

Genia Schönbaumsfeld, Southampton

How Threatening are Local Sceptical Scenarios?*

Abstract: In this paper I distinguish between ‘local’ and ‘global’ forms of ‘envatment’. I show that recent envatment arguments (the ‘local’ variety) work similarly to arguments from perceptual illusion and that neither of them are able, by themselves, to get us ‘global’ scepticism. Consequently, motivating the radical sceptical idea that all of our perceptual beliefs might be false is harder than it looks.

1. Introduction

In this paper I distinguish between three different kinds of envatment used to motivate ‘external world’ scepticism: recent envatment, lifelong envatment, and ‘nothing but envatment’. I will show that recent envatment scenarios (RE) are a kind of ‘local’ sceptical scenario, whereas lifelong envatment (LE) and ‘nothing but envatment’ (NBE) are global or radical sceptical scenarios that attack the very idea that anyone has ever had perceptual contact with an ‘external’ thing. This distinction is important, as one of the main claims that I seek to defend in this paper is that local sceptical scenarios by themselves are not sufficient for motivating the ‘global’ variety, that is to say, the thought that *all* our perceptual beliefs might be false.

I will proceed as follows. In the next section I show that RE scenarios and arguments from perceptual illusion work in a similar way in the sense that both must presuppose the existence of an ‘external world’ in order to function. Hence, these scenarios are not powerful enough to make radical scepticism (uncertainty about the existence of anything ‘external’) seem compulsory. I then go on to expose the picture that really motivates global sceptical scenarios – a Cartesian conception of our evidential situation – and show that this is something we can avoid. Following Wittgenstein, the aim is not to provide a refutation of global scepticism, but a way of preventing ourselves from getting into its grip in the first place – a way of resisting, in other words, the decisive, but misleadingly innocent, movement in the radical sceptic’s conjuring trick (PI 2009: §308). I conclude by showing that the perceived need for a ‘global validation’ of our

* Included in N. Venturinha (ed.), Special Section on Wittgenstein and Applied Epistemology.

epistemic practices is a product of a prior endorsement of the Cartesian picture, not the natural and intuitive starting-point of all epistemic investigation.

2. Local sceptical scenarios and arguments from perceptual illusion

Let us start by considering ‘local’ envatment scenarios. What is a local envatment scenario? Well, similarly, to other local sceptical scenarios, such as Dretske’s zebra example (of which more later), a local envatment scenario casts doubt on particular, ‘local’ perceptual knowledge claims, rather than attacking the very idea that perceptual knowledge is ever possible. Thus, local sceptical scenarios give one grounds for doubt about particular perceptual cases (or larger swathes of them), while leaving intact the thought that perceptual knowledge is sometimes available.¹ Applied to brain-in-a-vat (BIV) scenarios this implies arguing for recent envatment, rather than lifelong envatment or ‘nothing but envatment’, as both of the latter scenarios attack the notion that perceptual knowledge is ever possible. In LE I have always been a BIV and have therefore never been in touch with an ‘external world’, but only with electrode stimulations, whereas in ‘nothing but envatment’ there may be nothing else at all except me and my vat.² This is a kind of metaphorical BIV scenario analogous to Descartes’ evil demon argument, where we are discounting the notion that someone must have set up the vat and is continuing to service it.

RE, on the other hand, involves taking the BIV scenario *au pied de la lettre*. It presupposes the following: that there is a world, for example, in which the vat containing my brain exists; that there are evil scientists (or robots or aliens or what-have-you) who have caused me to become envatted (or who have ‘bred’ ‘me’ – my brain – to be envatted); that it is scientifically possible to separate brains from bodies without killing off the brain; that I might find out about my previous envatted state by, for instance, having my brain ‘reinserted’ into my body; that I ‘am’ my brain and so on. This, as it were, literal fleshing-out of the sceptical scenario is precisely what turns it into a ‘local’ case: if I imagine that I might be the victim of such a predicament, this does not imply that anyone else is, or, indeed, that the ‘external world’ as such does not exist (since, as we have just seen, a world containing vats, brains and evil scientists is in fact pre-

¹ Of course the sceptic may want to use local sceptical arguments in order to motivate ‘global’ sceptical scenarios, and it is this move that I will be attacking later in the paper.

² I attack NBE directly in chapter 3 of Schönbaumsfeld 2016.

supposed by this form of scepticism). All it implies is that, for as long as I am a BIV, many of my current perceptual beliefs about the world are false.

It is important to note here, however, that not *all* of my beliefs about the world will be false. Since I have formerly interacted with the world in the normal way, many beliefs I hold about the state of this world will survive envatment. For example, it remains true that London is the capital of England, that Great Britain is a small island off the Continent, that cats are mammals etc. Consequently, my beliefs about these things will remain true post-envatment also (even if I'm now a BIV on Alpha Centauri, as these beliefs don't depend on what I'm currently perceiving). RE will therefore mainly pose a threat to my current perceptual beliefs – e.g. about there being a computer in front of me, or a tree, or a person etc. It is important to stress this, as this fact poses a major stumbling block to any attempt at generalizing local sceptical arguments to get one the 'global' conclusion that if one can sometimes – or very often – be wrong, it is possible that one could always be wrong, and, hence that one cannot be certain that there is an 'external world' at all.

