



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Immunity to error through misidentification and the epistemology of de se thought

Citation for published version:

McGlynn, A 2016, Immunity to error through misidentification and the epistemology of de se thought. in M García-Carpintero & S Torre (eds), *About Oneself: De Se Thought and Communication*. Oxford University Press, Oxford, pp. 25-55. <https://doi.org/10.1093/acprof:oso/9780198713265.003.0002>

Digital Object Identifier (DOI):

[10.1093/acprof:oso/9780198713265.003.0002](https://doi.org/10.1093/acprof:oso/9780198713265.003.0002)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

About Oneself

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Immunity to Error Through Misidentification and the Epistemology of De Se Thought¹

‘Are you in pain, dear mother?’

‘I think there’s pain somewhere in the room,’ said Mrs Gradgrind, ‘but I couldn’t positively say that I have got it.’

(Charles Dickens, *Hard Times*)²

1. Immunity to Error Through Misidentification and De Se Thought

When essaying a singular judgment of the form [*a* is *F*] there seem, in principle at least, to be two distinct ways to fall into error. First, one might be wrong about whether the object *a* really is *F*. Second, one might be wrong about whether it is *a* that is *F*; one reaches for the right predicate, but misidentifies the object in question when applying it. Suppose, for example, that I am out walking in the park, when across the pond I see a woman wearing an enormous, bright red hat (Coliva 2006, Wright 2012). Mistakenly taking her for my Aunt Lillian, I judge

¹ A very early version of some of this paper was presented at the Northern Institute of Philosophy at a workshop in honour of Crispin Wright in December 2012. My thanks to all of the participants in that discussion, particularly Patrick Greenough, Marcus Rossberg, Robbie Williams, and Elia Zardini. An early version of the discussion of Pryor was presented at Institut Jean-Nicod at the first of two joint workshops with the Northern Institute of Philosophy on de se thought and self-knowledge in November 2013. I learned a tremendous amount from the talks, questions, and discussions throughout those two events, and would like to particularly thank Marie Guillot, Andy Hamilton, Guy Longworth, Giovanni Merlo, Daniel Morgan, Kristina Musholt, and Stephan Torre. Special thanks are due to Crispin Wright and Annalisa Coliva for inspiration and discussion, to Jan Pieter Maes, Bryan Pickel and Duncan Pritchard for discussion, to Ted Parent and an anonymous reviewer for extremely useful comments that led to a great many improvements, and to the editors of this volume for their comments, support, and patience.

² Thanks to Catherine Wright for drawing my attention to this passage; it is also mentioned by Hamilton (2013: 38fn3).

that [Aunt Lillian is wearing an extraordinary hat]. I'm right that the woman I see across the pond is wearing an extraordinary hat, the error in my judgment that [Aunt Lillian is wearing an extraordinary hat] being due to my having misidentified that woman as my Aunt Lillian. Let us call the first kind of error an *error of mispredication*, and the second an *error through misidentification*.

Certain classes of judgments have traditionally been thought to possess immunity to the first kind of error. In particular, self-ascriptions of one's occurrent mental states have been taken to be incorrigible in this sense. According to this familiar Cartesian thought about psychological self-knowledge, I cannot be mistaken if I judge that [I am in pain] or that [I believe that all human beings are created equal], assuming that I form those judgments in the usual first-personal way, rather than on the testimony of my doctor or psychoanalyst. Few philosophers today hold such a strong incorrigibility thesis about our own mental states, but more plausible descendants of the Cartesian picture are still in currency.

In a well-known passage in *The Blue Book* (1958: 66-7), Wittgenstein observed that certain judgments seem to have the property of being immune to the second kind of error distinguished above: the property of *immunity to error through misidentification*, to adopt now standard terminology introduced by Sydney Shoemaker (1968). Suppose that I judge that [I see a barn]. Now, this judgment isn't immune to mispredication; I might be looking at a well-crafted façade rather than a genuine barn (compare Shoemaker 1968: 557). However, my judgment doesn't seem to be vulnerable to error through misidentification, when formed

in the normal way. There doesn't seem to be much sense to be made of the suggestion that I've misidentified the person who is seeing a barn as myself; that sort of error doesn't seem to have been 'provided for', as Wittgenstein put the point in the passage mentioned above.

As Shoemaker and others have stressed (e.g. Shoemaker 1968: 557 and Evans 1982: 218-9), a judgment can be immune to error through misidentification when based on certain grounds, while being vulnerable to such error when based on others. For example, my judgment that [I believe that all human beings are created equal] might be immune to error through misidentification when made on the usual grounds I have for such a self-ascription, but not when made on the basis of reading a report prepared by my psychoanalyst.

What is the philosophical significance of the (apparent) fact that certain judgments display immunity to error through misidentification? Two tempting answers are suggested by reflection on two salient features of the example just given (and, perhaps tellingly, the examples offered by Wittgenstein himself: see Evans 1982: 218-20). First, the judgment that [I see a barn] is a *de se* judgment. This judgment is not merely about myself, as the judgment that [AM sees a barn] would be even if a bout of amnesia had caused me to forget that I am AM, but involves thinking of myself in a first-personal way. Second, the judgment involves self-attribution of a mental state. Given this, it is natural to wonder how immunity to error through misidentification relates to *de se* thought and psychological self-knowledge. In particular, we might ask to what extent the apparent immunity to error through misidentification of such judgments is

implicated in philosophers' tendency towards metaphysical, epistemological, and semantic extravagance when engaged in reflection about de se thought and talk and about avowals of psychological states. Here I'll be particularly interested in the phenomenon's bearing on the de se.

An immediate concern with the suggestion that immunity to error through misidentification bears any intimate relation to de se thought is that it doesn't seem to be at all *distinctive* of de se thought. On the one hand, many de se judgments are vulnerable to error through misidentification. Consider, for example, my judgment that [I am making a mess], made upon realizing that this is the best explanation of why I can't catch up with the person whose sugar bag is punctured and leaving a trail of sugar throughout the supermarket. This seems no less first-personal than the judgment that [I see a barn]—indeed, we may say that it is something of a paradigm of de se thought (Perry 1979)—and yet it's clearly quite possible that I have *misidentified* the mess-maker as myself (compare Wright 2012: 259). On the other hand, there seem to be wide classes of judgments that display immunity to error through misidentification despite not involving the de se. The standard examples involve perceptual demonstratives (e.g. Shoemaker 1968: 558, Evans 1982: 179-91). The judgment that [that thing is coming straight towards us], made on the basis of vision, is vulnerable to errors of mispredication, since the demonstrated object may be going to miss our location by a comfortable margin. But there seems to be no sense to be made of the suggestion that I'm right that the demonstrated object is coming straight

towards us, but I've gone wrong in misidentifying it as *that thing*.³ Likewise, when made on certain kinds of grounds, judgments that we would express using the second- and third-person pronouns, or using the indexicals 'here' and 'now' seem to be immune to error through misidentification (Wright 2012: 254). We seem to have a unified phenomenon here across the different kinds of examples, and that phenomenon doesn't seem to be distinctive of the *de se*.

So much by way of background. The principal topic of this paper is the idea, associated largely with pivotal work by Shoemaker (1970), that there is a significant distinction to be drawn within the class of judgments/grounds pairings that display immunity to error through misidentification, typically drawn with the labels 'logical' and '*de facto*' immunity to error through misidentification. Roughly, Shoemaker's thought was that certain judgments, in particular judgments about one's past based on memory, are immune to error through misidentification, but only contingently so. Other judgments, including all of the examples we have considered so far, are in some contrasting sense *guaranteed* to be immune to error through misidentification.

There are two main reasons to be interested in this topic, from the point of view of one who is interested in *de se* thought (though I will focus on the first here). First, in recent work Herman Cappelen and Josh Dever (2013) have argued that the *de se* is a philosophically 'shallow' phenomenon, and as part of their case for

³The same kind of example is standardly taken to show that immunity to error through misidentification should not be confused with the phenomenon of *guaranteed referential success*, highlighted in Anscombe 1975. My uses of 'I' may be guaranteed to refer (and to refer, in particular, to myself), but my uses of 'that' aren't guaranteed to refer to anything, even when I'm in possession of grounds that enable me to make judgments that are immune to error through misidentification.

this claim they attack the idea that *de se* judgments can be epistemically privileged. As such, they have asserted that no judgments are immune to error through misidentification, and they do so largely because of the kinds of cases taken to motivate a distinction between logical and *de facto* immunity (2013: 131-2). Having discussed these examples, they claim that ‘no satisfactory account has been given of a philosophically interesting epistemic phenomenon that survives the problem cases’ (2013: 132). Now, I don’t think Cappelen and Dever’s brief discussion justifies this claim. As we will see, the two main proposals for accommodating the problem cases in the recent literature come from James Pryor (1999) and Annalisa Coliva (2006), and Cappelen and Dever offer only a superficial discussion of the former, while failing to mention the latter at all. That said, I’ll argue at length below that neither Pryor nor Coliva’s proposals is satisfactory, and so there is a genuine challenge here for anyone who wants to defend the thesis that immunity marks a respect in which *de se* thought is epistemically privileged. Here, I’ll try to improve on these earlier proposals in order to meet this challenge.

Second, the debate over a distinction between logical and *de facto* immunity may have a bearing on whether studying the phenomenon of immunity can shed any light on the nature of *de se* thought. The observation that immunity is not distinctive of the *de se* might be taken to suggest that it can not (see, for example, McGlynn 2011, Morgan 2012, Wright 2012, Sainsbury and Tye 2012: 146-9 and Cappelen and Dever 2013: 133-9). That’s too quick, if there are stronger and weaker varieties of immunity to take into account. One possibility is that ‘logical’ immunity might prove to be distinctive of *de se* thought, even if immunity *per se*

is not. Unfortunately, there seems to be no plausibility whatsoever in this suggestion; indeed, *all* of the examples given above of non-de se judgments that possess immunity when based on the right kinds of grounds are naturally taken to fall on the logical side of Shoemaker's distinction, while the examples that motivate acknowledging something like a phenomenon of *de facto* immunity all involve de se thoughts.

