

INTROSPECTION AS A (LIMITING) CASE OF PERCEPTION

Giovanni Merlo

University of Geneva

© 2025, Giovanni Merlo

*This work is licensed under a Creative Commons
Attribution-NonCommercial-NoDerivatives 4.0 License*

doi.org/10.3998/phimp.4539

1. The puzzle

It is a beautiful summer day, and you are lying on the beach, looking at the blue sky above you. As your mind drifts lazily from one thought to another, you are aware of two simple facts, even if you may not be actively attending to them. On the one hand, looking at the sky, you know that:

(1) The sky is clear

This you know by perception: you're *looking* at the sky, and you see that it is clear. But just as much as you know that the sky is clear, you also know *that you see* that the sky is clear. Or, as you would put it:

(2) I see that the sky is clear

You know (1) by looking at the sky. But how do you know (2)? What kind of faculty, method, or way of knowing things allows you to gain knowledge of the fact that you see that the sky is clear?¹

Here lies a puzzle. There are only two possible ways of explaining your knowledge of (2): either you know (2) in the same way you know (1)—that is to say, by perception—or you know (2) in some other way. But even if these are the only options, neither seems the right one.

On the one side, there are good reasons to think that you do *not* know (2) by perception. For one thing, it does not ring true to say that, as you're lying on the beach looking at the sky, you 'see' that you see that the sky is clear. It is not as if you were looking at a CCTV camera tape of yourself looking at a visibly clear sky. Indeed, the more one thinks about this situation, the clearer it becomes that no perceptual verb is appropriate to describing your knowledge of (2) in the situation we're imagining—you don't 'hear', 'smell', 'taste', or otherwise 'feel' that you see that the sky is clear. One could say that you 'introspect'

1. Here and in what follows, when I say that you know (2), I mean that you have knowledge of the fact that (2) expresses in the relevant context, and that you do so under the particular mode of presentation that (2) encodes. Similarly for (1) and other sentences.

this—but, as Ryle famously pointed out, ‘introspection’ is a term of art (1949, 152); if we use it, we should be able to provide an account of what faculty, method, or way of knowing it stands for, and the need to provide such an account will simply bring us back to square one.

But the idea that you know (2) by perception can also be criticized on different, more principled grounds. Perception can be hindered by all sorts of occlusions and malfunctions. For one to be able to see that the sky is clear it is not enough that the sky actually be clear. It is not even enough that one directs one’s eyes towards the sky. One’s eyes must be working properly, and there must be no objects obstructing the view. But if you see that the sky is clear, you are almost ‘automatically’ in a position to know that you see that the sky is clear.² There is no additional question, here, of your sensory organs working properly, nor can any object—inner or outer—prevents you from acquiring the relevant bit of knowledge. Describing your knowledge of (2) as perceptual seems to fly in the face of the glaring asymmetry between the two cases.

And yet, on the other side, there are also good reasons to think that you *do* know (2) by perception. Your knowledge of (2) is a piece of a posteriori, non-inferential knowledge of a certain type of mind-independent state of affairs.³ It is not obvious how we may categorize knowledge of this sort if not as perceptual—especially when, in addition, phenomenology seems so closely involved in the process of acquiring it. Discussing people’s knowledge of their own beliefs, Moran notes ‘there is nothing it is like to have the belief that Wagner died happy or to be introspectively aware that this is one’s belief’ (Moran 2001, 14). But when you knowingly see that the sky is clear, there most certainly *is* something it is like for you to be in the mental state that

you are in, and this is—at least, on the face of it—an important part of the reason why you know that you are.

And then, again, the idea that you know (2) by perception is supported by other, more principled considerations. Note that, insofar as you believe (2), your belief in (2) is going to support your belief in (1). If you believe that you *see* that the sky is clear, that is going to give you a very good reason to believe that the sky *is* clear. But note also that your belief in (1) is already supported by experience. After all, we are assuming that you know (1) by perception. Now, suppose we were to deny that your knowledge of (2) is also perceptual. Then it seems we would be obliged to see your belief in (1) as epistemically overdetermined. For you would then be justified to hold that belief ‘twice over’: on the basis of experience and on the basis of your non-perceptual belief that you see that the sky is clear. And since (as we’ve seen earlier) whenever one sees that the sky is clear, one is almost ‘automatically’ in a position to know that one does, this kind of epistemic overdetermination would be quite widespread: pretty much anyone seeing that the sky is clear would be justified to hold the belief that the sky is clear ‘twice over’. Surely, that is a very strange result. But it is a result that we seem to be cornered into accepting once we grant that the source of one belief (i.e., perception) is not also the source of the other.⁴

And so, we have a puzzle. Either you know (2) by perception, or you don’t. But we have good reasons to think that you do and equally good reasons to think that you don’t. What should we do?

2. If Williamson’s (1996) is right about luminosity, conditionals of this kind call for qualification. See Conee (2005), DeRose (2002), Hawthorne (2005), and Merlo (forthcoming) for qualified formulations compatible with the anti-luminosity argument.

3. One may, of course, call these features into question. My claim, here, is only that *on the face of it* your knowledge of (2) exhibits these features.

4. Byrne notes that the result is more than strange. For example, ‘when one sees a red spot and believes both that this spot is red and that one sees this red spot, it is not a possibility that two spots are in play’—yet the identity of the spots *would* be in question if the beliefs had different sources (Byrne 2018, 132-3). I agree with this observation, though—as will become clear in due course—the lesson Byrne draws from the no-epistemic overdetermination argument is different from mine.

2. A conservative approach

Our puzzle is a particular instance of a more general problem. People ordinarily know many of their mental states, and they mostly do this despite the absence of CCTV cameras, mirrors, external witnesses, and other special props. The *Epistemological Problem of Self-knowledge* is the problem of providing a satisfactory account of knowledge of this 'introspective' kind.⁵ Now, many philosophers think that, in dealing with this problem, we shouldn't be guided by any superficial similarity between introspection and perception. Shoemaker, for example, has vigorously advocated the view that 'the relation of an "introspective" belief to the state of affairs it is about is altogether different from the relation of a perceptual belief to the state of affairs it is about' (1994, 289). But others are of the opposite opinion. For instance, both in *Mind and World* and at various places throughout his writings, McDowell has repeatedly suggested that 'we should [...] think about "inner sense" in parallel with "outer sense" to the fullest extent that is possible' (1994, 21-2).⁶ Which one is it, then? Should we think of introspection as 'altogether different' from perception, or should we rather think of it in parallel with perception 'to the fullest extent that is possible'?

Let us return to our puzzle. I propose we take this case—an ordinary case of self-knowledge of what one sees—as the starting point of our inquiry and see just how much solving (or trying to solve) this case may teach us about the general problem it exemplifies. The approach I want to explore is a 'conservative' one: since the puzzle already presents us with good reasons for treating your knowledge that you see that the sky is clear as perceptual, I suggest we concentrate our efforts on defusing the reasons against the perceptual option rather trying to come up with an alternative non-perceptual interpretation. One significant benefit of adopting this approach is that, by treating your

self-knowledge as perceptual, we will turn the problem of explaining it into a special case of a problem that philosophers need to tackle anyway—that of explaining perceptual knowledge in general.