If this is right, then RE scenarios cannot accomplish much more than arguments from perceptual illusion, which are similarly unable to get one from local to global error. That is to say, just as RE does not imply LE or NBE, so the possibility of perceptual error or illusion does not imply systematic or 'global' perceptual unreliability. For example, I am only able to determine that when I look at a square tower from a distance, it will appear round, because I can trust my perception that from close-by it looks square, and there is a scientific explanation available that can tell me why it nevertheless appears round from some way off. If perception in general were deceptive, I could not make the judgment that perceptual appearances are sometimes misleading. All I could do would be to report, for instance, that at time t1 I have the impression that 'thing 1' that I see is square, while at time t2 I have the impression that 'thing 2' that I see is round. And since 'thing 1' and 'thing 2' might, for all I know, be different things, I could not even conclude that one perceptual experience might be an accurate representation of the way things are, while the other might not. But if I cannot make this judgement, I am similarly unable to conclude that at time t2 I am being misled, for my perceptual experience at t2 would only be misleading if it were an experience of the same thing that I encountered at t1, so that my reports at t1 and t2 would turn out to be in conflict with each other. As long as I have no reason for assuming that my perceptual experiences of 'thing 1' and 'thing 2' are in fact experiences of the same physical object, however – which, of course, I would not if perception were generally defective (for then I would have no grounds for trusting one report more than the other; I would rather have to assume they are equally misleading) – there is no way of determining

that a perceptual ‘illusion’ has in fact occurred. All I could say is that I am having different perceptions at different times, but this, of course, does not suffice to allow me to infer that at time t_2 I was misled, and, hence, that perceptual errors are possible.

What if someone objected, however, that at t_1 , I have the impression that thing 1 that I see is square, at t_2 I have the impression that thing 2 that I see is round, and between t_1 and t_2 I have the impression that thing 1 hasn’t changed shape or been replaced – wouldn’t we then have to conclude that a perceptual illusion has in fact occurred? We might indeed conclude in such a case that a perceptual illusion has occurred (as we might also in the other scenarios), but the relevant point is that we can only draw this conclusion because we are taking our impression that thing 1 has not changed between t_1 and t_2 (or been replaced) to be veridical, which confirms what I was trying to show: namely, that we cannot conclude that a perceptual error has occurred unless we are willing to grant that some of our sense-perceptions can be veridical. If we don’t, then we certainly cannot grant that we can take our impression that thing 1 hasn’t changed between t_1 and t_2 at face value (and hence we could not conclude that a perceptual error has occurred). Consequently, far from showing that perception, in general, is defective, the possibility of perceptual error actually presupposes that perception is generally in good working order. So one might say that ‘aggregate arguments’ based on the possibility of perceptual error fail: I am not entitled to infer that because I can sometimes be wrong (‘local case’), it is possible that I could always be wrong (‘global case’). In other words, one cannot ‘aggregate’ local sceptical scenarios in order to reach a radical sceptical conclusion. That is to say, just as we need to presuppose the existence of an ‘external world’ that is broadly like our own in order to get the RE scenario off the ground, so perception must generally be taken to be reliable if an ‘argument from illusion’ is to be constructed.

It is worth pointing out at this juncture that this has nothing to do with the question of whether or not local error is always detectable. It may not be detectable in any given case. The point is rather that if I start with the assumption that *all* perceptions are unreliable, then I can *never* formulate an argument from illusion, as such an argument presupposes that some of our perceptions are reliable (e.g. that the tower has not changed shape and now looks square). Hence, I cannot use an argument from illusion in order to undermine the very thing the argument itself presupposes: the reliability of at least some of our sense-perceptions. But this is not to endorse the false principle that if it is not possible to tell that P, it is not possible that P.

If this is right, then what makes local sceptical scenarios possible is the very thing that they are, ironically, drafted in to undermine: a background of gener-

ally veridical perceptual experience. In this respect, RE is just a souped-up, ‘sci-fi’ version of the more pedestrian arguments from illusion and cannot, for this reason, achieve more than they can: to show that perceptual error is possible (that human beings are fallible). Since we have already seen, however, that ‘aggregate arguments’ are fallacious, an appeal to fallibility alone is not enough to get one the conclusion that everyone could be wrong about everything *all* of the time.

Nevertheless, one might, perhaps, wish to object that even if the foregoing is correct, it is at least *logically* possible that I am currently envatted and only vainly imagining that I am, say, looking out of the window (or, to speak with Descartes, that I am currently dreaming that I am doing this). Although this is true, it is important to bear in mind that the mere fact that I am able to conceive of such a state of envatment does not give me a real reason to suppose that I might *actually* be the victim of such a situation. Why not? Because, in the absence of a general argument designed to undermine the possibility of knowledge *per se*, I have no real grounds for thinking that the imagined scenario might be the actual one. For example, the state of current science is not yet advanced enough to make such scenarios even empirically possible; there is no evidence of alien or robot activity etc. In short, the thought that, despite all the evidence to the contrary, I might nevertheless be a BIV, purely because this is logically conceivable, is not yet to offer a reason *in favour* of this scenario: logical conceivability alone is not a form of evidence. Rather, if something is not logically possible, then it need not even be investigated, but this does not imply that logical conceivability counts, by itself, as a reason that speaks in favour of a given scenario.