There is, however, a subtler possibility. It may be that there's a unified account of why immunity arises where it does which, given the scope of the phenomenon, will presumably make no essential reference to distinctive features of de se thought. However, this doesn't entail that the explanations of, say, why some de se thoughts and why some demonstrative thoughts display *logical* immunity will also be the same, or that they will not need to advert to distinctive features of those judgments. This possibility has been pointed out and explored in recent work by Coliva (2012), which builds directly on her account of the logical/*de facto* distinction (2006). An interesting further project, then, is to assess the plausibility of her view of the relationship between immunity and the de se, and in particular to reconsider it in light of the present study's conclusions about how to draw the required distinction. This further project will have to wait for another occasion, but I do hope to accomplish some of the groundwork for that discussion here.

The remainder of this paper divides into seven sections. In section 2, I will introduce the kind of examples that suggest that there's an important distinction to be drawn within the class of judgments that are immune to error through

misidentification, and I'll offer a brief overview of the well-known debate between Shoemaker and Gareth Evans concerning the import of these examples. In section 3 I will criticize two proposals for how to draw the needed distinction found in Shoemaker's discussion, drawing two morals for the discussion in subsequent sections. Section 4 looks at Pryor's distinction between two varieties of immunity to misidentification and argues that it fails to offer a satisfying treatment of the problem cases, either in Pryor's own hands or as appropriated by Coliva, while section 5 criticizes Coliva's proposal concerning how to draw the logical/*de facto* distinction. Section 6 examines a recent discussion of the problem cases due to François Recanati, and section 7 briefly turns to Crispin Wright's position, according to which the problem cases show that fewer kinds of *de se* judgments are immune to error through misidentification than we might otherwise suppose.⁴ These critical discussions set the stage for my own positive proposal in section 8, which is that the problem cases call for a better understanding of the modality implicit in the notion of immunity to error. I will offer a particular account of immunity along these lines, showing how it enables us to draw a version of the distinction between logical and *de facto* immunity, arguing that it improves on the proposals criticized in earlier sections, and discussing some pressing objections to it.

2. *Quasi-Memory*

⁴ Bermúdez 2013 offers an account of when memory-based judgments about one's own past are immune, but he much too quickly sets aside the problem that animates the present discussion (2012: 212fn2), and so I don't consider his proposal here.

Consider a judgment about an event in my own past, made on the basis of remembering that event 'from the inside' (Shoemaker 1970: 273). To borrow Coliva's example (2006), suppose I judge that [I was in that remote part of Scotland five years ago], based (let's agree for now) on my apparent memories of having been there. Is my judgment, based on these grounds, immune to error through misidentification? Shoemaker (1970) argues that there's an important sense in which it is not, appealing to the metaphysical possibility of quasi-remembering another's past. The standard characterisation of quasi-remembering in the literature comes from Evans:

A subject *q*-remembers an event *e* if and only if (i) he has an apparent memory of such an event, and (ii) that apparent memory in fact embodies information deriving from the perception of that event by a person who is not necessarily himself.' (1982: 247-8; compare Shoemaker 1970: 271)

Quasi-remembering involves having an apparent memory of some event 'from the inside', as if it were part of one's own past, but where that memory impression may derive either from one's own past or the past of someone else. So it includes as sub-cases normal cases of remembering one's own past, as well as cases in which one has memory impressions of another's past (as when Rachel in Ridley Scott's *Bladerunner* 'remembers' witnessing the unenviable fate of the orange and green spider in the bush), and it excludes cases in which one has a memory *illusion*: a memory impression which derives from no one's past.

As mentioned, what puts pressure on the claim that our memory-based judgments are immune to error through misidentification is the possibility that one is quasi-remembering the past of another (see Pryor 1999: 289 for an

example involving fission). Similarly, the possibility of ‘cross-wiring’ cases, in which one receives proprioceptive information about the position of another’s body, suggest that my judgment that my legs are crossed isn’t immune even when based on proprioception rather than vision (see O’Brien 2007: 38-43 and Recanati 2007: 155-60). Some have even considered the possibility that perceptual examples, like our initial example of an immune judgment [I see a barn], might be vulnerable to error through misidentification due to cases of ‘quasi-perception’ (see Evans 1982: 184 and Recanati 2007: 155-6). In this paper, we’ll largely focus on the possibility of quasi-memory.

Let’s concede, as seems plausible, that it is metaphysically possible to quasi-remember past events in another’s life.⁵ In order for this possibility to threaten the immunity to error through misidentification of our actual memory-based judgments, we have to buy into certain assumptions about the kinds of grounds that memory can serve up. Most participants in the debate seem to assume that remembering is subjectively indistinguishable from quasi-remembering, and so that the grounds for judgments about events in one’s own past are normally the same whether based on memory or quasi-memory: as we presupposed above, they are one’s ‘apparent memories’ (e.g. Pryor 1999: 289, Coliva 2006: 418, and Wright 2012: 272). For example, Coliva suggests that my grounds for the judgment [I was in that remote part of Scotland five years ago] will be that I seem to remember being in that remote part of Scotland five years ago (2006: 419). In effect, what’s being presupposed here is a kind of internalist picture of one’s grounds, on which one’s grounds supervene on one’s non-factive mental

⁵ Due to limitations of space, I’ll follow most of the literature here in simply taking the possibility of quasi-remembering as a premise, though I note that this is a substantial assumption; see Roache 2006 for a nuanced defense and Hamilton 2013: chapter 2 for a recent denial.

states. It's a picture that I'm sympathetic to, but we should note that there are externalist alternatives; for example, on some views my actual grounds for the judgment [I was in that remote part of Scotland five years ago] might be the factive state that *I remember* being in that remote part of Scotland five years ago, not merely that I *seem* to remember this. Moreover, the arguments for this kind of internalism are rather underwhelming (McGlynn 2014: 66-72).

However, it doesn't look like one can simply sidestep the issues raised by quasi-memory by adopting an externalism account of the grounds available for the judgments in question. My judgment that I was in that remote part of Scotland five years ago is not incorrigible. Let's suppose that it is mistaken; I was actually there six years ago. My grounds in such a case cannot be that I remember being in that remote part of Scotland five years ago, since I was not, and yet most externalists will not want to say that my judgment was groundless. A natural suggestion is that my grounds in this case are that I seem to remember that I was in that remote part of Scotland five years ago (compare Williamson 2000: 198). Since my judgment is plausibly still immune to error through misidentification when made on these grounds, we can reconstruct the problem that this paper is directed at for the externalist too. So while there's no doubt wriggle room here, a natural way to develop externalism about grounds fails to address the issues raised by quasi-memory. In the remainder of this paper I'll grant the internalist description of the grounds typically available to us presupposed by Pryor and Coliva, but we should bear in mind that an externalist has reason to be invested in our discussion too.

Given this assumption about the kinds of grounds offered by memory, the possibility that I might quasi-remember another's past and judge that [I was in that remote part of Scotland five years ago] on that basis shows that it is metaphysically possible for this judgment to be in error through misidentification when based on the relevant kinds of grounds. Hence, Shoemaker infers, judgments like this are not immune to error through misidentification even when made on the relevant grounds. However, since cases like the one Pryor describes don't arise in the actual world, such judgments enjoy a *de facto* immune.

On the other hand, Evans stresses that the grounds that memory serves up for this judgment are *de se*, as is the resulting judgment. Given this, my judgment will not be based on a bridging identity claim (to the effect that I am the person whose past I seem to remember now); in Evans's terminology, my grounds will be 'identification-free' (1982: 180). So my judgment, based as it is on such identification-free grounds, is not vulnerable to error through misidentification. This immunity is not merely *de facto*; it's secured by the kinds of grounds available for the judgments in question, and not by contingencies about whose past my apparent memories derive from.

Coliva (2006: 407) suggests that we're close to having a paradox on our hands here, since both Shoemaker and Evans seem to be making solid points, and yet they reach diametrically opposed conclusions. On this basis, Coliva suggests that we should look for an evenhanded reconciliation, one that dissolves the disagreement while explaining what Shoemaker and Evans each get right. And a natural hope—one pursued in this paper—is that a more satisfactory

understanding of Shoemaker's distinction between logical and *de facto* immunity, or some suitable replacement, will allow us to make room for this kind of reconciliation.

3. Two Morals

In this section I want to consider two possible ways of trying to understand the needed distinction that can be extracted from Shoemaker's discussion. They are not particularly promising, at least as they stand, and so we shall be brief. Nonetheless, this section will suggest two morals for the discussion to follow.

First, it's worth asking what's wrong with Shoemaker's official treatment of the problem cases, as outlined in the previous section. As we saw, Shoemaker holds that judgments about one's own past aren't really immune to error through misidentification, as is shown by the possibility of quasi-remembering another's past. However, such errors will not in fact arise, since this possibility is not actually realized, and in this limited sense, such judgments are immune. The main problem with this account is that Shoemaker's notion of *de facto* immunity doesn't really seem like a kind of immunity at all. Indeed, it's not obvious that there's much more to it than the claim that, as Cappelen and Dever put it, 'our beliefs, in normal circumstances, track the world' (2013: 133), which they point out isn't enough to show that there's anything epistemically privileged about judgments that have this status; the modal force implicit in the notion of immunity—the topic of section 8 of this paper—seems to have been all but lost.⁶

⁶ For a different objection see Coliva 2006: 423.

