Knowledge and Perceptual Consciousness

Alasdair Craig

Lady Margaret Hall, University of Oxford

D.Phil Philosophy

Trinity Term 2022

Contents

Acknowledgements

Introduction

Chapter 1. Reasons outside thought

Chapter 2. Asking questions, giving instructions

Chapter 3. Thoroughgoing externalism

Chapter 4. Defeat

Chapter 5. Easy knowledge

Bibliography

Acknowledgements

I have been privileged to be supervised by Timothy Williamson and Mike Martin. I could not have asked for two more supportive, encouraging, generous, inspiring or insightful supervisors. I am extremely fortunate to have had the opportunity to study under their guidance. I am also eternally grateful to my family, and to friends who have supported and cheered me throughout by studies. Alesia Preite, Sapna Sinha, Vânia Pinto and many others have been constant sources of fun, energy and conversation. I gratefully acknowledge the Clarendon Scholarship for three years of funding.

Introduction

Perception is where thought, consisting of person-level propositional attitudes, confronts and is influenced by parts of the non-mental world that it has not antecedently embraced. When I come to believe that Robin is now in New York on the basis of my knowledge that the flight takes six hours and that she took off seven hours ago, the movement of my thought is guided or influenced in a rationally intelligible way by worldly facts that hold independently of my thoughts. But these are facts that I have already come to know and so already embraced in thought. They are therefore reasons that I possess and am in a position to base my beliefs on. In contrast, when I see a tomato on the table and form a perceptual belief that there is a tomato on the table, this movement of thought is influenced by a way the world is that I did not antecedently know or believe to be so.[1] Furthermore, when we introspect on our perceptual beliefs, it does not seem to us that they rest on prior beliefs that we have. At least on the face of it, there is no antecedent thought which serves as my reason for my perceptual belief that there is a tomato on the table. If this is right then perceptual beliefs and knowledge are non-inferential, since inferential belief- or knowledge-formation involves a rational transition between beliefs or states of knowledge.

[1] I'll be focusing on visual perception throughout.

Some qualifications and clarifications are in order. First, even if perceptual knowledge is non-inferential, antecedent thoughts may still play various roles in perceptual belief-formation. For example, knowledge of certain facts might ground the possession of concepts that I need in order both to see the tomato as a tomato and to form the belief that there is a tomato on the table. But these thoughts play the role of background enabling conditions rather than reasons on which I base my perceptual belief. Insofar as my perceptual belief is influenced by whatever aspect of the external world perception puts me in touch with, it is influenced by something that thought has not yet taken in, at least until the point at which I have formed the perceptual belief in question and thereby embraced in thought the relevant aspect of the external world.

Second, on some theories of perceptual knowledge – such as some classical foundationalist theories – our perceptual knowledge of the external world is inferred from knowledge of appearances or sense-data or some other mental entities. There is thus a sense in which, on such a picture, perceptual knowledge is inferential. But here the beliefs most immediately influenced by perception are not about the external world but rather about something like appearances. *These* beliefs are non-inferentially formed, since the relationship that holds between them and pre-cognitive perception is clearly not inferential.

Third, it should be admitted that it could be the case that when the non-mental world influences our thought through perception it does so by working in tandem with beliefs that we already have, and which do serve as reasons, so that perceptual beliefs do have antecedent beliefs as reasons. If this is right, then there is a sense in which perceptual knowledge is inferential after all. In fact, though, this is not how perceptual belief-formation strikes us. Pre-theoretically, it is quite unclear what antecedent beliefs our perceptual beliefs might be based on, which makes it difficult to see how such antecedent beliefs could be *our own* reasons for our perceptual beliefs. In any case, even if this is, despite appearances, how perceptual belief-formation works, it is still the case that some portion of the world is, through perception, influencing our thought in a way that cannot be assimilated to how the world can influence thought via inference, where thoughts lead to other thoughts by being reasons for them.

With these qualifications in place, what I have said is pretty uncontroversial. I have not committed to any claims about the structure of justification: coherentists, foundationalists and infinitists can accept that perceptual beliefs are, at least, *causally* influenced by aspects of the world that one has not yet embraced in one's thinking. But it is difficult to go much further in characterising the perception-thought interface without encountering problems and controversies. A central problem is to accurately characterise the perceptually mediated relationship between world and thought. In particular, is this relationship *rational*, in

something like the way that the relationship between inferential knowledge and the knowledge that the inference is drawn from is rational?

According to BonJour, introspection makes it quite clear that perceptual belief-formation is 'cognitively spontaneous':

> I do not *infer* that there is a red book on the desk, nor does the belief result from any other sort of deliberative or ratiocinative process, whether explicit or implicit. Rather it simply occurs to me, 'strikes me', in a manner which is both involuntary and quite coercive; such a belief is, I will say, *cognitively spontaneous*. It is cognitive spontaneity which marks the belief as putatively observational.[2]

Compare also Austin. If I look at a pig, he says,

> There is no […] question of collecting evidence; its coming into view doesn't provide me with […] *evidence* that it's a pig. I can now just *see* that it is, the question is settled.[3]

These passages illustrate something of the challenge we face in specifying the perceptually mediated relationship between mind and world. There

[2] BonJour 1985: 117.

[3] Austin 1962: 115.

is surely something right in both passages. There is a *sense* in which perception coerces belief; a sense in which perceptual beliefs simply 'strike' one; a sense in which I 'just see' that things are so. But there is also a sense in which these terms are inapt. My perceptual beliefs do not simply arise, as if from nowhere, in the way that thoughts sometimes spontaneously occur to me. The point is made vivid by considering the phenomenon of blindsight: patients claim to be blind, yet if forced to guess give reliably true answers to questions about objects placed in their blind field. To the blindsighter, these answers just strike them, as though out of nowhere. Normal perceptual knowledge-formation is not like that. It comes with a phenomenal character – there is something it is like to perceive. And this phenomenal character does not seem to us to be rationally inert or incidental; it is not random or unrelated to what we see and believe. But we struggle to articulate its role, rational or otherwise, in belief-formation. The most we can say is that it seems to consist in the objects and properties that we see.

Furthermore, there is also a sense in which we find ourselves coerced into having non-perceptual, *inferential* beliefs; and there is a sense in which inferential beliefs just strike us. Presented with a simple *modus ponens* argument, involving premisses I believe, I can't help but come to believe the conclusion. So it's not clear how far the passage cited from BonJour really goes towards distinguishing perceptual from inferential knowledge.

In Chapter 1 of this dissertation, I will describe some reasons for thinking that perceptual knowledge stands in a *rational* (and not merely a causal) relation to those features of the external world which it is about and which, prior to one's coming to have the relevant perceptual knowledge, one had not yet embraced in thought. This makes available an obvious way to differentiate between inferential and perceptual knowledge. The key difference is that inferential knowledge is rationally based on a fact which one had antecedently embraced in thought by coming to know, whereas perceptual knowledge is rationally based on a fact or some other kind of worldly entity – perhaps simply the objects and properties that we see – which one had not antecedently embraced in thought. That is, perceptual knowledge is based on what I call 'reasons outside thought'. Chapter 1 describes John McDowell's motivation for and development of this idea about perceptual knowledge, and it describes a popular line of argument according to which this idea must be wrong. According to this line of argument, one can only *possess* a reason for belief, such that one is in a position to form beliefs on the basis of this reason, by embracing the reason in thought. That is, the reasons one possesses are the facts that one knows or, perhaps, the contents of one's beliefs. In Chapters 2 and 3 I defend McDowell against this line of argument. While the picture of perceptual knowledge that I endorse is not exactly McDowell's, it is pretty close. In a nutshell, the main point that I make is that if we want to make plausible the idea of perceptual knowledge being based on reasons that lie outside thought, we must embrace what I call a thoroughgoing externalism. If we want to make

sense of reasons figuring in our perspective as the reasons they are, ready for us to respond to as such, but without being the contents of our thoughts, then we need to allow that the external world itself can constitute portions of our perspective by constituting the phenomenal character of our perceptual experiences. This implies that popular contemporary views about perceptual knowledge and justification, according to which perceptual experiences are non-factive, non-doxastic reasons or evidence for our perceptual beliefs cannot be right.

Chapters 4 and 5 defend the thoroughgoing externalism about reasons and evidence that I endorse from two objections, and provide a further positive argument for it. The message of this dissertation is thus that the phenomenal character of pre-doxastic perceptual experience is constituted by those worldly features which our perceptual knowledge is knowledge of, and that this enables it to provide a rational constraint on thought from outside thought.

Chapter 1

Reasons outside thought

To repeat, perception is where thought meets aspects of the world that it has not yet embraced. When things go well, this meeting leads to thought incorporating new aspects of the world, apparently without the need for the rational work that goes on in inference. There are two ideas about perceptual knowledge which I want to vindicate over the next three chapters. First, perceptual knowledge-formation is a rational process and involves the subject responding to *reasons* for belief. Second, the process is non-inferential, so these reasons for belief are *non-doxastic*: one possesses them and is in a position to respond to them as such in virtue of non-doxastic, perceptual-experiential mental states that one is in.

One reason for thinking that perceptual knowledge must be non-inferential is the traditional regress problem about justification. Assume that only justified or knowledgeable beliefs can justify further beliefs. Then if all justified or knowledgeable beliefs are justified by other beliefs, so that all justified and knowledgeable beliefs are inferential, there must either be an infinite regress of justified or knowledgeable beliefs or a closed circle of justified or knowledgeable beliefs.[4] Infinitism and

[4] See Smithies 2014 for an overview of the regress problem and the foundationalist solution.

coherentism bite these respective bullets. Foundationalism rejects the claim that all justified and knowledgeable beliefs are inferential. On the assumption that foundationalism is the most attractive of these three positions, we need to give an account of non-inferentially justified or knowledgeable beliefs. Since perception is where thought meets the world, perceptual knowledge seems a good candidate to be non-inferential.

This is by no means the only reason for thinking perceptual knowledge (and perceptually justified belief) is non-inferential. Much of the recent debate about whether perceptual knowledge or perceptually justified belief is inferential or non-inferential has coalesced around two positions labelled 'Liberalism' and 'Conservatism'. Formulating these positions is not straightforward but, roughly, according to Liberals one can have propositional perceptual justification to believe an external-world proposition purely in virtue of undergoing a non-doxastic perceptual experience; the resulting propositional perceptual justification is *pure* in the sense that it is constituted solely by the justification provided by the perceptual experience, all by itself. Put another way, perceptual experiences provide *immediate*, non-inferential justification to believe their contents. Conversely, according to Conservatives whenever a perceptual experience makes one have propositional perceptual justification to believe an external-world proposition, this justification is *impure*, in the sense that part of what constitutes it is one's propositional justification from a non-perceptual source to believe the negations of relevant

sceptical hypotheses (such as: 'you are now dreaming').[5] It's not straightforward to say which propositions count as relevant sceptical hypotheses; the key point is just that the Conservative thinks perceptually justified beliefs are based not solely on one's perceptual experience but also on one's antecedently justified belief that one is not hallucinating, dreaming, a brain-in-a-vat, and so on. Most of the participants in the debate between Conservatives and Liberals are fallibilists, who think that perceptual justification for beliefs about the external world is non-conclusive and that perceptual experiences are non-factive, not essentially world-involving mental states or episodes. Liberalism has been defended on a variety of grounds. According to Pryor, it is simply the obvious default position:

> An experience as of there being hands seems to justify one in believing there are hands in a perfectly straightforward and immediate way. When asked, 'What justifies you in believing there are hands?' one is likely to respond, 'I can *simply see* that there are hands.' One might be wrong: one might not really be seeing a hand. But it seems like *the mere fact* that one has a visual experience of that phenomenological sort is enough to make it reasonable for one to believe that there are hands. No *premises* about the character of one's experience […] seem to be needed.[6]

[5] Wright 2002; Wright 2007; Wright 2008.

[6] Pryor 2000: 536. Pryor's emphases.

Furthermore, it is surely psychologically unrealistic to think that our perceptual beliefs are based on prior beliefs in the negations of sceptical hypotheses or, indeed, on prior beliefs about anything. Therefore, as Silins says, 'A major advantage of Liberalism is that it is psychologically undemanding and straightforwardly compatible with the well-founded status of perceptual beliefs'.[7] Finally, also according to Pryor, Liberalism undercuts Cartesian scepticism.[8]

While all three of these considerations seem to me to be good reasons to think that perceptual knowledge (and perceptually justified belief) is non-inferential, it's less clear to me that they are good reasons to accept Liberalism (though in fact I think Liberalism is true). The same goes for the regress problem. One can think that perceptual knowledge is non-inferential without thinking that it is based on perceptual experiences, conceived of as reasons or evidence for perceptual beliefs. The fact that I say 'I can simply see that I have hands' doesn't show that I'm *basing* my belief that I have hands on my perceptual experience – if anything, 'simply see' suggests a lack of basing. A non-basing view would be equally or even more psychologically undemanding, and would forestall Cartesian scepticism at least as effectively as Liberalism. So why should we think that non-doxastic perceptual experiences justify and rationalise beliefs?

[7] Silins 2007: 119.

[8] Pryor 2000; Pryor 2004.

McDowell suggests an answer. According to him, we must be able to make sense of the idea that thought is normatively and rationally constrained by features of the world that lie *outside thought* in order to make sense of the idea that our thoughts are about that world.[9] In short, he thinks that the intelligibility of intentionality relies on the idea of rational constraint from outside thought. And he thinks that perception is where rational constraint on thought by aspects of the world which lie outside thought takes place.[10] His account of perceptual knowledge tries to make sense of this idea. Central to this account is the idea that perceptual experiences, conceived of as factive, world-involving mental states which lie outside thought, justify perceptual beliefs, or rationally ground perceptual knowledge.

McDowell's view is thus that perceptual knowledge is both non-inferential and based on experiential reasons. In this chapter I will first (sections 1.1 and 1.2) sketch his picture of perceptual knowledge. He develops this picture in reaction to what he takes the required rational constraint of the world on thought to involve, so I will also sketch his development of this notion of rational constraint. Although my own picture of perceptual knowledge will differ in several details from that of McDowell, it is helpful to have his view on the table, since the view I

[9] McDowell 1996: 8. Cf. McDowell 1986; Brewer 1999. For critical discussion of exactly what kind of rational constraint from 'outside thought' is required for making empirical content intelligible see, e.g., McCulloch 2002; McGinn 2009; Ginsborg 2018; and the debate between McDowell and Travis over a series of papers (the latest exchange is in Gersel et. al. 2018).

[10] McDowell 1996: xii.

endorse is similar to his. Furthermore, his view has been subjected to criticisms which any view of the kind I want to endorse must reckon with. Since these criticisms have been directed specifically at McDowell, to engage smoothly with them it will be helpful to have his view on the table.

It will quickly become clear that the central difficulty in sustaining his picture of perceptual knowledge is his combination of a kind of internalism about reasons with his insistence that we can possess reasons that lie outside thought. In sections 1.3–1.5 I will set out Ginsborg's argument that it is only by believing p that one is in a position to form beliefs for the reason that p. Her central contention is that believing p is what makes p figure in one's perspective as a reason for belief, such that one can respond to it as the reason that it is. Consequently, rationalising explanations of a subject's beliefs and actions must appeal to their beliefs.

My main aim in this chapter is simply to explicate McDowell's view and Ginsborg's critique of it. I will, however, point out some of the weaknesses in Ginsborg's own positive position. In particular, we will see in section 1.3 that her own account of what is involved in a reason figuring as such in a subject's perspective has important limitations. This prepares the way for Chapter 2, where I offer my own account of what it means for a reason to figure as a reason in one's perspective. Furthermore, we will see that Ginsborg assumes that only non-factive

mental states can figure in rationalising explanations of beliefs and actions. This point prepares the way for Chapter 3, where I argue that this assumption is what leads to the difficulty in conceiving of reasons outside thought figuring as such in a subject's perspective. I turn now to my explication of McDowell's picture of rational constraint and perceptual knowledge.

1.1. McDowell on rational constraint and perceptual knowledge

To say that the world rationally constraints thought is in part to say that the world makes certain thoughts appropriate or inappropriate, so it is in part to say that the world *normatively* constrains thought. McDowell's claim is that thoughts must be made appropriate or inappropriate by whatever they are about. Beliefs are a case in point (and are the thought-kind we'll focus on[11]): 'thinking that aims at judgement, or at the fixation of belief, is answerable to the world – to how things are – for whether or not it is correctly executed'.[12] So for my belief that you are tall to be about you and your height requires that you and your height are the measure of my belief's appropriateness.

[11] Following McDowell, I concentrate on judgement, belief and knowledge. As he says, 'If we can make sense of judgement or belief as directed towards the world […], other kinds of content-bearing postures or stances should easily fall into place.' Ibid., xii.

[12] Ibid., xii.

But there is more than this to the idea of rational constraint as it contributes to making intentionality intelligible, according to McDowell. In particular, this rational constraint must come in the first place from reasons that lie outside thought; and yet, at the same time, these reasons which lie outside thought must also lie inside the subject's perspective, as the reasons that they are. There is an obvious tension here, between seeing these reasons as external to thought yet internal to the subject's perspective. It's this tension that these first three chapters of the thesis are most concerned with. I'll now describe how McDowell develops the idea of rational constraint of thought from outside thought, with a focus on what he means by saying that perceptual reasons are external to thought but also internal to the thinker's perspective.

Rational constraint 1: The primacy of the 'good' case

To say that our beliefs are rationally constrained by the world is partly to say that the world really does influence, in a rationally intelligible way, the beliefs we have about it. We can put this point by saying that McDowell holds that, for intentionality to be intelligible, we must think of the world's normative import as genuinely operative in guiding or controlling our beliefs. Of course thinkers can make mistakes, sometimes forming beliefs that are not properly responsive to the world. But McDowell is suggesting that we make sense of the thinker's beliefs'

intentionality, whether these beliefs are correctly or incorrectly executed, in terms of the 'good' cases where she does successfully respond to the world's normative import by forming appropriate, *knowledgable* beliefs.[13] But the central question for McDowell is how this responsiveness to the world's normative import gets going. This takes us to the picture of perceptual reasons as both external and internal. First, what does it mean to say that perceptual reasons lie outside thought?

Rational constraint 2: Reasons Outside Thought

For McDowell, responsiveness to the world's normative import first gets going in perception and perceptual knowledge, construed as non-inferential. When I come to believe that it rained for the reason that, as I know, the streets are wet, this inferential movement of thought is normatively and rationally constrained by the world, and specifically by the fact that the streets are wet. However, McDowell thinks that this inferential rational constraint can only be understood in terms of 'inferential linkages' to non-inferential knowledge (inferential knowledge has 'mediately empirical content', as he puts it). We might put the point by saying that the world first gets into our thought through non-inferential rational constraint.[14]

[13] On this point, see McDowell 2018: 90.

[14] McDowell 1996b: 296. Cf. McDowell 2018: 90.

Non-inferential rational constraint is fundamental and prior to inferential rational constraint because in order to understand thought being about the world we need to be able to see that 'when we use our concepts in judgement, our freedom – our spontaneity in the exercise of our understanding – is constrained from *outside thought*, and constrained in a way that we can appeal to in displaying the judgements as justified.'[15] The fundamental problem with coherentism, according to which (in Davidson's slogan) 'nothing can count as a reason for holding a belief except another belief'[16], is that it doesn't make room for the idea of reasons which lie outside thought and yet which, at the same time, thought can be responsive to as such. To the extent that normative reasons lie outside thought, for the coherentist, they are not motivating or explanatory reasons. The result is thought as a series of 'moves in a self-contained game' – thought not constrained by anything outside of it.[17] This is supposed to prevent us from understanding how thought can be about the world outside thought.

Why exactly does our understanding of intentionality require us to see ourselves as having reasons that lie outside thought? It's difficult to find a

---

[15] McDowell 1996: 8. My emphasis.

[16] Davidson 1983/2001: 141. See also Davidson's afterword, in which he disclaims any endorsement of standard coherence theories of truth and says: 'My emphasis on coherence was properly just a way of making a negative point, that "all that counts as evidence or justification for a belief must come from the same totality of belief to which it belongs".' Davidson 1983/2001: 155.

[17] McDowell 1996: 5.

straightforward answer to this question in what McDowell says.[18] However, the problem, as far as I can see, is that if a subject could never find any reasons for belief beyond the contents of her own beliefs and states of knowledge then she could not appreciate that her reasons are facts (or some other kind of thing) which obtain whether or not she believes them to be so, and so she could not really appreciate them as having normative import with respect to what she should believe about the external world. That would render genuine rational constraint by the world on thought impossible, and this would render the subject unable to form beliefs about the world.

This way of putting the point suggests that 'reasons outside thought' must be reasons that are not the contents of our own thoughts, and yet which we can respond to as the reasons that they are. We can formulate this requirement as follows. For thought about a thought-independent world to be intelligible, we must be able to see the following condition as being met for at least some reasons for beliefs that we possess:

[18] See especially ibid., 16–18 and footnote 14 on p. 17 for arguments which I find difficult to fully understand. See Byrne 1996 for an expression of puzzlement about what McDowell's argument is meant to be here, and McDowell 1996b for a reply to Byrne. In any case, I find the argument that I sketch here forceful, whether or not it is McDowell's own.

> **Reasons Outside Thought**: Whatever propositional attitude (i.e. whatever thought) S does or does not adopt towards R, R counts as a reason she possesses.[19]

This is where non-inferential rational constraint comes in. If the Reasons Outside Thought condition is ever to be met, we must have some non-doxastic way of possessing reasons. Whatever this is, any beliefs that we form on the basis of these non-doxastic reasons will be non-inferential and non-inferentially rationally constrained. McDowell's suggestion is that it is by undergoing perceptual experience that we possess these reasons non-doxastically.

Rational constraint 3: Internalism, part one

The tension in this picture is clear. We must have reasons that lie outside our thinking, but they must be reasons which we *possess*. A standard way of understanding what it means to come to possess a reason is to say that the reason is incorporated into one's thought. If the reason is a fact, one might come to possess it as a reason by coming to know it. But reasons that lie outside thought cannot be possessed in this way. And yet McDowell insists that we must give positive content to the idea that they

[19] This formulation might be thought to imply that reasons must be propositions. That is not what is intended. Reasons may, for example, be everyday worldly objects. If S could come to possess such a reason then it would automatically fulfil Reasons Outside Thought, since it is not the kind of thing that S could have a propositional attitude towards.

are possessed. As he puts it, rational constraint by the world on thought 'require[s] the status as a rational constraint of what is responded to to be in the view of the responder'.[20] Rational constraint requires that the normative significance of features of the world figure 'in the view of the responder' as such, and in such a way that the thinker can respond to it by forming the appropriate beliefs. This is McDowell's internalism.

We can introduce two pieces of vocabulary in order to talk about (i) features in the world being such that they can figure in the perspective of a person *as* having normative significance, and (ii) the subject's capacity to recognise and respond to the normative significance of features in the world as such. Following Raz, let's say that 'our *rational powers* are a general capacity to recognise and respond to facts that make certain responses appropriate, and such facts are [normative] *reasons* because they can be recognised and responded to by our rational powers'.[21] McDowell, like Raz, thinks that only facts can be normative reasons, but it may be that other kinds of thing can be normative reasons too.

In any case, the first part of McDowell's internalist requirement can now be stated as the requirement that for a subject to form a belief in a way that is rationally constrained by the world requires the subject to form the belief *for* a reason, and this in turn requires the subject to *know* what

[20] McDowell 1996b: 293.

[21] Raz 2011: 86. My emphases.

their reason is for believing as they do.[22] Someone who knows what their reason for believing p is can cite that reason if prompted. Since reasons just are aspects of the world that make certain beliefs appropriate, someone's knowing what their reason is for believing p involves their knowing that this reason makes appropriate believing p: they recognise the reason *as* a reason, and thus as something that makes their belief appropriate.

The point can be illustrated in terms of inferential belief-formation. Suppose that I come to know that it recently rained on the basis of my knowledge that the streets are wet. What does it take for my knowledge that the streets are wet to count as normatively constraining my belief that it recently rained? According to McDowell's internalist requirement as described so far, I must *recognise* (and thus *know*, perhaps just tacitly) the normative significance of the fact that the streets are wet. I must, having formed my belief, be disposed if challenged to *cite* the fact that the streets are wet as my reason for believing that it recently rained. In short, I must know what my reason for so believing is.

Notice that so far the internalist requirement kicks in only *after* one has formed a belief for a reason. *Once* one believes p for a reason, one must know what one's reason for believing p is, and one must be able to cite it. But if we are to understand one's belief as a *response* to a reason, as *constrained by* that reason, then we surely need to specify some sense in

[22] 'Someone who makes a judgement knows, at least implicitly, not only what she judges but on what ground she judges it.' McDowell 2018: 91.

which that reason figured as a reason in one's perspective before one formed the belief. Furthermore, we cannot appeal to the idea that one knows what one's reason is in order to give content to the idea that one possesses a reason that lies outside thought, since knowing what one's reason is involves having incorporated that reason into one's thought. Thus McDowell expands his internalist requirement.

Rational constraint 4: Internalism, part two

We need to be able to see ourselves as possessing and responding to reasons which figure as such in our perspective on the world without yet being embraced by thought. The internalist requirement as so far stated does not secure this, since it kicks in only at the point at which we already have embraced the relevant reason in thought by coming to believe that it is our reason for such-and-such a belief. We therefore need to add to the internalist requirement to give content to the idea that we possess reasons *before* they enter into our thought, such that they can be *our* reasons for belief.

The first part of the internalist requirement said that for a subject to form a belief in a way that is rationally constrained by the world requires the subject to form the belief *for* a reason, and this in turn requires the subject to *know* what their reason is for believing as they do. Suppose that the relevant belief is that it recently rained, and the reason is that the

streets are wet. We can now extend the internalist requirement in the following way. For the subject's belief that it recently rained to count as rationally constrained by the relevant feature in the world *also* requires that had they *not* formed the belief, they would still have been *in a position to*

> (i) form the belief that it recently rained for the reason that the streets are wet
>
> and
>
> (ii) recognise the normative significance of the fact that the streets are wet, and in particular that this fact is a reason to believe that it recently rained.

We can say that to be in such a position is to possess a reason to believe that it recently rained. That is, one possesses a reason R to believe p only if one is in a position to recognise R's normative import with respect to p and to believe p on the basis of R.

There doesn't seem to be any difficulty in meeting this further internalist condition in cases of inference. Had I not formed the belief that it recently rained, I would still know that the streets are wet, and this knowledge would ensure that I am in a position to believe that it recently

rained for the reason that the streets are wet, and to recognise the normative significance of the fact that the streets are wet.

What about the non-inferential case? If my non-inferential belief is not based on a prior belief or piece of knowledge, then in virtue of what can I count as possessing a reason for that non-inferential belief in a case in which I do not form the non-inferential belief? This is where perceptual consciousness or experience comes into the picture. It is here that humans 'confront the world' in the first place, and so this is the obvious place to look for non-inferential rational constraint.[23] Furthermore, perceptual consciousness constitutes part of our perspective on the world, and yet does not seem to be part of our thought about the world, but seems rather to lie outside thought. Thus this is the obvious place to look for the possession of reasons outside thought. McDowell thus suggests that perceptual experience is where the world affords us reasons for belief by presenting itself to us as recognisably having normative import for what we should believe. He calls this '"a minimal empiricism": the idea that experience must constitute a tribunal, mediating the way our thinking is answerable to how things are'.[24]

If McDowell is right then we need to conceive of perceptual experience in such a way that having a perceptual experience puts one in possession of a reason that lies outside thought. Experience must itself thus be non-

[23] McDowell 1996: xii.

[24] Ibid.

doxastic, and it needs to constitute or contain reasons that the subject is in a position to form beliefs on the basis of. This is a potentially tricky balancing act: experiential reasons must lie outside thought yet also be internal to the subject's perspective in such a way that she can count as possessing them. To meet the full internalist requirement, experience must put the subject in a position to recognise experiential reasons as the reasons they are, and to form beliefs on their basis. But to fulfil the Reasons Outside Thought condition, experiences must do this without requiring the subject to take any view with respect to the truth of the facts that constitute the relevant reasons.

McDowell's view is that perceptual experiences can pull this off. For, first, they are indeed non-doxastic: they do lie outside thought. But they also put us in possession of reasons, fulfilling both parts of his internalism. They fulfil the first part of his internalism because, he says, we do cite our experiences as reasons for belief in everyday conversation:

> Someone who makes a judgement knows, at least implicitly, not only what she judges but on what ground she judges it. […] In the case of grounds by virtue of which judgements are *perceptually* knowledgeable, the point comes out in this: if someone who has expressed a bit of perceptual knowledge is asked 'What is your ground for that?', or, equivalently, 'How do you know that?', she

> can answer straight off, without further investigation, 'I perceive (have perceived, can perceive) that it is so'.[25]

So, according to McDowell, when one perceptually knows that there is a tomato on the table, one's reason for belief is *that I perceive that there is a red tomato on the table* or, as he also puts it, 'that [my] experience makes it manifest to [me] that [there is a tomato on the table].'[26] When one comes to perceptually know that there is a tomato on the table, one *also* comes to know both that one is perceiving or seeing that there is a tomato on the table, and that this is one's reason for believing that there is a tomato on the table. This way of putting things creates a problem for McDowell's view, often called 'the basis problem'.[27] The problem is that seeing that p is a way of knowing that p.[28] Thus if one's perceptual reason for believing that p is the fact that one sees that p, then one's

---

[25] McDowell 2018: 91. McDowell's emphasis.

