship’. If having an identity as a background presupposition makes for vulnerability to misidentification, as it does on the definition of IEM they propose, then ‘I was on a ship’, based on memory, is not IEM after all.

Coliva draws this conclusion in relation to her example of a memory judgment:

So the presence of identification components featuring merely as background presuppositions allows for the possibility of judgments – like "I was in that remote part of Scotland five years ago" when based on one's memories – which, despite being grounded in irreducibly first personal contents, and having no identification component among their grounds at all from the subject's point of view, may in principle presuppose a false identification (2006, 419)¹⁷

Wright draws a similar conclusion in relation to proprioceptive judgments:

‘Consider the proprioceptively based judgment that my legs are crossed. ... this judgment is grounded in a pure experience... But however that may be, the judgment does nevertheless rest on an identification, namely that it is my body — my legs — that are the source of the proprioceptive sensations that I am having — or perhaps better: the person whose arrangement of limbs is the causal source of my current proprioceptive experience is myself... One consequence is that such judgments are not IEM.’ (2012, 271)

Coliva and Wright’s focus on background presuppositions might initially appear like an ingenious way of capturing a dependence of a particularly subtle kind on an identity. In fact though, the strategy overshoots. The attempt to craft a definition of IEM that rules out memory or proprioceptive judgments as IEM results in a notion of IEM that captures an uninteresting epistemic status, one that no judgments have. The root problem is that background presuppositions are so promiscuous.

¹⁷ Coliva says here that the identity here is a background presupposition but not ‘among the grounds at all from the subject’s point of view’. I have been putting this point by saying that the identity is not part of the *basis* the subject relies on. Coliva sets out how she is using the phrase ‘the subject’s own grounds’ at (2016, 416).

Take an arbitrary judgment of the form ‘I am F’.¹⁸ And now consider a scenario in which I come to doubt whether I really am the person whose being F is responsible for my judging that I am F. In that scenario, I would of course come to doubt whether I am justified in judging ‘I am F’. It follows that ‘I really am the person whose being F is responsible for my judging that I am F’, an identity, is a background presupposition of my judgment. Since ‘I am F’ was arbitrarily chosen it follows that, on this definition, no judgments are IEM.

That argument seems very quick. It might be objected that, for some values of F, there will be no possible scenario in which I could be made to doubt whether I really am the person whose being F is responsible for my judging that I am F.

This seems wrong. It is certainly wrong for fallible judgments and it also seems wrong for infallible judgments. I will take the two cases separately.

Suppose ‘I am F’ is fallible. Then there will be possible cases in which I judge ‘I am F’ even though I am not F. If there are such cases, then a subset of the cases will be ones in which there is someone else who really is F, and where that person’s being F contributes to bringing about whatever combination of circumstances lead me to judge ‘I am F’. If such cases are possible, then there is no reason why I couldn’t be given reason to suspect that they obtain. So, no fallible judgment will count as IEM.

Perhaps surprisingly, the same is true for infallible judgments, if there are any of those. The key point is that even if a judgment is infallible, I may be persuaded that it is fallible. E.g. even if ‘I am in pain’, based on introspection, is infallible, I may be convinced by clever but ultimately unsound anti-luminosity arguments that it is fallible. I may further be convinced that, in a particular case, my inclination to judge ‘I am in pain’ is in error and moreover there is someone else whose being in pain has contributed

¹⁸ Or, indeed, an arbitrary judgment of the form ‘*a* is F’. *a* need not be ‘I’, though that is the case relevant here.

to bringing about my inclination to judge that I am in pain. E.g. suppose I am convinced that there is someone else whose being in pain regularly causes me to have an itch that is nearly but not quite intense enough to count as a pain, and that I am apt to mistake for a pain. I am further convinced, on a given occasion, that this has actually happened. In that case, I ought to withdraw from 'I am in pain' because of my doubt about the identity 'I am the person whose being in pain is responsible for causing me to be inclined to judge "I am in pain"'. This is despite the fact that, we are supposing, 'I am in pain' is in fact infallible. So, even infallible judgments will count as vulnerable to error through misidentification, on the definition being considered. That definition is unacceptable as a definition of IEM, or as a way of fleshing out the notion of thought about oneself *as subject*, because it captures an epistemic status that none of our judgments have.

One might offer a test for being a background presupposition stronger than the conditional one that I, following Coliva and Wright, have applied. But, first, if advocates of the definition had an appropriate stronger test up their sleeve, they presumably would have offered it. Second, and more importantly, Coliva and Wright's arguments that memory and proprioceptive judgments have identities as background presuppositions did rely on the conditional test for being a background presupposition being the right test. We can strengthen the notion of background presupposition to make it less promiscuous if we like. But we have no reason to expect that, on the new notion this yields, first-person bodily judgments will still count as always involving background presuppositions of identity.

(§V) Conclusion

I've discussed three different epistemic statuses – immunity to error through false belief misidentification, immunity to error through *wh*-misidentification, and immunity to error through false belief-or-background-presupposition misidentification.

In relation to the first status, immunity to error through false belief misidentification, I've argued, in agreement with everyone else in the literature, that bodily first-person judgments can have it.