Second, having noted that a judgment may be immune to error through misidentification when justified by some grounds and vulnerable when justified by others, it is natural to wonder whether any judgments are immune when justified by *any* possible grounds one could have for them. Let us call such judgments, if there are any, *absolutely* immune to error through misidentification (Pryor 1999: 279, drawing on Shoemaker 1968: 557). Might the logical/*de facto* distinction give way to a distinction between those judgments that are absolutely immune, and those that are only immune relative to certain grounds?

The main difficulty with this proposal is that virtually no judgments seem to be absolutely immune; as Pryor argues, even *cogito*-like judgments seem to be vulnerable when made on some possible grounds (1999: 300n19; compare Recanati 2007: 149-50). Few, if any, of the judgments we wanted to contrast with memory-based judgments about one's own past are plausibly construed as absolutely immune, and so the distinction offers no insight into the contrast we wish to understand.

A more general point suggests itself. Our discussion has just uncovered a cheap sense in which pretty much any judgment that is immune to error through misidentification is only contingently so; the same judgments made on different grounds may not be so immune. The moral I want to draw from this is that, on the face of it at least, what needs explaining (or, perhaps, explaining away) is a contingency that relates to pairings of judgments and grounds.

4. Wh-Misidentification

In an influential paper (1999), Pryor distinguishes two varieties of misidentification, and two corresponding notions of immunity, arguing convincingly that it's often not clear which is intended in the literature (1999: 286-8). First, we have what Pryor calls *de re* misidentification. Paradigm cases of *de re* misidentification have the following structure:

- i. One has grounds that justify one in believing, of some particular object *b*, that it is *F*.
- ii. One mistakenly (though perhaps justifiably) believes that $a = b$.
- iii. On this basis one judges that *a is F*.

It is important to note that this is intended as a reconstruction of the rational grounds of one's judgment—the causal or, better, explanatory basis of one's judgment—rather than as an attempt to describe a piece of conscious reasoning that one goes through in order to reach that judgment. Given this characterisation, our paradigm example of error through misidentification from the introduction above is a case of *de re* misidentification. One has grounds for believing that the elderly lady one can see is wearing an extraordinary hat, and (mistakenly) judges that [Aunt Lillian is wearing an extraordinary hat] because one has misidentified that woman as one's Aunt Lillian.⁷

⁷ Notice, however, that despite the title of Pryor's paper, his characterisation of *de re* misidentification is not, and is not intended to be, a characterisation of *error* through *de re* misidentification (and a parallel point holds for his characterisation of *wh*-misidentification). Nothing in Pryor's three conditions requires that one's belief that *a is F* itself is mistaken, but only that it rest on a mistaken identification. This is quite intentional on Pryor's part (1999: 298n9); he—in my view, correctly—sees error through misidentification (whether of the *de re* or *wh*- variety) as one species of a broader phenomenon of misidentification. This won't matter for present purposes, though I'll follow Pryor's preferred terminology in this section.

The contrast is with what Pryor calls *which-object*-misidentification, or wh-misidentification for short. Pryor's characterisation of this is rather more complex than his characterisation of de re misidentification. Pryor first characterises what it is for grounds G to offer one knowledge that *p*. Grounds G offer one knowledge that *p* if, in virtue of having G, one meets all of the conditions for knowing *p* —with two possible exceptions. First, G can offer one knowledge that *p* even if one does not believe *p*. Second, G can offer one knowledge that *p* even if one possesses an undefeated defeater for *p*, rendering *p* all things considered unjustified for one. So:

Roughly speaking, grounds G will offer you knowledge that *p* in those cases where, if you were to believe that *p* on the basis of G, *and you possessed no defeating evidence*, then your belief would constitute knowledge. (1999: 281-2, emphasis in original)

Pryor then defines wh-misidentification in terms of grounds offering one knowledge. In these terms, a case of wh-misidentification has the following structure:

- i'. One has grounds *G* that offer one knowledge of the existential generalization that something is *F*, where this is independent of any justification one has for believing of some particular object that it is *F*.
- ii'. Partly on the basis of *G*, one has justification (or believes oneself to have justification) for believing of some particular object *a* that it is *F*.

iii'. However, *a* is not *F*. Some distinct object *b* is *F*, and it is because the grounds *G* derive from the fact that *b* is *F* 'in the right way' that *G* offers one knowledge that something is *F*.

Pryor illustrates this structure with the following example:

I smell a skunky odor, and see several animals rummaging around in my garden. None of them has the characteristic white stripes of a skunk, but I believe that some skunks lack these stripes. Approaching closer and sniffing, I form a belief, of the smallest of these animals, that it is a skunk in my garden. This belief is mistaken. There are several skunks in my garden, but none of them is the small animal I see.' (Pryor 1999: 281)

Several comments on Pryor's characterisation of wh-misidentification are necessary. First, the qualification that one's grounds for the existential must be independent of any justification one has for any relevant de re belief isn't explicit in Pryor's official formulation of the conditions for wh-misidentification. However, it's suggested by his example and by other statements he makes about wh-misidentification in the surrounding text (in particular, see 1999: 282-3). Moreover, the qualification is necessary to cleanly distinguish wh-misidentification from de re misidentification.⁸

Second, and more importantly, we should take stock of the shift from talk of justification, as in Pryor's characterisation of de re misidentification, to talk of grounds offering knowledge. Notice first that this shift is only partial, since (ii') still makes reference to justification. It's this that allows this to be a

⁸ I'm grateful here to discussion with Giovanni Merlo.

characterisation of a kind of *misidentification*; Pryor's account of what it takes for grounds to offer one knowledge, together with the factivity of knowledge, rule out the possibility that any grounds can offer one knowledge of a *false* claim of the form *a is F*. Presumably Pryor is operating with a non-factive notion of justification, and so one's grounds *can* justify one in accepting a falsehood.

The shift from justification to offering knowledge between (i) and (i') has some intuitively unwelcome results. Consider a modified version of Pryor's skunk example in which I have understandably mistaken the smell of burning rubber for the smell of a skunk. I see an animal in my garden, and after getting closer and sniffing (for motives that we will leave unclear), I judge it to be the source of the skunky odour, and so mistakenly judge the animal to be a skunk in my garden. In fact, there isn't a skunk in my garden at all. This looks on the face of it like a case of wh-misidentification, but Pryor's characterisation rules this out; clause (i) isn't satisfied, since my grounds do not offer me knowledge that there is a skunk in my garden (since this isn't even true).

However, although Pryor's account of wh-misidentification has the result that the case just described is not an example, his characterisation of vulnerability to wh-misidentification arguably delivers the correct result that the judgement is vulnerable when made on the grounds in question. A judgment that [*a is F*] made on grounds *G* is vulnerable to wh-misidentification, according to Pryor (1999: 284), just in case it's possible for one's grounds *G* to be defeated by undercutting evidence—evidence that 'discredits' the justification *G* provides for the judgment that [*a is F*] without providing justification for thinking that *a* is in fact not-*F*—and the combination of *G* plus the undercutting evidence no longer justifies one

in judging that *a* is *F* but *could* offer one knowledge that something is *F*. In an attached note, Pryor writes (1999: 302n28, emphasis in original):

Some justification is such that it *could* offer you knowledge that *p* iff it's intrinsically such that, in some nearby possible world *where p is true*, justification of that intrinsic sort does offer you knowledge that *p*.

Pryor has in mind here precisely the kind of case I offered above: a case in which nothing is *F* and so one's grounds can perhaps justify one in believing that something is *F*, but cannot offer one knowledge that something is *F*. In the burning rubber case, for example, my judgment that the animal I see is a skunk in my garden, made on the basis of the grounds described, is plausibly enough *not* immune to misidentification by Pryor's lights. The reason is that we can perhaps say that there's a nearby world in which there is a skunk in my garden and my smelling a skunky smell offers me knowledge of this existential, and so my grounds *could* offer me knowledge of the existential, given how Pryor understands this condition. The upshot is that even if Pryor's characterisation of wh-misidentification yields implausible results, due to the employment of a factive epistemic status in its characterisation, this is corrected for by the time we reach his characterisation of vulnerability and immunity to wh-misidentification. Still, this all seems rather gratuitously complicated.

Why does Pryor characterise wh-misidentification as he does, introducing the awkward discrepancies with his characterisations of de re misidentification and immunity to de re misidentification that I have noted? That's rather unclear.⁹

⁹ I suspect that a clue here is provided by Pryor's appeal to his characterization when replying to an objection from Evans (1999: 294-5), but I won't press the point here.

Pryor says only that framing the characterisation in terms of grounds offering knowledge 'will greatly ease exposition' (1999: 281). This is dubious; rather, in the absence of any principled motivation, Pryor's own characterisation seem gratuitously complicated, restrictive, and unlike his characterisation of de re misidentification, leading to further complications when it comes time to characterise immunity to wh-misidentification.

Setting aside these puzzling issues about Pryor's characterisation of wh-misidentification, we have two notions of misidentification on the table, and two corresponding varieties of immunity to misidentification. Pryor takes the immediate interest of this distinction to be the vindication it offers Shoemaker over Evans in the debate reviewed in section 2. In order to lay out Pryor's position, we need to first understand how he views the relationship between de re misidentification and wh-misidentification, and between the associated notions of immunity.

First, Pryor holds that neither form of misidentification is sufficient for the other. Pryor takes the possibility of wh-misidentification without de re misidentification to be shown by his skunk example (1999: 285). As for the possibility of de re misidentification without wh-misidentification, Pryor asks us to consider an example in which I see a blue-coated man carrying a gun, and mistakenly believing him to be Sam, judge of Sam that he is carrying a gun (1999: 275). Here we have de re misidentification without wh-misidentification (1999: 285).