The plan is the following. I will begin by describing a familiar phenomenon involving beliefs acquired through perception—a phenomenon that I shall call *perceptual free-riding* (§3). I will then suggest that your knowledge of the fact that you see that the sky is clear can be plausibly construed as involving exactly this phenomenon (§4). The same approach applies to many other instances of self-knowledge, though not all of them, as I will explain in due course (§5). One kind of self-knowledge to which I would like to be able to apply the perceptual free-riding approach is self-knowledge of pain and other sensations. This application raises a special difficulty, having to do with the fact that our introspective awareness of pain and other sensations does not seem to leave any room for the kind of mismatch between appearance and reality that is typically made possible by perceptual awareness. It is in this connection that I will invoke the notion of a 'limiting case' which gives to this paper its title (§6).

The overall picture of self-knowledge I will be advocating is one on which—at least for the range of cases where the perceptual free-riding approach applies—no dedicated faculty, method, or way of knowing things is involved in our gaining knowledge of our own mental states. Instead, such knowledge turns out to be gained through the operation of whatever faculties, methods, and ways of knowing things are involved in our gaining perceptual knowledge of the outer world. In the last section (§7), I will highlight similarities and differences between this approach and others, of which it is, or may seem, a close relative.

3. Perceptual free-riding

I want to begin by considering another, apparently very different case. You are sitting at your kitchen table and lying in front of you are five eggs, carefully arranged, in a row, according to their colour. The first, leftmost egg is almost brown, and it is the darkest of the set. The last, rightmost egg is almost white, and it is the brightest one. The eggs in

5. For a classical presentation of the problem, see Wright (1998).

6. As it happens, McDowell is sceptical about the existence of a philosophical 'problem' in this area (see McDowell 1998, §3), but I take his position to be compatible with the minimal claim that articulating a satisfactory epistemological explanation of ordinary self-knowledge is not a trivial task.

the middle exhibit decreasing degrees of darkness, going from left to right.

As you contemplate the eggs, you are aware of two simple facts, even if you may not be actively attending to them. On the one hand, you know that:

- (3) Egg 1 is brown, and Egg 5 is white

This much you know by perception: you are looking at the eggs, and the first one to your left (Egg 1) visually strikes you as brown, while the last one to your right (Egg 5) visually strikes you as white. But just as much as you know that Egg 1 is brown and Egg 5 is white, you also know that:

- (4) Egg 1 is darker than Egg 5

Now, imagine someone were to ask how you know (4). Surely, we would like to be able to answer that your knowledge of (4) is *also*, in a broad sense, perceptual. Whenever you see one object as dark (say, brown) and another object as bright (say, white), you also—almost ‘automatically’—see the first object as darker than the other. It is not like you have to infer (4) from (3). Nor do you have any special faculty, method, or way of knowing things through which you know (4)-like facts. Rather, you know both kinds of facts by perception—and, interestingly, one bit of perceptual knowledge comes ‘for free’ relative to the other, given your understanding of what it is for one thing to be darker than another. It is this phenomenon that I am proposing to call ‘perceptual free-riding’.

One natural way of characterizing perceptual free-riding is in the context of the view that experiences justify belief. By ‘experiences’, here, I mean any conscious perceptual state and, by ‘justifying belief’, I mean making certain beliefs epistemically appropriate (possibly, only defeasibly so). On the view that experiences justify belief, when you see a red object in front of you, you are in a certain conscious perceptual state—an experience—which makes it epistemically appropriate for you to believe that there is a red object in front of you. In the gen-

eral case, whether or not you go on to form the belief in question is a separate matter (the kind of justification that the experience provides you with is, as this is usually put, ‘propositional’ rather than ‘doxastic’). The key point is that having the experience gives you an opportunity to form the relevant justified belief and, when all goes well, to acquire perceptual knowledge. This opportunity to acquire perceptual knowledge is what I shall hereafter call ‘perceptual justification’.

The thesis that experiences justify belief, though not uncontroversial, is in principle compatible with different views of the nature of conscious perceptual states. One view—the *straightforward view*, as one may dub it—holds that experiences are content-loaded states, and that one is perceptually justified to believe that *p* if and only if one has an experience that has the proposition that *p* as its content. This position implies that experiences have the same kind of content that beliefs have—i.e., conceptual, propositional content.⁷ Arguably, however, one can combine the thesis that experiences justify belief with other views—for instance, that experiences have conceptual content of a different sort than beliefs, that they have non-conceptual content, or even they have no content at all.⁸

I suggested that, insofar as we can characterize your knowledge of the fact that Egg 1 is darker than Egg 5 as perceptual, this knowledge comes ‘for free’ relative to your equally perceptual knowledge of the fact that Egg 1 is brown and Egg 5 is white. It is natural to see this phenomenon as a reflection of the fact that your justification to believe that Egg 1 is darker than Egg 5 comes ‘for free’ relative to your justification to believe that Egg 1 is brown and Egg 5 is white. And by this, I mean that:

[Free-riding eggs] Necessarily, if (i) you have a visual experience that justifies you believe that Egg 1 is brown and Egg 5 is white and (ii) you possess the concept ... *is darker than* ..., then you

7. A version of the straightforward view is defended by McDowell (1994) who later came to distance himself from it (cf. McDowell 2009, §§3-5).

8. See Silins (2024, §1).

also have perceptual justification to believe that Egg 1 is darker than Egg 5

On what I called the straightforward view, Free-riding eggs implies that, necessarily, anyone who, first, has an experience that justifies him or her to believe that Egg 1 is brown and Egg 5 is white and, second, possesses the concept ... *is darker than* ..., also has another experience, whose content is that Egg 1 is darker than Egg 5. The idea, here, is that possessing the concept ... *is darker than* ... allows one to have more experiences than one would otherwise.⁹ On other, less-than-straightforward views, Free-riding eggs doesn't have this implication. It is compatible with those views and Free-riding eggs to hold that possessing the concept ... *is darker than* ... merely allows one to, so to speak, extract more justifications from the experiences one has anyway.

My proposal is that your knowledge of (4) should be seen as a case of perceptual free-riding: because Free-riding eggs holds true and you satisfy (i) and (ii), some experience gives you propositional justification to believe that one egg is darker than the other. You then do whatever you normally do to turn that kind of justification into perceptual knowledge. The proposal assumes that experiences justify, but it is neutral among competing accounts of how they do that. And for what concerns the transformation of perceptual justification into perceptual knowledge, the idea is that the same story will work for (3) and (4) alike.¹⁰

An altogether different account of what goes on with (4) is possible. One may suggest that, instead of describing your knowledge that Egg 1 is darker than Egg 5 as perceptual, we should see it as inferential. Perception informs you about the colours of the eggs—brown and white. Then, knowing that brown things are darker than white things, you

go on to use that information to infer that Egg 1 is darker than Egg 5 (or, more neutrally, to form a belief that Egg 1 is darker than Egg 5 based on your belief about the colours of the eggs). But this alternative, inferential account is unconvincing, for several reasons.

First, it commits us to thinking of (3) as epistemically prior to (4), and it is not clear why we should accept this priority thesis. At first sight, your belief in (3) and your belief in (4) have equal claims to be treated as perceptual. It is plausible to think that (3) conceptually entails (4)—but it should certainly not be assumed without argument that two beliefs cannot both count as perceptual if the content of one conceptually entails the content of the other.