Klein provides a good example to illustrate this point. In order to have sufficient evidential support for the claim that a particular person that we see is Publius, for instance, we don’t first need to be able to rule out an incredulous bystander’s idea that it might not be Publius, but rather Magicus’s dog, whom Magicus has transformed into an exact duplicate of Publius, unless we already have independent grounds for thinking that Magicus has that “ability, an opportunity and a motive” (Klein 2004: 174). If this is right, then we analogously don’t need to accept the sceptic’s claim that we need to be able to rule out RE *before* we are entitled to appeal to any background knowledge, since the sceptic has so far given us no good reason, apart from bare logical possibility, for thinking that RE might be true.

In other words, we can accept the possibility of perceptual knowledge unless we are given a good reason not to, and not the other way round. We don’t, it seems to me, need to bother with a sceptic who wants an a priori demonstration of perceptual reliability without giving us a good reason first for indicting it: we

should espouse the principle ‘innocent until proven guilty’,³ not ‘guilty until proven innocent’. Not only is the former principle independently plausible, but it is made so much the more so given that, as I have shown in this section, both RE scenarios and arguments from perceptual illusion require a generally reliable background in order even to be capable of formulation. Ironically, it is that very background, however, that also gives us reason to resist the sceptical conclusion. If I really were a BIV, I’d be in a pretty bad epistemic state, but no compelling reason has so far been given for me to believe that I might *actually* be one.

So, contrary to what is commonly assumed, there is no ‘local’ route to a ‘global’ scepticism. For as we have seen, neither arguments from perceptual illusion nor RE scenarios are by themselves sufficient to entail the conclusion that perceptual experience in general is never to be trusted. If sceptical scenarios, therefore, amounted to no more than the local variety, one wouldn’t have much reason to be concerned about them.

3. Global sceptical scenarios and the picture that holds us captive

The sceptic has more weapons in his armoury, however. For what about LE or NBE? Doesn’t the mere raising of such scenarios imply that one must find a response to them (given that there can’t actually be any empirical evidence in favour of them)? But how would one do that? The approach I’m going to take is to show that what makes LE and NBE seem compelling are not independently plausible arguments that show why we should take seriously the idea that we could be victims of such scenarios, but rather a prior endorsement of a certain picture of our evidential situation which is non-compulsory. I contend that it is this picture which has the radical sceptical problem built into it from the very start, and it is therefore its hold on us that we need to break (PI 2009: §115).

Let us start by considering the concern that LE and NBE scenarios give expression to. For, clearly, nobody entertains such scenarios because they strike one as probable or in any way as evidentially motivated (as *ex hypothesi* there cannot be any empirical evidence either for or against them). Rather, such scenarios seem compelling because they express a deep-seated anxiety human beings appear to have about their relation to the world: how do we know that we are ever in touch with something genuinely ‘external’ and not just locked into

³ Compare Thorpe (forthcoming).

the cage of our own minds? Thomas Nagel gives potent expression to the fear that the very idea of possessing a point of view on the world, which is inherently limited by our particular perceptual and intellectual constitution, condemns us to the possibility that nothing in this world might be as we imagine it to be: “Objectivity and scepticism are closely related: both develop from the idea that there is a real world in which we are contained, and that appearances result from our interaction with the rest of it” (Nagel 1986: 68). Why does the fact that there is a real world open the door to radical scepticism? It is because, Nagel thinks, “it seems to follow that the most objective view we can achieve will have to rest on an unexamined subjective base, and that since we can never abandon our own point of view, but can only alter it, the idea that we are coming closer to the reality outside it with each successive step has no foundation” (ibid.). In other words, unless we somehow manage to discard our own point of view in favour of what Nagel calls ‘a view from nowhere’ (ibid.: 67), we can never really know about the world as it is ‘outside’ of our own minds; including if there is one. Stroud raises similar concerns:

We are confined at best to what Descartes calls ‘ideas’ of things around us, representations of things or states of affairs which, for all we can know, might or might not have something corresponding to them in reality. We are in a sense imprisoned within those representations, at least with respect to our knowledge. Any attempt to go beyond them to try and tell whether the world really is as they represent it to be can yield only more representations, more deliverances of sense experience which themselves are compatible with reality’s being very different from the way we take it to be on the basis of our sensory experiences. There is a gap, then, between the most that we can ever find out on the basis of our sensory experience and the way things really are. In knowing the one we do not thereby know the other. (Stroud 1984: 32)

The thought that we are ‘imprisoned’ by or locked into our own minds or perceptual experiences is, first and foremost, an expression of (what Pritchard calls) epistemic angst.⁴ Epistemic angst is the fear of epistemic limitation, which in the guise just sketched, takes the form of anxiety over the very idea of having a point of view, or of being a creature with a mind capable of taking up a perspective on a (putatively) objective world. The underlying idea seems to be not only that our point of view is limited in all sorts of contingently fallible ways and may therefore never convey the full picture of what is ‘out there’, it also expresses the full-blown angst that our ‘perspective’ might not so much as provide a view on *anything* – that we might be trapped in a mental cage with no beyond at all. It is this angst that fuels the Cartesian picture of our epistemic situation

⁴ This term was first used by Duncan Pritchard in his (2014).

which conceives of the ‘external world’ not as something that is, in principle, directly present to the senses, but rather as something that we can only know about by dint of making inferences from our mental states or apparent ‘perceptual’ experiences. And once such a picture is in place, it indeed becomes hard to see how we could so much as know about this world beyond the ‘veil of appearances’ (compare Stroud 1984), including whether it exists.