Despite this, Pryor argues that if a judgment is immune to wh-misidentification when based on grounds *G*, then that judgment is also immune to de re

misidentification when based on *G*. Crucially, however, Pryor thinks that the converse fails:

In [the skunk example], my belief that *that* animal is a skunk does not rest on any identity assumption. There is no particular animal I antecedently know to be a skunk, which I could reidentify as the small animal I see. So this belief is immune to *de re* misidentification, when based on the grounds described in [the example]. But in [the example], my belief *does* exemplify *wh*-misidentification. (1999: 285-6)

Pryor concludes on this basis that immunity to *wh*-misidentification is the more 'basic' (1999: 286) and 'interesting' (1999: 272, 287) epistemic status. This inference seems suspect. Knowing that *p* entails that one has a justified belief that *p* but not vice versa (at least according to orthodoxy: for a defence see McGlynn 2014: 39-42), but no one thinks that by itself this suggests that knowledge is more basic or interesting; indeed, as Williamson notes (2000: 3), there has been an tendency to think that the necessary but insufficient conditions on knowing are more basic and interesting. We will return to this point shortly.

Up to a point, Pryor's treatment of quasi-memory is simply an endorsement of Shoemaker's. According to Shoemaker, recall, memory-based judgments aren't really immune to misidentification, as the possibility of quasi-remembering events from another's past shows, but since this possibility isn't actually realised, one's judgments enjoy a kind of *de facto* immunity. The innovation in Pryor's treatment is that he takes these claims to be best understood as concerning immunity to *wh*-misidentification. Pryor advertises two advantages of this way

of recasting Shoemaker's position. First, as we've seen, he takes immunity to wh-misidentification to be the more basic and interesting variety of immunity. Second, he argues that some of Evans's objections to Shoemaker's position at best show that the judgments in question are immune to de re misidentification; they don't have any power to show that they are not vulnerable to wh-identification. So, 'with respect to the most interesting sort of immunity, Shoemaker is right: our first-person memory-based beliefs do *not* have that sort of immunity' (1999: 272).

An immediate worry with Pryor's Shoemaker-inspired position is that, as we have already seen, Pryor's claim that immunity to wh-misidentification is more basic and interesting has not been well motivated. However, it's not clear that this is really a problem. Coliva (2006: 408-9) has argued that Pryor has mis-advertised his own position as a vindication of Shoemaker over Evans, when in fact it offers an even-handed resolution of the conflict between them. On the envisaged resolution, Evans is right concerning immunity to de re misidentification, while Shoemaker is right concerning immunity to wh-misidentification. This appropriation of Pryor's distinction avoids the worry that he fails to motivate treating immunity to wh-misidentification as more basic; indeed, given the revised proposal's aspirations to be even-handed, it's an advantage that neither side of the distinction has proved more fundamental or interesting.

However, Coliva also contends that we cannot give a satisfactory response to the puzzle by appealing to Pryor's distinction, offering two objections. First, she suggests that the claim that memory-based judgments about one's own past are

vulnerable to wh-misidentification rests on a commitment to the grounds underwriting such judgments being ‘identity-neutral’ (2006: 411): that is, neutral on whose past is being represented.¹⁰ However, she takes one of Evans’s basic insights to be that this commitment is implausible:

So the proposal in effect merely invites a restatement of Evans’s original objection—that it simply distorts the epistemology involved to suppose that memory provides such identity-neutral grounds. (2006: 411)

Coliva’s second, ‘deeper’ worry is that Pryor’s distinction between the two varieties of misidentification isn’t as deep or as significant as he supposes. Consider the skunk example again. As Coliva reads Pryor, this is a case of wh-misidentification because ‘it involves identification of one object rather than another as the instance that verifies the existential quantification for which the thinker has independent grounds’ (Coliva 2006: 411-2). But what is one’s justification for that identification? Pryor doesn’t say, not does his template for wh-misidentification place any constraints on this. Now, Coliva contends that ‘the rational grounds underlying the identification’ will look like this:

I smell a skunky odor and this gives me grounds for:

(1) The animal (in my garden) which is actually responsible for this odor I can smell is a skunk.

¹⁰ It’s important not to confuse the claim that one’s grounds are identity-neutral with the claim that they are identification-free. One’s grounds are identification-free just in case no identity claim figures in them. Such grounds will *not* be identity-neutral if, for example, they are first personal in the way that we have been supposing the grounds offered by memory to be. Thanks to Ted Parent for prompting this clarification.

I go closer to one of the animals and sniff, while I can still smell the skunky odor. I then judge “This animal is a skunk’ *because* I believe:

(2) This animal (I can now see) = the animal (in my garden) which is actually responsible for this odor I can smell. (2006: 412)

On this reconstruction, the judgment that [this animal (I can now see) is a skunk] rests on grounds with a similar structure to paradigm cases of *de re* misidentification. Since in Pryor’s example the animal is not a skunk, that judgment is in error, and this is an error through misidentification; the mistaken judgment is justified by a false identity claim, namely (2). The crucial difference between Pryor’s example and examples of *de re* misidentification doesn’t lie in the structure of the grounds underwriting the judgments (any differences in structure being, on reflection, superficial) but rather in the kinds of *concepts* that figure in the identity claim figuring in those grounds:

An identification component may involve either *de re* concepts—that is, concepts which are grounded in identifying knowledge of their unique instances, like the concepts “That person” and “Aunt Miriam” in the recognition-based example I started with. Or else, it may involve non-*de re* but still singular concepts (hereafter, just “singular concepts”) that are not grounded in identifying knowledge of their unique instances, like “The animal (in my garden) which is actually responsible for this skunky odor I can smell.” (2006: 413)

So wh-misidentification is distinguished from *de re* misidentification only by the relative complexity of its structure, and by the kinds of 'conceptual ingredients'

that can figure in the mistaken identity claims; in the skunk case as Coliva reconstructs it, for example, the concepts that figure in the identity claim (2) are singular rather than *de re*. In this sense the notion of *wh*-misidentification is 'spurious' (2006: 414).

Unfortunately, neither objection looks entirely persuasive. Pryor had already offered a reply to the first kind of worry Coliva expresses in his paper. There he argues that his proposed treatment of the problem cases does *not* commit him to holding that one's grounds are identity-neutral. His proposal is that 'memory gives one a "package deal" of justification: it tells me that *someone* performed a certain action, and it also tells me that *I* am the person who performed that action' (1999: 296). He continues:

So long as these two pieces of justification always arrive as a package, Evans will be right to say that memory is not *silent* about whose past life it is reporting. But so long as these two pieces of justification are also, in principle, separable—so long as one can *undercut* the one piece of justification without undercutting the other—then it is possible for one's apparent memories to offer one knowledge that someone performed a certain action, even when one did not perform that action oneself. (1999: 296-7, emphasis in original)

Pryor doesn't offer much by way of positive motivation for this 'package deal' account (though he does offer a suggestive analogy to visual experience: see 1999: 297), but his main point is just that nothing in Evans's claim that one's grounds are not identity-neutral rules it out. Coliva's first objection merely invites a restatement of this reply to Evans's original objection.

What about Coliva's second worry with Pryor's position, namely her contention that *wh*-misidentification is spurious? Wright has observed that even if Coliva's proposed reconstruction of the grounds available to the subject in the skunk example works well enough, the recipe suggested by that reconstruction seems to misfire in other cases of *wh*-misidentification:

Another case: I am lost in sandy desert and, attempting to walk out, come across footprints which I misidentify as my own, concluding somewhat desperately, "I am going round in circles." Here the footprints give me reason to think that someone (maybe with feet about my size) has passed this way already; and I then misidentify—mistake myself for—the witness of that true existential claim. (Wright 2012: 256)

How might we reconstruct the grounds in this case, in line with the model suggested by Coliva's treatment of the skunk example? The first step is to identify the analogue of (1), the parent judgment that, together with my belief in a mistaken identity claim, leads me to mistakenly judge that I am the witness for the existential claim that someone has passed this way already. Coliva's recipe, as Wright reads her, involves identifying a singular concept by 'compos[ing] a definite description on the model: the unique object that caused—or is perhaps in some other way suitably related to—the thinker's grounds, *G*, for taking it that something is *F* (that the relevant existential is true)' (Wright 2012: 258). In the skunk example, my smelling the skunky odor gives me grounds for the existential claim that there is a skunk in my garden. The associated definite description is 'the animal (in my garden) which is actually responsible for this odor I can smell', and this does seem to fit the template Wright attributes to

Coliva. In the second example, my grounds for the existential claim that someone has passed this way already are the footprints, and so the definite description will be 'the person who caused these footprints in the sand'. Filling out the rest in the obvious way, we reach the following reconstruction:

(1') The person who caused these footprints in the sand has passed this way already.

(2') the person who caused these footprints in the sand = me

So:

(3') I have passed this way already.

However, Wright maintains that this reconstruction of my rational grounds in the example is implausible, since (1') is 'near enough, a tautology' (2012: 258). The thought is that Coliva's recipe for constructing a suitable definite description to figure in the parent judgment misfires when the predicate 'is, or approximates, the predicate: *the very cause of the grounds G*' (2012: 258). In Wright's example, that someone has already passed this way is the very cause of the footprints, and so constructing a definite description on Coliva's model, 'the unique object that caused the grounds G' leads to a reconstructed parent judgment that is tantamount to the claim that the unique object that caused grounds G is the very cause of grounds G. As Wright's wording suggests, such judgments need not be tautologies, but they are sufficiently wanting in

informational content to render the claim that they play a major role in the justification of judgments like (3') dubious.¹¹

Since Coliva's attempt to show that cases of *wh*-misidentification all involve essentially the same structure found in paradigm cases of *de re* misidentification looks to be in trouble, Wright concludes that *wh*-misidentification is 'a bona fide discrete phenomenon' (2012: 259), meriting further investigation. If that's right, Coliva's second objection to an appeal to Pryor's distinction is also problematic.