In relation to the second status, immunity to error through *wh*-misidentification, I've argued that it really is a different, and more demanding, status than immunity to error through false belief misidentification (agreeing with Pryor). But I've argued that it is a status that bodily first-person judgments can have (disagreeing with Pryor). They won't have it, if we understand immunity to error through *wh*-misidentification in terms of undermining defeaters, as Pryor proposes we do. But we shouldn't do that, since doing that causes the notion of immunity to error through *wh*-misidentification to collapse into that of infallibility.

In relation to the third status, immunity to error through false belief-or-background-presupposition misidentification, I've argued that bodily first-person judgments cannot have it (agreeing with Coliva and Wright). But that is only because the third status turns out to be one that no judgments have. We might avoid that undesirable result by modifying the definition of what it is for something to be a background presupposition. But we have no reason to expect that, on the new definition, first-person bodily judgments will still count as always involving background presuppositions of identity.

My conclusion is that we have no evidence against the *Parity Thesis*. On each of the two notions of IEM that succeeds in capturing an interesting epistemic status, immunity to error through false belief misidentification and immunity to error through

wh–misidentification, the *Parity Thesis* is true. Of course, it doesn't follow that we should endorse the *Parity Thesis*. There *might* turn out to be some interesting epistemic status that some mental first-person judgments have, that no bodily first-person judgments have, and that relates sufficiently closely to the kind of examples used to introduce the idea of immunity to error through misidentification that that status would count as a new kind of immunity to error through misidentification. If one thinks that there *must* be such a status, then the moral of this paper is just that we need to keep on looking for it. But the failure to find it, after considerable effort, naturally makes pressing the question of what reason there is to think that there must be such a status. The alternative possibility is that when I think about, e.g., the orientation of my limbs, or my location relative to the things in my environment, or my skin being covered by small beads of sweat, I think about myself as a bodily thing, and simultaneously as a subject – so I think about myself as a bodily subject. I've argued that the discussion of immunity to error through misidentification, and in particular of quasi-cases, does nothing to undermine this possibility. Is there any independent reason to assign it a low probability?¹⁹

¹⁹ For discussion and comments, thanks to Annalisa Coliva, Manuel García-Carpintero, Rory Madden, Michael Martin, Lucy O' Brien, Michele Palmira, James Pryor, Léa Salje, Carlota Serrahima and three anonymous referees.

References

- Ball, Brian and Blome-Tillmann, Michael. 2014. "Counter Closure and Knowledge Despite Falsehood", *Philosophical Quarterly* 64 (257): 552-568.
- Campbell, John. 1999. "Immunity to Error through Misidentification and the Meaning of a Referring Term," *Philosophical Topics* 26 (1&2): 89–104.
- Campbell, John. 1997. "Sense, Reference and Selective Attention", *Proceedings of the Aristotelian Society Supplementary Volume* (71): 55-98
- Campbell, John. 1999. "Schizophrenia, the Space of Reasons and Thinking as a Motor Process," *The Monist* 82 (4): 609-25.
- Cappelen, Herman and Dever, Josh. 2013. *The Inessential Indexical*. Oxford: Oxford University Press.
- Cassam, Quassim. 1997. *Self and World*. Oxford: Oxford University Press.
- Coliva, Annalisa. 2006. "Error Through Misidentification: Some Varieties," *Journal of Philosophy* 103 (8): 407-425.
- Annalisa Coliva, "Which 'key to all mythologies' about the self? – A note on where the illusions of transcendence come from and how to resist them," in Prosser and Recanati (eds) (2012).
- Evans, Gareth. *The Varieties of Reference*. New York: Oxford.
- Lane, Timothy and Liang, Caleb. 2011. "Self-Consciousness and Immunity," *Journal of Philosophy* 108 (2): 78-99
- O' Brien, Lucy. 2007. *Self-Knowing Agents*. Oxford: Oxford University Press.

Parfit, Derik. 1984. *Reasons and Persons*. Oxford: Oxford University Press.

Peacocke, Christopher. 1983. *Sense and Content: Experience, Thought and their Relations*. Oxford: Oxford University Press.

Prosser, Simon and Recanati, Francois. 2012. "Immunity to Error Through Misidentification: New Essays". Cambridge University Press.

Pryor, James. 1999. "Immunity to Error through Misidentification," *Philosophical Topics*, 26 (1–2)

Pryor, James "Problems for Credulism" in Chris Tucker (ed.) *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. Oxford, 2013

Recanati, Francois (2012). "Immunity to Error through Misidentification: What it is and where it comes from" in Prosser and Recanati (eds) (2012).

Shoemaker, Sydney. 1968. "Self-Reference and Self-Awareness." *The Journal of Philosophy* 65 (19): 555–67.

Smith, Joel. 2006. "Which Immunity to Error?" *Philosophical Studies* 130 (2): 273-83.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

Wittgenstein, Ludwig. 1958. *The Blue and Brown Books*. New York: Oxford.

Wright, Crispin. "Reflections on François Recanati's, 'Immunity to error through misidentification: what it is and where it comes from'" in Prosser and Recanati (eds) (2012).