That is to say, LE⁵ and NBE scenarios give expression to the fear that there might be a systematic mismatch between everything we collectively believe to be the case about the world (including if there is one), and the way the world actually is (if indeed there is one). And, if this were in fact so (if there were indeed a systematic mismatch), then, even if it seemed otherwise, we would never have knowledge of anything; nor could we ever find out (even in principle) whether such a scenario obtained, as any form of evidence one could appeal to would itself be part of the ‘grand illusion’.

Now what, apart from epistemic angst, is it that so much as gives us the idea that all of our beliefs taken together could fail to ‘match’ the way the world is given that, as we have just seen, we are not entitled to infer this conclusion from local sceptical scenarios? Following McDowell (1998a, 1998b, 2009), I locate this in a mistaken picture of our perceptual reasons – what McDowell calls the ‘highest common factor’ conception, and I call the Reasons Identity Thesis (RIT):⁶ the view that in both the good and the bad cases the perceptual reasons that we have access to must be the same. According to this conception – widespread in contemporary epistemology⁷ – the warrant one can have for one’s perceptual claims is always constituted by the ‘highest common factor’ of the good case and the bad case, which implies that our perceptual grounds can never be better than what the bad case provides.

There is an intimate connexion between this view of perceptual reasons and radical scepticism. Consider the following argument:

⁵ LE is less radical than NBE, where we are supposing there is no ‘external world’ at all. LE leaves intact the notion that there is an ‘external world’ in which my envatted brain exists, but I have never had any access to this world. These scenarios are often not distinguished, but I think that we should distinguish them. LE is a radicalized version of RE, but NBE is, in a sense, merely a contemporary version of the Cartesian sceptical scenario where there is nothing but my mind and its impressions, and there is no way of determining where these impressions come from (if, indeed, they come from anywhere).

⁶ See Schönbaumsfeld 2016.

⁷ This view is endorsed, *inter alia*, by Bennett (1971), Burge (2003), Conee (2007), Coliva (2012), McGinn (1984), Millar (1991), Nagel (1986), Pollock (1974), Pryor (2000), Stroud (1984), White (2014), Wright (2002, 2014).

The Indistinguishability Argument for 'External World' Scepticism

P1. In the bad case, the supporting reasons for one's perceptual beliefs can only consist of the way the world appears to one.

P2. The good and bad cases are phenomenologically indistinguishable.

C1. So, the supporting reasons for one's perceptual beliefs in the good case can be no better than in the bad case. (From (P2))

C2. So, the supporting reasons for one's perceptual beliefs can only consist of the way the world appears to one. (From (P1), (C1))⁸

C3. If the supporting reasons for one's perceptual beliefs can only consist of the way the world appears to one, one can never know more than how the world appears. (From (C2))

C4. If one can never know more than how the world appears, one can only know appearances. (From (C3))

C5. If one can only know appearances, the 'external world' lies 'behind' a 'veil of appearances'. (From (C4))

C6. If the 'external world' lies 'behind' a veil of appearances, knowledge of the 'external world' is necessarily 'indirect' – i.e. rests on an inference from these appearances to an 'external reality'. (From (C5))

C7. But such inferences will only be valid if one already knows *independently* (of relying on those appearances) that appearances are a good guide to the 'external world'. (From (C6))

C8. But one cannot know this, precisely because one only has access to appearances and their testimony is inconclusive. (From (C3), (C4), (C5), (C6), (C7))

C9. So, there is no non-circular way of demonstrating (from appearances) that the 'external world' exists and knowledge of it is possible. (From (C8))

⁸ I have borrowed this part of the argument from Pritchard (2008: 294).

C10. But such a demonstration is nevertheless necessary, as the hypothesis that there is an ‘external world’ could be false, for one might be in an LE or NBE scenario. (From (C9))

C11. So, scepticism.⁹

The notorious step in this argument is clearly the move from (P2) to (C1). If we grant this move, we grant – at the same time – that we cannot know that there is an ‘external world’. What a strong epistemological conclusion from an innocuous-seeming claim about the subjective indistinguishability of the good and the bad case! This prompts the question: is the move compulsory? That is to say, is (C1) actually entailed by (P1) and (P2)? It seems not.

One of the first philosophers (apart from McDowell) to notice this was Pritchard (2009, 2012). “It is the collective failure amongst contemporary epistemologists” to recognize that (P1) and (P2) fail to entail (C1), Pritchard claims, “that has led them to succumb to the sceptical problem in the first place” (Pritchard 2009: 472). In other words, Pritchard thinks that McDowell has shown that it is a mistake to suppose that phenomenological indistinguishability by itself implies the Reasons Identity Thesis (RIT): that from the fact that I might not, in a given case, be able perceptually to discriminate between the good and the bad scenario, it follows that the reasons available to me in both the good and the bad cases must be the same.