The real problem with appealing to Pryor's distinction to accommodate cases of quasi-memory is that it's not clear that his characterisation of vulnerability to *wh*-misidentification really delivers the result that memory-based judgments are vulnerable. If that's right, it's not clear how the proposal can be said to accommodate Shoemaker's key point, namely that it seems to be possible to make a mistaken judgement about one's own past because one has quasi-remembered an event in someone else's past, and that this seems to involve an error through misidentification.

The basic argument here is due to Evans.¹² Suppose again that I have an apparent memory of having been in that remote part of Scotland five years ago, and I form a judgment on that basis. Does the possibility that I might be quasi-remembering the past of another demonstrate that memory-based judgments

¹¹ Wright (2012: 258-9) distinguishes a second class of examples that create trouble for Coliva's treatment of *wh*-misidentification, but I lack space to discuss those here.

¹² I have tweaked Evans's argument since (as already noted) Pryor is explicit that vulnerability to *wh*-misidentification is to be characterized in terms of undercutting defeaters (1999: 284), and Evans's example makes a play with overriding defeaters instead (1982: 244-5). Strangely, Pryor spends two pages (1999: 294-5) discussing Evans's example without commenting on this. For the record, I don't share Pryor's 'intuition' (1999: 245) that in a case in which I have an apparent memory of being *F* but learn that I was not *F*, my grounds could still offer me knowledge that someone was *F*.

are vulnerable to wh-misidentification? Given Pryor's characterisation, to answer this question we need to consider whether it's possible for my grounds to be defeated by undercutting evidence in such a way that those grounds together with the defeating evidence no longer justify me in believing that I was in that remote part of Scotland five years ago, but could nonetheless offer me knowledge that someone was in that remote part of Scotland five years ago. Now, Evans's point is that if I were to fallback on the existential generalisation in such circumstances, this could be nothing more than a guess, and so I could not be said to have knowledge or justification (1982: 244-45). Once I can no longer take my apparent memory to be evidence concerning my past, I cannot take it to be evidence for the existential claim either.

On the face of it, this looks compelling. However, Pryor thinks that we can easily describe undercutting evidence that would defeat my judgment that I was in that remote part of Scotland five years ago while leaving those grounds capable of offering me knowledge of the existential generalisation (1999: 295-6). Suppose that a trustworthy source tells me that i. some of my apparent memories are due to my quasi-remembering someone else's past and ii. none of my apparent memories of having been in various locations in Scotland derive from events in my own past. This is an undercutting defeater for my judgment, since it doesn't give me evidence that I *wasn't* in that remote part of Scotland five years ago; it just undermines the evidential force of my apparent memory of being there. Moreover, it leaves open the possibility that my apparent memory of being in that remote part of Scotland five years ago is due to a memory illusion rather than due to quasi-remembering someone else's past.

The problem with this is one that Pryor gestures at when contending that it is undercutting rather than overriding defeaters that we need to appeal to in characterising immunity and vulnerability to wh-misidentification. In a parenthetical remark, he writes (1999: 284):

If you have evidence for believing p , and God appears to you and says, “ p isn’t true, but q is,” his testimony might defeat your original evidence and give you reason to believe q , but we don’t want to conclude that part of your original evidence therefore already justified you in believing q .

The moral of the example is that one’s original evidence together with a defeater may offer one knowledge of a proposition, not because a relevant part of one’s original evidence is unaffected by the defeater, but because the defeater itself offers one knowledge of that proposition.¹³ And the worry about the defeater Pryor describes for my memory-based judgment that I was in that remote part of Scotland five years ago is that it falls in the wrong side of this distinction. It seems undeniable that given my evidence plus the undercutting defeater as Pryor describes it, I would be justified in judging that someone was in that remote part of Scotland five years ago. However, the natural explanation for this is that the undercutting evidence includes the information that some of my apparent memories derive from events in another’s past. It’s true, as noted above, that this information leaves open the possibility that my memory impression of having been in that remote part of Scotland five years ago is just a

¹³ It should be stressed that it’s not clear that this is a moral that Pryor himself wants to draw. The placement of the quoted remark suggests that Pryor takes the example to motivate framing his characterization of vulnerability to wh-misidentification in terms of undercutting defeaters rather than overriding defeaters. If so, this strikes me as curious, since the point seems to stand even if we imagine that God says “Your grounds for p are worthless, but q is true”; this too seems to ‘defeat your original evidence and give you reason to believe q ’ without tempting us to suppose that your original evidence already provided justification for q .

memory illusion. Still, in light of my background knowledge that such illusions are relatively rare, it seems to be perfectly natural and reasonable for me to conclude that the reason I have a memory impression of being in that remote part of Scotland five years ago which does not derive from my own past is that this is one of my memory impressions which derives from someone else's past. If we suppose that I lack the relevant background knowledge—that for all I know, vivid memory illusions are in fact very commonplace—then I have considerably less inclination to think that my grounds justify me in the fallback existential.¹⁴

So I'm inclined to think that Evans's point is a good one; once my judgment that I was in that remote part of Scotland five years ago have been undercut, my original grounds do not justify the judgment or offer knowledge that someone was in that remote part of Scotland five years ago. The undercutting evidence that Pryor describes doesn't successfully undermine this point, since I have argued that he fails to show that it's really my original grounds that provide the justification I have for that existential after receiving the defeating evidence; rather, I contend, it's the combination of that defeating evidence and my background knowledge concerning the likelihood that I'm subject to a memory illusion.¹⁵ In light of Evan's objection, I find Pryor's contention that memory-based judgments meet his conditions for vulnerability to wh-misidentification implausible. Far from accommodating Shoemaker's insight, Pryor's proposal renders it elusive.

¹⁴ The point also undermines Smith's case against immunity to wh-misidentification (2006: 278-9). He in fact recognizes the worry (2006: 279-80), but his response simply takes for granted a number of disputed claims about the kinds of grounds that memory makes available and the kind of epistemic architecture characteristic of cases of wh-misidentification. I lack space to expand on this point here, but intend to do so elsewhere.

¹⁵ Pryor (1999: 290-7) extracts a further three arguments from Evans's discussion, but I'm happy to concede Pryor's responses to those. Cappelen and Dever (2013: 132-3) also object to Evans's argument discussed in the text, but the argument is very undeveloped.

5. Presuppositions

Consider an interesting variant Wright offers of our paradigm example of error through misidentification (which Wright sets in Washington Square Park):

Revisit the Washington Square Park scenario but this time let my thought that *Aunt Lillian is wearing an extraordinary hat today* be based on a close encounter, so that the judgment is based on directly observing my aunt and her extraordinary tither a few feet in front of me. Surely this is not, on any plausible construal, an *inferential* judgment – unless almost all routine observations are to count as inferential. But still, the judgment is not immune to error through misidentification. The person before me might be not my esteemed Aunt Lillian but her twin sister, whom I have never met or heard of, or some other look-alike. In that case my judgment would be wrong, and wrong as the result of a misidentification, but it is not inferential. (Wright 2012: 267)

Following Wright, let's refer to this as the 'close-by aunt' example. Since this is a non-inferential judgment, one's *grounds* will not contain the kind of complex structure needed for error through misidentification. Still, the judgment, made on those grounds, doesn't seem to be a plausible candidate to be regarded as immune to error through misidentification. This kind of case seems to refute Evans's claim that it suffices for a judgment to be immune to error through misidentification that it be based on grounds that are identification-free.

Here's one natural way to accommodate Wright's example. We give up on Evans's claim, and acknowledge the possibility that a judgment J justified by grounds G can be vulnerable to de re misidentification even if no identity claim is part of G because an identity claim is instead a *background presupposition* of J when justified by G. A presupposition of a judgment, made on particular grounds, is a proposition such that 'were it somehow to fall into question, a (rational and appropriately conceptually equipped) subject would be prepared to withdraw from the judgment' (Coliva 2006: 416). In the close-by aunt case, Wright's judgment's vulnerability to error through misidentification is plausibly due to it being a presupposition of that judgment, made on the grounds in question, that the woman he can see is his Aunt Lillian (Wright 2012: 270).¹⁶

We might make a similar proposal about memory-based judgments, such as [I was in that remote part of Scotland five years ago]. According to such a proposal, even though no identity claim typically figures in the grounds for such judgments, an identity—'I = the person whose past I seem to remember now' (Coliva 2006: 406)—figures as a presupposition of such judgments when they are based on the usual grounds, since doubt about that identity would rationally commit me to doubting my judgment that I was in that remote part of Scotland five years ago when made on the basis of memory.

Coliva (2006) suggests two ways that we might appeal to this proposal to offer an even-handed resolution of the disagreement between Evans and Shoemaker.

¹⁶ See also Pryor 1999: 291.

According to the first, there are two notions of error through misidentification, giving rise to two notions of immunity (2006: 420). On the one hand we have *error through misidentification relative to the subject's own rational grounds*, which involves making an erroneous judgment on the basis of grounds that include a mistaken identity claim. On the other, we have *error through misidentification relative to background presuppositions*, which involves making an erroneous judgment because a mistaken identity is a presupposition of that judgment when justified by the relevant grounds. A judgment made on particular grounds might be immune to either or both of these kinds of errors.

It's easy to see how this proposal is conciliatory. In effect, it characterizes a notion of immunity for which Evans's otherwise dubious claim that identification-free grounds is sufficient is (rather trivially) correct, and it places this on a more or less equal footing with a notion of immunity for which Evans's claim fails. Each party to the dispute—Evans and Shoemaker—comes out as correct about memory-based judgments about one's past on one of these understandings of immunity. Such judgments are immune to error through misidentification relative to the subject's own rational grounds, since one's grounds are first-personal and so identification-free. But they are vulnerable to error through misidentification relative to background presuppositions, as the possibility of quasi-remembering another's past shows.