Second, the inferential account commits us to thinking that undermining the justification for (3) would leave the subject without any justification to believe (4). But, again, it is not clear that this would be so. Suppose you were told that the lighting conditions are unusual, so that the exact colours of the five eggs are not what you would take them to be based on your visual experience. Presumably, you would still be in a position justifiedly to judge that the colour of Egg 1 (whatever it is) is darker than the colour of Egg 5 (whatever it is). And this suggests that your justification for (4) is not based on your justification for (3), contrary to what the inferential account invites us to think.

Third, even setting these difficulties aside, it seems that not all purported cases of perceptual free-riding very easily lend themselves to inferentialist reinterpretations. Consider:

(5) Egg 1 and Egg 5 are spatially related to me

(5) is another fact you can plausibly be described as knowing as you contemplate the eggs in front of you. Indeed, it seems that knowledge of (5) can be obtained 'for free' relative to knowledge of (3) (or (4)) by anyone possessing the concept ... *are spatially related to me*. But, of course, (5) in no way follows—logically, conceptually, or otherwise—from (3) (or (4)). So, in this case, the suggestion that your knowledge is inferential is hopeless. And if it is hopeless here, why take it seriously with respect to your knowledge of (4), when your knowledge appears to be

9. At least, assuming that experiences do not have multiple contents.

10. What that story may be is, of course, a difficult question, but not one that needs to be settled in this paper. For a classical account of how perceptual justification gets turned into perceptual belief (and, when all goes well, knowledge) see Millar (1991, § 4.3).

of the same kind in the two cases?

Saying that your knowledge of (4) is perceptual rather than inferential is compatible with thinking that your knowledge of (4) is, in certain respects, less basic than your knowledge of (3) (though, to be clear, Free-riding_{eggs} has no such implication). One might, for instance, hold the view that, while both your belief in (3) and your belief in (4) are perceptually justified, your belief in (3) is more 'strictly' perceptual, because the proposition that one egg is darker than the other is not part of the content of any experience you are having, whereas the proposition about the colours of the eggs is. On this non-straightforward view, not everything we are perceptually justified to believe is something that we are perceptually justified to believe in virtue of having an experience with exactly that content.¹¹ It is important that even a view of this sort could in principle discriminate between bona fide perceptual beliefs (of the more or less strict sort) and beliefs that are only 'impurely' perceptual. Arguably, if you have a perceptual experience that justifies you to believe that Egg 1 is brown and Egg 5 is white, and you possess the relevant mathematical concepts, you are going to be justified to believe:

(3*) Egg 1 is brown, Egg 5 is white, and $2+2=4$

But, since the justification from (3*) comes partly from the experience and partly from your understanding of the relevant mathematical concepts, a belief formed on this basis is going to be only impurely perceptual. By contrast, a proponent of the view I have in mind would say that your justification for (4) comes entirely from your experience concerning the colours of the eggs. In this case—they would say—possession of the concept ... *is darker than* ... plays a merely enabling (rather than justificatory) role: it allows you to extract justification from the experi-

ence (rather than providing you with a separate piece of justification). So, the following position is perfectly coherent and compatible with Free-riding_{eggs} (even if not implied by it): your knowledge of (3) is more strictly perceptual than your knowledge of (4), but both pieces of knowledge are purely perceptual, unlike, e.g., your knowledge of (3*).

'Perceptual free-riding' is a technical label for what I take to be a relatively familiar phenomenon: sometimes, when we learn certain things by perception, there are other, additional things that we can learn 'for free', and still by perception, provided that we possess the requisite concepts. Explaining a piece of knowledge as an instance of perceptual free-riding is not a matter of explaining it in terms of a sui generis faculty, method, or way of knowing things. It is simply a matter of reminding ourselves that we can explain it as an instance of perceptual knowledge: knowledge acquired by exploiting, in the usual ways, a kind of justification that comes from perceptual experience. This means that the piece of justification that acts as a 'base' for perceptual free-riding is not at all an epistemic ground for the free-riding belief—it is merely something of which the justification for that belief can be seen as a necessary accompaniment (at least in any subject possessing the concept ... *is darker than* ...). Let us see how this way of seeing things can be applied to our example involving self-knowledge of what one sees.

4. Introspection through perceptual free-riding

Recall our puzzle. You're lying on the beach looking at the blue sky above you. And, as you do this, there are two simple facts you are aware of, namely:

- (1) The sky is clear
- (2) I see that the sky is clear

You know that the sky is clear by seeing that it is. But how do you know that you see that the sky is clear?

I propose we describe your knowledge that you see that the sky is clear as perceptual knowledge obtained 'for free' relative to your

11. Cf. McDowell's view that 'some concepts that figure in knowledge afforded by an experience can be excluded from the content of the experience itself' (McDowell 2009, 260). Cf. also Byrne's view, in Byrne and Siegel (2017) that certain perceptual beliefs have contents that go beyond the deliverances of the perceptual module.

equally perceptual knowledge that the sky is clear. The idea is that your self-knowledge involves another instance of perceptual free-riding, which—in the context of the view that experiences justify belief—can be spelled out as follows:

[Free-riding _{seeing}] Necessarily, if (i) you have a visual experience which justifies you to believe that the sky is clear and (ii) you possess the concept *I see that...*, then you also have perceptual justification to believe that you see that the sky is clear.

As before, the proposal is compatible with different accounts of the nature of conscious perceptual states. On the straightforward view, the claim is going to be that having the concept *I see that...* allows one to have an experience with the conceptual content *that one sees that the sky is clear* whenever one has an experience with the content *that the sky is clear*. On other accounts, this multiplication of experiences is only optional. The general idea is to explain your knowledge that you see that the sky is clear as purely (and more or less strictly) perceptual—much as we did with your knowledge that Egg 1 is darker than Egg 5.

The account involves no commitment to the idea of a dedicated faculty of ‘inner sense’. Even on a view where the experience justifying one to believe that one sees that the sky is clear is distinguished from the experience justifying one to believe that the sky is clear, we can still think of the former experience as a by-product of the operations of the very same ‘outward looking’ faculty whose operations are also responsible for the occurrence of the latter experience. Indeed, we *must* see it as such a by-product. For, insofar as we want to account for that experience in terms of Free-riding _{seeing}, we have to conceive of it as a necessary accompaniment to every experience justifying one to believe that the sky is clear (at least, in any subject possessing the concept *I see that...*). And it would hardly be a *necessary* accompaniment if its occurrence were imputable to a distinct perceptual faculty—or even just to a separate operation of the same one. One could say that, on the proposal under consideration, the two experiences (if they are two) are

simply different aspects, or components, of a single piece of ‘testimony of the senses’.

This last point nicely dovetails with the no-epistemic-overdetermination argument advanced, in §1, against treating your knowledge of (2) as non-perceptual. The argument was that, if we were to describe your knowledge of (2) as non-perceptual, we would have to say that you are justified to believe (1) ‘twice over’: on the basis of experience and on the basis of your non-perceptual belief in (2). But, on the perceptual free-riding account, that is not so. Since—in individuals possessing the concept *I see that...*—the experience providing justification to believe (2) is, as we’ve just seen, a necessary accompaniment to the experience providing justification to believe (1), your belief in (1) and your belief in (2) turn out to have a common source. And so, there is no sense in which the former provides you with additional, independent justification to hold the latter.