For if one resists the non-mandatory move to C1, one can hold, against RIT, that ‘good case’ perception is factive: in the good case (where the environment is epistemically friendly) seeing that P can provide factive epistemic support for the belief that P, while in the corresponding bad case – where the environment is epistemically unfriendly and one only seems to see that P – one does not have factive epistemic support for one’s belief. And this disjunction holds, according to McDowell, whether or not one can phenomenologically distinguish between the two cases. That is to say, in the good case, seeing that there is a chair over there entails that there is a chair over there, which means that one doesn’t, in order to be epistemically entitled to this proposition, first need to be able to rule out that one is not radically deceived.

Of course, one might object here that McDowell’s disjunctive conception only works as a way of blocking scepticism if one already knows that one is in the good case. Given that everyone agrees that one cannot introspectively dis-

⁹ Notice how (C3)–(C11) reproduce the structure of Wright’s I–II–III argument for scepticism (Wright 2002: 338–9).

tinguish between the good and the bad case, however, since one's experiences would be the same, phenomenologically speaking, in both, it follows that perceptual reasons cannot, after all, be factive.¹⁰ The idea behind this argument is that one's perceptual reasons can only consist of what is common to the good case and the bad case – that is to say, their highest common factor. But although McDowell would accept (P1) and (P2) in the Indistinguishability Argument, he thinks that it would be a mistake to let our epistemic standing in the good case be dictated by our epistemic standing in the bad case (i.e. he would reject RIT). Consequently, this counter-argument is simply a question-begging attempt to insist that (C1) follows from (P2).

That Crispin Wright (2002, 2008, 2014), for example, believes, by contrast, that one needs to demonstrate that one is not dreaming (or not a BIV) *before* one is epistemically entitled to any perceptual knowledge claims already reflects his antecedent commitment to RIT (the HCF conception), since only this forces one to accept the thought that the epistemic grounds¹¹ one has access to are the highest common factor of the good case and the bad case, which can only be non-factive. Consider, for example, Wright's 'I-II-III' characterization of what gives rise to the sceptical problem:

Type-I Proposition: It seems to S as if she has two hands.

Type-II Proposition: S has two hands.

Type-III Proposition: S is not a brain in a vat (BIV).

According to Wright, S's evidential position is captured by the type-I, not the type-II proposition. The problem with this is that type-I propositions merely provide *prima facie* evidence for belief in type-II propositions, while only type-II propositions entail type-III propositions. For if S is in fact a handless BIV, then the type-I proposition – its seeming to S as if she has two hands – will provide no good reason for her belief in the type-II proposition that she has two hands. Hence, one is only entitled to the type-II proposition if one can already rule out that one is a BIV – that is to say, if one already has independent grounds for asserting the truth of the type-III proposition, which seems impossible.

On McDowell's alternative view, however, perceptual reasons do not, in the good case, fall short of the environmental facts, and consequently a lack of 'ex-

¹⁰ See, for instance, Wright 2002: 346.

¹¹ I am using 'grounds' and 'reasons' interchangeably.

periential markers' that would allow one introspectively to distinguish the good from the bad case is irrelevant. Consequently, LE and NBE will only pose a threat to one's knowledge of the world if one *already* agrees that one's perceptual reasons can only consist of what is common to both the good and the bad cases, and to assume that is just what, from the perspective of the disjunctive conception, is to be proved (i.e. it would be question-begging).

Perhaps it is hard to see this, because many philosophers just implicitly take 'the highest common factor thesis' for granted, and therefore believe that RIT is sufficiently motivated. The 'highest common factor thesis' is the view that the only facts that are reflectively available to an agent in the good case are facts that would also be available to one's physical duplicate in a bad case. But this is precisely the view that disjunctivists reject: according to disjunctivism, the good and the bad cases are asymmetrical and consequently provide us with different perceptual epistemic grounds. In the good case, these reasons give us access to the facts themselves (and if we have access to these facts, we don't first need to prove that the 'external world' exists); in the bad case they don't.

That is not to say that there is *nothing* that a good case agent shares with her bad case duplicate, for they may both be epistemically *blameless*. But that doesn't imply, as Pritchard correctly points out, that therefore the two agents enjoy equal levels of epistemic support or that a subject's epistemic support is constituted *solely* by properties that she has in common with her recently envatted physical duplicate (Pritchard 2012: 42). For why should one accept that one's perceptual epistemic support depends on introspective distinguishability *alone*?¹² Why shouldn't access to the facts in the good case also make some difference? Unless I have already accepted the radical sceptical conclusion of the Indistinguishability Argument, I cannot assume that my epistemic grounds are exhausted by knowledge of what is currently present to my consciousness.