[26] Ibid. See also McDowell 2006, where he clarifies this point. Some philosophers have interpreted McDowell's view to be that one's reason for one's perceptual belief that p is p itself. In this reply to Dancy, McDowell sets the record straight: for one of Dancy's arguments to work, says McDowell, it would have to be the case that 'when one forms a belief on the basis of the kind of experience that constitutes being perceptually open to a fact, it is the fact itself that is one's reason for believing what one does. […] Dancy attributes just that claim to me, including a denial that one's reason is – what would be a different fact – that one experiences the fact that makes one's belief true. But this is simply wrong. When I believe, on the basis of my visual experience, that there is a hummingbird at the feeder, my reason for believing is precisely that I see it to be so – that I experience, visually, its being so – or at least seem to. If I am asked to give my reason for believing that there is a hummingbird at the feeder, it would be absurd to respond by simply restating what I believe, and the absurdity is not lessened if in doing so I would be stating a fact' (134).

[27] The label is from Pritchard 2012.

[28] See Dretske 1969; Williamson 2000; French 2016.

perceptual reason for believing p presupposes that one (perceptually) knows that p. Clearly, this can't be the reason that grounds one's perceptual knowledge, since the reason itself just is the fact that one has perceptual knowledge. One alternative is to replace talk of seeing-that with talk of what French calls 'thing-seeing', or the kind of state we appeal to when we say things like 'I saw the dog on the armchair'.[29] Thing-seeing is distinct from seeing-that: I might see the dog on the armchair without seeing that the dog is on the armchair. A potential drawback of appealing to thing-seeing in this context is that one might see an object in such a way that one is unable to recognise it for what it is; this complicates how we should think about the rational import of thing-seeing. So there is a difficulty for anyone who wants to argue that perceptual knowledge is based on non-doxastic perceptual-experiential reasons, which is that it's not straightforward to specify a non-doxastic perceptual state which we refer to as a reason in everyday life. Let's side-step this problem for the moment, and just grant McDowell that there is 'seeing-that' state which is non-doxastic, and in particular does not involve knowing (I come back to this issue in Chapter 3, where I suggest that we replace McDowell's seeing-that state with a relational, object-seeing state).

So much for the first part of McDowell's internalism. The second part of his internalism was that it must be the case that had one not formed the relevant belief (in, say, p), one would still possess the relevant reason,

[29] See French 2016; Dretske 1969.

where possession of the reason involves being in a position to come to believe p *for* that reason, and being in a position to recognise the normative import of that reason, which is just to come to recognise that reason as a reason to believe p. McDowell says that this second part of the internalist requirement is also met by perceptual experience:

> The rational significance of an experience of perceiving is that it is potentially a ground for perceptually knowledgeable judgements […]. Suppose an experience has as part of its rational significance that it *could* ground a knowledgeable judgement that things are some way; even if that potential for grounding a judgement is not actualised in knowledge that things are that way, the availability for knowledge in self-consciousness that characterises an act, even a partial act, of the capacity for knowledge through perception, as the act of rationality it is, includes a potential for knowing that aspect of the experience's rational significance. […] Even if the potential for grounding some knowledgeable judgement that is part of the rational significance of an experience of perceiving is not actualised, the subject of the experience still has it in her power, just in undergoing the experience, to know that it has a rational significance that includes making it manifest to her that things are as she would judge them to be in that judgement.[30]

[30] McDowell 2018: 91–2. Cf. McDowell 2002: 98; McDowell 2003: 681.

Consider the case where you see a tomato on the table in good viewing conditions and have a perceptual experience which presents to you the fact that the tomato is on the table. Suppose you do not come to believe that there is a tomato on the table, nor that you perceive that there is a tomato on the table (nor that the fact that there is a tomato on the table is manifest to you). Perhaps you think you're hallucinating, or you've just been convinced by a sceptic that knowledge is impossible, or you're not paying attention to what you see as you daydream about what's for lunch. McDowell's claim in the quoted passage is that your perceptual experience still puts you in a position to believe that there is a tomato on the table for the reason that you see that there is a tomato on the table, and to know that this experience (or the fact that you're having it) is a reason to believe that there is a tomato on the table. So experience fulfils the second part of the internalist requirement.

What is it about experiential reasons that enable them to fulfil both parts of the internalist requirement while being non-doxastic? First, for McDowell it is essential that they have conceptual contents, which in 'good' cases are simply worldly facts.[31] Second, in 'good' cases of normal

[31] McDowell holds that only states or episodes with conceptual contents can be appreciated by a subject as having rational import. In my view we would do better, when it comes to understanding the rational role of experiences, to conceive of experiences as lacking contents, but let's go along with McDowell to see what we can make of his view. For arguments that only states with conceptual content can be or contain reasons see Brewer 2005; Ginsborg 2011; Hawthorne and Magidor 2018; cf. also Williamson 2000 for the view that evidence is propositional. See Sylvan 2016 for an overview of recent discussion of the ontology of reasons. See Kalderon 2011; Brewer 2011; Brewer 2018; Travis 2013; Logue 2018 for views on which experiences are or contain reasons while lacking conceptual content.

perception, experiences are *factive*: they are a kind of mental state which can only have facts as contents. A corollary of this is that perceptual reasons are *conclusive* reasons, which put the subject in a position not just to form beliefs on their basis, but to *know*. This aspect of McDowell's thinking is partly motivated by epistemological considerations: he denies the possibility of knowledge that is based on less than conclusive reasons. He thinks that the idea that perceptual experiences are less-than-conclusive reasons thus leads to scepticism.[32] However, the factivity of perceptual experiences also plays a more fundamental role in McDowell's thinking: in Chapter 3 I will argue that it is a crucial part of why experiences can be reasons that lie outside thought and yet simultaneously reasons that the subject possesses and is in a position to form beliefs on the basis of.

That concludes my interpretation of McDowell's picture of perceptual knowledge and his motivation of this picture in terms of the idea that we must make sense of the rational constraint on thought of a world that lies outside thought in order to make the intentionality of thought intelligible. So far in this chapter, I have introduced the idea that perceptual knowledge is non-inferential and yet based on non-doxastic reasons. The idea that it is non-inferential is attractive because it solves the regress problem, it is truer to how we apparently form our perceptual beliefs than inferential pictures of perceptual knowledge, and inferential approaches may lead to scepticism. The idea that perceptual knowledge

[32] See, e.g., McDowell 1982; McDowell 2008.

is, despite being non-inferential, based on reasons is attractive because it allows us to see ourselves as rationally responsive to reasons that lie outside thought, and so to see our thought as being *about* a thought-independent world.

McDowell's picture of perceptual knowledge tries to vindicate the idea of non-inferential, yet reason-based perceptual knowledge, and we have now seen many of the central features of this picture. In particular, we have seen that he treads a delicate line in displaying perceptual-experiential reasons as lying outside thought and yet as figuring in the thinker's perspective as the reasons they are, such that the thinker can respond to them as such. Almost every detail of the McDowellian picture of rational constraint in perception can and has been challenged, and there are several difficulties of interpretation which I have passed over. My aim has not been to motivate his picture from the ground up, but rather to present its basic rationale and to describe how its various pieces hang together. This seems to be the spirit in which it is meant to be taken. McDowell presents it as common sense, not in need of argument: 'no one would dream of abandoning [it], if abandoning it did not seem to be required by a powerful philosophical argument.'[33]

It is now time to consider whether or not there is in fact a powerful philosophical argument for abandoning this picture of perceptual knowledge. Many philosophers argue that it is pointless to try to identify

[33] McDowell 1996b: 286.

reasons which fulfil the Reasons Outside Thought condition. They think that only doxastic states can be or contain reasons for belief (and action), because only doxastic states can make it the case that reasons figure as such in a subject's perspective so that the thinker can respond to them as the reasons that they are.[34] In other words, the worry is that the tension in McDowell's account between his internalist requirement and his Reasons Outside Thought condition leads to incoherence. While this line of thought is common (see note 34), it has been developed in most detail by Hannah Ginsborg, and it is her argument that I will focus on. First, however, I want to acknowledge some puzzles about McDowell's own view. I do this through the lens of Clayton Littlejohn's statement of the view that only propositions which one believes can be one's reasons.

1.2. Two challenges for McDowell's internalism

Consider the following passage, in which Littlejohn is criticising McDowell's claim that there are non-doxastic experiential reasons (Littlejohn here uses McDowell's own account of believing- and acting-for-a-reason against him):

[34] BonJour 1978; Heck 2000; Hawthorne and Magidor 2018. The point is made specifically as a criticism of McDowell by Davidson 1999; Ginsborg 2006; Ginsborg 2018; Glüer 2004; Glüer 2009; Littlejohn 2018, and is hinted at in Stroud 2002. In several of these papers the point is simply stated as a truism; Ginsborg's papers provide the fullest development of the argument.

> To possess a reason or piece of evidence, you have to be in a state where you can believe, feel or do something for the reason that p without first having to change your attitude towards p. If your 'attitude' towards p consists of experience but not belief, you do not take it that p is true. You aren't so minded that you'd be right about p if p and wrong about p if p were false. What would it be to believe, feel or do something for the reason that p? McDowell (1978) gives us an answer. Your reasons for V-ing are, as he puts it, the light in which you took there to be something good, favourable, appropriate, or sensible about V-ing. If you don't believe p, whether it's because you're agnostic about whether p or because you believe not-p, p couldn't be your answer to the question, 'What's good about V-ing?' As such, it couldn't be the light in which you V. As such, it couldn't be your reason for anything at all. If we're working with the standard account of motiving reasons that we get from McDowell, the argument just sketched shows that the possession of reasons requires belief in the proposition that constitutes the reason[.][35]

Littlejohn's argument, compressed as it is, seems to me inconclusive. I'll go through the points that he makes.

[35] Littlejohn 2018: 123.

He first asserts that to possess p as a reason you must be in a position to believe or act for the reason that p without first having to change your attitude towards p. McDowell might respond that this is perfectly consistent with the idea of non-doxastic experiential reasons. According to McDowell, you can believe p for the reason that you see that p (recall: he construes seeing-that as non-doxastic and factive) without *first* changing your attitude towards the proposition that you see that p. For on his account of perceptual belief-formation, when you judge p for the reason that you see that p, you do so self-consciously, and so you also come to know that you see that p and that this is your reason for judging p. Neither perceptual knowledge of p nor knowledge of your reason for judging p is prior to the other. So the picture is not that you *first* form the belief or come to know that you see that p, and then on this basis judge p. The perceiving subject

> does not infer her knowledge that o is F from prior knowledge that she sees that o is F; but she comes to know that she sees that o is F in the very act in which she comes to know that o is F; and it is *because* that act provides her with knowledge that she sees that o is F that it, that very same act, provides her with knowledge that o is F.[36]

So, strictly speaking, Littlejohn's opening assertion is compatible with McDowell's view. In coming to perceptually know p, one *does* change

[36] McDowell 2019: 340.

one's attitude towards the proposition that one sees that p, but this is part of, rather than prior to, perceptually judging p.

Admittedly, this aspect of McDowell's view is difficult to understand. In particular, it is difficult to see what explanatory role the subject's knowledge that she sees that o is F is meant to be playing here, given that it is not playing the role of a premise in an inference. Furthermore, I share the puzzlement voiced by Brewer:

> I do not myself fully understand the thesis that these two judgements issue from a single cognitive act. They have quite different contents, one that o is F, the other that my current experience guarantees that o is F. Although the latter content entails the former, the former does not entail the latter. So they are by that light at least quite distinct judgements.[37]

If McDowell were forced to admit that there are two different cognitive acts here, one of coming to know p and one of coming to know what one's reason for believing p is, then the explanatory role of coming to know what one's reason for believing p is would be even more obscure.

So we might put the following challenge to McDowell, loosely inspired by Littlejohn's opening assertion. If merely seeing that p (which is, again, meant to be a non-doxastic state, lying outside thought) is sufficient for

[37] Brewer 2019: 371.

the fact that one sees that p to figure in one's perspective *as* a reason for believing p, to which one can rationally respond, then why does one *also* need to come to *know* that one sees that p in order to come to perceptually know p itself? In other words, what is the motivation for the *first part* of McDowell's internalism, given his claim that the requirement stated in the *second part* of his internalism is fulfilled by perceptual experience? Why does the subject need to embrace in thought the fact that she sees that p, in order for this to be *her* reason for believing p, when the pre-cognitive state of merely seeing that p is sufficient to put her in possession of the fact that she sees that p *as* a reason for belief? McDowell seems to want to have it both ways: he wants to hold, like Littlejohn, that only a fact which one knows can be one's reason for belief, while also holding, against Littlejohn, that an unknown fact can, through experience, be presented in the subject's perspective in such a way that the subject can form a belief on its basis. This is a challenge for rather than an objection to McDowell's view.

Littlejohn's second suggestion is that one must be able to cite one's reason for belief: to possess a reason to V, one must be in a position to cite this reason in response to the question 'What's good about V-ing?'[38] The thought here seems to be that to cite a reason is to assert a proposition which is a fact, and one can only assert a proposition which one believes to be true. As with Littlejohn's opening assertion, this suggestion seems to be compatible with McDowell's view, since

[38] Cf. Stroud 2002: 89.

McDowell agrees that you must be able to cite your reason for belief. On his account, someone who perceptually knows p also knows their reason for believing p and can cite it. To be sure, before one comes to perceptually know p, one cannot cite the relevant reason, but it's unclear why this should be a problem. Why should one be able to cite one's reason for believing p *before* one believes p? And, anyway, insofar as one is *in a position to know* that one's seeing that p is a reason one possesses to believe p, one might be said to be *in a position to* cite this fact as a reason to believe p.

Finally, Littlejohn claims that a reason for believing p must be the 'light in which you took there to be something good' about believing p or performing some action. Again, it's not obvious that this causes trouble for McDowell. That I see that p entails p, and when I judge p I also judge that I see that p and I recognise this to make appropriate my belief in p. Doesn't this give us everything Littlejohn wants?

Littlejohn could push back on both of these last two responses. The replies I've suggested on McDowell's behalf appeal to the fact that McDowell accepts that perceptually knowing p requires me to know, and thus embrace in thought, that I see that p. I can only cite the fact that I see that p as my reason for believing p in virtue of the fact that I know that I see that p. And I apparently only recognise my seeing that p to make appropriate my believing p in virtue of the fact that I know that I see that p. Littlejohn might say that what he is demanding is that the fact

that I see that p figures in my perspective as a reason to believe p *before* I come to believe p. Littlejohn could make this demand because this seems to be the only way that the fact that I see that p can figure in a rationalising *explanation* of why I believe p. This amounts to the complaint that the second part of McDowell's internalism, according to which possessing a reason requires only that the subject is in a position to form a belief for that reason and to recognise its normative significance, does not really specify a genuine way in which one can *possess* a reason outside thought, such that one's possession of it can figure in a rationalising explanation.

This leaves us with McDowell facing two initial challenges. The first has to do with his motivation for the first part of his internalism. Why must a subject *know* what their reason is, and so embrace their reason in thought, when McDowell is trying to show us how a subject can possess a reason without embracing it in thought? We might endow this challenge with a subtext: McDowell illegitimately gets us comfortable with the idea that merely having a factive perceptual experience is sufficient for reason-possession by smuggling the usual trappings of reason-possession into the cognitive act of judging oneself to have a reason. This takes us to the second challenge. For Littlejohn, McDowell's account of reason-possession outside thought, in terms of being in a position to know p and to know what one's reason for judging p is, does not really amount to a picture of genuine reason-possession. The usual trappings of reason-possession – being able to cite the reason, and taking it to be a reason –

are absent. And it's difficult to see how the reason figures in the subject's perspective *as* a reason in such a way that we could appeal to it in a rationalising explanation of the subject's perceptual belief.

The question, then, is whether McDowell can give some account of how an experiential reason to believe p figures *as* a reason in one's perspective *before* one forms one's perceptual belief on the basis of this reason, so that the reason can play a genuinely explanatory-rationalising role in perceptual belief-formation. A satisfactory answer to this question might also yield a principled reason for sustaining the first part of McDowell's internalism alongside the second part. Ginsborg (2006) and (2018) has given a detailed argument which tries to show that McDowell *cannot* give any account of how a non-doxastic experiential reason figures as a reason in one's perspective. She agrees with Littlejohn that reasons are doxastic. Let's now see what Ginsborg's argument is before considering possible responses to these challenges, and to the challenges Ginsborg develops.

1.3. Reasons$_1$, reasons$_2$ and the Reason-Taking Claim

Ginsborg begins her criticism of McDowell by making an assumption that he would reject out of hand. She assumes that rationalising explanations of a subject's beliefs and actions can properly appeal only

to the subject's non-factive mental states. As I will explain in more detail in Chapter 3, that is anathema to McDowell, for whom perceptual knowledge is to be explained in terms of the subject's rational responsiveness to a factive perceptual experience.

Ginsborg does not acknowledge that in making this assumption she is assuming something that McDowell would deny, and she does not offer a detailed argument for the assumption. This suggests that she does not take the assumption to be doing any special work in her argument against McDowell. Part of what I want to suggest is that that's not right: the assumption is absolutely crucial if Ginsborg's argument is to work. That doesn't mean that her argument against McDowell is uninteresting, however. It is simply misdirected. It is an argument that works well against those who think that perceptual experiences are non-doxastic reasons for belief, and who think of perceptual experiences as non-factive – and this is arguably the majority position in contemporary epistemology of perception. Her argument therefore brings home the fact that if we want to make room for the fulfilment of Reasons Outside Thought, we need to think of perceptual experience as factive.

In this section I will set out the details of Ginsborg's assumption that only non-factive states figure in rationalising explanations. This will also involve setting out Ginsborg's assumptions about what is involved in a reason figuring as a reason in a subject's perspective. These latter assumptions constitute critical premisses in her argument against the

idea that perceptual beliefs can be based on reasons outside thought. For this argument involves claiming that McDowell has no way of specifying a satisfactory sense in which reasons outside thought figure as reasons in the subject's perspective. By way of preparation for a more detailed critique of Ginsborg's argument, I will suggest in this section that her own account of what is involved in a reason figuring as a reason in a subject's perspective is in important ways limited and unclear.

Ginsborg first distinguishes what she calls reasons$_1$ and reasons$_2$, which are meant to correspond to two ways in which we use the word 'reason' in everyday life. 'Reasons$_1$' are the things or facts (for Ginsborg they are always facts) which, from a first-person perspective of deliberating about what to believe or do, we recognise as favouring or making appropriate a belief or action. For example, if I say 'I believe it's recently rained because the streets are wet' then I am citing my reason$_1$ – that the streets are wet – for believing that it's recently rained. Ginsborg's reasons$_1$ are what are often called normative reasons, since they make certain responses appropriate, and they are 'reasons *for* belief' (as opposed to, say, reasons *why* one believes as one does) because they are reasons that the subject herself is disposed to cite in support of her belief.

'Reasons$_2$' are the reasons we cite from a third-person perspective to provide rationalising explanations of other people's beliefs (although, of course, we can adopt this 'external' perspective on our own beliefs, too). For example, if I say 'Hannah believes it's recently rained because she

believes that the streets are wet', then I am citing the rationalising reason$_2$ – her psychological state of believing that the streets are wet – why Hannah believes it has recently rained. Reasons$_2$ are thus what are often called explanatory reasons, but they are not just any old explanatory reasons. Rather, they are specifically *rationalising* explanatory reasons: to cite them is to make rationally intelligible a subject's believing as she does, given her perspective on the world.

As Ginsborg recognises, from within the first-person perspective of deliberating about what to believe or trying to justify one's beliefs, we seek and try to attend to aspects of the world itself which favour or make appropriate certain beliefs. We generally treat as irrelevant our own non-factive mental states, including our beliefs about the world. So when one cites one's reason for believing p, one is either citing one's reason$_1$, a worldly fact or aspect of the world that makes appropriate believing p, *or* one is mistakenly citing a *supposed* worldly fact which one mistakenly takes to be a reason$_1$.

However, according to Ginsborg, from the perspective of third-person rationalising explanations of a subject S's belief, whether the subject

actually possesses a reason$_1$ is 'irrelevant'.[39] What matters is how things *seem* to the subject, not how they really are. That is, what matters is that from S's perspective it is to her as though she has a reason$_1$. This is what the reason$_2$ makes so. For example, when I want to provide a rationalising explanation of Hannah's belief that it recently rained, I'm uninterested in whether or not the streets are actually wet; instead, what I cite is Hannah's belief, true or false[40], that the streets are wet. These supposed observations constitute Ginsborg's argument for thinking that factive mental states do not play any rationalising-explanatory role over and above the role played by non-factive states.

According to Ginsborg, then, reasons$_2$ can do their rationalising-explanatory work whether or not the subject has a reason$_1$. Having a reason$_2$ for believing p does not entail having a reason$_1$ for believing p. S's having a reason$_2$ does, though, entail that from S's perspective it is as

[39] Ginsborg 2006: 290. Note that this is compatible with holding that the fact which is a subject's reason$_1$ and also the subject's knowledge of that fact may be explanatorily relevant with respect to S's belief and behaviour. For discussion of this point (not specifically in relation to Ginsborg, but in relation to a similar view) see Roessler 2014. As Roessler says, 'If almost anything can be a factor influencing human conduct [and, we might add, belief], why should not the agent's knowledge? Arguments to the effect that explanations in terms of knowledge can be correct and illuminating, insofar as they sustain distinctive sets of counterfactuals (Williamson 2000), are unlikely to cut any ice with a [theorist like Ginsborg]. For such arguments leave open whether explanations in terms of knowledge form part of the explanatory schema in terms of which we find intentional actions intelligible as such.' (348)

[40] A belief can be true or false and still be a reason$_2$, according to Ginsborg, but she doesn't say anything about whether or not it needs to at least be justified. She might say that it depends what kind of rationalising explanation we're trying to give. That is, she might say that there's a kind of rationalising explanation that can be given only in terms of justified beliefs or knowledge, but also that there's another kind that can be given in terms of any old belief.

though she has a reason$_1$, and in particular it entails that S 'takes herself to have a reason$_1$, which, [at least] on the face of it, will be the fact or supposed fact represented by the belief or other psychological state which serves as her reason$_2$.'[41] This claim is the centre of Ginsborg's critique of McDowell. Call it the Reason-Taking Claim:

> **Reason-Taking Claim**: Having a reason$_2$ to believe q involves taking p to be a reason$_1$ to believe q.

The Reason-Taking Claim is a close cousin of what Boghossian calls the Taking Condition on inference between belief-states:

> **Taking Condition**: 'Inferring necessarily involves the thinker *taking* his premisses to support his conclusion and drawing his conclusion *because* of that fact.'[42]

Whereas the Taking Condition is supposed to be a condition on inference between beliefs, Ginsborg intends to Reason-Taking Claim to apply to responding to epistemic reasons in general. Ultimately, though, since Ginsborg thinks all reasons are doxastic, all doxastic responses to reasons are for her inferential.

---

[41] Ginsborg 2006: 291.

[42] Boghossian 2014: 5.

Boghossian says that we must endorse the Taking Condition on inference if we are to understand inference as something that thinkers do with the aim of 'figuring out what follows or is supported by other things one beliefs', rather than as something that happens to or in them.[43] The Reason-Taking Claim is motivated by similar considerations. It constitutes Ginsborg's account of how something figures as an apparent $reason_1$ within one's perspective: something figures as an apparent $reason_1$ within one's perspective because one takes it to be a $reason_1$. And this, she says, is what we're interested in when we give rationalising explanations of subjects' beliefs, or when we try to explain a subject's belief in terms of how things seem to the subject.

Ginsborg offers little in the way of elaboration on the Reason-Taking Claim, and this leaves it uncertain what exactly she means by it. One thing that she does say is that to 'take' p to be a $reason_1$ to believe q involves being disposed to *cite* p *as* a $reason_1$ to believe q (see section 1.5 below). But that still leaves it unclear what kind of state 'taking' is. The most natural way to understand 'taking' is as a way of *believing*, and some things that Ginsborg says suggest that this is how she means it to be understood.[44] If that is what she means then the Reason-Taking Claim should be understood as follows:

[43] Ibid.

[44] See, e.g., Ginsborg 2006: 304–5. Here Ginsborg talks about having a $reason_2$ to believe p as involving one 'regarding' something to be a $reason_1$ to believe p.

> **Reason-Taking Claim (Doxastic Construal)**: Having a reason$_2$ to believe q involves believing p to be a reason$_1$ to believe q.

However, there are other ways of understanding what 'taking' means in this context. Ginsborg uses various formulations to state what I'm calling the Reason-Taking Claim, not all of which involve reference to 'taking'. For example, she says that 'a psychological state counts as a reason$_2$ for a subject's belief only if her being in that state involves her representing some fact [or, presumably, a supposed fact, although here Ginsborg omits this possibility] as a reason$_1$ for that belief'.[45] Boghossian describes several different ways in which his Taking Condition might be understood. Again, 'taking' might just mean believing. But, he admits, this is psychologically unrealistic, since children can perform inferences but might lack the concepts necessary to form the belief specified by the Taking Condition. Furthermore, he says, on the doxastic approach to 'taking' it's difficult to avoid having to say that the role of the 'taking'-belief is to be a further premise in the inference, and this leads to the familiar vicious regress described by Lewis Carroll.[46] To avoid this we might say that the 'taking'-belief is an enabling condition for the inference rather than a premise, but it's obscure how the 'taking'-belief 'enables' the inference if not by being a premise in it.[47] A doxastic construal of the Reason-Taking Claim faces these same problems.

---

[45] Ginsborg 2006: 291.

[46] Carroll 1895.

[47] Boghossian 2014: 6–9.

But, as Boghossian shows, non-doxastic construals of 'taking' are problematic too. The main alternative with respect to the Taking Condition for inference is one on which inferring involves following rules of inference. For example, one might infer in accordance with the rule of *modus ponens*. In so doing, says Boghossian, one might be thought to be 'taking' (without believing, at least not explicitly) the conjunction of the premisses to be a reason to believe the conclusion. However, as soon as we try to say what rule-following in inference amounts to we get into difficulties. In particular, how are we to explicate the difference between forming beliefs according to patterns that happen to conform to a rule, and actually *following* or being guided by a rule in forming beliefs? Boghossian suggests that the obvious move to make is to 'postulate the existence of an intentional state that represents the rule'. Following or being guided by the rule would then involve 'grasping what the rule requires, forming a view to the effect that its trigger conditions are satisfied, and drawing the conclusion that I must now perform the act required by its consequent. In other words, […] rule-following requires inference'. But if rule-following involves inference, and inference involves rule-following, we get another vicious regress.[48]

Instead, then, we might try to explicate rule-following without reference to an intentional state which represents the rule. The problem is that this tends to lead to accounts of rule-following which are unsuited to spelling

[48] Ibid., 14. Cf. Boghossian 2008. The route to regress is spelled out at Wright 2014b: 30.

out what's involved in fulfilling the Taking Condition on inference. For example, one might hold that inferential rules are, as Quilty-Dunn and Mandelbaum argue, 'built into the architecture' of our cognitive systems, which is just to say that 'whenever a mental representation is tokened that satisfies the antecedent of the rule, then, *ceteris paribus*, the system will token a representation that satisfies the consequent of the rule'.[49] This makes inference an automatic procedure, in which the subject themselves need have no appreciation of the premisses of their inference as making appropriate belief in the conclusion. Philosophers like Boghossian and Ginsborg, who want to do justice to the idea of the subject actively engaging in reasoning, will worry that by making inference automatic in this way, and by excising any essential role for person-level appreciation of the relevant rule and its application to the inference in question, we excise the subject's perspective from the process and make it explanatorily irrelevant. Similar remarks apply to the Reason-Taking Claim and rule-following accounts of believing for a reason.

There is a further interpretative difficulty about the Reason-Taking Claim. Ginsborg seems to think that a belief in p or knowledge of p by itself constitutes a reason$_2$ to believe whatever p entails or makes sufficiently probable. For example, Hannah's reason$_2$ for believing that it recently rained is her belief that the streets are wet (plus any other beliefs which might serve as premisses in the inductive inference).[50] But this

[49] Quilty-Dunn and Mandelbaum 2017: 7.

[50] See, e.g., Ginsborg 2006: 302 and 311.

doesn't seem right given how we've been trying to understand the Reason-Taking Claim. Suppose that p entails q. S might believe or know p but not notice the fact that p entails q; therefore S is not disposed to cite p as a reason$_1$ to believe q; therefore, given the Reason-Taking Claim, S does not have a reason$_2$ to believe q. This suggests that simply believing p is insufficient for having a reason$_2$ to believe all of p's consequences. More strongly, it suggests that a belief does not by itself *constitute* a reason$_2$ to believe anything, or, in other terms, one does not have a reason$_2$ to believe anything merely in virtue of having a belief (or merely in virtue of having a state of knowledge).[51] So, at least as Ginsborg encourages us to understand the Reason-Taking Claim, merely believing or knowing p does not amount to having a reason$_2$ to believe anything. Yet this is also contrary to Ginsborg's official line, according to which beliefs are reasons$_2$ *par excellence.*

[51] The difference between saying that belief is insufficient for reason$_2$-possession and saying that belief cannot constitute reason$_2$-possession can be illustrated in the following way. According to some philosophers, having a perceptual experience as of a tomato on the table gives one justification to believe that there is a tomato on the table. However, these philosophers say, having such an experience is not sufficient for having justification to believe that there is a tomato on the table, because it must also be the case that one lacks defeaters for this justification (for example: one must not have reason to think that one is hallucinating). On this kind of view, even though having the experience is not sufficient for having justification, one can have justification *in virtue of* having the experience, and the experience can *constitute* one's justifier. In general, it might be that p obtains in virtue of a set of conditions C, and yet C is not sufficient for p to obtain. See Chudnoff 2011b for discussion. The point I am making here is that Ginsborg's position seems to lead to the conclusion not only that belief is insufficient for reason$_2$-possession, say because some further negative conditions (like a no-defeater condition) need to obtain, but also that a belief cannot, by itself, *constitute* a reason$_2$.