This leaves us with the two arguments against treating your knowledge of (2) as perceptual. One was that it does not ring true to say that you ‘see’ that you see that the sky is clear (and similarly for other perceptual verbs). The other was that perception is subject to all sorts of occlusions and malfunctions whose possibility appear to be excluded in the case of self-knowledge.

Let us deal with the second argument first. A proponent of the perceptual free-riding account can plausibly say that it lies in the nature of perceptual free-riding to rule out any occlusions and malfunctions which, leaving justification for the ‘basic’ belief untouched, hinder justification for the ‘free-riding’ belief. Recall the egg scenario. If you are looking at the five eggs, no occlusion or malfunction can leave you with justification to believe that Egg 1 is brown and Egg 5 is white, while hindering justification to believe that one egg is darker than the other. The same can be said about your justification to believe that you see that the sky is clear at a time when you do.

It may be objected that some kind of ‘obstructions’ are compatible with perceptual free-riding. Suppose an expert persuades you to accept

a deviant theory according to which Egg 1 need *not* be darker than Egg 5 even if Egg 1 is brown and Egg 5 is white. Assuming you can be described as possessing the concept ... *is darker than*... despite your acceptance of the deviant view, one can ask whether looking at the eggs still provides you with propositional justification to believe that one egg is darker than the other. Free-riding_{eggs} implies that it does. But having that justification and being able to take advantage of it are different things. Plausibly, the deviant theory stands in the way of your forming a justified belief to the effect that Egg 1 is darker than Egg 5—it acts as some kind of ‘rational obstruction’. And so, in the situation we are imagining, an epistemic opportunity that your experiences provide you with cannot be taken advantage of.

I agree that perceptual free-riding does not rule out the possibility of ‘rational obstructions’. But I submit that the same possibility arises with respect to beliefs about what one sees. If Free-riding_{seeing} is true and you possess the concept *I see that*..., seeing that the sky is clear is sufficient for you to have justification to believe that you see that the sky is clear. But having justification to believe the proposition that one sees that the sky is clear is not yet having a justified belief in (let alone knowledge of) that proposition. Someone can persuade you to accept a deviant theory of what it takes for you to see that the sky is clear. Or, more simply, they can persuade you that your eyes are not working properly, and you are not *seeing* that the sky is clear. Misleading testimony *can* sabotage self-knowledge of what one sees just as much as it can sabotage knowledge of the outer world. In allowing ‘obstructions’ of this sort while disallowing obstructions of the more concrete kind, the perceptual free-riding account makes the right prediction.¹²

What about the other argument against treating self-knowledge as perceptual, based on the observation that perceptual vocabulary is a poor fit for cases involving such knowledge? I think that the observa-

tion is true but can be accounted for without giving up the perceptual free-riding account.

It seems plausible to think that, when we say that you ‘see’ that *p*, this is because there is some distinctive worldly appearance such that you are able to tell that *p* based on recognizing that appearance. It is appropriate to say that you ‘see’ that the sky is clear because there is a distinctive appearance the sky has when it is clear and, based on recognizing that appearance, you are able to tell that it is. Now, there certainly *isn’t* any distinctive appearance of oneself (or the sky) on the basis of which one can tell whether one is seeing that the sky is clear. And that’s why, when you are lying on the beach looking at the blue sky above you, we don’t say that you ‘see’ that you see that the sky is clear (and similarly for other perceptual verbs). But this, by itself, is no reason to think that your knowledge that you see that the sky is clear is not perceptual. Consider again another case of perceptual free-riding I briefly discussed in the last section:

(5) Egg 1 and Egg 5 are spatially related to me

I suggested that you can know (5) by perception. But I don’t think it is very natural to say that you ‘see’ that the two eggs are spatially related to you. After all, there isn’t any distinctive appearance of oneself (or the eggs) on the basis of which one can tell that the eggs are spatially related to oneself. Now, one can accept this claim while holding on to the thought that it is *by looking at the eggs* that one comes to know that the eggs are spatially related to oneself. Mutatis mutandis, it is the same combination of commitments that I want to recommend in the case of your knowledge that you see that the sky is clear. You don’t ‘see’ that you see that the sky is clear. But it is *by seeing that the sky* is

12. For the notion of rational obstruction invoked in this and the preceding paragraph, see Pryor (2004, 364). The cases can also be explained as involving epistemic defeat—see, e.g., Conee (2005, 448–9).

clear that you come to know this.^{13,14}

5. Extensions

Let us broaden our focus. The case we looked at so far involves self-knowledge that one sees that so-and-so is the case. How far can the perceptual free-riding account be extended beyond cases of this kind? While it falls beyond the scope of this paper to provide a full-scale solution to all instances of the Epistemological Problem of Self-knowledge, I want to argue that some extensions are at least initially promising.

The extension to other perceptual modalities follows naturally. If you are alone in your room and hear the noise of some footsteps coming from the corridor, you may think to yourself:

(6) Someone is approaching

This you know by perception: you are listening to the noise coming from the corridor and you hear that someone is approaching. But, as you hear that someone is approaching, you are (almost ‘automatically’) in a position to know that you hear that someone is approaching. Or, as you would put it:

(7) I hear that someone is approaching

It is natural to think that we can explain your knowledge of (7) as a case of perceptual free-riding relative to your auditory-perceptual knowledge of (6), given your possession of the concept *I hear that...*

13. An alternative, pragmatic explanation may be available of the fact that we don’t very often say ‘I see that the eggs are spatially related to me’ or ‘I see that I see that the sky is clear’: assuming these are things we see, almost ‘automatically’, every time we see something else, it would be redundant to point out that we do. Thanks to an anonymous referee for this suggestion.

14. I owe the point that seeing that *p* involves the exercise of visual-recognitional capacities, whereas knowing that one that one sees that *p* doesn’t to Millar (2019, 112). There is also significant consonance between the perceptual free-riding account and Millar’s suggestion that “when I look at the rose [...] I tell that I see it to be a rose [...] in an act of judgement that is an immediate response to its looking that way to me” (2019, 114).

And, plausibly, what goes for vision and audition also goes for the other senses.

A second, fairly natural extension is to self-knowledge of other kinds of perceptual states. Seeing that the sky is clear and hearing that someone is approaching are both factive perceptual states. But perceptual free-riding can also explain one’s knowledge of less committal facts like:

(8) It visually seems to me that the sky is clear

(9) It auditorily seems to me that someone is approaching

More precisely, it is plausible to think that we could explain your knowledge of (8) in terms of perceptual-free riding relative to your visual-perceptual justification to believe (1), and your knowledge of (9) in terms of perceptual free-riding relative to your auditory-perceptual justification to believe (6).¹⁵ Of course, the peculiarity of non-committal facts like (8) and (9) is that one can ordinarily know them even in cases where one’s perceptual justification to believe, respectively, (1) and (6) has been defeated—cases of known illusion, for instance. But, once we have accustomed ourselves to the idea that, in normal cases, a single faculty provides us with two pieces of perceptual justification (either through a single experience, or through two distinct-yet-correlated ones), we can easily envisage the possibility that, sometimes, one of these pieces of justification—the one that we would want to describe as a ‘base’ for free-riding—may get defeated while the other one doesn’t. Knowledge of (8) and (9) in cases of known illusion can be seen as involving exactly this phenomenon.¹⁶

The perceptual free-riding account can also be applied to object-directed perceptual states. If you see that Egg 1 is white, perceptual

15. In both cases, your attainment of the relevant bit of knowledge would be conditional on your possessing the requisite concepts. Hereafter, I shall often leave this specification implicit.