If this is so, however, then the Indistinguishability Argument is unsound, for, unless supplemented by the contentious 'highest common factor thesis', (C1) will not go through. The benefits of realizing that we have encountered a roadblock to RIT are great: not only are we preventing global sceptical scenarios from arising, we are also, at the same time, undermining the traditional epistemological picture according to which one is locked into an 'inner' world of appearances (or of 'sensory experiences') which gives one, at best, inconclusive,

¹² Of course epistemic externalists would reject this thesis, but they tend to go too far the other way by denying that one can have reflective access to one's epistemic grounds at all, which is a high price to pay (for more on this, see Pritchard 2012).

inferential knowledge of an ‘outer’ external world. So, McDowell is right that the ‘prop’ that radical scepticism relies on is the HCF conception – the thought that one’s perceptual reasons can never be better than what is *common* to the good case and the bad case (and, hence, that the reasons both scenarios give one access to are identical) – since it is this that fuels the idea that, for all one knows, one might be completely cut off from reality. Given that, on this view, all that one has perceptual access to is what is currently present to consciousness, one can never know non-inferentially and hence ‘directly’ what lies beyond it.

What might, at first sight, therefore seem only to be an epistemological thesis, (C1), actually has strong *metaphysical* implications—namely, the inaccessibility to direct perception of the ‘external world’. That is to say, endorsing (C1) in the Indistinguishability Argument not only commits one to the epistemic notion that perceptual reasons are inconclusive, but also to the metaphysical thesis that all one can have perceptual access to is to appearances *as of there being physical objects*, never to the physical objects themselves, as even in the best possible case, perceptual experience always falls short of providing knowledge of the facts. For if one can never know non-inferentially that one is in the good case, since what one has access to is only the appearance *as of there being a hand*, and one therefore has to make an inference from this appearance to the nature of reality, then what one, in fact, has ‘direct contact’ with is merely the *appearance* of hands, cookies and so forth, not with the actual hands, cookies etc. themselves.

If this is right, then the following line of defence is not open to proponents of RIT – namely, to grant that there *could* be such a thing as a good case, where one *would* enjoy direct perceptual access to physical objects, but to deny that one can ever know non-inferentially, via perception, that one *is* in the good case.¹³ For example, one might contend that one is ‘directly’ perceptually experiencing something that looks like a barn, say, while in fact being presented only by a barn-façade.¹⁴ In such a case, the objection goes, one would have ‘direct perceptual access’ to an object, but not be able to know, merely by looking at the object, that this object was a barn.¹⁵ But a proponent of epistemological disjuncti-

¹³ So, *pace* Byrne 2014, RIT can accommodate neither direct perceptual access to physical objects (what Byrne calls object-DPA), nor ‘inference DPA’ – the thesis that “we are in a position to non-inferentially know external world propositions by perception” (Byrne 2014: 284).

¹⁴ Compare Wright 2008 and Conee 2007.

¹⁵ This is a variation of Dretske’s ‘cleverly disguised mule’ example – that is to say, I might have ‘direct perceptual access’ to something that looks like a zebra, but if it is in fact a cleverly disguised mule, then I cannot come to know, merely by ‘looking’ at this thing, that it is a zebra. See the discussion of Pritchard below.

vism would simply deny that this sort of case counts as a good case. Hence, although one has direct perceptual access to a physical object in such a scenario, it is not to a *barn*, but to something that merely looks like one (i.e. to a ‘barn-façade’). It is only if one is actually perceiving a *barn*, however, (not a barn-façade, or what-have-you) that seeing that there is a barn in front of one puts one in a position to know that there is a barn in front of one. Furthermore, it must be the case that one’s cognitive faculties are operating normally and that no (misleading) defeaters are present.¹⁶ That one might not be able introspectively to distinguish, merely by attending to the putative ‘barn-experience’ itself, whether it is the experience of a real barn or of a barn-façade is not an objection to the view, just one more question-begging attempt to insist that phenomenological indistinguishability implies the non-facticity (and hence the inconclusiveness) of one’s perceptual reasons.

For example, all other things being equal, I would not normally have any reason to think that I am travelling through ‘barn-façade’ country, and, hence, I would generally have no good grounds for distrusting my perceptual experiences of barns. If, however, it were known, for instance, that along a certain stretch of countryside along the M4 motorway, there are barn-shaped objects that are not barns, put there as an artistic installation to remind travelling folk of their rural origins, then I ought not to trust that my perception of a barn will, in these circumstances, be veridical (but, again, this does not imply that it is *never* veridical – no ‘aggregate arguments’!). In other words, if the ‘highest com-

16 For this reason, Pritchard (2012: 29–34) distinguishes between ‘good’ and ‘good+’ cases, where seeing that *p* only entails knowing that *p* in the latter, not the former case, i.e. in situations that are both ‘subjectively’ as well as ‘objectively’ epistemically good (cases where one’s perceptual capacities are functioning normally and there is nothing present in the environment – such as, e.g., ‘barn façades’ – that ‘intervenes’ to prevent veridical perception). This taxonomy can also deal with cases, such as the ‘fake fake barn’ scenario discussed by White (2014: 314), where ‘misleading defeaters’ are present (i.e. situations where I am, for example, told by a reliable informant that there are fake barns about, although this is in fact false and I am looking at real barns). For such cases, although objectively epistemically good (i.e. I am seeing that there is a barn there), are not ‘subjectively epistemically good’ (since defeaters are present that undermine knowledge, so that I could only claim to know that *p* in such a situation if I made sure first that the defeaters did not obtain; as long as I am ignorant of this, I cannot know that *p* even though I see that *p*), and, hence, one is not, in such a situation, able to acquire knowledge that there is a barn in front of one, even though there is and one sees that there is. Hence, the ‘entailment thesis’ only holds in good(+) cases: seeing that *p* only entails knowing that *p* if the environment is epistemically friendly, my cognitive faculties are operating normally and no defeaters (misleading or otherwise) are present (also see next footnote). Although I think this is correct, I am not going to adopt Pritchard’s terminology. I will continue to use the ‘good case’ to mean what Pritchard means by the ‘good+ case’.