Thus, as we've been understanding the Reason-Taking Claim so far, a belief in p can at most amount to part of what constitutes a reason$_2$ to believe some further proposition. One must also 'take' p to be a reason$_1$ to believe the relevant proposition. If this 'taking' is doxastic, then a single reason$_2$ seems to have to be constituted by two beliefs – a belief in p and a belief that p is a reason$_1$ for some further belief. If this 'taking' is a matter of following a rule, then it's unclear what reasons$_2$ are – it depends on what following a rule amounts to. But here too, it seems, a belief or piece of knowledge by itself would not be a reason$_2$. Thus rationalising explanations which refer only to such a belief – as in 'Hannah believes it recently rained because she believes the streets are wet' – must be elliptical for more complete explanations which refer to the full reason$_2$.

We must admit, then, that it's not entirely clear what Ginsborg means by 'taking p to be a reason$_1$', and by extension it's not clear what she thinks reasons$_2$ are constituted by – although she does seem to want to hold that a mere belief in p constitutes a reason$_2$ all by itself, ready-made to be cited in rationalising explanations. There are thus substantial tasks involved in both further characterising the 'taking'-state mentioned in the Reason-Taking Claim and in specifying what explanatory and rational role it plays in belief-formation. Ginsborg's attempt to say what it is for a supposed reason$_1$ to figure in a thinker's perspective *as* an apparent reason$_1$ in such a way that they can base beliefs on it is thus, insofar as it appeals to the Reason-Taking Claim without further

explication, quite limited in its descriptive ambitions. That in itself needn't be a problem. However, as we will see, a central aspect of Ginsborg's criticism of McDowell is that his account of experiential, non-doxastic reasons cannot make sense of the idea that these reasons figure as such in the subject's perspective. Given the limitations of her own account of supposed reasons$_1$ figuring as such in a subject's perspective, defenders of McDowell are being set quite a low bar for producing a satisfactory response.

In Chapter 2 I will sketch my own suggestion for how best to understand the constitution of the rationalising mental states – what Ginsborg calls reasons$_2$ – that we appeal to in third-person-perspective rationalising explanations of inferential beliefs. In particular, I will suggest that simply believing or knowing p *does* by itself constitute what Ginsborg would call a reason$_2$. My suggestion will be that merely believing or knowing p makes p figure in one's perspective as a supposed reason$_1$, even if one doesn't explicitly *believe* it to be a reason$_1$, and my suggestion will make no essential reference to rule-following on the part of the subject.

Let's take stock. The aim in this section was to set out the assumptions Ginsborg makes in arguing, contra McDowell, that all reasons must be doxastic. She says that reasons$_1$ are the normative reasons that, from a first-person perspective, we seek and attend to in deciding what to believe and do. Reasons$_1$ are thus naturally understood to be worldly facts, or perhaps just worldly entities, which make appropriate or inappropriate

certain responses. Reasons$_2$ are the mental states we cite when, from a third-person perspective, we wish to give rationalising explanations of subjects' beliefs and actions, which is to say when we wish to explain why the subject believes or does as they do in terms of how things seem to them. From this perspective, says Ginsborg, we cite *non-factive* mental states rather than worldly facts (or factive states incorporating such worldly facts), since we are interested in the subject's perspective on the world, rather than the world itself. Reasons$_1$ thus play no rationalising-explanatory role. This is a major departure from McDowell's approach to reasons, but Ginsborg does not acknowledge this, suggesting, as I said, that she thinks her argument would work even if she allowed that reasons$_2$ can be factive states.

As well as this assumption about the non-factivity of reasons$_2$, the other major claim that this section considered was Ginsborg's Reason-Taking Claim. This constitutes Ginsborg's attempt to say what it means for a supposed reason$_1$ to figure as an apparent reason$_1$ in the thinker's perspective, in such a way that an external interpreter can rationalise and explain the subject's subsequent beliefs in terms of this feature of her perspective. Reasons$_2$ are the mental states that make the subject 'take' something to be a reason$_1$ for belief. We saw that there is a degree of interpretative difficulty in understanding how Ginsborg conceives her key notion of 'taking p to be a reason$_1$', and we saw that various standard ways of fleshing out similar 'taking' conditions come with well-

known problems. Despite this unclarity, the core of Ginsborg's argument against McDowell can now be set out.

1.4. Ginsborg's interpretation of McDowell

The first phase of Ginsborg's argument involves claiming that when McDowell says that experiences are reasons for belief, he must mean that they are reasons$_2$, and not merely that the fact that one is undergoing the experience in question is a reason$_1$. Or, in the terms in which McDowell often puts matters, McDowell's view cannot just be that the fact that one sees that p is a reason$_1$ for belief. He must also conceive of the experience (the state of seeing-that) itself as a reason$_2$.

Ginsborg argues for this claim in the following way. If McDowell's view were just that the fact that one sees that p is one's reason$_1$ to believe p then this would be entirely consistent with the coherentism that it is his purpose to reject. For the coherentist could agree and elaborate by insisting that one's reason$_2$ is one's belief or knowledge that one sees that p. Perceptual knowledge would then be based on one's knowledge about one's experiences – it would be inferential. This picture would not incorporate the idea of reasons which one possesses but which have not been embraced in one's thought – it does not make room for the

fulfilment of Reasons Outside Thought. Furthermore, for Ginsborg reasons$_1$ don't by themselves play any rationalising-explanatory role in belief-formation, and so if McDowell wants experiential reasons to be genuine *constraints* on thought then they must be the kind of reasons we appeal to in rationalising explanations, namely reasons$_2$. She concludes that McDowellian perceptual experiences must be reasons$_2$, which is to say that they must be (non-doxastic) mental states that figure in rationalising explanations of beliefs and (given the Reason-Taking Claim) make one 'take' something to be a reason$_1$ for belief.

The Reasons Outside Thought condition as I formulated it in Section 1.2 foregrounds what Ginsborg would call reasons$_1$: it says that there must be some reasons$_1$ which the subject possesses without their being the contents of doxastic states. Ginsborg's claim is that, really, what McDowell requires is that the following condition be fulfilled:

> **Reasons$_2$ Outside Thought**: R is a reason$_2$, and R is neither a thought nor a content of a thought (in particular, R is not a belief or state of knowledge, nor a content of a belief or state of knowledge).

If perceptual experiences are to fulfil this, then they must be reasons$_2$. They must therefore make the subject 'take' something to be a reason$_1$. Ginsborg identifies two candidates to be this supposed reason$_1$. First, the reason$_1$ might be the content of the experience, namely some content

which represents the external world as being some way. This would be analogous to how Ginsborg thinks beliefs make their contents into supposed reasons$_1$. Alternatively the supposed reason$_1$ could be the proposition that one sees that p.

We've seen that the second option is the one that McDowell would seem to favour, since he says one's reason for perceptually believing p is that one sees that p – and this is the reason one cites from the first-person perspective. In presenting this option Ginsborg assumes that the only way an experience could present to the subject the fact that they are having the experience is by the experience involving the subject recognising – and so believing – that they are having it.[52] This would make experiential reasons$_2$ doxastic after all, so Ginsborg rejects this option on McDowell's behalf. Instead, then, Ginsborg says that McDowell's position can only be coherent if he endorses the first option, on which seeing that p is an experiencer's reason$_2$ which makes it the case that the experiencer takes p to be a reason$_1$ for believing p. This option is of great interest since while McDowell may not endorse it, many other philosophers do.[53]

[52] Ginsborg 2006: 303 and Ginsborg 2018: 112.

[53] See, e.g., Pollock 1974; Brewer 1999; Pryor 2000; Huemer 2007; Chudnoff 2011; Smithies 2019. This is just a small sample. Note that many of those who endorse this option conceive of perceptual experience as non-factive – especially so-called 'dogmatists' and 'phenomenal conservatives'.

1.5. Ginsborg's argument

Ginsborg now frames her challenge to McDowell in the following way (using 'see that p', like McDowell, to refer to a non-doxastic perceptual-experiential state):

> If I see that *p*, but without believing that *p*, do I nonetheless take *p* to be a reason$_1$ for believing that *p*? That is, do I take myself to have, in *p*, a reason$_1$ for believing that *p*?[54]

If non-doxastic perceptual experiences are reasons$_2$ then, in light of the Reason-Taking Claim, they must make the subject 'take' something to be a reason$_1$. Here we're considering the suggestion that experiences make the subject 'take' their contents, which are about the external world, to be reasons$_1$ for beliefs about the external world. Ginsborg is asking whether non-doxastic perceptual experiences can do this, and her answer is that they cannot. For, she thinks, I must believe p in order to 'take' p to be a reason$_1$ to believe something and thereby have a reason$_2$.

What is Ginsborg's argument for this claim? One point she makes is that simply entertaining a content does not make one 'take' it to be a reason$_1$ for belief in anything, even if one recognises that the content stands in certain logical relations to other contents. As an example, she notes that

[54] Ginsborg 2006: 304.

we all know that *the moon is made of cheese* entails that *the moon is edible*, but we do not 'take' *the moon is made of cheese* to be a reason$_1$ for believing the moon is edible. Thus, 'for any propositions *p* and *q*, there is more involved in taking *p* as a reason for believing *q* than simply taking one to imply the other.'[55] Everyone will agree with Ginsborg here. What her point shows is that to argue that perceptual experiences are reasons$_2$, we need to do more than simply say that they have conceptual contents. The mere fact, if it is a fact, that experiential contents stand in relations like entailment to belief-contents doesn't show that experiences or their contents are reasons$_2$. However, this is at most a challenge, rather than an objection, to the view that non-doxastic experiences are or contain reasons$_2$. When McDowell and others emphasise their claim that experiences have conceptual content, they do so because they take this to be a necessary condition on experiences being reasons, not because they think that it is the only feature of experiences that makes them into reasons.

However, Ginsborg adds that 'intuitively […] the reason that I do not take the moon's being made of cheese as a reason$_1$ for believing that the moon is edible is that I do not believe that the moon is made of cheese'.[56] In order to explain this intuition, Ginsborg appeals to an idea that we already saw in Littlejohn's criticism of McDowell (see Section 1.2): 'Taking something as a reason$_1$ for belief is a matter of being

[55] Ibid., 304.

[56] Ibid.

prepared to cite it in defence of the belief, as a consideration which counts in its favour; and I cannot cite the moon's being made of cheese in defence of anything unless I believe that the moon is made of cheese.'[57] So while we don't have a full account of what 'taking' p to be a reason$_1$ to believe q is, we *are* being told that it involves bring disposed to cite p as a reason$_1$ to believe q. And now the line of thought seems to be, again, that being disposed to cite a proposition as a reason$_1$ to believe something involves being disposed to assert it, and being disposed to assert it involves believing it. So 'taking' p to be a reason$_1$ involves believing p. And Ginsborg agrees with Littlejohn that one must be disposed to cite p as a reason$_1$ to believe q *in advance* of coming to believe q, in order to be able to come to believe q for the reason that p. So one must believe p *in advance* of believing q in order to be able to believe q for the reason that p. But that is just what McDowell has to deny. So McDowell has to hold that one can 'take' p to be a reason$_1$ to believe q even though one is not disposed to cite p as a reason$_1$ to believe q. That is, thinks Ginsborg, counterintuitive.

Further, this produces the same challenge for McDowell that we saw Littlejohn as posing him: he needs to specify a sense in which a reason$_1$, or a supposed reason$_1$, figures in the perspective of the subject *as* an apparent reason$_1$, such that she can form a belief on its basis. And he must do this without appealing to the idea that the reason$_1$ figures as the

[57] Ibid.

content of a belief, and so without appealing to the idea that the subject is disposed to cite the reason$_1$ as a reason$_1$.

Ginsborg adds a consideration which seems to show that McDowell will have difficulty in meeting this challenge. She notes that McDowell, as she interprets him, thinks that when one has an experience with content p, one 'takes' p itself to be a reason$_1$ for perceptually believing p. The truth of p is clearly a *conclusive* reason$_1$ for believing p. Therefore, 'It makes no sense to suppose that I could take [the supposed reason$_1$] to be a reason for believing that [p], yet refrain from forming that belief because of other considerations which counterbalance or outweigh it'.[58] And yet we do sometimes refuse to believe our eyes, if we think we're hallucinating for example. Ginsborg's point here seems to be that there is no psychological room to 'take' p (or, for that matter, the fact that one sees that p, which would also be a conclusive reason$_1$ to believe p) to be a reason$_1$ to believe p while also refusing to believe p. But there is psychological room to enjoy a perceptual experience with content p while refusing to believe p. So merely enjoying a perceptual experience with content p does not make one 'take' p (or the fact that one is seeing that p) to be a reason$_1$ to believe p. So to merely enjoy a perceptual experience with content p does not amount to having a reason$_2$ to believe p. This concludes Ginsborg's argument against McDowell's position.

[58] Ibid.

1.6. Conclusion

This chapter began with the idea that perceptual knowledge is based on non-doxastic, experiential reasons. Several considerations for thinking that perceptual knowledge is non-inferential were described: it might solve the problem of the regress of justification; it is psychologically unrealistic to claim that our perceptual beliefs are inferentially formed; and it may have anti-sceptical force. The main consideration for thinking that perceptual knowledge is not only non-inferential but also based on experiential reasons was that, as McDowell argues, the intelligibility of intentionality depends on our ability to make sense of the idea that some of our knowledge is based on, and so rationally constrained or influenced by, reasons that lie 'outside thought'. Furthermore, if we deny that perceptual knowledge is based on experiential reasons, and we hold that it is non-inferential, we would apparently lack any explanation of how perceptual beliefs can be rational or justified. So, at least *prima facie*, there are good reasons to try to vindicate the idea of non-doxastic experiential reasons for belief.

We briefly noted that there is a division between those try to vindicate this idea in terms of a conception according to which perceptual experiences are non-factive mental states, and those who try to vindicate this idea in terms of a factive conception of perceptual experiences. We

explored in detail McDowell's version of the factive conception, and we saw that there is potentially a tension in it between an internalist conception of reason-possession and his insistence that perceptual reasons must lie outside thought.

This tension is exploited by his critics. They lay down a challenge: how can we make sense of the idea that a reason figures *as* a reason in one's perspective, such that one can rationally respond to it, without one *representing* it as a reason, and thereby embracing it in thought? Standard ways of thinking about reason-possession – in terms of 'taking' something to be a reason, and in terms of being disposed to cite something as a reason – seem to require the putative reason to be the content of a belief. If McDowell does away with these standard ways of thinking about reason-possession, then he needs to replace them with some other account.

Littlejohn and Ginsborg are sceptical that this will be possible. What seems to drive their scepticism is that the apparent point of the idea of non-doxastic perceptual experiences is that the subject themselves is passive and noncommittal in having such an experience. But if the subject is passive and noncommittal then how can the subject count as *recognising* or otherwise *registering* the rational import of the experience *in advance* of forming a belief on its basis? The subject only seems to get involved when it comes to forming the perceptual belief in question – but that's too late for the purposes of rationalising explanation.

Littlejohn and Ginsborg seem to assume that the only way to make sense of believing for a reason is in terms of the idea that the subject *first* (in some hard-to-specify sense) registers a reason *as* a reason and then *because* of this registration forms a belief on the basis of this registered reason.

What I'm calling registration is really what Ginsborg means by 'taking' – although as we saw, the notion of 'taking' is quite problematic. In particular, it's difficult to say what kind of mental state 'taking' is or involves; and then it's also difficult to say what kind of rational role 'taking' plays. Vicious regresses loom. Furthermore, it's difficult to discern Ginsborg's own positions on these matters. 'Taking' p to be a $reason_1$ to believe q seems to be something additional to simply believing p, since one might believe p without noticing that it's a $reason_1$ to believe q. Yet simply believing p seems to count as a $reason_2$, according to Ginsborg.

In the end, it could be that her point is really that believing p is a necessary (rather than a sufficient) condition for 'taking' p to be a $reason_1$. It is a necessary condition because, first, 'taking' p to be a $reason_1$ involves being disposed to cite it as such. And, second, the status of something *as* a $reason_1$ depends on its being a fact or a worldly entity, and so it must figure as a supposed fact in the subject's perspective for the subject to be in a position to respond to it as a $reason_1$. Ginsborg assumes that the only way it can so figure is by being believed.

In setting out this argument, we noted in passing two puzzles. First, McDowell himself, in the first part of his internalism, seems to *endorse* the idea that having a reason involves being disposed to cite something as a reason, at least once one has formed the belief which it's a reason for. Can McDowell motivate that claim while also insisting that one can possess a reason without being disposed to cite it *before* one has formed the belief it is a reason for?

Second, Ginsborg puts her critique in terms of a conception of rationalising explanation according to which only non-factive mental states are genuinely *explanatory*, in a rationalising way, of a subject's beliefs and actions. What explains why someone believes or acts as they do, she says, is that it is to them *as though* they have a reason to do so. The mental state which makes it as though they have such a reason she calls a reason$_2$; the actual reasons for which we try to believe and act are reasons$_1$, but these worldly facts never actually explain our believing or doing anything. Although Ginsborg explicitly sets out this picture, she does not note that it contradicts McDowell's own view, according to which factive mental states – and thereby worldly facts or reasons$_1$ – are genuinely explanatory. Further, Ginsborg does not say anything about the role, if any, that her non-factivism plays in her argument against McDowell. Indeed, at first glance, it seems not to play any role. In Chapter 3 I will argue that, in fact, her argument has no purchase against those who endorse McDowell's factivism. At the same time, her argument continues to have force against those who claim that

experiences, conceived non-factively, are non-doxastic reasons. First, however, in Chapter 2 I will take a closer look at the 'taking' condition as it figures in the Reason-Taking Claim.

Chapter 2

Asking questions, giving instructions

According to Ginsborg, non-doxastic perceptual experiences can't be reasons$_2$ because they don't make one 'take' a content to be a reason$_1$. In general, she thinks, to form a belief for a reason is to form a belief in response to one's 'taking' oneself to have a reason$_1$ for that belief. So 'taking' something to be a reason$_1$ is what is involved in having something figure *as* an apparent reason$_1$ in one's perspective, such that one can rationally form a belief for this reason$_1$. Her claim is that non-doxastic perceptual experiences don't make anything figure as an apparent reason$_1$ in one's perspective. Beliefs, on the other hand, do make their contents figure as apparent reasons$_1$ in one's perspective. Therefore, beliefs are what external interpreters properly appeal to in rationalising explanations of a subject's beliefs and actions. Thus, reasons$_2$ are always beliefs, and all reason-based beliefs and knowledge are inferential.

In the previous chapter we noted that the notion of 'taking' is unclear. It is unclear what 'taking' *is*, what it *involves*, and what its *rational role* is. And it is unclear how Ginsborg in particular conceives of 'taking'. This is a

crucial issue: the critical claim made against McDowell is that whatever 'taking' p to be a reason$_1$ is, it requires one to believe p. Thus undergoing a non-doxastic perceptual experience is not sufficient for 'taking' something to be a reason$_1$. It's hard to assess this claim without a clear notion of 'taking' p to be a reason$_1$.

In this chapter I provide what I take to be an attractive account of the nature and role of 'taking' in inference. More precisely, I give an account of the sense in which a supposed reason$_1$ needs to register *as* a reason$_1$ in one's perspective *in advance* of one forming an inferential belief on its basis. Whether this registration deserves the name 'taking' isn't clear. I'm inclined to think that it does not, because I think 'taking' is really just another word for 'believing', and I deny that one needs, in advance of inferring q on the basis of p, to *believe* that p is a reason$_1$ to believe q. Nevertheless, I'll continue to talk about 'taking', since it's common in the literature to use this word without necessarily meaning 'believing', and I want to engage with this literature and in particular with Ginsborg.

My hope is that once we get clear on the nature and role of 'taking' in inference, we can then consider whether non-doxastic perceptual experience might not after all make the subject 'take' something to be a reason$_1$. Chapter 3 argues that in the light of a realistic picture of how believing p makes one 'take' p to be a reason$_1$, we can see that non-doxastic perceptual experience can do something similar. So my account of 'taking' undermines Ginsborg's claim that only believing p can make

one 'take' oneself to have a reason$_1$, and so it will undermine her claim that only beliefs can be reasons$_2$.

The plan for this chapter is as follows. First I argue that when it comes to inference, a mere belief or state of knowledge is a reason$_2$, which can rationally motivate an inferential belief. Thus, when one infers from one's belief in p to a belief in q, one's belief in p, all by itself, makes one 'take' p to be a reason$_1$ in a way that is sufficient to rationally motivate you to believe q. 'Taking' is not a substantial further condition, over and above the belief that constitutes the reason$_2$.

I then sketch an account of how simply believing or knowing p makes you 'take' p to be a reason$_1$ in a way that is sufficient for you to be in a position to rationally believe q on the basis of p. This is my question-asking and instruction-giving account of inference. In developing this account, I describe ways in which inferential knowledge differs from perceptual knowledge and ways in which the two are more similar than is sometimes appreciated. In particular, while there is an element of passivity in perceptual knowledge-formation that is absent in inference, there is also an element of passivity that is common to perceptual and inferential knowledge. By recognising that inference involves an element of passivity, we can see that belief based on reasons can involve an element of passivity, and this makes it easier to accept that perceptual knowledge is based on reasons.

I then try to quell worries that my picture of inference (and by extension my picture of reason-based belief and knowledge in general) fails to assign enough of a role, in the rational motivation of the subject's inferential belief-formation, to the subject's antecedent person-level appreciation of the rational force of her (supposed[59]) $reasons_1$. I argue that my account can accommodate a substantial sense in which the subject appreciates a $reason_1$ as a $reason_1$ both in advance of and as a result of her performance of the inference.

On the account of inference that I give, one need not, in advance of inferring q on the basis of p, believe that p is a $reason_1$ to believe q. Further, one need not, in advance of inferring q on the basis of p, be disposed to cite p as a $reason_1$ to believe q. Finally, then, I point out that this account of inference makes it seem less heretical to claim that perceptual knowledge can be based on reasons which one possesses without having embraced them in thought. The overall message is that a realistic picture of inference is not so very different from the picture of perceptual knowledge that McDowell offers us; I take this to cast McDowell's picture of perceptual knowledge in better light.

[59] I include qualifications like this to take into account Ginsborg's insistence that one needn't actually have a $reason_1$ to form a belief that can be rationalised by an external interpreter in terms of a $reason_2$.

2.1. 'Taking' is not a substantial additional condition

My account of 'taking' might help make sense of some of what Ginsborg says. For, as we saw in Chapter 1, she seems to hold that simply believing p constitutes having a reason$_2$, and thus that simply believing p is sufficient for 'taking' p to be a reason$_1$. But we saw in section 1.3 that this view might seem difficult to sustain in the face of cases where one believes p but does not notice that the truth of p is a reason$_1$ to believe q (which, we can suppose, p entails). How, in such a case, can one's belief in p constitute one's 'taking' p to be a reason$_1$ for believing q? Something additional to simply believing p seems to be required. For suppose that a subject, S, infers from p to q, and thereby comes to know q. Our rationalising explanation of S's belief in q would, it seems, have to refer to something over and above simply S's knowledge of p, since S might have known p without noticing that p entails q. This seems to suggest that our rationalising explanation would have to refer, further, to the fact that S 'took' p to entail q. And this suggests that the 'taking' condition is something over and above the belief in p – a substantial additional condition. If so, then having a reason$_2$ is a more complex matter than simply believing p: it must involve both believing p *and* this substantial additional condition. If this is the right way to interpret Ginsborg's view then her argument against McDowell must be that believing p is a necessary *part* of having a reason$_2$ to believe propositions that p entails or makes sufficiently probable.

More to the point of this chapter, this would seem to suggest that we need to flesh out 'taking' in terms of the doxastic or rule-following approaches (see section 1.3), and take our chances in doing so. We saw that both of these approaches have well-documented difficulties. These difficulties, together with the fact that Ginsborg's official position seems to be that merely believing p is sufficient for having a reason$_2$, are reasons to seek an account on which 'taking' is not a substantial condition additional to believing or knowing p. I'll now add further considerations counting against the idea that believing p is not sufficient for 'taking' p to be a reason$_1$. That is, I will argue that, in the case of inference, the 'taking' condition is not a substantial condition additional to the belief or knowledge that constitutes the reason$_2$ of the inference.

Suppose that 'taking' p to be a reason$_1$ *is* something additional to believing p. We've seen that it's difficult to say what this substantial condition amounts to, given the problems with doxastic and rule-following accounts. But we might just try to live with this. That is, we might say that what 'taking' p to be a reason$_1$ amounts to is mysterious – just another philosophical puzzle, which we might one day solve. There are two reasons why this is not a position that offers even temporary relief.

First, the putative mystery would itself be mysterious. Whatever 'taking' p to be a reason$_1$ amounts to, it is meant to be a person-level matter: it is

meant to be an aspect of the thinker's *perspective*, contributing to how things seem to them. For it is a *rationalising*-explanatory factor, and rationalising factors are supposed to be features of the subject's perspective. But then it surely shouldn't be difficult to discover what this rationalising-explanatory factor is. For surely we should be able to discover it by introspecting on our own first-person perspective when we undertake our own reasoning. It is therefore mysterious how there could be a mystery about what this further rationalising factor amounts to.

This is perhaps an inconclusive worry. For if 'taking' p to be a reason$_1$ is a substantial condition, additional to believing p, then it is probably a mental state – perhaps another belief. But subjects may have beliefs and other mental states that they don't know they have, and are not in a position to know that they have.[60] So perhaps the mysterious further rationalising-explanatory factor which we're calling 'taking p to be a reason$_1$' is an aspect of our perspective that we're not generally in a position to know about.

All the same, it should leave *some* trace in the thinker's perspective. When I infer from my belief in p to belief in q, I may not know that I *believe* p, but I do know that it's p I'm inferring from. What else enters my perspective as part of what I infer *from*? If 'taking' is a substantial additional condition, we should expect there to be an obvious answer to

[60] See the anti-luminosity argument in Williamson 2000 and Srinivasan 2015; see also Gopnik 1983; Carruthers 2011; Schwitzgebel 2011 for empirical considerations supporting Williamson's *a priori* argument.

this question, but there doesn't seem to be any such obvious answer. We can make suggestions, of course. For example, perhaps *the fact that p favours believing q* enters my perspective as part of what I infer from. But it's not entirely obvious that it does, or that it needs to. But if this is a crucial part of my inferential basis for belief, surely it *should* be self-evident to me, as an inferrer, that this is so. The fact that it is not (and that nothing else is self-evidently part of my inferential basis, besides p itself) suggests that there is no essential substantial additional rationalising condition in play when I infer.

Secondly, if we accept that a subject S's 'taking' p to be a reason$_1$ is an essential part of what, from a third-person perspective, rationalises S's believing (say) q, while also holding that it is mysterious what this 'taking' amounts to, we seem to have to admit that our rationalising explanations of subjects' beliefs are incomplete. That is, a putative rationalising explanation of S's belief in q will actually fail to display S's belief in q as rational or as making sense given S's perspective. For if we don't know what it is about S's perspective that makes her belief in q rational, then we haven't really rationalised her belief in q.[61]

[61] Note that Dennett, discussing how we make sense of the actions of others from the 'intentional stance', allows that an entity with a belief or a desire may not be in a position to know that it has a belief or a desire (since it may not be a 'second-order' intentional system, capable of beliefs about beliefs, or it may be a second-order intentional system that nevertheless lacks access to its whole system of beliefs), but holds that an *interpreter* of such a system must, in order to make sense of one of its actions or beliefs, know about all of the mental states which rationalise those actions or beliefs. I am making the same point about the requirements of interpretation. See Dennett 2009, especially pp. 345–6.

There are therefore reasons to want to avoid having to say both that 'taking' p to be a reason$_1$ is essential for having a reason$_2$ and that what 'taking' p to be a reason$_1$ amounts to is mysterious or is a philosophical puzzle that we need to work to resolve. It is thus advisable to take Ginsborg at her word when she represents a belief, by itself, as constituting a reason$_2$. This allows us to sidestep the difficulties associated with the doxastic and rule-following approaches to 'taking' and to take rationalising explanations in terms of beliefs to be complete rather than mysteriously elliptical or incomplete. Doing this, however, requires us to articulate a sense in which simply believing p makes it the case that p figures as a reason$_1$ in one's perspective, so that we might say that one 'takes' p to be a reason$_1$.[62] This is what my minimalist account of 'taking' p to be a reason$_1$ will try to do. In this respect, then, I suggest that it is friendly to Ginsborg, since it allows us to take her at her world when she represents beliefs, by themselves, as constituting reasons$_2$, and it lets us avoid the difficulties associated with not doing so.

2.2. Treating a (supposed) fact as a reason$_1$ in inference

How might believing p involve 'taking' p to be a reason$_1$? My starting point is that the normative significance of a fact which is a reason$_1$ for

[62] I should emphasise that I am not saying that the only way to 'take' p to be a reason$_1$ is to believe p. Indeed, I am preparing the way for showing that this is not the case.

belief is grounded in its being a fact. It's plausible that all facts, just in virtue of being facts, have normative significance with respect to what we should believe. If that's right it raises the following possibility. Perhaps believing p – representing p as being a fact – involves treating p as the kind of thing which has normative significance with respect to what one should believe, namely a reason$_1$. I don't mean that believing p somehow involves a further belief that p is a reason$_1$. Rather, my suggestion is that part of what it is to believe p is to *treat it* as having normative import for what to believe.

What is it to treat something as having normative import for what to believe? Surely it is to treat it as the kind of thing from which one can learn further facts about the world. Here is a first-pass suggestion about how one might do this. Suppose that one knows that the streets are wet, and one is thus treating this fact as having normative import for what to believe. One does this, according to this initial suggestion that I am making, by being disposed when enquiring about the world to ask oneself questions like, 'What follows from the fact that the streets are wet?', and 'What can I learn about the world, given that the streets are wet?', and 'Does the belief I'm inclined to form conflict with the fact that the streets are wet?', and so on. In general, one treats the fact (or supposed fact) in question as a consideration that must be taken into account and that can be a resource in one's enquiries about the world. In this way, even if one has not yet noticed that p is a reason$_1$ for any

particular belief, one can still treat p as having normative import with respect to what to believe in general, and so treat it as a reason$_1$.