16. Note that this narrative does not oblige us to call into question the truth of Free-riding *seeing*, which does not imply any form of counterfactual or epistemic dependence of the free-rider on the ‘base’.

free-riding may allow you to know that you see that Egg 1 is white. But it may equally allow you to know (as you would put it):

(10) I see Egg 1

Similarly, hearing the footsteps coming from the corridor, you are going to be able to know that you are hearing some footsteps. That knowledge, too, can be explained as the result of perceptual free-riding relative, once again, to your auditory-perceptual knowledge of (6).

The third extension is to certain corresponding pieces of negative self-knowledge. Suppose that, by looking at the blue sky above you, you notice that:

(11) There aren't any clouds in the sky

Again, this is a piece of perceptual knowledge. Knowing (11) by perception, you are going to be in a position to know that you don't see any clouds in the sky—or, as you would put it:

(12) I do not see any clouds in the sky

I submit that your knowledge of (12) can be explained in terms of perceptual free-riding relative to your perceptual knowledge of (11). If this is correct, perceptual free-riding can not only explain how you know that you see some clouds in the sky (when you do). It can also explain how you know that you *don't* see any clouds in the sky (when you don't). And that means that we can use perceptual free-riding to explain your general ability to tell *whether* you see any clouds in the sky. Again, the point generalizes to other perceptual modalities.

The last (and, for reasons I shall come to shortly, most controversial) application of the perceptual free-riding account is to self-knowledge of one's own sensations. Suppose you feel pain in your right foot. According to what may be called the *Perceptual Theory of Pain*, feeling pain in one's foot is a matter of experiencing that something is wrong with one's foot by exercising a distinctive perceptual faculty—often called 'nociception'—specialized for delivering information about bodily dis-

turbances of various sorts.¹⁷ Now, feeling pain in your right foot, you may be able to know:

(13) Something is wrong with my right foot

And if the perceptual theory of pain is right, we can say that you know (13) by nociception. But just as much as you know (13), you also know:

(14) I feel pain in my right foot

A proponent of the perceptual free-riding account will describe your knowledge of (14) as obtained 'for free' relative to your nociceptual knowledge of (13), given your possession of the concept *I feel pain*. Abstracting away from the details of the case, the general idea is that even self-knowledge of our 'innermost' mental states can be seen as a byproduct of the operations of certain perceptual faculties—those specialized for delivering information about one's own body.

The reliance on some independently characterizable species of perceptual knowledge—a trait that is common to all four extensions—constitutes, at the same time, the cornerstone and the source of the main inherent limitation of the perceptual free-riding account. In effect, what the account does is taking the phenomenon of perceptual justification for granted and inviting us to identify, for each of various bits of self-knowledge, some instances of such justification which can plausibly act as a 'base' for free-riding. But there are many instances of self-knowledge for which this strategy doesn't work. Your self-knowledge that you believe that $2+2=4$ cannot be explained in terms of perceptual free-riding—for there are no conscious perceptual states (hence no perceptual justifications) that can plausibly be argued to correlate, in the requisite way, with self-knowledge that one believes that $2+2=4$. And the same problem arises for self-knowledge of intentions, conjectures, hopes, desires, and many other propositional attitudes.

This limitation need not be seen as a black mark against the credibility of the perceptual free-riding account. According to an indepen-

17. See Byrne (2018, 147).

dently plausible view, a fundamental difference holds, not only between 'passive' (or 'receptive') mental states and 'active' (or 'spontaneous') ones, but also between self-knowledge of the first and self-knowledge of the second.¹⁸ The present proposal fits naturally with that view. The 'passive' (or 'receptive') mental states could be identified as 'experiences' in the sense of §3. And since experiences are, by definition, states conferring perceptual justification, self-knowledge concerning them, and only them, would be a most natural explanatory target for the perceptual free-riding account.

This way of setting things up, however, makes it crucial to the tenability of the account that at least *all* the extensions outlined above can be plausibly defended. If it turned out that the perceptual free-riding account can only explain self-knowledge of some 'passive' (or 'receptive') mental states and not others, this would—at least, potentially—constitute a blow to its credibility. It is in this connection, then, that an important worry must be addressed about the fourth and last extension, concerning self-knowledge of pain and other sensations.

6. Introspection as a *limiting case of perception*

The worry starts from the observation that obstructions and malfunctions aren't the only kind of perceptual mishaps that can put us out of touch with the outer world. Besides cases in which we fail to register what is going on with external objects and states of affairs, there are also cases in which *it seems to us* that we are registering certain types of external objects and states of affairs when, in fact, we are not. These are, as one may put it, cases of misleading (rather than missing) appearances. You look at an object under unusual lighting conditions and you take it to be red even if it is white. A straight stick looks bent to you because it is half immersed in transparent water. The examples abound and are exceedingly familiar.

Now, it may be suggested that the possibility of giving rise to *errors*

from misleading appearances (as I am going to refer to them hereafter) is integral to any kind of genuinely perceptual justification. For example, Burge writes that 'it is a fundamental feature of perceptual warrant [...] that it allows that an individual can be fooled while retaining warrant' (2003, 536). I take this to imply that no kind of justification (or 'warrant') ruling out the possibility of errors from misleading appearances (i.e., cases in which one is 'fooled while retaining the warrant') can be properly described as perceptual.

But it may also be suggested that, in the case of our introspective awareness of our own sensations, no errors from misleading appearances are possible. As Kripke pointed out in a much-quoted passage of *Naming and Necessity*:

Someone can be in the same epistemic situation as he would be if there were heat, even in the absence of heat, simply by feeling the sensation of heat [...]. No such possibility exists in the case of pain [...]. To be in the same epistemic situation that would obtain if one had a pain is to have a pain. (1980, 152)

I take it that being 'in the same epistemic situation that would obtain if one had pain' is having the kind of justification to believe that one is in pain which is provided by a pain sensation. If Kripke is right that having this kind of justification implies being in pain, one cannot have this kind of justification and be wrong.¹⁹

And so, we have a tension. Perceptual justification implies the possibility of errors from misleading appearances (that is Burge's point). But the kind of justification we have to believe that we are in pain when we have a pain sensation rules out that possibility (that is Kripke's point). It seems we should conclude that the kind of justification we have to believe that we are in pain when we have a pain sensation is *not* of the perceptual kind. And this means that the perceptual free-riding proposal is ill-suited to account for ordinary self-knowledge of pain (and

18. Cf. Boyle 2009. For a defense of a 'pluralist' approach to self-knowledge, see Coliva (2016).

19. Cf. the phenomenon of 'immunity to error through misascription' (Christofidou 2000 and Bar-On 2004, ch. 6).

of any other sensations to which Kripke's point applies).