mon factor thesis' is rejected, I am entitled to rely on some 'background evidence' (of the sort, say, that the Department of Transport doesn't generally go to the trouble of putting up barn-façades); I need not confine myself only to knowledge I have gleaned from my present 'mental state' or perceptual experience.¹⁷

Similarly, unless I have specific reasons for thinking that I might be in a hoaxing zoo (desperate 'austerity' measures, say), seeing that there is a zebra in front of me, when there is, is ordinarily sufficient to know that there is. That is to say, I am entitled to appeal to my 'background knowledge' that zoos don't generally go to the trouble of producing credible fakes; I don't need to confine myself to introspection about the 'zebra-experience' alone which of course could not by itself 'tell' me whether I am confronted by a real zebra or not. (But I could, of course, if needs be, go to the trouble of making special checks, such as checking for paint or other clues of deception.)

4. Conclusion: 'global validation' and radical doubt

In this paper I have shown that once we realize that there are no good independent grounds for accepting the move from (P2) to (C1), we no longer have to concede that RIT is the right way of conceiving, in epistemology, of perceptual reasons. This means that we can accept the thought that, unless misled 'locally', we are in direct touch with the world, and so a 'global' scepticism that suggests that even in the good case our perceptual experiences might *always* fall short of the facts, can no longer get off the ground.

I believe that this is also what Wittgenstein has in mind in a remark in *On Certainty* that often escapes commentators' notice: "So one might grant that Moore was right, if he is interpreted like this: a proposition saying that here is a physical object may have the same logical status as one saying that here is a red patch" (OC 1969: §53). In other words, Wittgenstein believes that physical object propositions and propositions about experiential seemings can be cognitively on a par: we can have the same kind of 'direct access' to physical objects as is commonly thought, on the Cartesian picture, to accrue only to one's experiential seemings (or 'sense data'). The problem with Moore's own account, of course, is that he wants to regard it as a refutation, not a dissolution, of radical

¹⁷ See also Pritchard's discussion of 'favouring' and 'discriminating' epistemic grounds in his 2012.

scepticism, and *qua* ‘refutation’ it appears question-begging. Wittgenstein explains:

The statement ‘I know that here is a hand’ may then be continued: ‘for it’s *my* hand that I’m looking at’. Then a reasonable man will not doubt that I know. – Nor will the idealist; rather he will say that he was not dealing with the practical doubt that is being dismissed, but that there is a further doubt *behind* that one. – That this is an *illusion* has to be shown in a different way (OC 1969: §19; emphasis in original).

The radical sceptic wants a ‘global validation’ of all of our epistemic practices taken together; she wants a proof that we are not wrong about everything all of the time. Contrary to what Moore is trying to do, however, namely, to play the same game as the sceptic, but to come down on the opposite side, Wittgenstein thinks that the notion of ‘wholesale validation’ is not a coherent idea. Since if we suppose, with the sceptic, that *all* of our ‘external world’ propositions taken together could be false, then the words we use to formulate this ‘doubt’, would lose their meaning too: “If you are not certain of any fact, you cannot be certain of the meaning of your words either” (OC 1969: §114).

That is to say, it is not just that, according to Wittgenstein, global scepticism is *unstateable* – but might, for all that, be true – it is rather that, if Wittgenstein is right, it is no longer clear what the sceptic’s ‘thought’ can really amount to (i.e. this is not merely a ‘semantic’ point). For, if, for example, it were conceivable that I were falsely ascribing pain to myself all the time, then it is just no longer clear what ‘pain’ really means.¹⁸ But if the meaning of ‘pain’ is under threat, what, then, would it mean to say, with the sceptic, that “I doubt that there is such a thing as pain”? It is for these reasons that Wittgenstein says that “if you tried to doubt everything you would not get as far as doubting anything. The game of doubting itself presupposes certainty” (OC 1969: §115).

If this is right, then it seems that the kind of justification that the sceptically-minded philosopher is after – a fully general evaluation of our entire rational system – is impossible. But this is not because human powers are unequal to the task. Rather, the very idea seems incoherent: “So is the *hypothesis* possible that all the things around us don’t exist? Would that not be like the hypothesis of our having miscalculated in all our calculations?” (OC 1969: §55) For a ‘miscalculation’ in *all* our calculations is not, as it were, just an ‘aggregated’ mistake, but rather implies that we have never *calculated* at all – that nothing that we have ever done counts as an instance of *calculating*. Similarly, we have never, if the sceptic is right, correctly identified a physical or material object either.