Of course, believing p doesn't involve constantly asking oneself what follows from the fact that p. The suggestion is rather that it involves being disposed to ask oneself such questions in the course of enquiry. One is not always enquiring, so one is not always asking such questions. Furthermore, in the course of enquiry one does not ask what follows from every one of one's beliefs. Rather, one's understanding of their contents enables one to select, with more or less success, the beliefs that are most relevant for the topic of the particular enquiry at hand. For example, if one is wondering about recent weather conditions then one will consult all of one's beliefs which seem to one to have a good chance of producing knowledge about this topic. There is no guarantee that one will consult the right beliefs: one might omit to ask what follows from a proposition which one believes and which is relevant to the enquiry at hand. Our enquiries are more or less thorough. But part of what enquiry involves is asking what follows from one's beliefs; and, I am suggesting, part of what it is to believe p is to be disposed to ask what follows from p in the course of one's enquiries. In this sense, if one believes p then p figures in one's perspective as a reason$_1$: as a fact which has implications for what one should believe about the world. This is so despite the fact that believing p is not sufficient for realising that p is a reason$_1$ for any particular belief. Let us explore the advantages of this suggestion before considering extensions and qualifications that might be required.

This suggestion offers us a way of characterising the figuring of something as a reason$_1$ in a subject's perspective in less demanding terms than the doxastic and rule-following approaches. We can avoid the psychologically unrealistic requirement that the subject must actually believe that they have a reason$_1$, and we can similarly avoid the difficulties associated with rule-following accounts. Furthermore, we don't need to worry about the fact that, obviously, one can believe p without noticing that p supports believing q or whatever other proposition one's belief in p in fact supports believing. For one needn't notice these support relations in order to 'take' p to be a reason$_1$ as I am conceiving of 'taking'. Rather, the subject simply needs to treat the relevant proposition as having a certain kind of relevance for enquiry about the world. This means that they are disposed to ask the right kinds of questions about p and what follows from it in the course of future enquiries. And it's plausible that treating the proposition in this way is involved in believing it. In a nutshell, the idea is that part of what it is to believe a proposition to be true is to treat it as having a certain kind of normative import – the kind of normative import that makes it a reason$_1$ for belief.

This suggestion recovers the idea that merely believing p is sufficient for having a reason$_2$. By citing S's belief in p as a reason$_2$ in a rationalising explanation of her belief in q, we explain the subject's belief in terms of the relevant supposed reason$_1$ figuring in S's perspective as a reason$_1$ – as

something that makes appropriate other beliefs in general. In particular, by citing S's belief in p we rationalise S's belief in q in terms of the idea that, in forming that belief, S was asking themselves what followed from p. For example, if I say 'Hannah believes it recently rained because she believes that the streets are wet', I am explaining Hannah's belief that it recently rained in terms of its being a good answer to a question like: 'What does the fact that the streets are wet indicate about recent weather conditions?' This rationalising explanation is not omitting mention of a special 'taking' state in which the subject somehow 'takes' the fact that the streets are wet to be a reason$_1$ to believe that it recently rained and on this basis comes to believe that it recently rained. Nor is it leaving out a special 'taking' state otherwise conceived, or left mysterious. All that we need to know, in order to make sense of Hannah's belief, is that she was treating the streets being wet as a reason$_1$. Being told that she believed that the streets are wet tells us this, and so it tells us enough.

Of course, there must be other factors involved in Hannah's formation of the belief that it recently rained. Since this is not a case of deductive inference, it's likely that further background beliefs are playing an inferential role, and no doubt we could articulate this role in our rationalising explanation and the subject herself could articulate further premisses. In this respect our rationalising explanation *is* elliptical, but there is no great mystery about what is being left tacit. However, even in a case in which our rationalising explanation fully articulates the premisses of an inference, there are other factors involved in the

production of the belief. In Hannah's case, she must somehow be granted a *good answer* to the question 'What does the fact that the streets are wet indicate about recent weather conditions?' Similarly, in a case of deductive inference, the subject must somehow come to *recognise* that 'p' is a correct answer to a question like 'what follows from q?'

The further factors involved in granting the subject a good answer are, I suggest, below-the-hood: they do not figure in the subject's perspective and are not among her reasons for belief. They are sub-personal causal factors, not the stuff of rationalising explanation.[63] For example, it may be that, as one prominent cognitive scientific approach to inference holds, inference involves the (unconscious, sub-personal) manipulation of mental models of the world. Alternatively, it may be that inference involves (unconscious, sub-personal) transitions between sub-personal representations governed by rules built into the architecture of our

[63] See Drayson 2014 for an overview of the various different senses in which 'subpersonal' has been used since Dennett first drew the personal/subpersonal distinction. I have in mind the usage she associates with the 'Pittsburgh school', according to which there is a 'personal level' of psychological explanation in which we make actions and mental states intelligible in terms of normative reasons, and a 'subpersonal level' of explanation in which we explain phenomena in terms of causation and natural laws.

cognitive systems.[64] My approach is compatible with different theories of the sub-personal mechanisms that make inference possible. The role of these mechanisms does not mean that the subject herself is not involved in answering the question: the subject remains active throughout the process of inference. Indeed, it is one of the virtues of the question-asking account of 'taking' p to be a reason$_1$ in cases of inference that it captures the right mixture between activity and passivity on the part of the inferring subject. Let's explore this mixture in more detail.

[64] For the mental models theory see Johnson-Laird 2006. The spirit of the mental models theory is captured in Charles Sanders Pierce's idea that 'Deduction is that mode of reasoning which examines the state of things asserted in the premisses, forms a diagram of that state of things, perceives in the parts of the diagram relations not explicitly mentioned in the premisses, satisfies itself by mental experiments upon the diagram that these relations would always subsist, or at least would do so in a certain proportion of cases, and concludes their necessary, or probable, truth.' Cited in Johnson-Laird 2012: 135.

For the rule-following approach see Quilty-Dunn and Mandelbaum 2017. Note that this rule-following approach does not involve the idea that the *subject* is *actively following* rules. Rather, the rules are built into the subject's cognitive architecture. As Quilty-Dunn and Mandelbaum note, it's not clear whether on this approach it's apt to speak of the relevant rules being 'followed' at all – whether by the thinker or by a subsystem of the thinker (see p. 13). To use a common example: rules are presumably similarly built into the architecture governing the digestive process, but it would not be apt to talk of the subject as following these rules, and perhaps it would not be apt to talk of their digestive system as following them either.

2.3. Activity and passivity in memorial, inferential and perceptual knowledge: getting the mixture right

In conceiving of the relationship between active investigation and questioning on the part of the subject and the sub-personal information processing that this investigating initiates and involves, I draw on Naomi Eilan's attempt to do justice to the 'passive' and 'active' elements in remembering and perceptual belief-formation. Here is her account of purposeful remembering:

> Consider the following kinds of questions one might ask oneself in the course of a normal working hour. Where did I put the keys? What did Mary ask me to get from the shop? What is that student's name? […] Sometimes the answer comes, as we say, immediately, sometimes with irritating delay and sometimes not at all. Much of what passes for mental activity consists in asking oneself such questions and waiting for longer or shorter periods for the answer. Now, in all such cases we may speak of the answer 'coming' to us [… and] it is appropriate to speak of the thinker being 'saddled with content'.[65]

[65] Eilan 1998: 188.

According to Eilan, the thinker is active in asking and answering these questions, but there is also a passive element in the process, captured metaphorically by talk of being 'saddled' with content. She suggests that

> [t]he way to do non-metaphorical justice to this passivity while holding on to the active subject is to treat information-processing theories of the mechanisms involved in remembering as descriptions of non-conscious ingredients in what the subject actively does when asking and answering her own questions. The fact that there is this non-conscious cognitive ingredient just is what introduces the passivity. So reference to such mechanisms is not an optional extra relative to our everyday conception of remembering, but, rather, an extension of them.[66]

As Eilan presents matters, the subject is active all the way through the process of remembering: it is the subject who both asks and answers the questions. All the same, it is not clear just from introspecting *how* we answer our questions, beyond the fact that we simply remember the answer. According to Eilan, our 'bafflement' about how we remember is a result of the fact that we rely on unconscious, sub-personal information processing to come upon the answers to our questions. Asking questions, like focusing attention on the environment that is presented by perception, is a way of 'actively selecting information to be processed'[67];

---

[66] Ibid., 188–189.

[67] Ibid., 193. Cf. Campbell 2000, a paper which Eilan is drawing on.

so there is commerce between the personal and sub-personal levels, even if the work of the sub-personal mechanisms doesn't enter into our first-person perspective.

What Eilan calls the 'passivity' in remembering recalls BonJour's description of perceptual belief as 'cognitively spontaneous' and of perceptual beliefs simply 'striking us' with 'coercive' force (see the Introduction). The question-asking account of 'taking' p to be a reason$_1$ provides new ways of thinking about what is right about such metaphors, and of contrasting perceptual passivity with inferential activity.

As BonJour uses them, these metaphors are supposed to highlight a way in which perceptual belief differs from inferential belief. There is certainly something right about this. The image of perceptual beliefs as spontaneous, as just striking us, quite properly draws our attention to the fact that perceptual beliefs are normatively constrained by something outside thought in a direct rather than mediate way.[68] That is, when we form perceptual beliefs we do not ask questions about the normative import of facts that we already know or believe in order to come to know *further* facts, as I have been suggesting we do in inference. Instead, insofar as questions are involved in the process of perceptual knowledge-formation, these questions are directly about our immediate environment itself. Eilan suggests that perceiving involves asking the question 'What is going on here?', and extensions of this question. Perceptual beliefs thus

[68] This idea is the idea that McDowell adverts to when he says that inferential beliefs have only 'mediately empirical content'. See section 1.1.

seem spontaneous because we, as subjects with perceptual beliefs, cannot point to or cite some prior fact out of which the content of our perceptual belief has emerged as an evident corollary. There is nothing that we have available in thought that came *before* our perceptual belief, as its rational support. Instead, the content of our perceptual belief just strikes us as being the case.

In contrast, there is a kind of cognitive *work* that the subject is doing in inference that is absent in perceptual (and, arguably, memorial) belief-formation. My suggested account of 'taking' p to be a reason$_1$ helps show what this work amounts to. In inference, the subject themselves is representing the fact or supposed fact which they are asking questions about – questions like 'What follows from this fact?' By believing or knowing p, the subject is representing p as something with respect to which it makes sense to ask such questions. Representing p in this way is something in which the subject is active.[69] Consequently, when one infers from one's belief in p to a new belief in q, the new belief in q is not spontaneous in the same way that perceptual belief is, since it has emerged out of the subject's activity – her believing p – in a way that a

[69] A belief is a *state*, not a process or an event, and belief is not voluntary (not something we can have or stop having at will). So how can believing be an *activity*? See Boyle 2009 on belief as active yet non-voluntary, and Boyle 2011 on belief as active yet stative. As Boyle says, 'my beliefs are in important respects analogous to my actions themselves, rather than to objects on which I act – not because I can believe whatever I will myself to believe, but because both believing and willing are exercises of a more generic power of rational self-determination, a power that lies at the basis of both theoretical and practical self-control' (Boyle 2009: 121). The sense of 'active' I have in mind is just this: that believing involves the exercise of rational powers.

perceptual belief has not. Of course, this isn't to say that the subject is totally passive in perceptual belief-formation: there too, the subject asks and answers questions. And in both inference and perceptual knowledge, the subject's answer to the question involves their forming a belief, and in this respect the subject is active, too. It is just that in perception the subject is not active in representing a fact about which a question is asked.

At the same time, I would argue that there is a kind of passivity which inferential belief-formation has in common with non-inferential, and in particular perceptual, belief-formation. There is a sense in which the contents of inferential beliefs simply strike us as being true, or in which these contents spontaneously occur to us. When, say, I make a *modus ponens* inference, I can *just see* that the conclusion is correct in the light of the premises. Of course I can articulate the premisses and I can see that they are the source of my supposed knowledge of the conclusion. And, at least according to the question-asking account of 'taking' p to be a reason$_1$ in inference, I initiated the inference by asking the right kind of question about the premisses, or about the supposed fact which is constituted by their conjunction. But I can't articulate any *further* reason for thinking that the conclusion follows from the premisses, and I may not be able to explain what it is about the premisses that make them guarantee the truth of the conclusion. Nor can I explain *how* I can tell that the conclusion follows from the premisses. There is thus a kind of passivity here. I attend to the premisses, ask what follows from them, and

the answer simply arrives – I do not know how. My suggestion, inspired by Eilan's account of the passivity in remembering and perceiving, is that we should do 'non-metaphorical justice to this passivity while holding on to the active subject' by treating theories of the mechanisms involved in inferring (such as the mental models theory) as 'descriptions of non-conscious ingredients in what the subject actively does when asking and answering her own questions' (the quoted words are from what Eilan says about remembering and perceiving, not inferring).

The account of 'taking' p to be a reason$_1$ (and so having a reason$_2$) in inference that I am offering fits well with the fact that inference involves this element of passivity. Less minimalist approaches to 'taking' p to be a reason$_1$, such as the doxastic and rule-following accounts previously discussed, seem partly to be driven by a desire to deny that our notion of inference involves the idea that inference essentially involves sub-personal mechanisms.[70] Boghossian, for example, says that his 'taking' condition on inference must be a substantial person-level state, additional to one's belief in or knowledge of the premise(s) – a further belief or an

[70] Compare what Eilan says about a way of thinking about remembering, according to which 'our everyday "personal level" concept [of remembering] is wholly independent of any reference to the [sub-personal] mechanisms that make remembering possible.' Eilan complains that this way of thinking of remembering is inadequate because it can't do justice to the element of passivity in remembering. Eilan 1998: 188. I am suggesting that there is a similar tendency to mistakenly think that the concept of *inference* is wholly independent of any reference to the sub-personal mechanisms that make it possible.

intuitional state[71] or a person-level following of a rule involving representation of the rule in question – because this is the only way that we can make sense of inference as 'person-level reasoning, reasoning as a mental action that a person *performs*, in which he is either *aware*, or can become aware, of why he is moving from some beliefs to others'.[72] Since merely believing p is, as we have seen, not constitutive of nor sufficient for the subject's recognising that p is a reason$_1$ to believe q (where q is a proposition that p entails), Boghossian assumes that there must be some further person-level state or activity which constitutes the subject's awareness of why moving from belief in p to belief in q is appropriate, and which helps explain why the subject makes this transition. I am denying that there must be such an additional person-level state or rule-following activity, constituting the subject's awareness of why moving from belief in p to belief in q is appropriate, which plays an explanatory role in the subject's actually transitioning to her new belief in q. Instead, I suggest, we should make do with the subject's treating p as a reason$_1$ by asking what follows from it, and then accepting the answer by coming to believe it. And I am suggesting that this story about inference is supported by the passivity that we find when we introspect on the inferential process.

---

[71] Above I discussed only the doxastic and rule-following approaches, but Boghossian also mentions the possibility of an 'intuitional construal of taking', according to which one has an 'intellectual seeming' that such-and-such a supposed fact supports believing such-and-such a proposition. See Boghossian 2014: 9–10.

[72] Boghossian 2014: 16. Boghossian's emphases.

Let's take stock. In the previous chapter we saw that Ginsborg emphasises that a rationalising explanation a subject's belief involves displaying that belief as rational in the light of the subject's perspective, or of how things seem to the subject. In this context, what matters for Ginsborg about the subject's perspective is that the subject 'take' something to be a reason$_1$ for the belief in question. This is supposed to rationally motivate them to form the relevant belief. But Ginsborg says little about what this 'taking' involves. What she does say is that the putative reason$_1$ must be the content of one of the subject's states of belief or knowledge. And she implies that the subject's having this state of belief or knowledge is sufficient for the subject 'taking' its content to be a reason$_1$ for the belief we're trying to explain.

But we noted (in sections 1.4 and 2.1) that merely believing p surely doesn't by itself make one 'take' p to be a reason$_1$ *for a specific further belief*. It seems that one can believe p without noticing that p's truth has a normative bearing on whether one should believe, say, q (where, in fact, p entails q). That made it difficult to see how Ginsborg can sustain her view that merely believing p makes one 'take' p to be a reason$_1$ for whichever belief the subject subsequently forms on its basis. And this seemed to suggest that, on Ginsborg's account of believing for a reason, a mere state of belief or knowledge of p can't, by itself, rationally motivate a subject to believe q; and, as a corollary, it seemed to suggest that a rationalising explanation of a subject's inferential belief in q must appeal to more than just the subject's belief in p. This leads to a

conception of 'taking' p to be a reason$_1$ for belief according to which this 'taking' is something additional to simply believing or knowing p. But, I argued, this leads to difficulties. Besides the well-documented problems with doxastic and rule-following accounts of 'taking', the fact is that we *do* employ rationalising explanations of beliefs (and actions) which, at least on the surface, appeal only to a subject's belief or knowledge of a premise (or a conjunction of premisses), as in 'Hannah believes it recently rained because she knows the streets are wet'. If these rationalising explanations are to be complete and not mysteriously elliptical, then believing or knowing p must be sufficient for making p figure as a reason$_1$ (or supposed reason$_1$) in one's perspective, such that believing or knowing p can rationally motivate one to form a belief in q (which, we're supposing, p entails). The challenge, then, is to articulate a 'taking' condition which merely believing or knowing p fulfils and which is sufficient to rationally motivate further beliefs via inference.

Section 2.2 took up this challenge. The suggestion there was that since believing or knowing p involves representing p as a fact or as true, and since being a fact is what grounds the normative import of a fact, believing or knowing p involves *treating* p as a fact, and in particular treating it as having normative import for what one should believe. Treating p as having normative import for what to believe can involve having a number of different dispositions. The one I focused on was the disposition to, when enquiring about the world, ask questions about what the truth of p makes it appropriate for one to believe. Insofar as believing

p endows one with such a disposition, it ensures that p figures in one's perspective as a reason$_1$ for belief – though not as a reason$_1$ for a *specific* belief. And, I claimed, this understanding of what it is for a supposed fact to figure as a reason$_1$ in a subject's perspective is sufficient for enabling this feature of a subject's perspective to play the role of reasons$_2$ in complete rationalising explanations of beliefs.

This characterisation of 'taking' p to be a reason$_1$ is quite minimalist. In particular, it does not involve the subject appreciating, prior to making the inference in question, that p's truth supports believing whatever proposition concludes the inference's argument. Consequently, there must be an explanatory role in inference for something that lies outside of the subject's perspective, which makes it the case that the subject, on asking the question, gets the right answer. Acknowledging this challenges us to map the boundaries lying between the person-level, active elements in inference, and the sub-personal, passive elements. Section 2.3 has made a start on doing this (I acknowledge that it is only a start). My main claim here has been that while inferential belief-formation is active in ways that non-inferential belief-formation is not, inferential belief-formation also shares passive elements with non-inferential belief-formation. And I have claimed that my minimalist characterisation of 'taking' p to be a reason$_1$, which requires us to acknowledge the role of sub-personal processes in inference, fits well with this element of passivity in inference. Next, we will extend the minimalist characterisation of 'taking', in order to better describe the role of 'taking' in inference.

2.4. Regress and instructions

I mentioned that there are various dispositions involved in treating p as having normative import with respect to what to believe: being disposed to ask questions of the sort we've been discussing is not the end of the matter. This is important because, on reflection, it's clear that the question-asking account of 'taking' p to be a reason$_1$ is not adequate to the task of characterising the role of reasons$_2$ in inference. On the account we've been working with so far, when one infers from p, one begins by asking what follows from the fact that p. Supposing that the answer this question yields is 'q', then, as I've been presenting matters, one simply forms a belief in q. But that can't be quite right. The question asks what follows from the fact that p, so strictly speaking the answer must be of the form 'q follows from the fact that p'. If, on my account of inference, one first comes to believe that *q follows from the fact that p*, and then on this basis infers *q*, then I am after all smuggling a substantial doxastic taking-state into the account, and vicious regress looms.[73]

[73] A vicious regress looms because we can ask how one comes to know q, having come to know that q follows from the fact that p. It seems that I would have to answer by saying that one asks oneself 'What follows from the fact that q follows from the fact that p?', to which the answer would be 'q follows from the fact that q follows from the fact that p' – and so the regress is launched. Cf. note 48 above.

There seem to me to be two avenues of response to this worry, one concessionary and involving an extension of my view, the other less concessionary and involving a clarification of my view. In the end, I think that the concessionary reply is required, but let's quickly consider the less concessionary reply, since it appeals to a feature of McDowell's picture of perceptual knowledge.

We saw in section 1.2 that McDowell holds that one comes to know that one sees that o is F in the 'very same act' that she comes to know that o is F . I said that if this is right then it is difficult to see how one's knowledge that one sees that o is F could be part of what explains one's knowledge that o is F; and I agreed with Brewer that it is difficult to see how one could come to have knowledge of two different contents in a single cognitive act. Nevertheless, it is tempting to say something similar about inferential knowledge as I have been describing it. The idea would be that when the answer, q, comes to one's question 'What follows from the fact that p?', one would come to know, as part of the same cognitive act, both that q follows from the fact that p *and* q itself. As I said about McDowell's version of this move, it makes it difficult to see how S's knowledge that q follows from the fact that p could be something that an interpreter could appeal to in a rationalising explanation of S's belief in *q*. But it's not clear that this matters in the present context: I have been defending the idea that a rationalising explanation of S's belief in q which appeals simply to S's belief in p (where p entails q) is a complete rationalising explanation of S's belief in q, not in need of

supplementation. S's knowledge that q follows from the fact that p might simply be a side-effect of her coming to know q in this inferential way: she knows that she asked what follows from the fact that p, and she knows that the answer to this question is q, and so she knows that q follows from the fact that p, but her purpose in asking the question was not primarily to find out about this connection between p and another fact, but rather to find out the further fact, q.

The problem remains, however, that it's simply not clear how a single cognitive act can produce knowledge of two different contents. Nor is it clear how a subject can rationally believe q in this way without making a further inferential step from her question's answer (that q follows from the fact that p), thus setting off the vicious regress.

Instead, then, a more concessionary response to the problem extends the account of 'taking' p to be a reason$_1$ in the following way. I suggest that, as well as being disposed to ask questions like 'What follows from the fact that p?', taking p to be a reason$_1$ involves the subject being disposed to *give instructions*. If p is the putative fact in question, the instruction might be: 'Extract further truths from the fact that p'. Note that 'p' might be the conjunction of several premisses: the subject can offer more or less information in the instruction. Furthermore the instruction can be more or less open-ended. For example, if the fact in question is that the streets are wet, the instruction might be 'Extract further truths from the fact

that the streets are wet', or it might be 'Extract further truths about the recent weather from the fact that the streets are wet'.

What or who is one giving such instructions *to*? Recall Eilan's claim that 'Much of what passes for mental activity consists in asking oneself […] questions and waiting for longer or shorter periods for the answer.' In the same way, I am suggesting, we can give ourselves instructions. The idea that inferring involves giving oneself instructions also parallels information-processing approaches to selective attention, which Eilan makes use of in her account of the role of attention in perception. As she puts it, the idea is that 'attention is the selection of information for further processing'; and she follows John Campbell in suggesting that, as with remembering, 'the information-processing accounts of the mechanisms involved in such selection and further processing [are] the non-conscious cognitive components of what the subject does when attending'.[74] My suggestion is that just as in attending one can select information for processing, so one can instruct oneself to extract information from facts that one already knows. And, further, I am suggesting that giving such instructions can be a way of treating p as a fact, and thus 'taking' it to be a reason$_1$, and that this is what one does when one makes inferences.

If these suggestions can be sustained then they solve the problem we confronted. Whereas the answers to the questions that I described would

[74] Eilan 1998: 192. Cf. Campbell 2000; Campbell 2002.

be of the form 'q is what follows from p', the outputs of the sub-personal process initiated by the kind of instruction I described would be of the form 'q is a fact'. There is thus no implication that in cases of inference one could only come rationally to believe q via a doxastic 'taking' state.

Note that I am not retreating from the idea that part of what is involved in 'taking' p to be a $reason_1$ (and thus part of what is involved in believing p) is being disposed to ask questions of the sort discussed earlier. What I am retreating from is the suggestion that such question-asking and question-answering is the subject's canonical contribution to inference; instead, in inference, a subject gives instructions and takes up the results of the processes thus initiated by forming beliefs. Still, part of being a reflective thinker is to ask questions like 'What follows from the fact that p?' By asking and answering such questions, a thinker can perform an inference in a more reflective way. For, in this way, they might come to know that (the fact that) q follows from (the fact that) p. The thinker can then conclude that q is true by instructing themselves to extract further truths from this fact (that is, from the fact that (the fact that) q follows from (the fact that) p).[75] Questing-asking of this sort can therefore play a role in inferential enquiry. But to avoid vicious regress instruction-giving must kick in at some point.

[75] I include 'the fact that', rather than writing simply 'q follows from p', to make explicit that the thinker is thinking with contents that she takes to be facts. One might know that q follows from p but deny that either of p or q are true.

2.5. Enriching the thinker's perspective

One might worry that the instruction-giving account of inferring simply outsources the work of inferring to sub-personal mechanisms. The subject's appreciation of the rational significance of the premisses of the inference might seem not to play enough of a role in rationally motivating their inferential belief. By removing any role for a doxastic 'taking'-state in rationally motivating the subject's uptake of whatever contents are produced by the sub-personal process, we seem to have removed any rational-explanatory role for the subject's grasp of the contents involved and consequent understanding of rational connections holding between them. Suppose, for example, that S instructs herself to 'Extract truths from the fact that p', and this instruction initiates a process whose output is 'q', which S thus comes to believe. According to this description of the inferential process, S herself has no idea whether the relationship between p and q is one of entailment or something else; so she has no idea whether she is making a deductive or an ampliative inference. Can we really attribute the inference to *her*, in this case? Can she really say that p is *her* reason_1 for believing q, and can outside interpreters really say that her belief in p is her reason_2 for believing q?

There are two parts to this worry. First, it seems that the subject's appreciation of the rational connection between premisses and

conclusion, such as it is, plays too little a role in rationally motivating the subject to believe the conclusion. A corollary of this is that outside interpreters cannot appeal, in rationalising explanations of inferential beliefs, to the subject's appreciation of rational connections holding between premisses and conclusion. Second, even if we sacrifice the idea that the subject's having a rich appreciation of such rational connections rationally motivates (and so explains in a rationalising way) her belief in the inference's conclusion, we might still want to insist that, *as part of* or *as a result of* her inference from p to q, she comes to appreciate that a rational connection holds between the truth of p and her belief in q. But the account of inference in terms of instruction-giving seems not to have the resources to explain how inferring q from p could give rise to any such appreciation. For, from the subject's perspective, q is simply the output of a process initiated by an instruction to extract truths from p: no further specification of the relationship between q and p figures in her perspective.

As an initial reply to the first part of the worry, I would re-emphasise that the account of inference in terms of instruction-giving, as it has been presented so far, does not totally deprive the subject of appreciation of the rational significance of the premisses of the inference, even before she has come to believe the conclusion. For by instructing that truths be extracted from the premisses (that is, from their conjunction), the subject treats the conjunction of the premisses as the kind of thing which makes appropriate certain further beliefs. In this sense, the subject does treat the

conjunction of the premisses as a reason$_1$, and it figures as such in her perspective. Below I'll say more about how its so figuring rationally motivates the subject to believe the conclusion of the inference.

As an initial reply to the second part of the worry, I deny that the instruction-giving account of inference lacks the resources to explain how inferring q from p could give rise to an appreciation of the fact that the truth of p supports believing q, and is one's own reason$_1$ for believing q. On the contrary, when S comes to believe q as a result of having instructed herself to extract a truth from the fact that p, this puts S in a position to recognise q *as* the output of the inferential process. The purpose and function of the inferential process is to produce a true content *in the light of* p's truth. S doesn't just find herself believing q as though out of nothing; rather, they find themselves believing q as the outcome of this process. Since they themselves initiated and specified the point of the process, they know its purpose. Thus, although on my account there is no doxastic 'taking' state which mediates, as a rational motivator, between the premisses and the conclusion-belief of the subject's inference, carrying out the inference does, at least in standard cases involving reflective adult thinkers, put the subject in a position to know that p is not just a reason$_1$ but *her* reason$_1$ *for believing q*. Actually, in my view, a subject can inferentially know q on the basis of p without knowing that p is her reason$_1$ for believing q; but it should now be clear that my account of inference can accommodate the view of those

philosophers who do hold that S must know that p is her reason$_1$ for believing q in order for her inferential belief in q to be knowledgeable.

These responses have the limitation of only giving the subject a very basic appreciation of the kind of rational connection obtaining between knowledge of p and belief in q. For all I've said, the subject has no idea, either before or after performing the inference, whether the inference is deductive or ampliative. However, the subject's appreciation of the rational connection in play can be made richer by building more into the instruction that they give. For example, instead of simply 'Extract truths from the fact that p', the subject might give the instruction 'Extract truths which are entailed by the fact that p', or 'Extract truths which are strongly indicated but not entailed by the fact that p'. A subject who gives the instruction which mentions entailment, say, is thereby specifying and directing the process in which they are engaged: they are consciously *deducing*. When their instruction produces a result, say q, this (at least in standard cases[76]) puts them in a position to know that p *entails* q, since they know that their instruction specified the production of an entailment. Again, however, it's not that they come to know that p entails q and infer q from this[77]; rather this is something they might come to

[76] Since I accept that no non-trivial condition is luminous, I do not rule out the possibility that there are possible cases where the subject infers q from p in this way but is not in a position to know that p is her reason$_1$ for believing q.