The problem is not about any predictions that the free-riding account may make about the circumstances in which one is perceptually justified to judge that one is in pain. Generalizing from the example discussed in the last section, a proponent of the account may say that you have perceptual justification to believe that you are in pain whenever you have nociceptual justification to believe that there is something wrong with your body. This claim doesn't commit one to thinking that a pain sensation can sometimes justify you to believe that you are in pain when you aren't. On the contrary, it is compatible with it to insist that having nociceptual justification to judge that there is something wrong with one's body *just is* having a pain sensation, which *just is* being in pain. The question is whether a kind of justification that allows for these 'equations' can be legitimately described as *perceptual*. If it can't (as Burge's point implies) the perceptual free-riding account is wrong in claiming that the problem of explaining ordinary self-knowledge of pain can be subsumed under the general problem of explaining perceptual knowledge.

In principle, one could react by denying that ordinary beliefs about one's own sensations are really immune to error from misleading appearances. Reflecting on cases of 'phenomenal continua' where one is in pain and the pain slowly subsides until it no longer exists, Fumerton concludes that 'I might surely have some level of justification for believing that I am in pain even when if I'm not, but where I am instead in a state that is only very similar to pain' (Fumerton 2010, 380). The familiar fraternity initiation scenario can be invoked in support of the same conclusion: if one is told that a particular spot on one's neck is about to be cut with a razor and then an ice cube is placed on that spot, one may experience a sensation providing one with justification to believe that one is in pain. But the justification in question would be misleading.²⁰

20. This is probably a descendant of Churchland's (1984, 77) spy-under-torture scenario.

These counterexamples, however, are highly controversial, and I think that advocates of the perceptual free-riding account would be well-advised not to make their proposal depend on rejecting Kripke's point on their basis.²¹ Instead, I suggest they should say that, if Kripke is right about pain and other sensation, the kind of justification we have to believe that we are in pain when we have a pain sensation is best regarded as a *limiting* (but still perfectly bona fide) case of perceptual justification. Let me explain.

In mathematics, a limiting case of an object *X* is a case where one or more components of *X* take on their most extreme values. A circle is a limiting case of a polygon because it is a polygon with infinitely many sides. A point is a limiting case of a circle because it is a circle with radius equal to 0. The kind of justification we have to believe that we are in pain when we have a pain sensation is a limiting case of perceptual justification because it is a case where a certain component present in ordinary cases of perceptual justification takes on its most extreme value. What component, and what value?

The notion of an error from misleading appearances is clearly premised on a *fallible* notion of perceptual justification (that perceptual justification leaves room for errors of this sort means, as we've seen, that one can be 'fooled' while retaining the warrant). Now, someone who operates with a fallible notion of perceptual justification is likely to see any given perceptually justified judgment as involving a certain degree of *risk*, with some judgments qualifying as riskier than others. Suppose you are looking at a red-hot iron poker. One proposition you may be perceptually justified to judge on the basis of your visual experience is:

- (15) The poker is hot

For someone who operates with the fallible notion, though, the perceptual justification you have for (15) is compatible with various 'adverse' scenarios. Maybe the poker is red but not hot (because someone

21. For discussion, see, e.g., Wright (2015).

painted it red). Or maybe it is not even red, but only looks red (because the lighting conditions are unusual). Insofar as it tacitly relies on these adverse conditions *not* obtaining, your judging (15) would be riskier than your judging:

(16) The poker is red

And, for analogous reasons, your judging (16) would be riskier than your judging:

(17) The poker looks red

I want to suggest that it is with respect to *this* notion of riskiness that the kind of justification we have to believe that we are in pain when we have a pain sensation—a justification that one has whenever one is nociceptually justified to believe that something is wrong with one's body—may be regarded as a limiting case of perceptual justification. Ordinary cases of perceptual justification involve a certain degree of risk (or so one may think if one operates with the fallible notion). The kind of justification we have to believe that we are in pain when we have a pain sensation doesn't (or so one may think if one agrees with Kripke). But this doesn't make the latter kind of justification fundamentally different from the former. Instead, it is possible to see it as the extreme end of a spectrum along which we could in principle place all other instances of perceptual justification—each ever so slightly less risky than the previous one, each involving less tacit reliance on various adverse conditions not obtaining.²² The conclusion, then, is the following: even if Kripke is right about our introspective awareness of pain and other sensations, the credibility of the free-riding account is not necessarily in danger.²³

22. Cf. Sosa (2003, 282).

23. The idea that introspective awareness of one's own sensations should be seen as a limiting case of perceptual awareness is also advocated by McDowell, but in a different form. For McDowell, the point is that, unlike outer objects, 'the objects of 'inner' awareness [...] have no existence independently of that awareness' (1994, 21; cf. also 1998b and 2011). For criticism of this thesis, see Finkelstein (2003, ch. 3) and Bain (2009).

7. Comparisons

Our discussion took off from an ordinary case of self-knowledge of what one sees. It turned out that, by invoking the phenomenon of perceptual free-riding, one can explain that kind of self-knowledge as perceptual (§4) and then apply the same strategy to self-knowledge of many other 'passive' (or 'receptive') mental states (§5)—including, most controversially, self-knowledge of pain and other sensations (§6). What we ended up with, then, is a distinctive approach to the Epistemological Problem of Self-knowledge. Let me conclude by comparing it to some other prominent approaches, focusing on those that can be regarded as its closest relatives.

As I briefly mentioned in §2, Shoemaker thinks we should not try to assimilate introspection to perception. Part of the reason why Shoemaker thinks this has to do with what he calls the 'independent condition' on perception—roughly, the thesis that our knowledge of a certain class of states and events cannot qualify as perceptual unless 'the existence of [those] states and events is independent of their being known' by us (1994, 271). On Shoemaker's own view, introspection doesn't satisfy the independent condition.²⁴ And, at bottom, this is because introspection delivers a kind of knowledge that is intimately tied to rationality. Being rational is compatible with lacking perceptual knowledge but requires possessing (a certain amount of) self-knowledge.

Obviously enough, the present proposal differs from Shoemaker's in that the former tries to explain a significant portion of self-knowledge in perceptual terms.²⁵ Importantly, though, the way it does this is not at all incompatible with Shoemaker's insistence on the intimate relationship between self-knowledge and rationality. It is plausible to think that concepts like *I see / don't see that..*, *I hear / don't hear*

24. That is, he says, the 'fundamental difference' between perception and introspection (1994, 289).

25. What of the independence condition? I think that, as Shoemaker formulates that condition, self-knowledge satisfies it. Nothing in the perceptual free-riding account implies a kind of omniscience whereby mental states cannot obtain unknown by their subject.

that..., etc., as well as their non-factive counterparts, have a rational relevance that concepts like ... *is darker than*... do not have. Indeed, a good case can be made that concepts in the former class are *rationally mandatory*—meaning that if you are a fully rational agent, you are required to possess those concepts.²⁶ Suppose this is right. Then the perceptual free-riding account has the following implication: if you are a fully rational agent, whenever you see that the sky is clear, you are perceptually justified to believe that you do. Similarly, if you are a fully rational agent and don't see any clouds in the sky, you are going to be perceptually justified to believe that you don't. And so on and so forth. In brief, it is compatible with the perceptual free-riding account to hold that being rational requires being equipped to 'see' (scary quotes) certain things for free when one sees (no scary quotes) certain other things. I think it is evident how consonant the resulting view would be with much of Shoemaker's own theorizing about self-knowledge—specifically, with his insistence that rational self-blindness (the condition of a rational agent lacking first-personal knowledge of his own mental states) is not possible.