So far, though, nothing has been done to say why any *significance* should be attached to the notion of immunity to error through misidentification relative to the subject's own grounds. Once we recognize that a judgment can rest on a false

identity claim without that identity figuring as part of its grounds, it's hard to see why we should care about a notion of immunity that's been designed to simply ignore that insight. It seems more principled to hold that Evans was simply wrong to assume that the grounds of a judgment being identification-free suffices for that judgment to be immune to error through misidentification. We can illustrate this worry with reference to the close-by aunt example. Wright's judgment *is* immune to error through misidentification relative to his own rational grounds, assuming that he is right to insist that it's noninferential. Nonetheless, the judgment doesn't seem to be immune to error through misidentification in any interesting sense, and in this it contrasts with memory-based judgments about one's own past.

There's more that could be said here, but in any case, Coliva prefers a somewhat different proposal according to which there is only one notion of immunity. Somewhat surprisingly, her suggestion is that this is immunity to error through misidentification relative to the subject's rational grounds, and so Evans's claim that memory-based judgments about one's own past are immune is vindicated. According to Coliva's preferred position, the fact that such judgments may have identity claims as part of their presuppositions doesn't immediately introduce vulnerability to error through misidentification. The significance of these identity presuppositions is rather that under some possible circumstances they shift from being presuppositions to being part of one's grounds. If, for example, I'm informed that there's a live possibility that my memory impressions in fact causally derive from someone else's past, then the identity $I = \text{the person whose past I seem to remember now}$ becomes something which I need evidence for if

my judgment that I was in that remote part of Scotland five years ago is to be justified by my apparent memories; the identity claim moves into the grounds of the judgment, and the judgment is not justified by the available grounds if I lack evidence for it.

Coliva now gives this a Shoemakerian twist. Given my actual grounds, which are identification-free, my memory-based judgment that [I was in that remote part of Scotland five year ago] is immune to error through misidentification, just as Evans would have it. However, 'this normal epistemic architecture is a contingency' (2006: 421); in the circumstances envisaged in the previous paragraph, for example, my grounds are no longer identification-free, and this introduces scope for error through misidentification.

Coliva goes on to draw the distinction between logical and *de facto* immunity in these terms. A judgment J based on by grounds G will be merely *de facto* immune when it has an a posteriori true identity amongst its background presuppositions, while it is logically immune if it either rests on no identity presuppositions at all, either in its grounds or background presuppositions, or if any identity claims that figure as presuppositions for it are recognizable a priori (2006: 423-4).¹⁷ Let us set aside judgments which don't rest on identity claims at all (such as, perhaps, demonstrative judgments like *that thing is moving really fast* (2006: 423-4)). Then the distinction concerns whether the truth of

¹⁷ Coliva in fact has 'contingent' where I have 'a posteriori' (2006: 423). The change I have made is in line with Coliva's own insistence that the logical/*de facto* distinction should be drawn in epistemological rather than metaphysical terms (2006: 423-4); moreover, in discussing whether psychological self-ascriptions are *de facto* or logically immune, she takes the crucial questions to be whether they 'they have identification components as part of their background presuppositions and whether, in such a case, it is an a priori true one' (2006: 424).

presupposed identity claims are recognizable a priori or not. The significance of this, according to Coliva, is that only in the latter case is there room for the kind of epistemically abnormal circumstances described in the previous paragraph: where a doubt is raised that has the effect of moving the identity out of the judgment's presuppositions and into its grounds, introducing vulnerability to error through misidentification.

Unfortunately, I don't think that we can rest content with Coliva's proposal. There are a number of problems. First, the reason it is meant to be significant that a presupposed identity claim is a priori is that this status means that it can 'not sensibly be questioned, or called into doubt' (2006: 424). However, a priori true propositions don't in general seem to be indubitable in anything like this sense.¹⁸ Perhaps this first worry can be avoided by characterizing logically immune judgments directly as those that either have no identity claims as presuppositions, or which have indubitable identity claims as presuppositions (compare Smith 2006: 276). Since there are deeper problems with Coliva's proposal, I won't consider how adequate this response to the first worry is.

The second problem is one that the proposal under consideration shares with Coliva's first, namely that it intuitively mishandles Wright's close-by aunt example. The present proposal entails that Wright's judgment is in fact immune to error through misidentification *simpliciter* (since it's immune relative to Wright's own rational grounds, and this is the only variety recognized); it would

¹⁸ This is based on a point that Robbie Williams raised in discussion.

only become vulnerable to error through misidentification were Wright's judgment to fall into doubt in the specified kind of way.

Related to this last point, the most important objection is that, like Pryor's proposal, Coliva's doesn't seem to really respect Shoemaker's insight. Here's how Coliva describes what she takes to be right about Shoemaker's claims:

Surely there is something right about Shoemaker's view. Once one accepts the notion of quasi-memory, it seems that I might have apparent memories as of, say, being in a particular remote part of Scotland five years ago, which in fact causally derive from someone else's past. And in such an event, according to Shoemaker, my judgment, "I was in that remote part of Scotland five years ago" would be affected by an error of misidentification. (Coliva 2006: 406)

Suppose that I judge that [I was in that remote part of Scotland five years ago], on the basis of quasi-remembering someone else's past. As Coliva affirms in the passage just quoted, my judgment seems to be in error through misidentification. However, if I've no reason to think that anything is amiss, her account doesn't actually deliver that result; there's no scope for error through misidentification until my judgment is *defeated* in the right kind of way, leading to a shift in the grounds I have for my judgment. Shoemaker's point here seems pretty clearly correct, as Coliva's own endorsement in the passage just quoted acknowledges, and so it's a significant cost of her account that it stands in tension with it. Of course, it's open to Coliva to respond that her proposal gets at the kernel of truth

in Shoemaker's claims about quasi-memory. However, that's only a good response if no viable proposal underwrites Shoemaker's claims as they stand.

6. Two Grades of Immunity

Recanati's recent work (2007, 2012) contains a sophisticated account of the nature of immunity to error through misidentification and its source, and his discussion contains some proposals that engage directly with the topic of this paper. These proposals will be the topic of this section. Recanati's account of immunity to error through misidentification is embedded within a larger account of *de se* thought, and perspectival thought more generally, and I can't hope to do justice to the details of his position here. Rather, I'll try to get just the relevant aspects of his view on the table, so that we can focus on his treatment of quasi-memory.

Recanati's central idea is that certain *de se* judgments will be immune to error through misidentification when they are based directly on experience, since the relevant experiences are intrinsically *de se*.¹⁹ However, Recanati has a very particular picture of what this involves. Such experiences have 'thetic'—subjectless—content; they are first-personal not in virtue of representing the experiencing subject, but rather because the 'mode' of the experience ensures that it must concern the subject of that experience. Recanati introduces the

¹⁹ Recanati now also recognizes a category of judgments that display 'derivative' immunity, having been 'based upon' judgments that are 'primitively' immune (see 2012: 183-4 and Wright 2012: 260-1 for discussion). This complication won't matter here.

notion of a 'mode' with an analogy to illocutionary force in speech act theory (2007: 125-9), but I find the following gloss from Wright more helpful:

The mode of an experience, in Recanati's understanding of the notion, is something along the axis: visual, auditory, tactual, olfactory, kinaesthetic, proprioceptive, etc. But more generally, a mode is a manner in which a content is present to or entertained by the mind, so that in addition to the modes of experience, states like belief, desire, intention, hope, etc., also count as modes. (2012: 262 fn18)

On Recanati's account, a *de se* judgment such that [my legs are crossed] does not have a classical proposition as its content, but rather a 'person-relative proposition', which can only be evaluated as true or false upon the provision of a subject. And the 'internal' mode—the proprioceptive and kinaesthetic mode (2007: 146)—can, as a matter of fact, only serve up information about oneself, and so a judgment based on information from that mode has to be evaluated as true or false with respect to oneself; however, the subject is not 'identified' in the content of either the judgment or the experience it is based on. That's why there's no scope here for error through misidentification.²⁰

As mentioned, there are several aspects of Recanati's discussion that speak to the topic of this paper. The first is that he draws a distinction between two 'grades' of immunity, and suggests in a footnote that this may correspond to what

²⁰ Recanati's revised position in his 2012 paper allows that a judgment may be immune to error through misidentification when its content explicitly represents the subject, so long as the experience it is based on has thetic content. This difference also won't matter here.

Shoemaker was after with his logical/*de facto* distinction (2007: 150fn61). The idea, applied to memory-based judgments about oneself, seems to be that in making such a judgment one implicitly attributes a certain internal state to oneself—having certain memory impressions, say—and one also implicitly ascribes to oneself the ‘world-involving’ property of having experienced the event represented by those memory impressions (2007: 152-3). Each of these kinds of self-ascription is associated with a grade of immunity; the self-ascription of the inner state has grade-1 immunity while the world-involving self-ascription has grade-2 immunity (2007: 152). Recanati suggests that the possibility of quasi-remembering events in another’s past (and likewise the possibilities of quasi-propriocepting the position of someone else’s body and quasi-perceiving someone else’s environment) shows that grade-1 immunity is ‘stronger’:

What happens in quasi-perception and quasi-memory can be described by saying that there is a *failure of grade-2 immunity*. Grade-1 immunity is retained: the subject is entitled to automatically self-ascribe the property of being in a certain perceptual/mnesic state, without any possibility of error. What the subject cannot rightly do, however, is automatically self-ascribe the property of standing in the right relation to the content of the state: for it may be that he does *not* possess the property of seeing a flower in front of him, or of having seen the remembered scene in the past. Someone else may be seeing, or someone else may have seen. (2007: 156, emphasis in original)

Contrary to Recanati's suggestion, the distinction between grade-1 and grade-2 immunity doesn't seem to map very cleanly on to the distinction between logical and *de facto* immunity. One way to bring this out is to appeal to demonstrative judgments; the judgment [that thing is coming straight for us] is standardly and plausibly classified as *logically* immune to error through misidentification, but it's not clear how (or indeed if) it fits into Recanati's scheme at all. Still, our principal concern here is not with drawing the distinction between logical and *de facto* immunity; that distinction is only of interest to the extent that it enables us to reconcile the immunity of memory-based judgments with the possibility of quasi-memory. Recanati's grades of immunity don't really help us to do this by themselves, but some other remarks he makes in the course of his discussion seem more promising.²¹

What quasi-remembering and quasi-proprioception show, according to Recanati, is that it's a contingency that the mode associated with memory can only deliver (first-personal) impressions about one's own past, and that the internal mode can only deliver information about one's own body (2007: 146). Why isn't this just to concede that the relevant judgments simply aren't immune to error through misidentification even when made in the relevant way? Recanati offers two suggestions. The first is that the error involved in the problem cases need not involve misidentification. My judgment still rests directly on an experience that is identification-free, as Evans says, but it also presupposes the 'normal functioning of the system' (Recanati 2007: 158). In one of the problem cases, in

²¹ Guillot (2014: 1807-8) suggests and criticises another way one might try to draw the logical/de facto distinction within Recanati's framework.

which I quasi-remembers another's past, my judgment is in error because that presupposition is false, and so I have the illusion of (say) having been in that remote part of Scotland five years ago. But there's no misidentification going on.