[77] Of course, nothing in my account precludes this from taking place: a subject might well instruct himself to produce an entailment of p, and, on receiving 'q' as the answer, come to believe that p entails q, and then make a further inference to belief in q. However, my account would analyse this process in terms of *two* instructions: the initial instruction to find entailments of p, and then a subsequent instruction to find entailments of $p \wedge (p \rightarrow q)$.

know as a result of making their inference. Thus, the subject's knowledge that p entails q is not part of what rationally motivates their belief in q.

However, by allowing that the subject can give richer and less rich instructions, we do also allow that both the subject's person-level grasp of concepts like *entailment* and their grasp of the content of the belief from which the are inferring can have roles in rationally motivating their inference. There are in fact two different ways in which this is so.

First, the subject may need to give more highly specified instructions in order to successfully perform an inference. They may, for example, need to try to extract entailments from a fact they already know, rather than merely vaguely instructing themselves to extract further truths from it. Similarly, they may need to specify, with more or less precision, the topic or subject-matter of the truths that they want to extract. Earlier, for example, we saw that the subject might instruct themselves not just to extract truths from the fact that the streets are wet, but rather to extract truths about the recent weather from the fact that the streets are wet. By giving more highly specified instructions, the subject offers a greater degree of direction to the sub-personal mechanisms involved in inference, and this may enable these mechanisms to produce results that less specific instructions would be unable to extract from them. The intuitive idea here is that by specifying in their instruction that they are looking specifically for, say, entailments or truths about a particular topic, they narrow down the scope of the enquiry and make the problem more

tractable. I take this to be a realistic picture of how person-level thinking interacts with sub-personal processing to produce new beliefs. Now, by giving these more highly specified instructions, the subject not only 'takes' p to be a fact which makes appropriate further beliefs, but also 'takes' p to be a fact which, say, *entails* other facts (in a case in which the instruction mentions entailment). And the subject's 'taking' p to entail further facts might be part of what explains her ability to draw the inference in question, since, as we saw, it may be that only a highly specific instruction can successfully produce the inference in question. In this way, then, the subject's appreciation of the rational import of p is a crucial part of the explanation of her performance of the inference. In general, the greater the degree of a subject's conceptual sophistication, the richer the instructions they can give; this explains why more conceptually sophisticated subjects may be able to make inferences that less conceptually sophisticated subjects cannot.

Second, the subject *believes* the conclusion of the inference because doing so is part of the inferential process which she set in motion by giving her instruction. She initiates the inferential process with the intention of coming to have a new belief, and this makes rational sense for her to do because she 'takes' herself to have a reason$_1$ – a fact which makes appropriate further beliefs. So part of what explains her forming her inferential belief is her having given her instruction, which is itself part of what it is for p to figure in her perspective as a reason$_1$ and (if the

instruction specifies that she seeks, say, an entailment) as a fact which entails other facts.

In a nutshell, then, in cases of inference the subject's appreciation (however rich) of the rational import of something she believes, constituted in part by her instruction-giving, explains both the successful performance of her sub-personal inferential mechanisms, and her forming a belief in the conclusion of the inference.

Note that I am not proposing a level of richness that a subject's instruction must reach for the subject to count as performing an inference. After all, it may be that subjects with quite low degrees of conceptual sophistication can make inferences. Rather, the point is that the instruction-giving account of inference allows for deeper degrees of appreciation on the part of the thinker of the kind of inference they are undertaking and the kind of rational connection that obtains between premisses and conclusions, and I am suggesting that we can make sense of this appreciation on the part of the thinker being part of what rationally motivates and enables the process and the concluding belief. In turn, we can, at least in principle, appeal to this aspect of the subject's perspective in our rationalising explanations of their inferential beliefs. In practice, it will often be difficult for an external interpreter to know exactly how rich the subject's appreciation is of the kind of inference they are performing. However, if the subject is a young child, for example, then the interpreter should be in a position to know that their

appreciation of the kind of inference they are performing is not very rich.

The worry was that the instruction-giving account of inferring gives the subject too minimal an appreciation of the rational connection between her premises and her conclusion-belief, and gives whatever appreciation she does have too minimal a role in rationally motivating her conclusion-belief. It should now be evident that this worry is misplaced on both counts.

A further worry concerns the question of how psychologically realistic the instruction-giving account of inference is. We rarely explicitly internally voice instructions or questions of the sort I've been discussing. Eilan addresses a similar worry about her own account of question-asking in remembering and perceiving: 'it does not seem right to say that one is incessantly interrogating the environment, and only has experiences as a consequence of such non-stop interrogation. Surely one just soaks up what is going on in a more relaxed way.'[78] In response to this worry she suggests that 'perceiving is something we learn to do', and that 'much of what this involves is the acquisition of routines of specific questions which become automatic or habitual extensions of the general question: what is going on here? (Somewhat in the manner in which components of a physical skill become automatic).'[79] I would give a

[78] Eilan 1998: 194.

[79] Ibid., 195.

parallel response to the worry as it applies to my account of inference. The idea is that inferring is something we learn to do, and that it thereby becomes relatively automatic and habitual, and thus does not involve the explicit or conscious voicing of instructions and questions. But these instructions and questions still structure the process of inference, and we can always slow down and proceed more reflectively. When we do this, I suggest, we do find ourselves giving instructions and asking questions of the kind I've been highlighting.

2.6. Two puzzles about McDowell's picture of perceptual knowledge

We can now compare my picture of what is involved in 'taking' p to be a $reason_1$ and forming an inferential belief on the basis of p, with some aspects of McDowell's picture of perceptual knowledge. In particular, Chapter 1 described two puzzles about his picture.

The first puzzle has to do with his claim that a perceiving subject 'comes to know that she sees that o is F in the very same act in which she comes to know that o is F'.[80] This is puzzling because it's unclear how a single

[80] McDowell 2019: 340.

cognitive act can produce these two different states of knowledge (see section 1.2).

Above (section 2.4) I mentioned that it's tempting to say something similar about inferential knowledge, but I admitted that it remains unclear how a single cognitive act can produce knowledge of two different contents. But it does seem to me that a watered-down version of this claim is available. On the instruction-giving account, one infers by instructing oneself to extract a truth from something one already knows or believes. It's plausible that an instruction like 'Extract a truth which is entailed by the fact that p' could be used both to infer q (where p entails q) *and* to recognise that p entails q, and that p is a reason$_1$ to believe q. That is, it seems plausible that one could give this instruction with the intention of recognising what p entails or what p is a reason$_1$ to believe, rather than or as well as with the intention of drawing an inference. Which intention one has in giving the instruction will determine what belief one forms at the conclusion of the process: one might simply believe q, or one might believe that p entails q or that p is a reason$_1$ to believe q. The content of the instruction is such that it would make sense to form any of these beliefs. Thus, the same instruction, and the same sub-personal process initiated by this instruction, can produce different rational beliefs. In particular, it can produce a rational belief in q, and it can also produce a rational belief that p is a reason$_1$ to believe q.

It wouldn't be quite right to say that these rational beliefs can therefore be formed in the same cognitive act. Rather, what seems right to say here is that they can be formed via the same capacity, in different acts. Furthermore, the beliefs in question are not quite inferential analogues of the beliefs McDowell says are formed in the same cognitive act. The beliefs I've suggested can be rationally formed using the same instruction are a belief in q (which is entailed by p) and a belief that p entails q or that p is a reason$_1$ to believe q. In contrast, McDowell suggested that the same cognitive act can induce knowledge of q (where q is the content of a perceptual experience) and knowledge that one sees that q. The fact that one sees that q is what McDowell thinks is one's reason for believing q. Despite this difference, the idea in both the inferential and the perceptual case is that one can come to know both q and can come to know something about one's rational basis for believing q, and that one can do so via the same capacity (on my view) or through the same cognitive act (on McDowell's view).

It's important for McDowell that one does come to know both that o is F and that one sees that o is F in the same cognitive act. This is partly because he thinks that a subject can only believe that o is F on the basis of the fact that she sees that o is F if she knows that the fact that she sees that o is F is her reason (what Ginsborg would call her reason$_1$) for believing that o is F. And he thinks that she can only know that the fact that she sees that o is F is her reason for believing that o is F if she knows that she sees that o is F. Consequently, she cannot perceptually know that

o is F before she knows that she sees that o is F. But, McDowell thinks, nor can she know that she sees that o is F before she knows that o is F. So her knowledge that she sees that o is F must arrive together with her knowledge that o is F. This is surely part of what leads him to make the puzzling claim that the perceiving subject comes to know both of these facts in the same cognitive act.

However, perhaps McDowell can drop this puzzling claim and replace it with a perceptual analogue of the watered down version of it that my account of inference makes available. For perhaps having a factive perceptual experience puts one in a position to know both that o is F and that one sees that o is F, and perhaps McDowell could say that one must simultaneously come to know both of these things in order to come to perceptually know either of them. This doesn't mean that one comes to know both of them in the same cognitive act; rather, one comes to know them in two simultaneous cognitive acts, using the same capacity to extract information from one's experience. And perhaps something similar happens in inference: one simultaneously comes to know both q and that one's reason for believing q is p. In my view, though, it is not necessary to make this claim, in either the perceptual or the inferential case. This is because I do not agree with McDowell that in order to believe q for the reason that p one must know that p is one's reason (what Ginsborg would call one's reason$_1$) for believing q. This brings us to the second puzzle in McDowell's view.

As we just saw, McDowell thinks that a subject can only believe q on the basis of p if she knows that p is her reason (what Ginsborg would call her reason$_1$) for believing q. The subject must, as it were, embrace in thought her reason$_1$, in order for it to be *her* reason$_1$ for belief. This is the first part of his internalism (see section 1.1). But, at the same time, according to the second part of his internalism, for a subject's belief in q to be a belief for the reason that p, it must be the case that had she not formed her belief in q she still would have been in a position to believe q for the reason that p and to recognise the normative significance of the fact that p. In the case of perceptual knowledge, McDowell holds that this second internalist condition is fulfilled in virtue of the subject's enjoying a non-doxastic factive perceptual experience of seeing-that-p. Pre-cognitive seeing-that-p is sufficient, he says, for possessing the fact that one sees that p as a reason (or a reason$_1$) for belief: the seeing-that state has its 'rational significance' for the subject regardless of whether she embraces it in thought, which is to say regardless of whether she comes to believe or know that she sees that p. The second puzzle about McDowell's view has to do with the apparent difficulty about motivation that (in my discussion of Littlejohn's criticism of McDowell in section 1.2) I suggested might be thought to afflict this two-part internalism.

The difficulty has to do with how to motivate the first part of the internalism given the assumption that McDowell makes in affirming that the condition in the second part of his internalism is fulfilled. He assumes that the subject can possess a reason for belief without

embracing the reason in thought, yet the first part of his internalism demands that the subject embrace the reason in thought in order to possess it as *her* reason for belief. Why is this necessary, given that a perceptual experience, lying outside thought, can make a reason one's own reason? Or, from the other direction, if McDowell accepts the more demanding requirements for reason-possession contained in the first part of his internalism, why is he satisfied that one can possess a reason for belief in virtue of meeting the less demanding requirement contained in the second part of his internalism?

I think that something like this worry is part of what animates Ginsborg and Littlejohn's doxasticism about reasons, according to which only beliefs or states of knowledge can be reasons (or reasons$_2$ in Ginsborg's parlance) for beliefs. It can seem that McDowell is forced into an unprincipled two-part internalism by the need that he feels to make room for reasons that lie outside thought – reasons that fulfil the Reasons Outside Thought condition. If we forsake the enterprise of locating reasons outside thought, we needn't claim that one can possess a reason without that reason being a fact or supposed fact that the subject knows or at least believes. That is what Ginsborg and Littlejohn recommend that we do.

However, depending on how it is fleshed out, my account of the role of 'taking' p to be a reason$_1$ in inference may generate a two-part internalism that is similar to McDowell's, and which faces a parallel

challenge of motivation. This suggests that the impetus for his two-part internalism, and the challenge this creates, is not his devotion to the project of describing reasons that lie outside thought, since my account of inference is an account of believing for reasons which lie *inside* thought.

To see that my account can face a parallel challenge, recall first that my account of 'taking' oneself to have a reason$_1$ does not include a doxastic 'taking' state, in which one believes or knows that one has a reason$_1$, or in which one believes or knows that p is a reason$_1$ to believe q. Consequently, on my view, no such doxastic 'taking' state is involved in rationally motivating one to believe the conclusion of the inference, and thus no such doxastic 'taking' state can be appealed to in rationalising explanations of the subject's inferential belief. In a word, on my account, one's belief in p is a reason$_2$ one possesses to believe q (where p entails q) – no *further* belief connecting p and q is required. Yet we might well hold that for one to believe q for the reason$_1$ that p one must at least implicitly recognise (and thus believe) that p is a reason$_1$ to believe q. That is, we might well hold that to have a rational inferential belief the subject must also have a doxastic 'taking' state of the kind that I have argued is absent in the rational motivation of the inferential belief. After all, Littlejohn, Ginsborg and McDowell all agree that to believe q for the reason$_1$ that p one must be disposed to cite p as one's reason$_1$ for believing q.[81] In order

[81] See Littlejohn 2018: 123; Ginsborg 2006: *passim*; Ginsborg 2011: 144–5.

to be so disposed, it might be thought, one surely needs to believe that p is one's reason$_1$ for believing q, and thus that p is *a* reason$_1$ to believe q.

If I were to endorse this claim then I would be committed to a two-part internalism. According to the first part of this internalism, for one to believe q for the reason$_1$ that p, one would need to know that p is one's reason$_1$ for believing q. According to the second part of this internalism, for one to believe q for the reason$_1$ that p, it must be the case that had one not formed one's belief in q, one would still be in a position to believe q for the reason that p. And I claim that this second internalist condition is fulfilled by the subject's simply believing or knowing p, which involves the subject 'taking' (in the non-doxastic sense of 'taking' I've been developing) p to be a reason$_1$, but does *not* involve a doxastic 'taking' state and does not involve the subject 'taking' p to be a reason$_1$ *to believe q*.

This picture faces a motivational challenge parallel to that facing McDowell. How is is that one can be in a position to believe q for the reason that p while *lacking* a doxastic 'taking' state in which one 'takes' p to favour believing q, and yet one can only actually believe q for the reason that p if one *has* such a doxastic 'taking' state? The parallel question for McDowell is: how is it that one can be in a position to (knowledgeably) believe q for the reason that p while *not* believing p, and yet one can only actually believe q for the reason that p if one *does* believe p?

One might push back on the claim that in order to be disposed to cite p as one's reason$_1$ for believing q, one must believe that p is one's reason$_1$ for believing q and that p is a reason$_1$ to believe q. Littlejohn says only that someone who believes q for the reason that p must be disposed to answer 'p' in response to questions like 'What's good about believing q?' or 'Why do you believe q?' Surely a subject can do this without possessing the concept *reason* and so without believing that p is their reason$_1$ to believe q. That seems right to me. But this doesn't mean that I can avoid two-part internalism. For on my view, one need not be disposed to answer 'p' in response to questions like 'What's good about believing q?', in order to 'take' p to be a reason$_1$ in the way required to have a reason$_2$ to believe q. On my view, one simply needs to be disposed to ask the right kinds of questions and give the right kinds of instructions in relation to p, and I've argued that simply believing or knowing p is sufficient (not necessary) to be so disposed. But believing or knowing p does not by itself dispose one to answer 'p' in response to the question 'What's good about believing q?' For one might believe or know p but not yet have any idea that p is a reason$_1$ to believe q. So if I were to agree with Littlejohn that to believe q for the reason that p one must be disposed to answer 'p' in response to the question 'What's good about believing q?', then I would be committed to a two-part internalism. The motivational challenge would now be to answer the following question: How is it that one can be in a position to believe q for the reason that p while *not* being disposed to answer 'p' in response to questions like

'What's good about believing q?', and yet one can only actually believe q for the reason that p if one *is* so disposed?

I won't offer a response to this motivational challenge here. For one thing, as I said in relation to the first puzzle about McDowell's view, I simply do not agree with him that in order to believe q for the reason$_1$ that p one must know that p is one's reason$_1$ for believing q. Nor am I convinced that in order to believe q for the reason$_1$ that p one must be disposed to cite p as one's reason$_1$ for believing q. It is one thing to perform some mental action, like an inference; it is another thing to form accurate higher-order beliefs about that mental action. One might be mistaken about one's own mental states and actions, and about their interrelations. So I do not endorse the first part of McDowell's internalism, and I do not want to commit myself to any of the various two-part internalisms just described.

But my main concern here is not to argue for or against McDowell's two-part internalism. I want to stay official neutral on that matter, so that everything I say can be accepted by the many philosophers who do think that believing q for the reason that p involves knowing p to one's reason$_1$ for believing q or at least involves being disposed to cite p in response to the question 'Why do you believe q?' Rather, my concern is to show that *if* one accepts one these claims – whether McDowell's full-blooded claim about knowing what one's reasons$_1$ are or Littlejohn's weaker claim about being disposed to cite one's reason$_1$ – then one will struggle to

avoid committing oneself to a two-part internalism, whether or not one seeks to identify reasons that fulfil Reasons Outside Thought. For I have argued in this chapter that even inferential beliefs have rational motivations that do not include either doxastic 'taking' states or the disposition to answer 'p' in response to the question about what's good about believing q. So unless one insists on a more demanding conception of having a reason$_2$, with all the difficulties that comes with, then one will be committed to a two-part internalism – unless one simply drops the first part of the internalism, as I have suggested may be the right thing to do. In a nutshell, my point is that there is no special problem here having to do with trying to identify reasons that fulfil Reasons Outside Thought, and which can be neatly solved by holding that only beliefs or knowledge can be reasons$_2$ for beliefs. So those philosophers who do not join me in simply rejecting the first part of McDowell's internalism do not here have an argument against the idea that perceptual reasons are reasons that lie outside thought.

To sum up, if the instruction-giving account of inference is roughly right, it provides us with lessons for thinking about non-inferential reason-based belief. First, the instruction-giving account vindicates a watered-down version of McDowell's thought that one can both come to know p and come to know what one's reason$_1$ for believing p is in the same cognitive act. For it suggests that we can use the same capacity both to infer the truth of q on the basis of p and to recognise that p's truth favours believing q. This doesn't mean that we can come to know both q

and that p favours believing q in the same cognitive act. But the fact that we can use the same capacity to come by both pieces of knowledge is significant, since it suggests that inferring p from q involves activating one's capacity to recognise that p's truth favours believing q, even if one doesn't actually come to believe this latter fact. Something similar may well go for perceptual knowledge: the capacity to come to perceptually know p may be the very same capacity that we use to come to know that we see that p. This goes at some but not all of the way towards McDowell's claim that one comes to perceptually know p and to know that one sees that p in the same cognitive act.

The second lesson is more significant. The key obstacle that Ginsborg, Littlejohn and others see to the idea that there can be reasons which fulfil the Reasons Outside Thought condition is that, as they see things, possessing p as a (supposed) reason_1 to believe q involves being disposed to *cite* p as favouring q, and perhaps involves *knowing* or *believing* that p is a reason_1 to believe q. And, crucially, they insist that to believe q for the (supposed) reason_1 that p, one must have been disposed to cite p as favouring believing q and one must have known or believed that p is a reason_1 to believe q *in advance of* forming the belief in q. From this perspective, McDowell's position is apt to seem strange and ad hoc. For, on the one hand, he agrees that possessing p as a reason_1 to believe q involves being disposed to cite p as favouring q; indeed he thinks that it involves knowing that p is a reason_1 to believe q. But, on the other hand, he denies that one must possess p as a reason_1 *in this way* in advance of

forming the belief in q on the basis of p. Rather, one *comes* to possess p as a reason$_1$ in this way *as part of* coming to know q. In advance of that, he thinks, one need only possess p as a reason$_1$ in a far less demanding way. This combination of views can seem puzzling, since it seems to combine an acknowledgement that genuine rational appreciation of a reason$_1$ *as* a reason$_1$ involves embracing it in thought with a claim that one's belief can be rationally motivated by a reason$_1$ that one has not yet embraced in thought. It is natural to suspect that McDowell is combining these views in an ad hoc attempt to make room for reasons outside thought.

However, what my account of inference suggests is that it's false, even in the case of inference, that to believe q for the (supposed) reason$_1$ that p, one must have been disposed to cite p as favouring believing q and one must have known or believed that p is a reason$_1$ to believe q *in advance of* forming the belief in q. Thus, anyone who insists that inferentially believing q on the basis of p requires these things will be forced into a McDowell-style two-part internalism. This suggests that if McDowell's view is ad hoc and strange, this is not because of his commitment to the fulfilment of Reasons Outside Thought. Rather, it is because of his commitment, shared by Ginsborg and Littlejohn, to the idea that believing q for the reason that p involves being disposed to cite p as one's reason$_1$ to believe q and knowing that p is one's reason$_1$ for believing q. The puzzle, if there is a puzzle, is a general one about believing for reasons, whether the believing in question is inferential, perceptual or of another kind. The puzzle is why, if one can rationally respond to a

reason$_1$ (or supposed reason$_1$) without antecedently fulfilling a substantial doxastic 'taking' condition, one must fulfil such a condition *as part of* or *as a result of* rationally responding to the (supposed) reason$_1$. As I've indicated, in my view one needn't fulfil such a condition even in this way, and so there is no puzzle at all – but it has not been my primary purpose to argue for this further claim.

2.7. Conclusion

This chapter has sketched an account of the nature and role of 'taking' in inference. That is, I have tried to describe the sense in which a (supposed) reason$_1$ needs to figure or register *as* an (apparent) reason$_1$ in a subject's perspective in order for her to be in a position to infer from it to a further belief. As I acknowledged at the outset, 'taking' may not be the best label here: perhaps 'figuring' or 'registering' is better. The central message of the chapter has been that we should embrace a fairly minimalist account of what this figuring, registering or taking comes to in inference, and that we should therefore give a relatively significant role to sub-personal processes in which the thinker is passive. By the same token, the fact that there is a degree of passivity in perceptual knowledge-formation, and the fact that non-doxastic perceptual experiences cannot provide for the fulfilment of a 'taking' condition,

construed doxastically or in terms of person-level rule-following, should not lead us to dismiss the idea that such experiences might be reasons$_2$. This is not to deny that there is an element of activity in inference that is absent in perceptual knowledge-formation; but, as the next chapter will argue, this difference does not indicate that perceptual knowledge cannot be reason-based.

Chapter 3

Thoroughgoing externalism

Inference is active in a way that the formation of perceptual knowledge is not.[82] The image of perceptual knowledge as spontaneous, as just striking us, rightly suggests that perception involves the external world exerting a degree of direct control over our mental lives. The process of forming a perceptual belief does not involve asking questions about the normative import of facts which we already know or believe, whereas (as I argued in Chapter 2) in inference it does. So in inference, the subject themselves has to represent as true the (supposed) fact which they ask questions about or give instructions in terms of. Only by doing this does the subject put themselves in a position to rationally ask questions and give instructions of the kind I suggested are at work in inference. Asking questions and giving such instructions would make no sense in a mental context in which the relevant content is merely entertained or is the content of a state of (for example) desire or fear. In contrast, if perceptual experience is non-doxastic then a subject undergoing a perceptual experience does not thereby represent anything as true. How, then, can simply undergoing a perceptual experience make it make sense for one to, as it were, ask questions about whatever the experience

[82] See section 2.3 above.

presents one with, with a view to finding out new facts about the external world?

That is, if we extend our question-asking and instruction-giving account of the subject's role in inference so that it applies more generally to the subject's role in reason-based belief-formation, it might seem that we have to make all reason-based belief inferential after all. For while I deny that a (supposed) reason$_1$, p, must figure in one's perspective *as* a reason$_1$ *to believe q*, for one to be in a position to believe q for the reason that p, I *have* been arguing that there is a sense in which p must figure in one's perspective *as* a reason$_1$ – as a fact which makes appropriate beliefs and so which should constrain and inform future enquiry. How can p figure in one's perspective in this way if one does not believe it? The challenge for McDowell is to specify a way in which, merely by enjoying a perceptual experience, it makes sense for the subject to ask questions of the experience or about the experience's content in order to learn how things are in the world.

I begin this chapter by considering McDowell's own response to this challenge. The rational significance of a perceptual experience is, he says, that it puts one in a position to know something about the external world. This offers a useful starting point, by highlighting the conclusive nature of perceptual-experiential reasons as McDowell conceives of them. However, the notion of being in a position to know is too vague to amount to a fully satisfactory conception of what is involved in

possessing a reason for belief. In the rest of the chapter I therefore try to flesh out McDowell's response, in order to provide a positive characterisation of perceptual experience, such that to enjoy a perceptual experience is to have a reason$_1$ figure as such in one's perspective. The key idea is that perceptual experience has a phenomenal character which is constituted by the very worldly entities which make true one's perceptually knowledgeable beliefs. These entities might be facts or they might be objects and properties. In arguing that this enables us to conceive of perceptual experience as constituting or containing reasons that fulfil Reasons Outside Thought, I also argue that accounts on which perceptual experience is not world-involving in this way cannot sustain such a conception. That is, I try to show that Ginsborg's argument constitutes a compelling argument against positions which are currently popular in the literature (such as Pryor's Dogmatism and Huemer's Phenomenal Conservatism[83]). First, though, let's consider in more detail McDowell's idea that perceptual experiences put one in a position to know worldly facts.

[83] Pryor 2000; Huemer 2007.

3.1. Being in a position to know

It is pretty clear what McDowell himself would say in response to this challenge. According to him, in a 'good' case of ordinary perception, perceptual experience is a factive mental state which puts the subject in a position to know both something about how things are in the external world and that she sees that they are that way. This is what makes the perceptual experience a reason that the subject possesses. Recall the passage quoted in section 1.1, in which he says that 'the rational significance of an experience of perceiving is that it is potentially a ground for perceptually knowledgeable judgements', and that

> [e]ven if the potential for grounding some knowledgeable judgement that is part of the rational significance of an experience of perceiving is not actualised, the subject of the experience still has it in her power, just in undergoing the experience, to know that it has a rational significance that includes making it manifest to her that things are as she would judge them to be in that judgement.[84]

In short, for McDowell an experience is what Ginsborg calls a reason$_2$ because it puts the subject in a position to know its own rational significance and to know something about the external world. Elsewhere

[84] McDowell 2018: 91–2.

he talks about an experience being an 'opportunity' to know, and an 'invitation' to believe, or as making knowledge available.[85] Does this amount to a satisfactory characterisation of the experience's making it the case that a reason$_1$ figures *as* a reason$_1$ in one's perspective, such that we may call the experience a reason$_2$?

Ginsborg thinks that it does not. She rejects this kind of response to the challenge we've seen her as posing on the grounds that it cannot do justice to 'the internalist character of rationality'.[86] The problem is, she says, that it does not characterise the experience's content as being the subject's *own* reason$_1$ for belief.

To illustrate this she gives two examples, which she says are cases in which a fact does not (as she puts it) 'register' in one's consciousness, and yet one can be aptly described as being in a position to know that fact. First she describes a case in which a subject sees someone wearing a brown sweater and has a normal, veridical visual experience, but fails to notice either the experience or the sweater because her attention is focused on something else (say, what to eat for lunch). Ginsborg says of this case that

> [e]ven though the fact does not register in her consciousness – if asked at a later time why she did not form the belief [about the

[85] McDowell 2002a: 278; McDowell 2003: 158–9. Ginsborg cites these passages at Ginsborg 2018: 113.

[86] Ginsborg 2018: 113.

> brown sweater], she is likely to say that she did not see that the sweater was brown – there is clearly a way in which it is visually available to her, namely that the sweater is there, manifesting its cooler, right there in her line of sight.[87]

Second, Ginsborg describes a case in which the subject does not even see the brown jumper but is in the same room as it and so could see it if only she moved her head in the right direction. Here, says Ginsborg, 'the fact is, in an intelligible sense, available to her: being in the same room, she is well placed to know that the sweater is brown, and if she had just turned her head, she would have acquired that knowledge.'[88]

To grasp the point of Ginsborg's examples we need to keep in mind that, as she interprets McDowell, his view is that a perceptual experience is a reason$_2$ which makes the subject possess its content – a worldly fact – as a reason$_1$ to believe that very content. And, on the development of this view that she is considering, an experience does this by putting the subject in a position to know the worldly fact that is its content. So the point of her first example is that it seems to show that one can have a factive perceptual experience with a certain content – in this case that there is a brown sweater before one – and yet this content does 'not register in [one's] consciousness', and so cannot figure in one's perspective as a reason$_1$. For surely, Ginsborg's thought is, a reason$_1$ must

[87] Ginsborg 2006: 307.

[88] Ibid., 307.

figure in one's consciousness if it is to figure in one's perspective as a reason$_1$. So even if the experience can truly be said to put one in a position to know something about the world, this doesn't mean that it makes one possess its content as a reason$_1$ (or, she might add, as an apparent reason$_1$). Although Ginsborg is here arguing against the idea that what makes experiences reasons$_2$ is that they put us in a position to know their contents, the example she gives and what she says about seem to cut more generally against the idea that perceptual experiences are reasons$_2$, whatever we might think about why or how they are reasons$_2$. For the example seems to show, at least as Ginsborg interprets it, that a perceptual experience simply doesn't by itself make its content figure as a reason$_1$ in the subject's perspective, since it doesn't by itself insert its content into the subject's consciousness. Let's come back to this point after considering her second example.