There are significant points of contact between the perceptual free-riding account and Peacocke's view of self-knowledge. According to Peacocke, 'you do not observe your perceptual states; nor do you know about them only by inference'; rather 'a perceptual experience's having a certain content makes reasonable, for one conceptually equipped to think it, the firstperson judgement that he is having an experience with a certain content' (1998, 83). This sounds a lot like saying that, with the requisite conceptual equipment in place, the subject of an experience can get a 'free-ride' to self-knowledge of the fact that he or she has that experience. But what is supposed to explain the 'free-ride' in this framework? The mere occurrence of the experience and its having the content it does are not sufficient for the subject of the experience to be justified in self-ascribing it. It is crucial to Peacocke's proposal that the experience also qualify as conscious (Peacocke 1998, §I). Now, this con-

sciousness requirement gives rise to a dilemma. If for the experience to be conscious is for the subject to be aware of it as the experience it is, Peacocke's account can be criticized for presupposing the very kind of self-knowledge (or self-awareness) it aims to explain. If, on the other hand, for the experience to be conscious is just for it to contribute to 'what it is like' for the subject who enjoys it—i.e., for it to be *phenomenally* conscious—it's unclear how satisfying the consciousness requirement could be sufficient for the experience to provide the subject with justification for the self-ascription.²⁷

The perceptual free-riding account provides a way out of this dilemma. On this account, self-knowledge of perceptual experiences can indeed be seen as involving, or presupposing, awareness of those experiences as the experiences they are—but this awareness is now explained as a matter of the subject's being perceptually justified to believe that he or she has that kind of experience (where this justification comes either from the perceptual experience itself or from an experience that necessarily accompanies it, depending on one's preferred account of how experiences justify). This way of escaping the dilemma, while in principle available to Peacocke, is not compatible with his explanatory ambitions. Peacocke aims to explain self-knowledge of occurrent beliefs, desires, intentions, etc. in the same way in which he explains self-knowledge of perceptual states (Peacocke 1998, §II), whereas I have argued that the perceptual free-riding account is inherently limited to self-knowledge of 'passive' (or 'receptive') mental states. Turning the perceptual free-riding account into an account of *all* species of self-knowledge would require renouncing the insight on which the whole account is based—namely that the kind of justification you have to believe that you see that the sky is clear is of the same (perceptual) sort as the justification you have to believe that the sky is

26. Cf. Merlo (2022).

27. For discussion of this dilemma see Coliva (2008).

clear.²⁸

Insofar as it gives ‘outward-looking’ knowledge explanatory (even if not epistemic) priority over ‘inward-looking’ one, the perceptual-free riding account may appear to resemble Byrne’s ‘transparency-based’ view of self-knowledge (and, of course, Evans’s (1982, 227-8) view, from which Byrne took inspiration). The impression of resemblance is reinforced by the fact that both accounts try to explain self-knowledge without positing dedicated faculties, methods, or ways of knowing things. But the similarities end here. The perceptual free-riding account differs from Byrne’s in that it does not treat self-knowledge of what one sees (and, more generally, self-knowledge concerning ‘passive’ mental states) as inferential in nature. For Byrne, when you know that you see that so-and-so is the case without the aid of mirrors, CCTV cameras, etc., this is because you have drawn an inference from a certain ‘visual proposition’ to the effect that so-and-so is the case to the conclusion that you see that so-and-so is the case (2018, ch. 6).²⁹ The inference is supposed to generate safe beliefs because—Byrne assumes—visual propositions are usually entertained as contents of the very same kinds of states that the beliefs in question ascribe to the subject—i.e., states of seeing.³⁰ But, leaving many subtleties aside, this position is well-known to suffer from the problem that its being the case that so-and-so is, in general, a terrible reason for one to believe that one is seeing that so-and-so (cf., e.g., Boyle 2011). Clearly, the present proposal doesn’t face this problem. The view I have been advocating is one on which,

when you know that you see that so-and-so is the case, your knowledge is, itself, perceptual in nature—it is knowledge acquired by exploiting, in the usual ways, a kind of justification that comes from perceptual experience. On this view, its being the case that so-and-so is not at all your reason to believe that you see that so-and-so is the case. It is merely something else that you see.

Other transparency-based accounts of self-knowledge do not construe people’s beliefs about their mental states as inferential. For example, according to Fernández’s ‘bypass model’, ‘we self-attribute beliefs on the basis of our grounds for those beliefs’ (2013, 50). This is a transparency-based account because, according to it, our grounds for self-attributing beliefs are often states that concern the outer world. But it is not an inferential account because, for all Fernández says about the basing relation, a self-attribution may be based on certain grounds without being inferred from them (2013, §2.2). Indeed, according to Fernández, that is exactly what happens when, e.g., one self-attributes a perceptual belief. When I form the perceptual belief that there is an apple in front of me, I self-attribute that belief on the basis of the very experience that justifies it—and I do this ‘immediately’ rather than inferentially. On the resulting account, my perceptual experience of the apple does double duty—it immediately justifies the perceptual belief, and it immediately justifies the self-attribution of that belief. Again, this looks a lot like a perceptual ‘free-ride’.

Interestingly, Fernández does not generalize his account to the kind of self-attributions of perceptual states that have been my primary focus in this paper. Maybe the reason for this is that applying the ‘bypass model’ to those self-attributions would risk introducing into the account the same kind of circularity that Peacocke’s proposal would suffer from on the first horn of the dilemma discussed above. As part of his analysis of the basing relation, Fernández requires that the state on which a certain belief is based should be epistemically available to the subject (Fernández 2013, §§2.2 and 2.5). The account one would obtain by applying the ‘bypass model’ to self-attributions of perceptual states would, therefore, presuppose (rather than explain) the epistemic

28. Cf. Smithies’s ‘simple theory of introspection’, according to which ‘introspective justification is fundamentally different in character from other kinds of justification, including perceptual justification’ (Smithies 2012, 261). Interestingly, Smithies (2019, 154-155) sees his account as an extension of Peacocke’s.

29. This is an extrapolation from Byrne’s own proposal, which is primarily concerned with object-directed perceptual states (e.g., seeing a hawk). The extrapolation will only account for certain basic cases of knowing that one sees that so-and-so is the case—at least if we assume, with Byrne, that visual propositions are concerned with a limited range of sensible qualities (2018, 137).

30. At least, that is supposed to be so ‘in creatures like ourselves’ (2018, 140).

availability of those states (what, in discussing Peacocke, we described as awareness of those states as the states they are). Again, one could perhaps see the perceptual free-riding account as filling a gap in the ‘bypass model’. But it matters that the two accounts are premised on very different conceptions of epistemic justification. Since, in many cases, one’s grounds for believing something are going to be a terrible reason to believe that one believes it, we can insist that those grounds justify the second-order belief only if we adopt something like a regularity account of epistemic support—and that is exactly what Fernández does.³¹ We do *not* need to do that when we explain your knowledge that you see that the sky is clear in terms of perceptual justification (or, for that matter, when we explain your knowledge that one egg is darker than the other in the egg scenario). So, not only are the two accounts—Fernández’s and the free-riding one—different in scopes, but the difference in scope reflects a deeper difference in their respective explanatory resources.