This is a little hard to stomach. The 'illusion' in question arises because I am receiving information about events in another's past as if I experienced them myself. In making judgments on that basis, I presuppose that I am the person whose past I seem to remember now, where, recall, presupposing this doesn't require me to entertain a thought with this identity claim as its content, let alone for my grounds to include this identity claim. In the problem cases, this identity presupposition is mistaken, and it seems like my judgment is in error because of this misidentification. It's not clear that Recanati has offered us any principled grounds on which to resist this simple and appealing Strawsonian line of reasoning to the conclusion that if I make false judgments about my own past having quasi-remembered the past of someone else, those judgments will be in error through misidentification. It's true that there may be other presuppositions that are false in the problem cases, but to say this much is not yet to show how to resist the conclusion that one's error is in part due to one's presupposing a mistaken identity claim.

Recanati's other thought is that our judgments about our own pasts are in fact based on memory rather than quasi-memory, and memory-based judgments of the relevant sort are immune to error through misidentification even if their possible counterparts based on quasi-memory are not:

Cases of wrong connection are ‘possible’ only in the sense that, in some possible world w , they would happen. We should clearly distinguish that possible world w , in which wrong connections happen, from the actual world, in which they do *not* happen and memory and perception *are* immune to error through misidentification. A case of ‘ordinary memory’, that is, a case in which the rememberer is the original experiencer, would be a special case of quasi-memory in w , but in our world it is not. One of the things that show this is precisely the fact that memory is immune to error through misidentification, while quasi-memory is not. (2007: 157, emphasis in original)

The final sentence of this passage is apt to strike as rather question-begging, given that it’s part of a reply to the worry that the possibility of quasi-remembering another’s past shows that memory-based judgment aren’t immune to error through misidentification after all. But the bigger problem here is that Recanati assumes that there is only a threat to the immunity of memory-based judgments if remembering turns out to be a subspecies of quasi-remembering in the actual world. However, as we saw in section 2 above, the crucial question is whether memory and quasi-memory can serve up the same kinds of grounds, and it’s not clear what Recanati’s view on this is. Either they make available different kinds of grounds, in which case Recanati’s proposal would just be a version of the externalist one that we considered above in section 2, or the grounds are the same, rendering it difficult to see the relevance of the point made in the passage just quoted.

So I don't think that Recanati's discussion helps us to reconcile the problem cases with the immunity of the relevant judgments. It needs emphasizing that nothing I've said in this section constitutes any kind of objection to Recanati's general framework or his treatment of de se thought and immunity to error through misidentification in particular. Moreover, I can't see any tension between Recanati's account and the proposal I will make in the next section. My claim for now is just that some such proposal is still needed.

7. Narrowing the Scope of Immunity Within De Se Thought

In light of the difficulties we have encountered, why not simply conclude that the possibility of quasi-remembering another's past shows that memory-based judgments *aren't* immune to error through misidentification (while resisting Cappelen and Dever's wholesale rejection)? This is the line suggested by Wright's most recent discussion (2012). Wright proposes that we take the close-by aunt case as our model for how to treat memory-based judgments and the like (2012: 272), along the lines discussed above in section 5, and so these will turn out to be cases in which one's judgment is vulnerable to error through misidentification, due to having an identity claim as a presupposition when made on the usual kind of grounds. The upshot is a dramatic 'narrowing of the scope of IEM I-thought' (2012: 271).

That certainly seems like a coherent view. However, it's liable to strike us as counterintuitive and defeatist, and it's hard not to feel like the proposed assimilation of the problem cases to the close-by aunt example is strained; it

feels like there's an important difference being overlooked. In the next section I'll try to articulate that difference, hopefully thereby showing that Wright's conclusion is indeed overly pessimistic.

8. Immunity and Safety

Our discussion of existing treatments of the possibility of quasi-remembering has suggested a number of morals or desiderata. First, we want to try to rescue Evans's claim that memory-based judgments about one's own past are immune to error through misidentification, while acknowledging the validity of Strawson's insight that cases of quasi-memory are metaphysically possible and that they involve an error through misidentification. Second, we want to do this while avoiding robbing the notion of immunity of all of its modal force. Third, we also need to avoid appealing to a contingency that involves a shift in one's grounds for the judgment in question. Fourth, if possible, we want to explain why memory-based judgments about one's own past and cases like Wright's close-by aunt seem to differ in respect of immunity to error through misidentification, despite the fact that both have identity claims as presuppositions when made on the usual kinds of grounds. I'll attempt to show that the proposal I will defend in this section will fare better with respect to these desiderata than the competitors that we have considered so far.

Pryor and Coliva both try to accommodate quasi-memory by refining our understanding of the kinds of *misidentification* that might be involved. My proposal is that we shelve that approach, and instead rethink the notion of

immunity to error. As we have seen, immunity to error through misidentification is standardly characterized in terms of the *impossibility* of being in error through misidentification when making the judgment in question on the relevant kind of grounds (e.g. Shoemaker 1968: 557, Pryor 1999: 279 and 283, Bar-On 2004: 57, Coliva 2006: 403, Recanati 2007: 146, Stanley 2011: 90, Bermúdez 2013: 212, and Cappelen and Dever 2013: 130). As our discussion to this point has confirmed, this makes it rather difficult to make room for anything like *de facto* immunity, except in the cheap sense involving a shift in the grounds for the judgment. My hope here is to table and to begin to motivate and defend an alternative conception of immunity to error that will enable us to fare better.

Let's start by considering what it is for a judgment to be *safely true*. A judgment J, based on grounds G, is safe in the relevant sense just in case there's no nearby world in which one forms J on the basis of G and yet that judgment fails to be true.²² Let us say that J is safely safely true when based on G just in case there's no nearby world in which there's a nearby world in which one forms J on the basis of G and yet J fails to be true. However, it is well known that safety does not iterate; one's judgment can be safely true without being safely safely true, and each further iteration is non-trivial, since a world y can be nearby to world x, and x nearby to w, without y being nearby to x—that's to say, the relevant relation between worlds *of being nearby to* is not transitive (see e.g. Sainsbury 1997 and Williamson 2000).

²² As is well-known, refinement is required to stop judgments of necessary truths and the like from trivially counting as safely true. Such complications won't matter here.

My suggestion is that we think of immunity to error as *sufficiently iterated safety from error*. That is, for a judgment J based on grounds G to enjoy immunity to error is for there to be no nearby world in which there's a nearby world in which there's a nearby world in which...one forms J on the basis of G, and J is in error. The worlds in which one forms J on the basis of G and J is in error are *very* modally remote. For J based on G to enjoy immunity to error *through misidentification*—to stress, whether of the de re or the wh- variety—is for there to be no nearby world in which there's a nearby world in which there's a nearby world in which...one forms J on the basis of G, and J is in error through misidentification.

An immediate red flag is raised by my use of both the weasel phrase '*sufficiently iterated*' and the associated 'magic dots'.²³ Just how many iterations are needed? It's hard to envisage a principled and satisfactory answer to that question being forthcoming, and so we'll need to decide if this is an acceptable kind and level of indeterminacy. I'll return to this important worry shortly; for now, let's continue to lay out the proposal.

Consider again judgments about one's own past, made on the basis of memory. Quasi-remembering another's past isn't merely something that doesn't happen in the actual world; rather, it's a modally remote possibility. Given that the possibility of such judgments being in error through misidentification when made on the kinds of grounds in question is so remote, such judgments are safely

²³ I have borrowed the phrase 'the magic dots' from Weir 1998: 768, though Weir uses it in a rather different context.

safely safely safely... free from error through misidentification. By the lights of the proposal just made, that's what it is for such judgments to be immune to such error. Nonetheless, we have clearly left room for the possibility that one might make such a judgment on the very same grounds that would *actually* secure immunity, and yet that judgment be in error through misidentification. In a world in which I quasi-remember another's past my judgment that [I was in that remote part of Scotland five years ago] may well be in error through misidentification, and my judgment will already be vulnerable to such errors if worlds in which I quasi-remember events from the lives of others are too close, even if I remain completely unaware of this possibility.

We can now offer a characterization of the distinction between logical and non-logical immunity to error through misidentification. For a judgment to be logically immune when based on certain grounds is for there to be no possible world in which one bases one's judgment on those grounds and that judgment is subject to error through misidentification; this is, of course, just to appropriate the standard characterization of immunity. For a judgment to be non-logically immune when based on certain grounds is for it to be immune to error through misidentification, in the sense that one's judgment is safely, safely, safely, safely... safe from such error, and for there be some possible world in which one bases one's judgment on those grounds and that judgment *is* subject to error through misidentification. I won't call the latter *de facto*, since there's an important sense in which non-logical immunity is never merely *de facto* on this account. For the sake of having a label, let us call this *contingent* immunity to error through misidentification.