The second example is more obviously directed against trying to utilise the notions of being in a position to know or making knowledge available in explicating the status of experiences as reasons$_2$. Although Ginsborg doesn't put it this way, the force of the example is surely to show that these notions are too vague to really help with the task of characterising a substantial sense in which one's having a perceptual experience can make a reason$_1$ figure as such in one's perspective. For it can be truly said that one is in a position to know all sorts of facts, even though one's perspective contains no evidence or reasons to believe these facts to be true. Indeed, there is a sense in which one is in a position to

know all knowable facts. So the notion of being in a position to know a fact does not offer a way of characterising the internalisation or partial internalisation of a fact or of reasons which might aid one in coming to know that fact. It therefore does not offer a way of characterising the figuring of a reason$_1$ or supposed reason$_1$ as such in one's perspective.

Do the subjects in these examples have reasons$_2$ to believe that there is a brown sweater nearby? We might try to vindicate Ginsborg's claim that they do not by showing that we cannot harvest the materials for a rationalising explanation of any imagined subsequent belief that there is a brown sweater from the mental states of the subjects in the examples. Consider the first example, where the subject enjoys what McDowell would think of as a factive visual experience in which the brown sweater is presented to her as such, but in which she does not attend to it. With this characterisation of a visual experience in mind, are we happy to appeal to such visual experiences in rationalising explanations of subjects' beliefs? Suppose that the subject continued to see the sweater and that eventually she attends to it and forms the belief 'there's a brown sweater'. Are we happy to offer a rationalising explanation of this belief solely in terms of her visual experience of the sweater, or is more needed? I take it that Ginsborg would hold that it is clear that more is needed, that the visual experience alone neither explains nor rationalises the subject's belief. That claim is much more obviously true with respect to the second case, where the subject is merely in the same room as the brown sweater. The subject's being in a position to know that there is a

brown sweater, as that position is constituted in this case, is clearly not a state that we could cite in a rationalising explanation of any subsequent belief on the part of the subject that there is a brown sweater in the room. We might conclude that the notion of being in a position to know must at least be refined and elaborated if it is to have any hope of being the kind of position we could appeal to as a reason$_2$ in rationalising explanations of beliefs.

There are two things to say in response to Ginsborg. The first is that whatever way we characterise the figuring of a reason$_1$ as such in one's perspective, we are going to have to allow for gradations of attentional access to the reason$_1$ as a reason$_1$. Put aside perception for a moment and consider inference again. Suppose S knows p, and suppose that p favours belief in q, and that S knows this too. That would suffice for S to have a reason$_2$ to believe q on Ginsborg's own account (I've included a doxastic 'taking' state in the description of the case to ensure this result). Still, it might be that p is currently far from S's attention. It might even be that S is considering whether or not q is true, and has so far not brought into her attention p and its status as a reason$_1$ for believing q. Still, on Ginsborg's own account, S's knowledge of p is a reason$_2$ which puts her in possession of p as a reason$_1$. Allowing that this is so does not start us on a slippery slope towards having to admit that all normative reasons in the world are reasons$_1$ that S possesses whether or not these reasons have ever entered her perspective, on the basis that if she only attended to the right parts of the world she would bring them to mind.

There is a principled difference between facts that S knows and facts that she does not know, even if among the facts that S knows there are variations in the immediacy with which she can or does attend to them.

Now reconsider Ginsborg's first example, which I said might be thought to cut against the very idea that experiences are reasons$_2$, whether or not that idea is explained in terms of experiences putting subjects in a position to know. The subject sees the brown sweater, and has a normal perceptual experience, but her attention is focused elsewhere and so the fact that there is a brown sweater before her does not 'register in her consciousness'. It's not clear what this means. Does Ginsborg have in mind a distinction between access- and phenomenal-consciousness?[89] If so, does she think that the subject is phenomenally conscious, in having her perceptual experience, but that the fact that there is a brown jumper before her is not access conscious – not immediately available to serve as a reason$_1$ for belief?[90] Or perhaps she has no access/phenomenal consciousness distinction in mind, and just thinks that in such a case one does have a conscious visual experience, but the worldly fact that it presents does not '*register*' in consciousness, which is just to say, again, that

[89] See Block 1995.

[90] See ibid., 231.

it is not immediately available to serve as a reason$_1$ for belief.[91] Either way, surely the proper reaction to this kind of case is to say that it doesn't obviously differ significantly from a case in which one knows p but is temporarily distracted from the fact that p. Yet we would still say that someone who knows p thereby has a reason$_2$. It's thus not clear to me why Ginsborg thinks a case in which one is distracted from one's perceptual experience with content p can't be a case in which one's experience is a reason$_2$.

In general, there are features in the world that are currently embraced by our perceptual consciousness, and features in the world that are currently outside our perceptual consciousness. The fact that the features that lie inside our perceptual consciousness may capture more or less of our attention is not obviously damaging to the intuitive line between what lies inside perceptual consciousness and what does not. Even if something we know is not something we are now attending to, it is still something we know; and even if something that has entered our perceptual consciousness is not something we are now attending to, it is still something that has entered into our perceptual consciousness. So the

[91] A third possibility is that when she says that the fact doesn't 'register' in the subject's consciousness, she means that the subject simply doesn't have any kind of conscious visual experience. If that's the correct description of the case then we can just say that the subject lacks a reason$_2$ here because she lacks a perceptual experience (and lacks anything else which might be a reason$_2$). The case then wouldn't offer any reason to think that perceptual experiences are not reasons$_2$. Instead, it would be better utilised in the same way that I suggest Ginsborg's second example might be utilised, in an attempt to show that the notion of being in a position to know is vague in a way that makes it useless for the purpose to which McDowell is putting it. I'll come to this point momentarily.

mere fact that one may fail to attend to whatever perceptual experience presents one with is not a reason to think that factive perceptual experiences are not reasons$_2$.

Nevertheless, the second thing to say in response to Ginsborg is to accept the point that we need to do more by way of giving a positive characterisation of perceptual experience in order to make it plausible that they constitute reasons$_2$. What her second example shows is that the notion of being in a position to know, at least if it not given further specification, has too wide an application for it to be able to give substance to the idea of the non-doxastic figuring of normative reasons as such within a subject's perspective.

So we need to offer a further positive characterisation of perceptual experience which makes it intelligible that perceptual experiences are reasons$_2$ – mental states or episodes which we can appeal to in rationalising explanations of subjects' beliefs, and which make it the case that a reason$_1$ or supposed reason$_1$ figures as such in a subject's perspective. That's the task I take up now.

3.2. Factivity and the relational view

Ginsborg suggests that McDowell simply overlooks the reasons$_1$/reasons$_2$ distinction.[92] In fact, though, McDowell is one of a number of theorists whose understanding of the relationship between normative and explanatory-motivating reasons is in direct competition with the reasons$_{1/2}$ distinction as Ginsborg draws it. Although I have alluded to this already, so far I have not explored in detail the differences between the two approaches. Now is the time to do so, since we have just seen that we need to give a positive characterisation of perceptual experiences which makes it plausible that they are reasons that we can cite in rationalising explanations (i.e. reasons$_2$), and in my view we cannot do that if we draw the reasons$_{1/2}$ distinction in the way that Ginsborg does. For, as Ginsborg draws the distinction, reasons$_2$ are non-factive, non-relational mental states; whereas, in my view, the factive or relational metaphysics of perceptual experiences (in the 'good cases') is a central part of what makes them reasons$_2$. In the next section I will argue that Ginsborg is wrong to cast reasons$_2$ as non-factive and non-relational. In this section I will briefly lay out the factive and relational views of perceptual experience.

McDowell can agree with Ginsborg to call reasons as they figure in rationalising explanations of subjects' beliefs from a third-person

[92] 'McDowell, at least as I read him, fails to appreciate the distinction between reasons$_1$ and reasons$_2$' (Ginsborg 2006: 311).

perspective 'reasons$_2$'. And he can agree with Ginsborg that rationalising explanations of a subject's beliefs or actions try to show how the subject's belief or action made sense or was appropriate in the light of her perspective on the world. So reasons$_2$ must be or involve some aspect of the subject's total mental state. However, where Ginsborg takes a step that McDowell would not take is in claiming that having a reason$_2$ to believe p does not entail having a corresponding reason$_1$ to believe p. As we saw in Chapter 1, Ginsborg claims that from the third-person perspective (from which we seek to make rationally intelligible a subject's beliefs and actions) whether or not the subject actually has a reason$_1$ is irrelevant. What matters is just that the subject 'takes' herself, perhaps mistakenly, to have a reason$_1$ for a certain belief. By denying that it matters for the purposes of rationalising explanation whether or not the subject actually has a reason$_1$, Ginsborg is denying that reasons$_2$ are factive, since she is denying that it matters whether their contents are true or false. She thereby rules out the possibility that the factivity of perceptual experiences is part of what makes them into reasons$_2$.

What is it to be a factive mental state? Propositional attitudes are factive if and only if they can only be taken up towards truths. Thus knowledge is factive: one cannot know p without p's being the case. This is how McDowell thinks of perceptual experience: in 'good' cases of fully functioning vision in normal environmental conditions, one enjoys a kind of perceptual-experiential mental state which can only present one with worldly *facts*.

Furthermore, factive mental states, as I will be understanding them (following McDowell and Williamson), are fully mental. To say that knowledge is a fully mental factive state is to say that there is (as Williamson puts it) 'a mental state being in which is necessary and sufficient for knowing p'.[93] Similarly, we might say, there is a mental state being in which is necessary and sufficient for having a factive perceptual experience. This implies that hallucinations and illusions are different kinds of mental state altogether. This view amounts to a *thoroughgoing externalism*.[94]

---

[93] Williamson 2000: 21.

[94] For defences of this thoroughgoing externalism as it pertains to knowledge see Williamson 2000: Chs 1–3. Thoroughgoing externalism with respect to perceptual experience is usually associated with naive realism and the associated disjunctivist treatment of veridical perceptions, illusions and hallucinations. Classic statements include Hinton 1973, Snowdon 1981, Snowdon 1990, McDowell 1982, Martin 1997, Martin 2004, Martin 2006, Brewer 2011. For epistemological applications, see McDowell 1982, McDowell 1995, McDowell 1996, McDowell 1998, Williamson 2000, Travis 2005, Kalderon 2011, Brewer 2018. Travis 2005 traces historical precedents for the relevant metaphysical picture to, for example, Cook Wilson 1929. See also Dutant 2015 and Le Morvan 2017 for similar examples. See also Antognazza 2015: 167, where she points out that the 'dominant reception in Western thought' of Plato's epistemology held that there is a 'sharp distinction of kind between knowledge and belief', and this same sharp distinction is found in Aristotle, the Stoics, Scholastic Aristotelians and Early Modern philosophers, too. This is strongly suggestive of thoroughgoing externalism. So although the view may be non-standard today, this may well be an historical anomaly. Besides its historical pedigree, it may also enjoy support from empirical evidence in psychology. Nagel 2013 presents evidence that children learn to use the concept of *knowledge* before learning to use the concept of *belief*, which undermines the idea that the concept of knowledge is composed of the concept of belief plus other concepts. If our concepts are a good guide to metaphysical reality, then this supports the idea that knowledge is a *sui generis* mental state rather than a composite of belief plus further conditions. See also Miracchi 2015 for a quite different argument that knowledge is a mental state.

Many philosophers endorse semantic externalism, according to which the fact that a subject is in a mental state with such-and-such a *content* is partly determined by or supervenient on how things are or have been or will be in the world beyond the subject's skin or brain or whatever other boundary might be thought to sunder mind from world.[95] But externalism about *attitudes* is less common.[96] Semantic externalism combined with internalism about attitudes holds that fully mental states (as opposed to composite states composed of mental states and the obtaining of non-mental conditions) are determined by the external environment only in virtue of their contents. Thus, on this view, a belief and a desire (or any other non-factive mental state) which share the same content are different kinds of mental state but do not differ in anything that is determined by the world outside of the skin or brain: they differ only in internal conditions. In general, the obtaining of all attitude-kinds is determined entirely by internal conditions. This view entails that all fully mental states are non-factive. Consequently, knowledge cannot be a fully mental state. This leads to a conception of knowledge as a metaphysical composite, consisting of a fully mental attitude (belief), plus further external conditions. Similarly, on this view veridical perceptual experiences must consist of a fully mental state, which is common to veridical perception, hallucination and illusion, plus the obtaining of external conditions. The thoroughgoing externalism that is externalist

[95] See Putnam 1975 and Burge 1979.

[96] See Williamson 2000: 50.

about both attitudes and contents rejects this conception of knowledge and perceptual experience. According to this thoroughgoing externalism, these states are at once factive and, in the sense defined above, fully mental.

Factivity is a property of representational states. McDowell holds that perceptual experience is representational, and so on his view it makes sense to say that perceptual experiences are factive. However many disjunctivists about perceptual experience deny that it is representational, and I think there are persuasive arguments for a non-representational, relational conception of perceptual experience.[97] If this is right then perceptual experience isn't the kind of thing which could be either factive or non-factive.

However, on a basic relational view of perceptual experience, 'good'-case experience involves the subject standing in a primitive acquaintance relation with worldly objects and their properties, and these worldly objects and their properties are such that certain propositions about the world are true. The phenomenal character of such experiences consists in this relation. Perceptual experience can, on this view, be seen as part of what we talk about when we say things like 'I saw the dog on the armchair' and 'I see the yellow ball', so the 'basis problem' for McDowell's view does not afflict this view, since seeing objects does not

[97] E.g. Martin 2002a; Travis 2004; Weir 2004; Brewer 2006; Hutto and Myin 2012 and see also Breckenridge 2007 for criticism of a common argument for the view that experience does have content.

involve knowing anything about them.[98] This raises the prospect that these experiences, or the worldly objects and properties that partly constitute them, are conclusive reasons (which the subject possesses in virtue of having the experiences) to believe the worldly propositions that they guarantee the truth of.[99] These reasons would be conclusive because they would be (or would at least be constituted by) the very things, or the very layout of the world, which make the relevant propositions true.

While I incline towards the relational conception of experience, I won't take an official stance with respect to the question of whether we should see experiences as representational and factive or as non-representational and relational. For the purposes of responding to the worries about the idea that there can be reasons outside thought that I have been considering, it does not matter whether experience is relational or factive. I will therefore proceed to defend the idea that part of what it is about perceptual experiences that makes them reasons$_2$ is either their factivity or their being relations to worldly objects and properties.

[98] For the 'basis problem' see above, section 1.1.

[99] The non-representational, relational conception of reason-constituting or reason-bearing sensory consciousness here adverted to is found in, for example, Martin 2002a; Kalderon 2011; Brewer 2011; French 2016; Brewer 2018 (cf. Brewer et. al. 2018 for an important modification of Brewer's view); Travis 2013; Logue 2018.

3.3. Rationalising explanations

I said that, given the relational view of perceptual experience, it is plausible that worldly objects and their properties constitute conclusive reasons for beliefs about the world. In saying this, do I mean that they plausibly constitute reasons$_1$ or reasons$_2$? Similarly, do the worldly facts which, according to factivism, perceptual experiences present us with constitute reasons$_1$ or reasons$_2$?

Well, if they are reasons at all then they both clearly constitute reasons$_1$, since they would be reasons which, from a first-person perspective, we would recognise as making appropriate forming certain beliefs about the world. But might they not also be reasons$_2$ – reasons we can cite from a third-person perspective in rationalising explanations of subjects' beliefs? The thoroughgoing externalism described in the previous section raises this possibility. We cite reasons$_2$ in order to make sense of a subject's beliefs or actions in terms of her perspective. That led Ginsborg to the view that reasons$_2$ must be non-factive, non-relational mental states. But this only follows if a subject's perspective is exhausted by non-factive, non-relational mental states. That is just what thoroughgoing externalism denies: there are, according to this view, fully mental factive and relational states, which help to constitute a subject's perspective on the world. If this is right then the reasons$_{1/2}$ distinction collapses, at least as Ginsborg presents it. The dispute between those who embrace

thoroughgoing externalism and those who reject it is multi-faceted and I cannot here engage with it in anything like an exhaustive way. I will, though, argue that the idea that reasons$_2$ are always non-factive mental states leads to serious difficulties and is not well-motivated. Genuinely normative reasons – reasons$_1$ – surely must sometimes be part of what we appeal to when we offer rationalising explanations of others' beliefs (and actions, but here we will concentrate on beliefs).

Although Ginsborg herself does not say much in defence of the idea that reasons$_2$ are always non-factive mental states, it is a familiar idea. It can motivated by general considerations against embracing thoroughgoing externalism. But it can also be motivated by considering our practice of giving rationalising explanations. We can give good rationalising explanations of a subject's believing p in terms of her false and unknowledgeable beliefs. 'Hannah believes it rained because she believes the streets are wet' is apparently a good rationalising explanation of Hannah's belief that it rained even if Hannah is wrong that the streets are wet. Consequently, there must be some sort of reason$_2$ in play in this kind of case, and the natural candidate is Hannah's *belief* that the streets are wet (or perhaps the *proposition* that the streets are wet[100]). But this reason$_2$ can't also be a reason$_1$ – a normative reason – since as the case is described there is no normative reason for Hannah to believe that it rained. Now we generalise. In a corresponding 'good' case, where Hannah knows that the streets are wet and knowledgeably infers that it

[100] For this anti-psychologising approach, see Dancy 2000.

rained, exactly the same rationalising explanation of her belief that it rained is apparently available. The fact that in this good case Hannah actually does have a reason$_1$ adds nothing in terms of rationalising-explanatory power.[101] So whether a subject knows or not, and whether their belief is true or not, is irrelevant with respect to rationalising explanations. What matters are non-factive states, and specifically beliefs.

This argument is not decisive. We can accept that there is a kind of rationalising explanation available in both the 'good' and 'bad' cases, while also holding that there is a further kind of rationalising explanation available in the good case which makes essential reference to the subject's normative reason, and so which is unavailable in the bad case. This is exactly what McDowell thinks: 'When the "better" disjunct holds [i.e. when one has a genuine, factive perceptual experience], there is not only the kind of [rational] intelligibility that there is when the "worse" disjunct holds [i.e. when one is merely hallucinating, say], but also a different kind of intelligibility, which is special to the "better" disjunct'.[102] That is, we can give third-person rationalising explanations of a subject's belief in terms of perceptual experiences conceived of as factive or relational states. Similarly, we can give third-person rationalising explanations of a subject's beliefs or actions in terms of their *knowledge*. In both cases, a kind of rational intelligibility would be lost if the explanation was re-cast in terms of non-factive, non-world-involving

[101] For a recent statement of this kind of claim see Dancy 2014: 84.

[102] McDowell 2013b: 26.

states (perceptual appearances or mere beliefs). On this view, it's true that the non-factive states do provide for a kind of rational intelligibility, but it is a different kind to that provisioned by factive states.

Furthermore, there are good reasons to think that McDowell is right about this. To begin with, we do in everyday life sometimes formulate rationalising explanations in terms which are at least compatible with the explanation making essential reference to a normative reason (a reason$_1$). Indeed, we don't usually appeal to a subject's beliefs in rationalising explanations. If you ask me why Hannah thinks it's recently rained, I'm more likely to reply with the factive construction, 'because the streets are wet' (or 'because of the wetness of the streets', a formulation which fits nicely with relational conceptions of experience), than to reply 'because she believes that the streets are wet'. The latter response carries the conversational implicature that Hannah is wrong that the streets are wet. This suggests that our primary way of rationalising subjects is in terms of their factive and relational mental states. This evidence is not decisive, though, since one might insist that the factive and relational constructions are really abbreviations for an explanation which involves reference to a non-factive mental state, or that the implication of the explanation that the streets really are wet is superfluous to the rationalising force of the explanation.[103]

[103] E.g. Davidson 1980. See Roessler 2014: 350 for critical discussion.

But a good reason to think that McDowell is right is that, as he argues, the rational intelligibility yielded by explanations that appeal to non-factive mental states must itself be derivative of a different kind of rational intelligibility yielded by explanations which involve essential reference to world-involving factive or relational mental states. This is, he says, 'especially clear in the case of beliefs grounded in experience':

> If one has an experience in which it appears that things are a certain way, it is, so far as that goes, rational for one to believe that things are that way. But that is because when one has such an experience, one's experience seems to reveal things to be a certain way: that is, one's belief seems to have the kind of rational status that characterises the "better" disjunct. The kind of intelligibility that is common to both disjuncts needs to be understood in terms of there seeming to be an instance of the kind of intelligibility that is distinctive of the "better" disjunct.[104]

Now, Ginsborg's own position suggests that she would concede that there is something right about what McDowell says here. Why does Ginsborg think that a reason$_2$ must make it the case that one *takes* oneself to have a reason$_1$? Surely because reasons$_1$ – normative reasons – are what make certain beliefs appropriate beliefs to have. Non-factive reasons$_2$ only make beliefs appropriate, to the extent that they do, derivatively of this fact. In responding to reasons$_2$, we are doing the best we can to respond

[104] McDowell 2013b: 26. Cf. Hornsby 2008.

to reasons$_1$. This point seems to be conceded by, for example, Stampe, when he says (in a passage Ginsborg quotes approvingly) that '[w]hether the belief that *p* may […] be described as a reason for something, depends on whether the *fact* that *p* would be a reason for that thing'.[105] Similarly, Ginsborg herself says that 'the notion of a reason$_1$ has priority over that of a reason$_2$: a psychological state counts as a reason$_2$ for a subject's belief only if her being in that state involves her representing some fact as a reason$_1$ for that belief.'[106] Conversely, the notion of a reason$_1$ can be grasped without reference to reasons$_2$.

What Ginsborg seems to be conceding that is that our understanding of rationality arises in the first place from our practices of deliberating about what to believe, and thus from the first-personal perspective on reasons. From this perspective, what we care about are reasons$_1$. And we do not, from this perspective, regard our beliefs as, in and of themselves, reasons$_1$. The mere fact that I believe that the streets are wet does not make it appropriate for me to believe that it's recently rained, from the first-person perspective. In general, the fact that I believe p is not, taken by itself, a source of any normative constraint (with some exceptions,

[105] Stampe 1987: 337, cited by Ginsborg 2006: 291.

[106] Ginsborg 2006: 291.

which really just clinch the point[107]), at least when the matter is considered from the first-person perspective. Of course, the fact that I believe p might be used as a premise in an inference to a conclusion about some further non-mental fact. The point is just that, by itself, it does not (so we think, when within the first-person perspective) generally make appropriate further beliefs about, for example, the propositions p entails.

Some philosophers might deny this last claim. They might point out that if, in the course of my deliberation about what to believe, I am reminded that I already believe p, and I recognise that p entails q, then I will

[107] For example, the fact that I believe p makes it appropriate for me to believe that I believe p. Perhaps, also, the fact that I believe that there are pink rats in my boots makes it appropriate for me to believe that I need medical help, and makes it appropriate for me to seek that help (although perhaps it only does so with the addition of further background premisses, such as that the idea of pink rats in boots is absurd). The pink rats example is adapted from Dancy 2000: 125.

thereby come to see believing q as, in some sense, appropriate.[108] That seems right, but this doesn't show that it is my belief in p, or the fact that I believe p, that I take to make appropriate believing q. It might instead be that I take (what I take to be) the *truth* of p to make appropriate believing q, and your reminder that I believe p has simply reminded me of (what I take to be) p's truth. For when you remind me that I believe p, I verify that I do indeed believe p by asking myself whether or not p is true and receiving an affirmative answer (as per 'transparency' approaches to self-knowledge[109]). If this is the right account of the

---

[108] The claim might be that the kind of appropriateness determined by a belief differs from the kind of appropriateness determined by the fact that the belief is directed at. Still, the thought is, both kinds of appropriateness are genuinely normative.
Kolodny 2005 provides an overview of and tries to explain the kind of appropriateness apparently in question here. He distinguishes between what he calls 'the "ought" of reasons' and 'the "ought" of rationality' (the latter is the putative kind of appropriateness we're discussing here). The 'ought' of reasons is what we are getting at when we say 'there is *reason* for you to have the attitude, because of some feature of your *actual situation*' (i.e. worldly environment). The 'ought' of rationality is what we are getting at when we say that 'it would be *irrational* of you not to have [such-and-such] attitude, because of some *other attitude* that you have' (509). For example it would be irrational not to believe the drink is safe to drink because you believe (whether truly or falsely) that it was poured by your trustworthy friend. This distinction roughly tracks Ginsborg's reason$_{1/2}$ distinction, since the attitudes in question are non-factive. Kolodny notes that there are those who (unlike Ginsborg) try to explain the 'ought' of reasons in terms of the 'ought' of rationality.
See also Pryor 2018's discussion of what he calls 'hypothetical' rational requirements, which roughly match Kolodny's 'ought' of rationality. I will touch on Pryor's 'hypothetical' rational requirements in Chapter 5.

[109] Evans wrote: '[I]n making a self-ascription of belief, one's eyes are, so to speak, or occasionally literally, directed outward – upon the world. If someone asks me, "Do you think there is going to be a third world war?", I must attend, in answering him, to precisely the same outward phenomena as I would attend to if I were answering the question "Will there be a third world war?" I get myself into the position to answer the question whether I believe that *p* by putting into operation whatever procedure I have for answering the question whether *p*.' Evans 1982: 225.

introspection of one's own beliefs, then we can't cleanly separate my learning that I believe p from my affirming to myself p's truth. So it's not at all clear that in this example I am taking my belief, as opposed to the truth of p itself, to make appropriate belief in q.

There is an additional reason to think that it is (what I take to be) the truth of p, and not my belief in p, that is what I treat as making appropriate belief in q, from the first-person perspective. Suppose that instead of reminding me that I believe p, you tell me that *you* believe p. Unless combined with some further premise that I have reason to believe, such as that the majority of your beliefs are true, the mere fact that you believe p gives me no reason$_1$ to believe p and, by the same token, no reason$_1$ to believe q (which p entails). But if *your* believing p is not a reason$_1$ for me to believe q (or p), then why should *my* believing p be a reason$_1$ for me to believe q (or p)? One suggestion is that coherence is a source of normative force: I should have coherent beliefs. But that doesn't help answer the question because it's unclear what principled reason there could be for limiting the terms of the coherence relation to

one's own beliefs.[110] I suggest, then, that cases in which we learn about the beliefs of others make vivid a general point: the fact that someone (including oneself) believes p is no $reason_1$ in and of itself to believe p or its consequences to be true. The fact that this point is less vivid when it comes to learning about one's own beliefs is easily explained by the fact that one's method for finding our or verifying that one does believe p proceeds via something like Evans's transparency routine. This makes it easy to confuse asking oneself about p's truth-value with asking about whether or not one believes p, and this in turn makes it easy to confuse the roles that the belief in p and the supposed fact p itself play in one's thinking. There is no danger of such confusion when it comes to learning about the beliefs of other people, which is why considering cases in which one learns about the beliefs of someone else makes vivid that facts about what someone believes are not facts which, by themselves, make it appropriate to believe in the consequences of those beliefs' contents. By the same reasoning, the psychological state of belief in p itself (as opposed to the fact that one has the belief) is not a $reason_1$ to believe p or

[110] One suggestion here is that the terms of the coherence relation should be limited to beliefs which are or have the potential to be inferentially related to one another, or reach some threshold of inferential integration. Perhaps your beliefs are not reasons for me because I can't base my beliefs on enough of your beliefs and you can't base your beliefs on enough of my beliefs, because we don't have enough access to one another's beliefs. But suppose that we engage in a long conversation, perhaps intermittently over many years, about some topic, in the process coming to have an extremely intimate and detailed knowledge of one another's belief-sets with respect to this topic and other topics in the immediate vicinity. It does not seem that at any point in this process there would come a point at which your beliefs become reasons for me, in and of themselves (of course, if I come to recognise that you are usually right then I might rationally infer from the fact that you believe p that p is true or likely to be true, but this is not the sense of beliefs-as-reasons which matters here).

p's consequences. In conclusion, then, it seems to be true that (at least from within the first-person perspective) we do not regard a belief in p, or the fact that one believes p, as by itself making appropriate believing p or believing p's consequences.

Furthermore, it's easy to see why we don't, from the first-person perspective, regard beliefs as such as our reasons. Beliefs, being non-factive, can occur whether or not their contents are true. It's not even obvious that anything in the nature of belief makes beliefs likely to be true. But whatever makes a certain belief appropriate is surely something that makes that belief at least likely to be true.[111] This is why beliefs do not make themselves appropriate, are not self-justifying. By the same token, they do not automatically justify believing the consequences of their contents. So the first-person perspective on what matters when it comes to deciding what to believe accords with the perspective of philosophers who are concerned with evaluating the epistemic status, as justified or not, of the beliefs both of others and of themselves. And the rationale that lies behind this accord is that beliefs are normatively constrained by truth, whereas beliefs themselves are non-factive. To sum up, then, from a first-person perspective, in which we deliberate about what to do and believe, we ignore our beliefs and concentrate only on worldly facts and perhaps factive and relational mental states.

[111] As Cohen says, most philosophers have simply 'presupposed the view that [epistemic] justification and truth are conceptually related – that there is an internal connection between a belief being justified and being true'. Cohen 1984: 279.

If it's accepted that our understanding of reasons and rationality arises in the first place from these first-person practices, a number of challenges arise for those who claim that from the third-person perspective of giving rationalising explanations we appeal to a subject's non-factive mental states and in particular their beliefs. In particular, it remains unclear how our third-person explanations of a subject's beliefs, in terms of their non-factive mental states, count as genuinely *rationalising* explanations, as opposed to simply brute causal explanations. In what does the rationality consist? From a first-person perspective, reasons are reasons because they have normative significance: they make a certain belief appropriate. And they have this normative significance precisely because they are external-world facts or are ways that the external world is: that's what grounds their normative import. But from the third-person perspective, reasons$_2$ can't be understood as having normative import in the same way. For a subject might have a reason$_2$ without having a reason$_1$, and so the sense in which a reason$_1$ would make a certain belief appropriate is simply absent. So in what sense are reasons$_2$ genuinely normative? The fact that, from a first-person perspective, beliefs are *not* reasons$_1$ suggests that reasons$_2$ aren't normative at all. So the question remains: in what sense are explanations of beliefs in terms of beliefs genuinely rationalising? I can't prove that there's no good answer in the offing, but this is certainly a major challenge for Ginsborg and those who share her view.[112]

[112] I am closely following Roessler 2014 in this paragraph and the next.