8. Conclusion

It is a beautiful summer day, and you are lying on the beach, looking at the blue sky above you. As your mind drifts lazily from one thought to another, there are many facts you have knowledge of: not only that you see that the sky is clear, but also that you hear seagulls in the distance, that you smell smoke from a barbecue nearby, and that you feel a slight sunburn sensation developing on your back. You know facts of this sort almost ‘automatically’ whenever they obtain, and this may suggest that your knowledge of them is not perceptual. But I argued that this suggestion can be resisted. On the view I put forward in this paper, your knowledge of those facts is a by-product of the operations of the very same ‘outward looking’ perceptual faculties whose operations are also responsible for their obtaining. You see that the sky is clear and, in the same way you know that, you also know that you do. If this view is right, introspection—or, at least, a significant portion of

what philosophers would call ‘introspection’—is nothing else than a (limiting) case of perception.³²

References

- Bain, D. (2009). “McDowell and the presentation of pains”. *Philosophical Topics* 37 (1): 1-24.
- Bar-On, D. (2004). *Speaking My Mind: Expression and Self-Knowledge*. Oxford: Oxford University Press.
- Boyle, M. (2009). “Two kinds of selfknowledge”. *Philosophy and Phenomenological Research* 78 (1): 133-164.
- Boyle, M. (2011). “Transparent self-knowledge”. *Aristotelian Society Supplementary Volume* 85 (1): 223-241.
- Burge, T. (2003). “Perceptual entitlement”. *Philosophy and Phenomenological Research* 67 (3): 503-48.
- Byrne, A. and S. Siegel. (2017). “Rich or thin?”. In Nanay, B. (ed.), *Current Controversies in Philosophy of Perception*. New York, Routledge: 59-80.
- Byrne, A. (2018). *Transparency and Self-Knowledge*. New York: Oxford University Press.
- Christofidou, A. (2000). “Self-consciousness and the double immunity”. *Philosophy* 75 (294): 539-569.
- Churchland, P. M. (1984). *Matter and Consciousness*. Cambridge, MA: MIT Press.
- Coliva, A. (2008). “Peacocke’s self-knowledge”. *Ratio* 21 (1): 13-27.
- Coliva, A. (2016). *The Varieties of Self-Knowledge*. London: Palgrave.
- Conee, E. (2005). “The comforts of home”. *Philosophy and Phenomenological Research* 70 (2): 444-451.

31. For criticism of this aspect of Fernández’s view, see Gallois (2015, 122-123).

32. The research leading to this paper received funding from the Swiss National Science Foundation (SNSF) through the Starting Grant project ‘Metaphysics We Can Believe In’ (TMSG11_218195). Earlier versions of the paper were presented at the Appearance & Reality Workshop (Ligerz, February 2022) at the ANU Philosophy of Mind seminar (Canberra, November 2022). I am grateful to the participants in these events for their critical feedback. A special thanks to Adrian Haddock, Alan Millar, and Crispin Wright.

- DeRose, K. (2002). "Review of *Knowledge and Its Limits* by Timothy Williamson". *The British Journal for the Philosophy of Science* 53 (4): 573-577.
- Evans, G. (1982). *The Varieties of Reference*. Oxford: Oxford University Press.
- Fernández, J. (2013). *Transparent Minds: A Study of Self-Knowledge*. Oxford: Oxford University Press.
- Finkelstein, D. (2003). *Expression and the Inner*. Cambridge, MA: Harvard University Press.
- Fumerton, R. (2010). "Poston on similarity and acquaintance". *Philosophical Studies* 147 (3): 379-386.
- Gallois, A. (2015). "Fernández on transparency". *Teorema* 34 (1): 121-134.
- Hawthorne, J. (2005). "Knowledge and evidence". *Philosophy and Phenomenological Research* 70 (2): 452-458.
- Kripke, S. (1980). *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- McDowell, J. (1994). *Mind and World*. Cambridge, MA: Harvard University Press.
- McDowell, J. (1998a). "Response to Crispin Wright". In Wright, C., Smith, B. C. & Macdonald, C. (eds.), *Knowing Our Own Minds*. Oxford, Oxford University Press: 47-62.
- McDowell, J. (1998b). "One strand in the Private Language Argument". In his *Mind, Value, and Reality*. Cambridge, MA., Harvard University Press: 279-96.
- McDowell, J. (2009). "Avoiding the Myth of the Given". In his *Having the World in View—Essays on Kant, Hegel and Sellars*. Cambridge, MA., Harvard University Press: 256-272.
- McDowell, J. (2011). "David Finkelstein on the inner". *Teorema* 30 (3): 15-24.
- Merlo, G. (2022). "Self-knowledge and the Paradox of Belief Revision". *Review of Philosophy and Psychology* 13 (1): 65-83.
- Merlo, G. (forthcoming). "Privileged access without luminosity". In Melis, G., G. Merlo, & C. Wright (eds.), *Self-knowledge and Knowledge A Priori*. Oxford, Oxford University Press.
- Millar, A. (1991). *Reasons and Experience*. Oxford: Oxford University Press.
- Millar, A. (2019). *Knowing by Perceiving*. Oxford: Oxford University Press.
- Moran, R. (2001). *Authority and Estrangement*. Princeton: Princeton University Press.
- Peacocke, C. (1998). "Conscious attitudes, attention, and self-knowledge". In Wright, C., Smith, B. C. & Macdonald, C. (eds.), *Knowing Our Own Minds*. Oxford, Oxford University Press: 63-98.
- Pryor, J. (2004). "What's wrong with Moore's argument?". *Philosophical Issues* 14: 349-378.
- Pryor, J. (2005). "There is immediate justification". In Steup, M. & E. Sosa (eds.), *Contemporary Debates in Epistemology*. Malden, MA., Wiley-Blackwell: 181-202.
- Ryle, G. (1949). *The Concept of Mind*. New York: Barnes & Noble.
- Shoemaker, S. (1994). "Self-knowledge and 'inner sense' - Lecture II: The Broad Perceptual Model". *Philosophy and Phenomenological Research* 54 (2): 271-90.
- Silins, N. (2024). "Perceptual experience and perceptual justification", *The Stanford Encyclopedia of Philosophy* (Fall 2024 Edition), E. N. Zalta & U. Nodelman (eds.), URL = <<https://plato.stanford.edu/archives/fall2024/entries/perception-justification/>>.
- Smithies, D. (2012). "A Simple Theory of Introspection". In Smithies, D. & D. Stoljar (eds.), *Introspection and Consciousness*. New York, Oxford University Press: 259-194.
- Snowdon, P. (2012). "How to think about phenomenal self-knowledge". In Coliva, A. (ed.), *The Self and Self-Knowledge*. Oxford, Oxford University Press: 243-62.
- Sosa, E. (2003). "Privileged access". In Smith, Q. & A. Jokic (eds.), *Consciousness: New Philosophical Perspectives*. New York, Oxford University Press: 238-251.

- Williamson, T. (1996). "Cognitive homelessness". *Journal of Philosophy* 93 (11): 554-573.
- Wright, C. (1998). "Self-knowledge: the Wittgensteinian legacy". In Wright, C, Smith, B. C. & Macdonald, C. (eds.), *Knowing Our Own*

- Minds*. Oxford, Oxford University Press: 101-122.
- Wright, C. (2015). "Self-knowledge: the reality of privileged access". In Goldberg, S (ed.), *Externalism, Self-Knowledge and Scepticism: New Essays*. Cambridge, Cambridge University Press: 49-74.