This is, as promised, a proposal that seems to do well with respect to the items on our list of desiderata. We can endorse Evans's claim that memory-based judgments about one's own past are immune to error through misidentification. And we can do so while both preserving much of the modal force naturally associated with the notion of immunity and accommodating the metaphysical possibility that one might fall into error through misidentification while keeping fixed the grounds on which the judgments in question are based. The fourth desideratum is a little trickier. How does the proposal explain the difference between memory-based judgments about oneself and the close-by aunt case?

In fact, consideration of Wright's example raises a complication for my proposal, which we will need to face squarely. In the previous section I suggested that Wright is correct that his example involves a judgment that is not immune to error through misidentification given its grounds. However, the view I've adopted in this section doesn't let me endorse that conclusion without qualification.

From our current perspective, Wright's example is rather underspecified. Whether or not Wright's judgment that his Aunt Lillian is wearing an extraordinary hat is immune to error through misidentification, given my proposal, depends on the modal proximity of worlds in which I might misidentify someone else as Aunt Lillian at such close quarters. So I can't hold that the judgment, based on the grounds as Wright describes them, is vulnerable to error through misidentification *tout court*. However, given natural assumptions about

the relevant modal facts, Wright's judgment is vulnerable partly in virtue of the fact that some of the ways in which Wright's judgment could be in error through misidentification aren't sufficiently modally remote. Wright mentions the possibilities that his aunt has an identical twin that has been kept secret from him and that she has a lookalike (2012: 267). Now, it may be that neither possibility is all that nearby; both possibilities will normally be sufficiently remote for Wright's judgment to count as safe, and so perhaps to count as a good candidate for an item of knowledge. But compared with the kind of science-fiction scenarios we need to look at in order to find cases of error through misidentification for memory-based judgments about one's past, such possibilities are relatively close. So on my account the difference between Wright's case, as most naturally fleshed out, and memory-based *de se* judgments is one of degree not kind, and it lies in the proximity of worlds in which one's judgment is in error through misidentification when made on the grounds in question.

The general picture is one on which whether a judgment is immune to error through misidentification depends not just on its content and on its underlying epistemic architecture, but also on certain modal facts. It's this that allows the proposal to comfortably accommodate contingent immunity to error through misidentification, but it also leads to a number of worries. Here I'll only have space to discuss the two that are most pressing. The principal one has already been flagged, and my proposed treatment of the examples will have thrown it into sharp relief again. I've suggested that my proposal vindicates the intuitively correct verdicts that Wright's close-by aunt example is one in which immunity is

absent (given a natural conception of the relevant modal facts), while memory-based judgments are immune. However, given the degree of indeterminacy in the proposal there is ample room to wonder whether the proposal in fact yields these results.

Here, we're running into the familiar phenomenon of vagueness, of a sort that is ubiquitous throughout epistemology (as well as pretty much everywhere else, of course). On a familiar fallibilist conception of epistemic justification, for a proposition to be justified by one's evidence is for it to be *sufficiently* likely to be true, given that evidence. Closer to our present concerns, whether or not a judgment is safely true turns on what's going on at *sufficiently* close possible worlds, and yet it's hard to say anything informative about what counts as sufficiently close.²⁴ Now, immunity, as I've suggested we conceive of it, inherits any vagueness introduced by safety, and it opens up a further dimension of indeterminacy, concerning how many iterations of safety are required. The question is whether this kind and degree of vagueness is tolerable, or whether it undermines any attempts to draw particular conclusions from the account of immunity to error through misidentification that I have offered. I concede that it's not obvious that pessimism is unwarranted here. But it's certainly not obvious that it *is* warranted, and that immunity, so conceived, is a less serviceable notion than the other vague epistemic notions mentioned. As for the particular verdicts about particular cases I have discussed above, I hope to have

²⁴ As Ted Parent has reminded me, many philosophers are happy to live with the vagueness in closeness when developing or endorsing theories of counterfactuals. To draw this comparison is not to commit to the idea that we should identify the notion of closeness appealed to in the characterization of safety with that appealed to in standard account of counterfactuals; see Hawthorne 2004: 56fn17 for critical discussion of that identification.

made those plausible, but I concede that I have at present no more principled basis I could appeal to in order to bring around a dissenter. I'm not sure that makes us any worse off than we are in other debates in epistemology, and in philosophy more generally, but the matter certainly deserves further investigation.

Finally, it might be objected that I've simply changed the subject: that the long list I provided of philosophers who understand immunity in terms of impossibility shows not that the entire literature has been overlooking a live theoretical option, but rather that that's just how immunity *should* be understood. I'm not sure what would justify this attitude, but in any case, for some purposes it wouldn't matter if I were forced to concede the point. In particular, it wouldn't diminish a claim to have offered a response to Cappelen and Dever's contention that *no* epistemically interesting phenomenon has been identified that can withstand the pressure exerted by the problem cases. Whether any such phenomenon deserves the label 'immunity to error through misidentification' seems at best a matter of secondary importance.

Conclusion

I have argued that possibilities such as quasi-remembering the past of another need not shrink or empty the class of *de se* judgments that display immunity to error through misidentification. After an extensive critical discussion of existing proposals in the same spirit, my own suggestion has been that we should rethink the notion of immunity to error, understanding it as sufficiently iterated safety

from error rather than as the impossibility of error. On this characterization, it's built in from the start that immunity can be merely contingent; it is logical immunity that emerges as a special case. I have conceded that the proposal introduces vagueness that goes beyond that already present in the notion of safety, and that there is a need for further specification, or further reassurance that the vagueness is of a sort and degree that can be tolerated, though I haven't attempt to offer either here. I hope to have said enough to have illuminated the nature of immunity to error through misidentification, and to have defended it from the kind of scepticism about the phenomenon we find in Cappelen and Dever.

References

Anscombe, G. E. M. 1975. The First Person. In Samuel Guttenplan, ed. *Mind and Language: Wolfson College Lectures 1974*. Oxford: Oxford University Press.

Bar-On, Dorit. 2004. *Speaking My Mind: Expression and Self-Knowledge*. Oxford: Oxford University Press.

Bermúdez, José Luis. 2013. Immunity to Error Through Misidentification and Past-Tense Memory Judgements. *Analysis* 73: 211-20.

Cappelen, Herman and Dever, Josh. 2013. *The Inessential Indexical: On the Philosophical Insignificance of Perspective and the First Person*. Oxford: Oxford University Press.

Coliva, Annalisa. 2006. Error through Misidentification: Some Varieties. *Journal of Philosophy* 103: 403-25.

Coliva, Annalisa. 2012. Which 'key to all mythologies' about the self? A note on where illusions of transcendence come from and how to resist them. In Simon Prosser and François Recanati, eds. *Immunity to Error Through Misidentification: New Essays*. Cambridge: Cambridge University Press.

Evans, Gareth. 1982. *The Varieties of Reference*. Oxford: Oxford University Press.

Guillot, Marie. 2013. The Limits of Selflessness: Semantic Relativism and the Epistemology of De Se Thoughts. *Synthese* 190: 1793-1816.

Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.

Hamilton, Andy. 2013. *The Self in Question*. Basingstoke: Palgrave Macmillan.

McGlynn, Aidan. 2011. Review of Anthony Hatzimoysis, ed., *Self-Knowledge*, Oxford University Press. *Notre Dame Philosophical Reviews*.

McGlynn, Aidan. 2014. *Knowledge First?* Basingstoke: Palgrave Macmillan.

Morgan, Daniel. 2012. Immunity to Error Through Misidentification: What Does it Tell Us About the *De Se*? In Simon Prosser and François Recanati, eds. *Immunity to Error Through Misidentification: New Essays*. Cambridge: Cambridge University Press.

O'Brien, Lucy. 2007. *Self-Knowing Agents*. Oxford: Oxford University Press.

Perry, John. 1979. The Problem of the Essential Indexical. *Noûs* 13: 3-21.

Pryor, James. 1999. Immunity to Error through Misidentification. *Philosophical Topics* 26: 271-304.

Recanati, François. 2007. *Perspectival Thought: A Plea for (Moderate) Relativism*. Oxford: Oxford University Press.

Recanati, François. 2012. Immunity to Error Through Misidentification: What It is and Where It Comes From. In Simon Prosser and François Recanati, eds. *Immunity to Error through Misidentification: New Essays*. Cambridge: Cambridge University Press.

Roache, Rebecca. 2006. A Defence of Quasi-Memory. *Philosophy* 81: 323-55.

Sainsbury, R. M. 1997. Easy Possibilities. *Philosophy and Phenomenological Research* 57: 907-19.

Sainsbury, R. M., and Tye, Michael. 2012. *Seven Puzzle of Thought and How to Solve Them: An Originalist Theory of Concepts*. Oxford: Oxford University Press.

Shoemaker, Sydney. 1968. Self-Reference and Self-Awareness. *Journal of Philosophy* 65: 555-67.

Shoemaker, Sydney. 1970. Persons and Their Pasts. *American Philosophical Quarterly* 7: 269-85.

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

Stanley, Jason. 2011. *Know How*. Oxford: Oxford University Press.

Weir, Alan. 1998. Naïve Set Theory is Innocent! *Mind* 107: 763-98.

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

Wright, Crispin. 2012. Reflections on François Recanati's 'Immunity to Error Through Misidentification: What It is and Where It Comes From'. In Simon Prosser and François Recanati, eds. *Immunity to Error through Misidentification: New Essays*. Cambridge: Cambridge University Press.