¹⁸ For a similar point see Thomas 2014.

But, if so, the sceptic owes us an explanation of what ‘calculation’ or ‘physical object’ really means, given that we have never, allegedly, seen a single instance of such things. Since such an explanation will either turn out to be question-begging¹⁹ or entirely revisionary,²⁰ it seems that radical sceptical doubt is not only insufficiently motivated (given that it cannot, as I have shown, be constructed out of local sceptical scenarios, and RIT by itself is not independently plausible), but also, if Wittgenstein is right, entirely self-undermining – the sceptic, by calling into question what she is simultaneously relying on, is sawing off the branch on which she is sitting.²¹ Consequently, we seem to have been given no good reason to accept the sceptic’s challenge that an a priori demonstration of perceptual reliability is necessary before good case perception can be taken at face value.

Bibliography

- Bennett, Jonathan: *Locke, Berkeley, Hume: Central Themes*, Oxford 1971.
- Burge, Tyler: *Perceptual Entitlement*, in: *Philosophy and Phenomenological Research*, 67 (2003), 503–548.
- Byrne, Alexander: *McDowell and Wright on Anti-Scepticism, etc.*, in: Dylan Dodd & Elia Zardini (eds.): *Scepticism and Perceptual Justification*, Oxford 2014, 275–297.
- Coliva, Annalisa: *Moore’s Proof, Liberals and Conservatives – Is There a (Wittgensteinian) Third Way?*, in: Annalisa Coliva (ed.): *Mind, Meaning and Knowledge*, Oxford 2012, 323–351.
- Conee, Earle: *Disjunctivism and Anti-Skepticism*, in: *Philosophical Issues*, 17 (2007), 16–36.
- Klein, Peter: *Closure Matters: Academic Scepticism and Easy Knowledge*, in: *Philosophical Issues (Epistemology)*, 14 (2004), 165–184.
- McDowell, John: *Criteria, Defeasibility and Knowledge*, in: *Meaning, Knowledge and Reality*, Cambridge, MA 1998a, 369–394.
- McDowell, John: *Knowledge and the Internal*, in: *Meaning, Knowledge and Reality*, Cambridge, MA 1998b, 395–413.
- McDowell, John: *The Disjunctive Conception of Experience as Material for a Transcendental Argument*, in: *The Engaged Intellect*, Cambridge, MA 2009, 225–242.
- McGinn, Colin: *Wittgenstein on Meaning*, Oxford 1984.
- Millar, Alan: *Reasons and Experience*, Oxford 1991.
- Nagel, Thomas: *The View from Nowhere*, Oxford 1986.
- Pollock, John: *Knowledge and Justification*, Princeton 1974.

19 E.g. along the lines of ‘a physical object’ is what you would be perceiving when looking at the table in front of you, if an a priori demonstration could first be given that you are not in an LE or NBE scenario.

20 I.e. the sceptic would be revising our ordinary concepts, say, in the direction of phenomenalism, for example. For more on this, see Strawson 1964.

21 For more on these themes, see Schönbaumsfeld 2016.

- Pritchard, Duncan: McDowellian Neo-Mooreanism, in: Adrian Haddock & Fiona Macpherson (eds.): *Disjunctivism: Perception, Action, Knowledge*, Oxford 2008, 283–310.
- Pritchard, Duncan: Wright contra McDowell on Perceptual Knowledge and Scepticism, in: *Synthese*, 171 (2009), 467–479.
- Pritchard, Duncan: *Epistemological Disjunctivism*, Oxford 2012.
- Pritchard, Duncan: Entitlement and the Groundlessness of our Believing, in: Dylan Dodd & Elia Zardini (eds.): *Scepticism and Perceptual Justification*, Oxford 2014, 190–212.
- Pryor, Jim: The Skeptic and the Dogmatist, in: *Noûs*, 43 (2000), 517–549.
- Schönbaumsfeld, Genia: *The Illusion of Doubt*, Oxford 2016.
- Strawson, Peter Frederick: *Individuals*, London 1964.
- Stroud, Barry: *The Significance of Philosophical Scepticism*, Oxford 1984.
- Thomas, Alan: McDowell on Transcendental Arguments, Scepticism and Error Theory, in: *International Journal for the Study of Skepticism*, 4(2) (2014), 109–124.
- Thorpe, Joshua Rowan: Closure Scepticism and the Vat Argument, in: *Mind* (forthcoming).
- White, Roger: What is my Evidence that here is a Hand?, in: Dylan Dodd & Elia Zardini (eds.): *Scepticism and Perceptual Justification*, Oxford 2014, 298–321.
- Wright, Crispin: Anti-Sceptics Simple and Subtle: G. E. Moore and John McDowell, in: *Philosophy and Phenomenological Research*, 65 (2002), 330–348.
- Wright, Crispin: Comment on John McDowell, in: Adrian Haddock & Fiona Macpherson (eds.): *Disjunctivism: Perception, Action, Knowledge*, Oxford 2008, 390–404.
- Wright, Crispin: On Epistemic Entitlement (II): Welfare State Epistemology, in: Dylan Dodd & Elia Zardini (eds.): *Scepticism and Perceptual Justification*, Oxford 2014, 213–247.