On the other hand, if we accept that normative reasons, or reasons$_1$, can be essential elements in genuinely explanatory and motivating reasons$_2$, then we can make sense of the idea that explanations in terms of beliefs can be rationalising explanations. For if reasons$_2$ can be factive or relational mental states, then there is no problem about how explanations in terms of reasons$_2$ can be rationalising: they are rationalising because they display the subject as responding properly to a genuine normative reason$_1$. They display the subject's belief as appropriate by appealing to the aspect of the world that makes it appropriate. This helps us answer the question of what is rationalising about explanations of beliefs in terms of beliefs (rather than in terms of factive states). For we can now say that although such explanations do not display the subject's belief as normatively appropriate, they do display it as the outcome of a process which is *relevantly similar* to a process which would produce a normatively appropriate belief. When a subject believes that it recently rained because they believe (wrongly) that the streets are wet, they are doing *something* right: they are recognisably utilising their capacity for knowledgeable inference, but their utilisation of the capacity is, as disjunctivists often say, 'defective'. As Roessler puts it, in cases in which subjects form beliefs on the basis of mere (non-knowledgeable) beliefs, we can think of what they are doing as 'defective (though typically excusable) exercises of the capacities operative in a "good" case, and it's characteristic of the latter that the exercise of the capacity is explained in terms of a normative reason.'[113] So whereas the view on which reasons$_2$

[113] Ibid., 354.

are always non-factive, non-relational mental states leaves us without any account of how reasons$_2$ actually rationalise a subject's beliefs, a view on which reasons$_2$ can be factive or relational provides a satisfying account of the rationalising force of rationalising explanations.

Let's take stock. We're trying to vindicate the idea that perceptual experiences can be reasons$_2$ by showing that they make something figure as a reason$_1$, or an apparent reason$_1$, in one's perspective. Chapter 2 suggested a fairly minimalist way of understanding what this figuring amounts to. But even on this fairly minimalist approach to 'taking' something to be a reason$_1$, we still need to vindicate the idea that perceptual experience makes it *make sense* for the subject to *treat* something as a reason$_1$, perhaps by asking the right kind of questions or giving the right kind of instructions. We began this chapter by noting that simply saying that perceptual experiences put subjects in a position to know something about the external world does not offer a fully satisfactory vindication of this idea, because the notion of being in a position to know applies too widely, and in particular applies to situations in which the subject cannot be said to really possess a reason$_1$ (or, for that matter, a supposed reason$_1$). I then suggested that the key to vindicating the idea of perceptual experiences being reasons$_2$ is appreciating that they are both factive or relational *and* fully mental. The factive or relational nature of the mental state in question can then be understood to be an aspect of the subject's own perspective, which makes it possible

to conceive of rationalising explanations that appeal essentially to factive or relational mental states.

In this latest section we have been considering the plausibility of this conception of rationalising explanations. We have considered (admittedly not exhaustively) some of the usual motivations for thinking that reasons$_2$ must be non-factive and found them wanting. Next, we found that thinking of reasons$_2$ as non-factive creates a mystery about how rationalising explanations are genuinely *rationalising*. And we found that if we conceive of reasons$_2$ as factive, at least in the 'good' cases, then that mystery dissipates, allowing us not only to see rationalising explanations in 'good' cases as genuinely rationalising, but also allowing us to see that there is a kind of rational intelligibility available in 'bad' cases, too. None of this is original – I am simply re-stating a standard disjunctivist line of argument. One point which I emphasised and argued for was that from a first-person perspective we do not regard our beliefs to be reasons$_1$ – normative reasons – for belief. That is partly what creates the mystery about the source of the rationalising force of reasons$_2$, conceived of as non-factive states like beliefs. Furthermore, what I now want to argue in the next section is that it also creates a mystery about how non-doxastic perceptual experiences, if conceived of as non-factive and non-relational states common to veridical perception and hallucination and illusion, could be reasons$_2$ for belief.

3.4. Assertoric force, presentational phenomenology, etc.

Chapter 1 briefly touched on the debate between Liberalism and Conservatism in the epistemology of perception, concerning whether or not perceptual justification for belief is immediate. Liberalism, recall, is the position that perceptual experiences, all by themselves, give subjects justification to believe external-world propositions. This justification is 'immediate' in that it does not rely on justification to believe any other propositions, such as 'I am not now hallucinating' or 'my experience is veridical'. The majority of those who defend this Liberal view agree with their Conservative opponents that perceptual experiences are non-factive and non-relational. Call these philosophers non-factivist Liberals. They conceive of experiences as non-doxastic, non-factive, non-relational immediate justifiers or reasons. This view is worth considering, since it promises to vindicate the idea that perceptual reasons fulfil Reasons Outside Thought; and its strategy is in competition with the one I am recommending. For my claim is that it is the factivity or relational nature of perceptual experiences that enables them to be reasons that lie outside thought while also being a subject's reasons for belief. Does non-factivist Liberalism constitute a viable alternative to my appeal to thoroughgoing externalism? In particular, does non-factivist Liberalism have a response to the concerns that we have seen Ginsborg raising about the very idea of non-doxastic reasons$_2$?

An approach taken by some non-factivist Liberals involves thinking of perceptual experience as belief-like. The strategy is clear: even if perceptual experiences are non-doxastic, perhaps they have enough of the features that beliefs have to be reasons$_2$ in the same way that beliefs are. For example, Heck says that perceptual experiences justify beliefs 'pretty much [in] the same way beliefs do'. This is so both in terms of what it is about a belief or experience that endows it with rationalising power, and in terms of the structure of the rationalising process: 'there is an inference or transition – call it whatever you like – from my perceptual state to some belief I recognise it to underwrite.'[114] Heck is explicit that experiences should be understood as non-factive content-bearing states (although he thinks their contents are non-conceptual), and regards their non-factivity as no obstacle to their being reasons, since beliefs too are non-factive reasons:

> The fact that not all of my beliefs are true gives me no reason not to infer other beliefs from them. So when the skeptic asks, 'How *can* you form beliefs on the basis of perceptions you know sometimes to be misleading?' I want […] to ask: Would you also suggest that we abandon the practice of inference just because some of our beliefs are false?[115]

[114] Heck 2000: 509, 511.

[115] Ibid., 509.

Heck is not here distinguishing between the third- and first-personal perspectives in the way that Ginsborg does. This suggests that Heck may think that we treat our beliefs as normative reasons from a first- as well as from a third-person perspective – a claim I argued against in the previous section. But let's put aside my general qualms about non-factive, non-relational pictures of rationality: I am here interested in whether Heck might have a way of responding to Ginsborg's criticism of the very idea of non-doxastic experiential reasons. The question for which I want to extract an answer is: what is it that perceptual experiences have in common with beliefs that allows them to provide reasons in the same way as beliefs do?

Heck's answer is that 'perceptions are attitudes, attitudes that are *like* beliefs insofar as to be in a perceptual state is to hold an *assertive* or *presentational* attitude towards a certain content'. Assertive attitudes 'purport to represent how the world is'.[116] Perceptual experiences have what Heck calls 'assertoric force' or a 'presentational aspect'. This label, 'assertoric force', echoes the distinction in the philosophy of language between the force and the content of an utterance.[117] The idea is that the same content can be uttered with different forces, whether assertoric or interrogative or some other kind of force. Only when we know both the force and the content of an utterance can we give the meaning of the utterance. Assertoric force is the force that an assertive utterance has. An

[116] All the quotations in this paragraph are from Heck 2000: 508–509.

[117] For a sceptical historical overview of this distinction see Baker and Hacker 1984: Chapters 2 and 3.

assertion presents its content as being true, or perhaps expresses the speaker's judgement that the content is true. Now Heck's idea is presumably that we can draw a distinction in the philosophy of mind between attitude and content that is analogous to the distinction between force and content in the philosophy of speech acts, and, further, that at least some mental attitudes parallel some illocutionary forces. So perceptual experiences are like beliefs in that they have assertoric force. Of course, Heck insists that they are *not* beliefs, for the subject is not engaged in, and their conceptual capacities are not exercised in, having a perceptual experience. But their assertoric force allows them to be reasons in the same way that beliefs are.

An immediate worry about Heck's approach has to do with the distinction between beliefs which are justified and beliefs which are unjustified. Heck, like Ginsborg, thinks that there is a sense in which beliefs justify and rationalise other beliefs which can be understood without reference to the idea of rationalisation and justification by factive or relational mental states or aspects of the world. Let's suppose that they are right, even though, as we saw in the last section, it is quite unclear how this claim can be sustained. On the factive-first, thoroughgoing externalist approach I've been recommending, there is a clear distinction between cases of genuine rationalisation and justification by factive or relational states, and cases in which rational capacities are exercised in a defective way. But even if this factive-first approach is rejected, there is still a distinction to be drawn between

justified and unjustified beliefs. Consider cases of wishful thinking, for example. Suppose that Alvin wishes for rain tomorrow, and this causes him to take the dramatic red sunset to be a portent of rain, and so he goes to bed happily believing that it will rain tomorrow. Or suppose that he receives strong evidence that it will be dry tomorrow and his desire causes him to ignore this evidence and retain his belief that it will rain. In a case like this, there is something irrational about Alvin's belief that it will rain tomorrow: it's not a justified belief. And it is usually assumed that this irrationality is infectious, in that were Alvin to believe that his crops will be watered tomorrow on the basis of his unjustified belief that it will rain tomorrow then this further belief would also be unjustified, even though there remains a sense in which it is rationally intelligible. Of course, on a factive-first approach we can explain this rational intelligibility in terms of the idea of the defective exercise of a rational capacity. Although Heck and Ginsborg are not factive-firsters, they can make a parallel move with respect to unjustified but rationally intelligible beliefs like Alvin's belief that his crops will be watered tomorrow. That is, they can say that it is rationally intelligible insofar as it is the product of a defective capacity to make good inferences from *justified* beliefs. But – and this is the crucial point – this move implies that we understand the rational intelligibility that beliefs based on beliefs have in terms of the idea of *justified* beliefs transmitting their justification to beliefs formed on their basis. That creates a challenge for Heck's claim that perceptual experiences rationalise and justify beliefs in much the same way that beliefs do, since perceptual experiences are arguably not the kind of state

that can be justified or unjustified. Even if experiences have assertoric force, they lack this further feature of beliefs, which seems to be essential to our understanding of how beliefs justify and rationalise further beliefs. This isn't a knock-down criticism, but it does highlight the fact that Heck has done nothing to confront an old question: how can there be mental states which justify beliefs without themselves needing to be justified? The question is especially urgent within the context of an approach to perceptual justification according to which perceptual experiences justify and rationalise beliefs in much the same way that beliefs do.

This suggests that a more promising approach for non-factivist Liberals is to appeal to features of experience which differentiate it from belief, rather than to look for features it has in common with belief. In this vein, a number of non-factivist Liberals appeal to the idea that perceptual experiences have 'presentational phenomenology', and have identified this as what makes these non-doxastic experiences into reasons that subjects possess for external-world beliefs. For example, according to Pryor, perceptual experiences have

a distinctive phenomenology: the feeling of *seeming to ascertain* that a given proposition is true. This is present when the way a mental episode represents its content makes it feel as though, *by* enjoying that episode, you can *thereby just tell* that that content obtains […]. When you have a perceptual experience of your hands, that experience makes it feel as though you *can just see* that hands are present. It feels as though hands are

being shown or revealed to you. […] Having the phenomenology of seeming to ascertain P is *what makes us* have prima facie justification to believe P. (You really need *to have* the phenomenology. It's not enough to think you do.)[118]

This description of presentational phenomenology is flawed. According to Pryor's parenthesis, 'You really need *to have* the phenomenology', and he says that to have the phenomenology it's not enough to 'think you do'. But the phenomenology in question is, as Pryor puts it, a feeling or a seeming that, just by having such-and-such an experience, you can just tell that there is a hand before you; or it's a feeling that hands are being revealed to you. Chudnoff gives a similar account of presentational phenomenology: a perceptual experience has presentational phenomenology with respect to a proposition P if and only if in having the experience 'you both seem to fact-perceive that P and seem to be sensorily item-aware of an item that makes it the case that P'.[119] However, Smithies points out that if we 'explain the distinctive phenomenal character of perceptual experience in terms of a "that"-clause that articulates the content of how things seem to you when you have the experience', then we run into the problem that 'you might judge such content to be true without currently having any visual experience at all.'[120] If you judge or believe p to be true, then surely p seems to be true to you. Thus, if you believe that you fact-perceive that P and you believe

[118] Pryor 2004: 357. Pryor's emphases.

[119] Chudnoff 2011: 320.

[120] Smithies 2019: 94.

that you are sensorily item-aware of an item that makes P true, then you have presentational phenomenology as Chudnoff describes it. This implies that one could enjoy presentational phenomenology without having any perceptual experience at all; or one could enjoy presentational phenomenology with respect to the proposition 'there is a cat on the mat' while enjoying a perceptual experience which has no such content and in which, in fact, it does not appear that there is a cat on the mat. Presentational phenomenology, if specified in the way that Pryor and Chudnoff specify it, thus seems not to be distinctive of perceptual experience at all. If one can have presentational phenomenology simply by having a certain belief (or even, we might add, a certain inclination to believe), then it's not clear how merely having presentational phenomenology gives one justification to have a certain belief about the world.

Smithies therefore suggests that we embrace an 'impure representationalism', according to which 'some phenomenal differences are differences in representational force, rather than differences in representational content'. Then we can talk about the 'presentational force' of an experience:

Presentational force is a kind of phenomenal character that we know ostensively by introspecting perceptual experience and contrasting it with other experiences. It is the kind of phenomenal character that your experience has when you perceive things, and when you hallucinate

them, but not when you merely visualize or make judgements about things. To rephrase what Pryor says, it is the kind of phenomenal character that your experience has *when* you can just see that hands are present.[121]

The best we can do, then, to articulate the nature of presentational 'force' is to gesture at it and to say that it has something to do with the way that the representational content of perceptual experiences figures in one's perspective. Indeed, Smithies' 'presentational force' sounds quite like Heck's assertoric force, except that Smithies is explicit that the presentational force of an experience is an aspect of the experience's phenomenal character, and it is something that beliefs don't have.

According to these non-factivist Liberals, then, perceptual experiences are non-doxastic, non-factive, non-relational mental states which constitute reasons that a subject possesses for beliefs about the external world. They are reasons – they have rational significance or justificatory power – in virtue of their presentational phenomenology or assertoric or presentational force. Heck's 'assertoric force' is something that experiences share with beliefs: both kinds of state 'purport to represent how the world is'. 'Presentational phenomenology' or 'presentational force', in contrast, is something perceptual experiences have but beliefs lack, and has to do with how the experience's content is presented in the experience. What should we make of these suggestions?

[121] Ibid.

It seems to me that if we try to pursue this approach to perceptual reasons then we run into a dilemma. Let's use the term 'presentational force' to refer to whatever it is about non-doxastic experiences in virtue of which they are reasons we possess – so this term is neutral between presentational phenomenology and assertoric force. First recall Ginsborg's reasons$_{1/2}$ distinction. One way to understand this is as a distinction between features of a subject S's perspective which S attends to and treats as her reasons for belief when deliberating about what to believe (reasons$_1$), and features of S's perspective which she does not attend to when deliberating about what to believe, but which somehow make something figure as a reason$_1$ in her perspective (reasons$_2$). A belief's *content* is thus a reason$_1$ because the deliberating subject attends to this when looking for reasons for belief. The belief itself is just a reason$_2$, because the deliberating subject doesn't attend to it, but it is what makes the content figure as an apparent reason$_1$ in her perspective. Now the dilemma can be stated in terms of the reason$_{1/2}$ distinction, so understood, as follows. Either presentational force is part of what the deliberating subject attends to and treats as a reason$_1$, or it is something which she does not attend to and treat as a reason$_1$ but which, like belief, makes something figure in her perspective as an apparent reason$_1$. These options are exhaustive, since if presentational phenomenology does not play one of these two roles then it is not part of what *rationalises* a subject's perceptual belief, at least if we stick to the assumption, made by Ginsborg and non-factivist Liberals, that only aspects of a subject's

perspective can play a rationalising role. In short then, either presentational force is an aspect of the subject's reason$_1$ or it is an aspect of the subject's reason$_2$.

Suppose first that presentational force is an aspect of the subject's reason$_2$ but not an aspect of their reason$_1$. This is how I interpret Heck's view. We might understand this view as simply covertly endorsing a doxastic theory of perceptual experience, so that what it is to undergo a perceptual experience with a content p is to believe p in a certain way. Experiencing would then be a way of believing. This is a capitulation because it gives up on the project of identifying reasons which fulfil Reasons Outside Thought.

If, instead, we try to understand this view as one according to which the subject is not doxastically committed in having a perceptual experience, then presentational force cannot be something that the subject themselves contributes to the experience. The experience's content cannot be something that the subject themselves represents as being true. Nor, on this horn of the dilemma, is presentational force part of or a feature of what the subject attends to in deliberating about what to believe – it is not an aspect of an apparent reason$_1$. The problem is that this makes it hard to see how the presentational force of a perceptual experience is an aspect of the subject's perspective at all. For it is not an aspect of the subject's *stance* towards an experience's content, because in order to avoid capitulation we've agreed that the subject has no stance

towards the experience's content simply in virtue of having the perceptual experience. Nor, we have stipulated, is it an aspect of what the experience puts the subject in a position to attend to and take into account when deliberating about what to believe. Or, more precisely, if it is something that the subject can attend to, it is not something that they would treat as part of or relevant to their having a reason$_1$. Thus, presentational force seems to play no part in the subject's perspective, at least with respect to rationalising perceptual beliefs. Perhaps we could think of presentational force as something which stands off-stage, causally influencing or grounding the subject's perspective in the same way that sub-personal processing might be thought to do. But this is not the kind of thing in virtue of which one has a reason$_2$ or in virtue of which something figures as a reason$_1$ in one's perspective. It's therefore hard to see how it can do the job that Heck wants it to do.

In short, then, on this first horn of the dilemma, it's hard to see how presentational force plays any role in the rationalisation of beliefs by perceptual experiences. Earlier in this section, we also saw that this approach faces a particularly acute version of the challenge of explaining how perceptual experiences can justify and rationalise beliefs without themselves being justified. This challenge is particularly acute since on this view perceptual experiences justify in a way that is analogous to how beliefs justify beliefs.

Now consider the other horn of the dilemma: suppose that presentational force is an aspect of the subject's reason$_1$ or supposed reason$_1$. On this horn of the dilemma, presentational force *is* an aspect of what the experience puts the subject in a position to attend to when deliberating about what to believe, and it is relevant, from the subject's perspective, to their having a reason$_1$ for belief. Perhaps, as in Smithies' description, an experience's content comes with a certain phenomenology, which is part of what the subject attends to when they attend to the content of the experience. Or perhaps the subject doesn't merely attend to the experience's content but also to the fact that it is the content of an experience, or to the metaphysical structure of the experience. All of these options and more are possible ways of articulating what presentational force is on this horn of the dilemma. I suggest that the same problem afflicts them all.

The lesson of our earlier considerations of the first-person perspective on reasons was that, as Ginsborg says, what we're interested in from the first-person perspective are worldly facts or ways the world is. Reasons$_1$ just are worldly facts or ways the world is, or they at least *include* worldly facts or ways the world is (as when a state of knowledge is treated as a reason$_1$ to believe the consequences of its content). That is why, as I've argued, we don't from the first-person perspective consider the fact that we *believe* p to be a reason$_1$ to believe p itself or any of p's consequences. (Again, compare the fact that we *know* p: this is something we'd be happy to treat as a reason$_1$ for further beliefs about the world). Note that those

who defend the idea that perceptual experiences are non-conclusive, non-doxastic reasons for belief tend to accept this. In recent defences of this position, Feeney & Schellenberg (2020) and Moretti (2015) emphasise that perceptual evidence should be expressed using propositions representing the external world as being some way: when I see my hand, my perceptual evidence should be indicated by saying simply 'here is a hand'. They emphasise that perceptual evidence is a content that is solely about the external world; it is *not* about one's own perceptual experiences or other mental states.[122] I take this to be a recognition that, at least from the first-person perspective, we take normative reasons to be (or include) ways the world is or worldly facts.

This creates a problem for any philosopher who doesn't embrace thoroughgoing externalism. The problem is that, according to such philosophers, what we most immediately confront when we deliberate about what to believe are aspects of our perspective which are independent of the world's being any particular way: they are features that could be present in one's perspective however things are in the world. That creates the worry that we can never attend directly to the facts or the ways the world are that we are interested in when we're

[122] Both Moretti and Feeney & Schellenberg (F&S) think that when a subject comes to rationally perceptually believe p, p should be modelled as having probability 1 for that subject. Nevertheless, they don't take perceptual experience to be factive. See Moretti 2015: 265; F&S 2020: 77. The main target of these papers is White 2006, who represents the non-factivist Liberal conception of perceptual evidence as being that perceptual evidence consists in a believed proposition about one's perceptual experience, rather than an experienced proposition about the external world.

deliberating about what to believe. Take, for example, a simple version of non-factivist Liberalism espoused by Moretti, according to which perceptual evidence or reasons are just *experienced content*: 'empirical evidence is a content *experienced* rather than believed'[123], and the contents in question are everyday propositions about the external world, like 'here are some hands'. Similarly, such a view would treat inferential evidence as just *believed* content. But from the first-person perspective, it's wrong to say that my evidence or my reason for belief is a *content* or a *representation*. I might, after all, merely entertain the content in question, but I wouldn't thereby think of it as a reason$_1$. But the non-factivist seems to be forced to say that, really, what I'm attending to when I treat something as a reason$_1$ is simply a content, which can be true or false. Consequently, the non-factivist has to explain how it is that this content comes to be something that figures as a reason$_1$ in my perspective. For nothing about the content itself makes it figure there as a reason$_1$.

The solution to this problem is to appeal to something which connects the content in question to *truth* in the subject's perspective. Standardly, this is the role of the subject's stance towards the content: this is the role of belief. But we are here supposing that a content can figure in one's perspective as a reason$_1$ by being experienced rather than being believed. The role of presentational force is to solve the same problem that belief solves – to make the experience's content figure as a reason$_1$ in the subject's perspective. How might presentational force achieve this?

[123] Moretti 2015: 271.

Presumably by doing the same thing that belief does, namely by connecting the experience's content to the truth in the subject's perspective. But, on this horn of the dilemma, we are supposing that presentational force achieves this connection to truth in a quite different way from belief. In particular, we supposing that presentational force achieves this by entering into the subject's perspective as an aspect of her apparent reason$_1$, rather than by playing the role of reason$_2$.

This produces a puzzle. As we have been continually noticing, belief cannot make its content figure as a reason$_1$ in a subject's perspective by *itself* being part of the reason$_1$, or by itself being part of what the subject attends to in trying to decide what to believe. Were the subject S to believe that 'I believe that p', this higher-order belief would not make p figure as an apparent reason$_1$ in the subject's perspective. But, on the horn of the dilemma under consideration, things must be different in the case of experience. For, on the approach under consideration, the presentational force that the experience attaches to its content is meant to be an aspect of the reason$_1$. To be sure, it is not, as Smithies emphasises, itself a *content*, but is rather an ineffable feeling that one has *when* one hallucinates or enjoys a veridical experience. But, on the approach we're considering, it does figure along with or as a property of the experience's content as part of what one attends to when deliberating about what to believe. The puzzle is how and why presentational force can do this, while belief cannot. That is, a belief cannot connect its content to the truth in the right kind of way to make that content figure

as a reason$_1$ in the subject's perspective by the belief's itself being part of what the subject attends to in deciding what to believe. But presentational force *can* do this. How? And why this difference between belief and presentational force?

As we saw in section 3.3, the reason that beliefs are not reasons$_1$ and do not enter into one's considerations when one is deliberating about what to do or believe is that beliefs are non-factive: they can occur whether or not their contents are true. The very problem that beliefs are meant to solve for those who deny thoroughgoing externalism – connecting contents to the truth – re-emerges for belief itself. For beliefs themselves can be true or false, and so when a subject attends to their own belief something needs to connect that belief to truth in order for it to figure as a reason$_1$ in the subject's perspective. When one deliberates about what to do or believe, one looks for *facts* or features of the world itself; mere beliefs don't cut it. But according to the non-factivist Liberal, presentational force is non-factive in just the same way that belief is: it can attach to contents whether or not they are true, since hallucinations and illusions can have it in just the same way that veridical experiences have it. It might be replied that presentational phenomenology is different from belief insofar as each instance of presentational phenomenology presents *itself as*, in the case in question, attaching to a true content. But it's not clear that belief does not do that: beliefs represent their content as being true, and so each belief might be said to present itself as attaching to a true content. Even if presentational

phenomenology and belief present themselves as attaching to true contents, it remains the case that, on the approach we're considering, both of them can attach to falsehoods. Since this seems to be what rules out belief as something we concern ourselves with from the first-person perspective of deliberating over what to believe about the external world, it presumably rules out presentational force too, so long as it is conceived in the way that we have been discussing. The puzzle was how and why presentational force can do something that belief cannot do – connect its content to the truth, in the right kind of way to make that content figure as a reason$_1$ in the subject's perspective, by itself being part of what the subject attends to in deciding what to believe. My response to the puzzle is that insofar as presentational phenomenology is, like belief, non-factive, it cannot do this after all. Thus, on the second horn of the dilemma, according to which presentational force is an aspect of the subject's reason$_1$ or supposed reason$_1$, it cannot do the job that it is meant to do. It cannot make the experience's content figure as a reason$_1$ in the subject's perspective.

So on neither horn of the dilemma can presentational force make an experience's content figure as a reason$_1$ in the subject's perspective. This section began with the question: does non-factivist Liberalism have a satisfying response to the concerns that we have seen Ginsborg raising about the very idea of non-doxastic reasons$_2$? I conclude that their uses of the notions of presentational phenomenology and assertoric force, at least, do not offer any such satisfying response.

3.5. Conclusion

Let's circle back to Ginsborg's argument against McDowell, to see what we can say about it now. As we saw in Chapter 1, she argues that if experiences are to genuinely rationally constrain thought then they must be reasons$_2$ – the kind of reasons we appeal to in rationalising explanations. And for experiences to be reasons$_2$, they need to make the subject 'take' something to be a reason$_1$. As Ginsborg conceives of this 'taking' condition, this means that simply in virtue of having an experience with p as its content, the subject must 'take' p to be a reason$_1$ to believe p. And, while it is not entirely clear what Ginsborg means by this, she seems to think that this 'taking' involves the subject actually believing that p is a reason$_1$ to believe p. Certainly, she says, it involves the subject being disposed to cite p as a reason$_1$ – and presumably, absurdly, as a reason$_1$ for believing p.[124] But experiences can't do this, she says. I must believe p in order to be disposed to assert it, as I must do when I cite it. Furthermore, if I did 'take' p to be a reason$_1$ for believing p, there would be no psychological room to do this while refusing to believe p, since the reason$_1$ would be conclusive.

In Chapter 2 I argued that Ginsborg builds too much into her 'taking' condition. In particular, I argued that for a subject to believe q on the basis of p, the subject need not have 'taken' p to be a reason$_1$ to believe q

[124] Ginsborg doesn't actually point out the absurdity of this. Brown 2018 uses this point to argue against McDowell.

before coming to believe q on the basis of p. And I argued that 'taking' should not be construed doxastically or in terms of rule-following. Instead, I argued, we should embrace a fairly minimalist construal of 'taking', or of what is involved in a reason$_1$ (or supposed reason$_1$) figuring as such in the subject's perspective. This led to an account of inference in which sub-personal processes play an essential role, and thus in which there is a degree of passivity as well as activity on the part of the thinker. My hope was that this would prepare the way for an account of how perceptual experiences can make reasons$_1$ figure as such in the subject's perspective, given that Ginsborg's criticism of McDowell boils down to a worry that if the subject is passive in having a perceptual experience then there is no sense in which, in having the experience, the subject can count as registering an apparent$_1$ reason *as* a reason$_1$.

Although Chapter 2 offered a more minimalist account of 'taking' than Ginsborg supplies, it still required reasons$_2$ to make something figure as a reason$_1$ in a subject's perspective. For, on my account of inference, the subject is not entirely passive: part of what it is to have a reason$_2$ is to treat something as a reason$_1$. In the case of inference this involved being disposed to give instructions and ask questions of the right kind. For it to make sense for the subject to give these instructions and ask these questions, the reason$_1$ had to figure as such in her perspective, and in particular it had to figure as a fact in her perspective. In inference, this is taken care of by the subject believing or knowing something to be the case: the content of their doxastic state figures as a fact (or supposed fact)

in virtue of being believed or known. This leaves us needing to find a way in which perceptual experience makes something figure as a fact, or supposed fact, in one's perspective.

Chapter 3 began by considering McDowell's own suggestion about the rational significance of experience: in 'good' cases, one's factive perceptual experience puts one in a position to know its content and to know that one is seeing that content to be so. While this may be a good starting-point for explaining how it is that perceptual experiences can make a reason$_1$ figure as such in one's perspective, it doesn't say enough. The notion of being in a position to know is too vague, applying to cases in which one clearly lacks any mental state which an interpreter could cite in a rationalising explanation of perceptual belief. It does, however, point the way towards a positive characterisation of what it is about perceptual experiences that makes them reasons$_2$ by highlighting that, as McDowell thinks of things, they offer *conclusive, factive* rational support. In section 3.2 I pointed out that Ginsborg's criticism of McDowell does not take into account this aspect of his view; Ginsborg herself assumes that only non-factive states can be reasons$_2$. Section 3.3 described the rationale for thinking that reasons$_2$ are indeed factive (or relational): in short, if we deny this then we struggle to make sense of the very idea of rationalising explanation.

Section 3.4 offered what amounts to a development and vindication of Ginsborg's argument against the possibility of non-doxastic reasons$_2$, so

long as we assume that only non-factive, non-relational states can be reasons$_2$. Given this assumption, our perspective is exhausted by non-factive states, so that whatever figures in our perspective is independent of the world's being any particular way. Thus, if anything is to figure as a reason$_1$ or an apparent reason$_1$ in one's perspective, it must somehow be connected to how things are in the external world, or to the truth. A content, for example, must somehow be assigned the truth-value *true*. Now, when we are dealing with reasons$_1$ which lie inside thought, this problem is apparently solved by the subject believing a content to be true. While I do not think that this solution ultimately works, it has a degree of plausibility because beliefs can achieve this mapping of contents to truth without figuring in the subject's perspective as a potential object of attention. Thus, the fact that from a first-person perspective beliefs are not reasons$_1$ can seem not to matter: reasons$_1$ and reasons$_2$ can be said to play different roles, as Ginsborg argues. But reasons outside thought are a different matter: their figuring as reasons$_1$ in the subject's perspective cannot depend on the subject adopting any attitude towards them. Instead, I argued, whatever makes them figure as reasons$_1$ must be something – call it 'presentational force' – which achieves this by figuring in the deliberating subject's perspective as a potential object of attention. But presentational force is itself non-factive. Thus, just as a belief cannot make its content into a reason$_1$ by itself being an object available to the deliberating subject's attention, surely presentational force cannot either.

In short, the problem is that, given the rejection of thoroughgoing externalism, whatever figures in a subject's perspective can only figure there as a reason$_1$ if something connects it to the truth. But whatever connects it to the truth cannot do so by itself figuring in the subject's perspective, since this further feature of the subject's perspective will also need to be connected to the truth, for the same reason as the original feature of the subject's perspective needed to be connected to the truth. So whatever does the truth-connecting needs to achieve this without figuring in the subject's perspective. At the same time, however, this truth-connector must be something that we can appeal to in a rationalising explanation of the subject's beliefs and actions, and so it must be an aspect of the subject's perspective, even if it is not something which figures *in* the subject's perspective. This leads to the idea that the truth-connector is an *attitude* of the subject – and in particular a belief. This, then, is what is right about Ginsborg's argument. If we assume, as she does, that our perspectives are exhausted by non-factive states, then there is good reason to think that only a belief can make something figure as a reason$_1$ or apparent reason$_1$ in a subject's perspective, and so that only a belief can be a reason$_2$.

If, on the other hand, we accept thoroughgoing externalism, how does this neutralise Ginsborg's argument? We need to show that perceptual experience can make something figure as a reason$_1$ in one's perspective, and I have argued for quite a minimalist construal of what this figuring amounts to. Since I have argued that this figuring need not involve the

subject *believing* something to be a reason$_1$, we need not worry about the lack of psychological room between 'taking' (i.e. believing) p to be a reason$_1$ to believe p and believing p. But we need to show that something can figure in one's perspective as a fact or as a way that the external world is, without the subject believing it to be a fact or a way the world is. The suggestion that the thoroughgoing externalist can make here is that experience does this by presenting one with a fact or a way that the world is. In virtue of one's perceptual experience, this worldly thing enters into one's perspective as the worldly thing that it is. There is no need for a belief to connect it to the truth because it is not something that can be false, or that can be present in one's perspective independently of the world's being some way. But, at the same time, it really is present in one's perspective. In particular, on the relational approach, perceptual experience is *constituted by* a relation between the subject and external-world objects and properties, and as a result of this its *phenomenal character* is partly determined by these external-world objects and properties. This claim provides the resources needed to expand on McDowell's account of the rational significance of perceptual experience, going beyond his claim that this significance is constituted by perceptual experience's putting subjects in a position to know things about the external world. For to this we can now add that the rational significance of perceptual experiences is constituted by the experience having a phenomenal character that is determined by the external-world facts or objects which make true one's perceptual judgements. It is this character which makes these facts or objects figure in one's perspective as

reasons$_1$. It does this not by connecting them to the truth, for no such connection is required. Instead, it does this simply by being the way that these facts or objects enter into one's perspective. The role of the phenomenal character of experience, that is, is not to connect an aspect of one's perspective to the truth, but is simply to be the figuring as such of a fact or object in one's perspective.

Of course, there is much more to say in order to flesh this picture out, and much of it has already been said.[125] We need to get clear about what exactly one's reasons$_1$ are: are they facts, objects, experiences, or something else? And we need to say more about what is required of a subject in order to come to have perceptual knowledge. In particular, do they need to know what their reason$_1$ is, and be disposed to cite it, in order to count as having perceptual knowledge? And, crucially, in what ways is the subject active in coming to have perceptual knowledge, and in what ways are they passive? As we saw in Chapter 2, Eilan has suggested that perceiving involves the subject actively interrogating the world by asking questions about it; but, equally, there are multiple ways in which the subject might be said to be passive in coming to have perceptual knowledge, and which have not received anything like full exploration in the last three chapters.

For the moment, however, I leave these issues. My aim has been to show that thoroughgoing externalism offers a way – perhaps the only way – of

[125] See Kalderon 2011; Brewer 2011; Brewer 2018; Travis 2013; Logue 2018 for ways of developing the picture I've been describing.

vindicating the idea that there are perceptual reasons for belief that fulfil the Reasons Outside Thought condition. By arguing, contra Ginsborg, that we should think of the figuring of reasons$_1$ as such in one's perspective in quite minimalist terms, I have made room for the idea that non-doxastic experiences can make reasons$_1$ figure as such in one's perspective simply by placing external-world entities in one's perspective.

Chapter 4

Defeat

My response to Ginsborg's argument that there can be no non-doxastic reasons involved appealing to the idea of factive or relational perceptual experiences which provide conclusive evidence or reasons (which are both reasons$_1$ and reasons$_2$) for perceptual beliefs. Perceptual reasons for belief, as I have suggested we understand them, are conclusive because they are constituted by those features of the world that make true the beliefs they are reasons for. Furthermore, they figure in one's perspective *as* the reasons that they are, simply in virtue of one's experience having the phenomenal character that it has. Consequently, as McDowell emphasises, they put one in a position to know the propositions that they are reasons to believe. In this way, they are not just truth-guaranteeing, but knowledge-guaranteeing too: perceptual beliefs that are properly formed on the basis of perceptual reasons are guaranteed to be knowledge.

In this chapter I will examine one kind of argument against the notion of knowledge-guaranteeing evidence or reasons. An increasingly common pattern of argument in epistemology is to suggest that a given epistemological theory cannot accommodate or explain a certain kind of *defeat*. The argument against knowledge-guaranteeing evidence or

reasons that I want to consider is of this kind. Defeat, roughly speaking, is the loss of an epistemic status like knowledge or justification in a particular kind of way. Knowledge can be lost by forgetting, by ceasing to believe, by losing evidence, by the known proposition becoming false, and so on. But knowledge-defeat cases are cases where the subject knows p, and then on receiving a defeater no longer knows p, but can continue to believe p on the same basis as before receiving the defeater. Even if the subject retains their belief in this way, it will no longer amount to knowledge, despite being true and based on evidence or a reason which previously sustained it as knowledge. Defeaters tend to be evidence that one has erred, and so that one does not know p. There is something puzzling about such cases when they are described in this way: if one is justified in believing p, or knows p, how can one later continue to believe p on exactly the same basis and yet fail now to be justified or to know p? After all, as theorists we appeal to the basis of the belief to explain why it is justified or amounts to knowledge. But these kinds of defeat case suggest that believing on the original basis was not sufficient for ensuring knowledge or justification.

These cases are particularly puzzling for philosophers who think that knowledge is based on conclusive evidence or reasons, which are such that beliefs properly based on them are guaranteed to knowledgeable. This infallibilist conception of reasons and evidence seems to predict that knowledge-defeat is not possible. Insofar as compelling cases of knowledge-defeat can be described, then, the infallibilist seems to make

false predictions. In this chapter I take up this challenge to the infallibilist conception of perceptual reasons that I argued for in Chapter 3. First I describe various kinds of defeat, and outline the options that infallibilists and fallibilists have in analysing and explaining such cases. Next, I describe the challenge to infallibilism in more detail. We see that the infallibilist either has to deny that knowledge-defeat ever takes place, or has to explain how defeaters disrupt the subject's ability to sustain their knowledgeable belief on the same basis as before receiving the defeater. Neither option seems palatable, on the face of it. Nevertheless, I argue, there is a viable path towards denying that knowledge defeat ever takes place. It seems likely that, in denying the reality of knowledge-defeat, the infallibilist will have to say different things about the rational effects of different kinds of defeater on different kinds of knowledge. Thus, I suggest that one popular way of denying that knowledge-defeat takes place, the Level-Splitting strategy, may provide an accurate analysis of some kinds of putative defeat case, but that it is less aptly applied to cases of putative defeat of specifically perceptual knowledge. Instead, then, I try to pinpoint a strategy that is better suited to the perceptual case.

4.1 Introducing defeat, fallibilism and infallibilism

Cases of defeat, as we're understanding them here, involve a subject coming to possess some new piece of evidence which destroys an epistemic status like knowledge or justification. The new evidence is

called a 'defeater'. Taxonomising these defeaters is an ongoing enterprise: Dutant & Littlejohn (2021) gather from the literature opposing (or rebutting), undercutting, reason-defeating, collective, pragmatic, identification, higher-order and negative self-appraisal defeaters.[126] We will be discussing opposing, undercutting and higher-order defeaters in this chapter.

Opposing defeaters for knowledge that p are evidence for not-p. For example, suppose you come to know that your neighbour is away by seeing that his car is not in his driveway. If you then notice that the lights are on in his house (and you know, say, that he almost never leaves the lights on when he's out), this might be an opposing defeater for your knowledge that he's away (yet, we can imagine, he *is* away, and he did leave the lights on this time because he left in an exceptional hurry).

Undercutting defeaters are typically described as defeaters which give one reason to think one's basis for believing p does not in fact support or justify believing p.[127] Suppose, for example, that you look at a red wall in normal viewing conditions and it looks red to you, and on this basis you come to know that the wall is red. But suppose that you are then misleadingly told by an informant (who you are justified in trusting) that

[126] Dutant and Littlejohn 2021: 223.

[127] E.g. Dutant and Littlejohn 2021. They are also sometimes described (e.g. in Kotzen 2019) as defeaters which do not directly support believing not-p but which nevertheless undermine the support your evidence gives you for believing p. But this description fails to provide a sufficient basis on which to distinguish them from the other kinds of non-opposing defeaters on Dutant and Littlejohn's list.

the wall is illuminated by red lights capable of tricking people into thinking white walls are red. Or suppose, instead, that you are misleadingly told by an informant (who you are justified in trusting) that you are currently hallucinating. Both of these new pieces of evidence are standardly classified as undercutting defeaters for your knowledge that the wall is red: neither of them directly support believing that the wall is not red, but both seem to upset the connection between your perceptual evidence or perceptual belief-formation process and the proposition that the wall is red.

Of course, in all of these cases it is likely that the subject would cease to have the relevant belief, and so in that way would cease to know, given that knowledge entails belief. But the point is that if the subject stubbornly ignored the defeaters and hung on to their belief on the same basis as before, they would still apparently cease to have knowledge. Giving up the belief is apparently the epistemically appropriate thing to do.

We can divide the various possible reactions to such cases into two broad families: infallibilist and fallibilist. As I'll understand these labels here, the infallibilist tries to give an account of bases, whether evidential or process-based or both, on which the basis of one's knowledge offers conclusive support for one's knowledgeable belief. That implies that putative defeat-cases in which the subject's basis is apparently left intact either do not involve genuine defeat after all or, despite appearances, do

not allow the subject to continue to believe on the same basis as before receiving the defeater. Here is how we should think of conclusive bases:

> Conclusive Basis: For a knowledge-basis B to be conclusive with respect to a proposition p is for it to guarantee that one knows p if one's belief that p has B as its basis.

If we identify one's basis with a piece of evidence or a reason, then the evidence or reason can be thought of as conclusive with respect to p if and only if believing p on the basis of that evidence guarantees that you

know p; likewise if we identify one's basis with a process or method.[128] As we will see in more detail below, the challenge for the infallibilist is to explain why cases which seem to involve defeat actually do not; or, where the infallibilist accepts that defeat takes place, they must explain why,

[128] There are three clarifications worth making with respect to this way of defining infallibilism and conclusive bases:

1. Talk of 'basis' should not be taken to imply that infallibilists are committed to the idea that knowledge is always *based on* a basis, where 'based on' indicates a substantive basing relation holding between one's knowledge and a basis. This is certainly one way that an infallibilist might think, especially those who think of bases in evidential terms. However, for those who think of bases as processes or methods, this is not a natural way of fleshing out their view: we don't base our knowledge on a method in the same way as we might base a belief on a piece of evidence. For these infallibilists, then, the idea that the basis of your knowledge that p is your method of forming the belief that p should not be understood to imply that you base your belief that p on your method of forming that belief in the way that someone might base a belief on evidence.

2. I say that a basis is conclusive iff a belief with that basis (i.e. based on that evidence or formed via that method or process) is guaranteed to be knowledge. Note that this should strictly speaking be distinguished from the idea that a basis is *factive* or *relational* since one could hold that, say, perceptual experiences are factive while denying that perceptual beliefs properly based on perceptual experiences are always knowledgeable. One might be motivated, in holding this view, by fake-barn country cases: suppose that you are travelling in fake-barn country, and happen to see the only genuine barn, and on this basis come to believe that there is a barn before you. Some philosophers have the intuition that your belief is not knowledgeable, despite being based on a genuine, factive perceptual experience. This is a tricky issue which I cannot adequately address here. Suffice to say that I agree with McDowell that the rational significance of perceptual experience is that it puts one in a position to know. If this is right then knowledge-grounding perceptual evidence is conclusive, in the sense I've been describing. The factivity or relational nature of experience is a central part of what enables it to provide conclusive evidence, but the fake-barn country case suggests that this may not be the end of the story.

3. 'Infallibilism' is usually thought of as the view that knowledge-grounding evidence must be conclusive; it does not involve the further anti-sceptical commitment to the claim that we do in fact have such conclusive evidence. However, here I am using 'infallibilism' to refer to a position which does include this further anti-sceptical claim.

contrary to appearances, the subject's original basis for believing is no longer available to them.

The fallibilist, in contrast, holds that the basis of one's knowledge of p can be inconclusive, which is to say that believing p on that basis does not guarantee that one knows p. One could continue to believe p on the same basis and yet no longer know p, as a result of some defeating factor. Fallibilists sometimes argue that this approach is more in-keeping with our judgements about defeat cases. But the fallibilist must say more about what further conditions, beyond the basis of one's belief, make one's belief into knowledge. A potential cost of fallibilism is that it can be hard to see how something can be a basis of one's knowledge (or doxastic justification) while failing to guarantee that one knows (or is doxastically justified).[129]

I've defined infallibilism and fallibilism in terms of knowledge, but there are also related positions with respect to doxastic justification. Here is a very cursory overview. An infallibilist about the bases of knowledge might also say that the bases of doxastic justification are knowledge-guaranteeing, so that all doxastically justified belief is knowledge. Alternatively, one might say that the bases of doxastic justification guarantee doxastic justification but not knowledge. A further alternative is to say that the bases of doxastic justification guarantee neither knowledge nor doxastic justification. In practice, infallibilists about the

[129] This point has been made forcefully by McDowell. See in particular McDowell 1982; McDowell 1998; McDowell 2011. See also Chapter 3 of this thesis.

bases of knowledge tend to pick one of the first two options with respect to justification, while reliabilist fallibilists (for whom bases are truth-conducive processes) pick the second option and evidentialist fallibilists (for whom bases are evidence or reasons) pick the third.

Fallibilists and infallibilists thus envisage different explanatory roles for bases. That one believes p on a certain conclusive basis provides a full explanation of how one knows p.[130] But that one believes p on a certain inconclusive basis can't be the full explanatory story about how one knows p, and nor can it be the full explanatory story about how one's belief is doxastically justified according to the evidential fallibilist. For fallibilists, then, there are at least some cases where the basis of one's belief is only part of what makes it knowledge.

4.2 White against infallibilism

As I said earlier, a common pattern of argument in epistemology is to suggest that a given epistemological theory cannot accommodate or explain a certain kind of defeat. This kind of argument has been particularly directed against infallibilists. Consider a recent argument made against infallibilism by Roger White. Let's suppose that

[130] This shouldn't be taken to imply that all infalliblists think that we can give a fully *satisfying and non-circular* explanation of how one knows p by appealing to the basis of one's knowledge. For some infallibilists hold that an analysis of the basis will involve reference to the fact that one knows p. See, for example, Williamson 2000.

infallibilism is true, so that the basis of perceptual knowledge is conclusive. White argues that if this is so then we cannot explain how defeaters can defeat perceptual knowledge. Consider the following case:

> **Suspected Hallucination**: You see your hand, and thereby come to know that you have a hand (*I have a hand* = 'H'). You then progressively gain more and more misleading evidence that you are currently hallucinating the hand (the evidence is misleading because, in fact, you're seeing the hand in a perfectly normal way). However, you are stubborn, and you continue to believe H in spite of the misleading evidence that you are hallucinating.[131]

White says that it's clear in this case that at a certain point the misleading evidence that you are acquiring acts as a defeater so that your belief in H ceases to be knowledge. However, it does not seem, says White, that the basis of your stubbornly maintained belief changes once you receive the defeating evidence. You seem to continue to believe H on the same basis as before. This suggests that your basis is inconclusive – that it is not knowledge-guaranteeing.

Indeed, says White, a fallibilist can easily explain the mechanism of defeat in such cases. What he calls the Cartesian Picture of perceptual evidence would give the following account. According to the Cartesian Picture, perceptual evidence is constituted by propositions like E (*that it*

[131] See White 2014: 313–314. I am paraphrasing his 'Suspected Hallucination' case.

appears to you that H). We can then explain the mechanism of defeat in Apparent Hallucination in probabilistic terms:

> Evidence E provides support for H to the extent that E is more to be expected given H than given not-H. The misleading evidence [that you are hallucinating] reduces the difference between these likelihoods […]. Evidence of hallucination *screens off* the evidential link between the appearance that H and H.[132]

White adds that we might alternatively explain the defeat 'in terms of explanation'.[133] What explains its appearing to me that there is a hand? He suggests that there are, *prima facie*, two good explanations: there being a hand, and my suffering from a hallucination or illusion as of a hand.[134] Evidence of hallucination favours the latter explanation.

Whether we take the probabilistic or the explanationist approach, we can explain the defeat while allowing that one continues to believe on the same basis as before. The role of the defeater is to weaken the supporting connection between the perceptual evidence and the proposition H, in a way that the subject can recognise.

---

[132] Ibid., 314. White's emphasis. I've altered the labelling to match the labels I've been using. To see how this kind of probabilistic approach to defeat might be developed in more detail see Kotzen 2019.

[133] Ibid.

[134] Actually, White describes the explanations a little differently, but I think his descriptions reduce to these kind of explanations, which I use for the sake of simplicity of presentation. See ibid., 314.

Of course, it's only if we think of perceptual evidence as inconclusive that we can conceive of this supporting connection being weakened or destroyed. If perceptual evidence, or the non-evidential basis of a perceptual belief, is conclusive then by definition no evidence that one later gets can defeat the support it offers for the relevant belief. Thus, goes White's line of thought, if the justifying bases of perceptual beliefs are conclusive then there are just two ways of thinking about what happens in Suspected Hallucination.

First, we might deny that any defeat of knowledge takes place here, so long as the subject stubbornly maintains their belief in H. White regards this as implausible. Second, we might say that the basis of the subject's belief in H somehow changes when she acquires the misleading defeating evidence. Referring to conclusive perceptual evidence as involving the 'direct apprehension' of a fact or of the presence of a hand, White suggests that the infallibilist taking this route would have to say something like the following:

> When I first look at my hand it is not only apparent to me that there appears to be a hand before me, I also directly apprehend its presence and hence have an additional source of justification for my conviction that here's a hand. As it becomes more likely that I am hallucinating I lose […] my direct access to the presence of the hand and I am left with just the fact that there appears to be a

> hand before me. And then my predicament is just as the Cartesian describes it: this evidence provides less support for there being a hand before me the more likely it becomes that I'm hallucinating.[135]

But, says White, this story is implausible too:

> [W]hy is it that misleading evidence that I'm hallucinating robs me of this extra evidence? Nothing appears to have changed about my perceptual state and connection to the world. Suppose I […] go on dogmatically thinking *here's a hand* despite the mounting evidence that I'm hallucinating. I am at least still *seeing my hand*. My eyes are still receiving accurate information from the world and my belief forming mechanisms are operating as usual in forming a true belief. Isn't that enough to have *direct perceptual access* to the layout of my environment?[136]

4.3. Level-Splitting

Let's consider the possibility, which White dismisses, that defeat does not take place in in Suspected Hallucination**.** For there is a growing literature

[135] Ibid.

[136] Ibid., 314–315.

defending the idea that genuine defeat of knowledge is less pervasive than at first sight it might appear.

An initial classificatory point: exactly how we categorise the putative defeater in Suspected Hallucination depends on whether we are fallibilists or infallibilists. At first sight, the infallibilist should classify the putative defeater as higher-order, and probably not as undercutting or rebutting (but note that defeaters may fall under more than one of these categories at the same time). Recall that the infallibilist holds that perceptual evidence is conclusive, which we said means that if one properly bases one's perceptual belief on it then one is guaranteed thereby to come to have perceptual knowledge. Mere hallucination is not and does not provide conclusive evidence about the external world, so evidence that you're merely hallucinating is evidence about what evidence you have; in particular, it is evidence that you lack the kind of perceptual evidence that the infallibilist says we have when we perceive normally. It is thus evidence about your first-order evidence: this makes it higher-order.[137]

[137] It is also evidence that you were never justified in the first place (assuming that hallucinations do not provide justification for external-world beliefs), and certainly that you never had knowledge, even before you received the putatively defeating evidence about hallucination. For it is evidence that you were hallucinating all along. It thus functions exactly as Lasonen-Aarnio says higher-order evidence does: 'defeat by higher-order evidence has a retrospective aspect, providing a subject with evidence that her belief was never rational, reasonable, or justified to start out with.' Lasonen-Aarnio 2014: 317.

On the other hand, the putative defeater is probably best classified as undercutting on a fallibilist conception of perceptual evidence. Undercutting defeaters in cases of perception add to your total evidence in such a way that your original perceptual evidence no longer supports what it supported before you received the defeater. But they do not alter or suggest anything about the prima facie evidential force of that perceptual evidence, and thus do not suggest that your doxastic response to that perceptual evidence was irrational before you received the undercutting defeater. Fallibilists about perceptual evidence generally think that you have the same perceptual evidence whether you're veridically perceiving or merely suffering from a subjectively indistinguishable hallucination or illusion. So evidence that you're merely hallucinating is not evidence about what evidence you have or even what that evidence, taken on its own, supports (that is, it suggests nothing about the *prima facie* evidential force of your perceptual evidence). It is thus not higher-order, in the way that we saw it is for the infallibilist. Rather, it is evidence about the etiology of your perceptual evidence: it suggests that your perceptual experience was not caused in the right kind of way by your environment to count as part of an episode of successful perception. Since the right kind of causal connection is what underpins perceptual experience's propensity to reveal truth, evidence that it doesn't obtain is evidence that your experience is not veridical. Its evidential force is thereby undercut. There is of course a sense in which the undercutting evidence here is higher-order, insofar as it concerns the etiology and, indirectly, the truth-value of your evidence and so concerns

your evidence, but there is also a clear distinction between the way that this undercutting evidence defeats (if it does defeat) and the way that higher-order evidence defeats (if it does defeat). The undercutter leaves you thinking that you were right about what your evidence was, but weakens the connection you thought obtained between that evidence and the truth; the higher-order defeater leaves you thinking you were wrong about what your evidence was.

In responding to White, then, the infallibilist should start from the assumption that the putative defeater is higher-order. This opens up the possibility that the infallibilist can apply what has become known as a 'Level-Splitting' analysis to the case, according to which the case does not involve genuine knowledge-defeat. Let's see whether this strategy works.

According to Level-Splitters, epistemic akrasia is sometimes the epistemically appropriate response to putatively defeating misleading higher-order evidence. A thinker is epistemically akratic if they believe p while believing that they ought not to believe p (or that they probably ought not to believe p, or they have high confidence that they ought not to believe p). Within the context of evidentialism, according to which what is rational to believe is determined by one's evidence, a thinker is paradigmatically epistemically akratic when they believe *p but my evidence doesn't support p* or when they believe *p* while also being highly confident

that they lack evidence supporting *p*.[138] Epistemic akrasia can be characterised in non-evidentialist terms (*p but I ought not to believe p* is a more neutral characterisation of an epistemically akratic belief-content), and the considerations motivating Level-Splitting views can be translated into non-evidentialist frameworks; here I continue to work in an evidentialist framework.[139]

The driving motivation for Level-Splitting views, stated in evidentialist terms, is that in order to fully respect both one's first-order and one's higher-order evidence, one may sometimes need to have beliefs at the first-order level which in some way conflict with one's beliefs at higher order levels. Level-Splitting is often defended by appeal to cases.[140] Consider Long Deduction:

> Long Deduction: You competently undertake a complex deductive inference, involving many steps, finally in this way coming to believe C, the conclusion of the inference. You know that you are

[138] See Lasonen-Aarnio 2010, Lasonen-Aarnio 2014, Lasonen-Aarnio 2015; Lasonen-Aarnio 2020; Lasonen-Aarnio 2021; Williamson 2014; Wedgwood 2012. If justified belief must be based on conclusive evidence, as some hold, then one cannot rationally believe *p but my evidence doesn't support p*, since one cannot know this. However, believing p while having high confidence that one's evidence does not support p cannot be ruled out as inevitably irrational in this way. See Lasonen-Aarnio 2020: 609.

[139] See Lasonen-Aarnio 2020: 602 for an argument that even if we replace evidentialism with other popular theories of rational belief and knowledge, there will remain an incompatibility between our theory of rational belief and knowledge and anti-akratic principles.

[140] See Lasonen-Aarnio 2008; Williamson 2011; Williamson 2014. The Long Deduction case is due to Williamson.

> fallible in making deductive inferences, and so you know that for each step in the long deduction there was a small risk that you made an error. Adding up the risk of error in each step, you realise that there is a very high probability that you made an error at some point in the long deduction. In fact, you made no error, and so you really do have evidence that entails C, and you now believe C on the basis of this evidence. But you also know that it's highly likely that you made an error in your deduction.

The Level-Splitter argues that to give up your belief in C would be to fail to do justice to your first-order evidence, which after all entails C, and which continues to do so even when you knowledgeably reflect on the high possibility of error. But, equally, you do have evidence that it's highly likely that you made an error and so that your evidence doesn't support C. So to do justice to this higher-order evidence you should also believe that you lack evidence supporting C, or that it's highly unlikely that you have evidence supporting C. It also seems that you should believe that it's highly unlikely that you know C. So, if you should always believe in accordance with all your evidence, then you should have an epistemically akratic belief. Other cases often discussed in the higher-order evidence literature, such as ones in which you receive misleading evidence that you have undertaken some reasoning under the brain-addling influence of hypoxia or sleepiness, have a similar structure. To complete a defence of Level-Splitting, one would need to explain why

epistemic akrasia is so often judged to be obviously irrational. In the next section we will consider an error-theory which purports to do this.

Before we try to apply the Level-Splitting analysis to Suspected Hallucination, note a complication. Your misleading evidence that you're hallucinating is both a putative higher-order defeater for your perceptual knowledge of H *and* a putative *rebutting* defeater for your knowledge (assuming you have such knowledge) that you are now seeing a hand, and that you are now seeing that there is a hand. Perhaps you don't know that you are now seeing a hand, but you are surely in a position to know this, assuming that this is a fairly standard case of perception. One way that an infallibilist might therefore try to deal with Suspected Hallucination is to insist that one has conclusive evidence that one is seeing a hand, and one can leverage this evidence to rebut the misleading evidence that one is merely hallucinating. With the misleading defeater itself defeated in this way, one can simply go on knowing both H and that one sees a hand. Call this the Steadfast strategy. The Steadfast strategy is attractive for infallibilists, because it saves them from having to explain how rebutting defeat of knowledge takes place, which is as difficult a challenge for infallibilism as explaining how undercutting and higher-order kinds of defeat take place. In short, then, Level-Splitting might seem not to be needed for the purposes of explaining why defeat does not take place in Suspected Hallucination, given the subject's possession of evidence which might pre-emptively

defeat the putative higher-order defeater. Let's put this suggestion to one side for now; I will come back to it in the next section.

According to the Level-Splitting analysis of Suspected Hallucination, the epistemically appropriate thing for you to believe in this situation is something like *H but I'm merely hallucinating so I don't have any evidence for H*. I will now argue that the problem with this Level-Splitting approach, as applied to Suspected Hallucination, is that the first- and higher-order levels cannot be kept apart, so that if one is epistemically licensed to have epistemically akratic beliefs at these different levels then one will be licensed to have straightforwardly inconsistent beliefs at the same level. But it's implausible that the epistemically appropriate position for the subject to take in Suspected Hallucination is to have straightforwardly and obviously inconsistent beliefs. Epistemic akrasia looks bad, but this apparent badness can perhaps be explained away (see the next section). But it's difficult to see how Level-Splitters could persuade us to live with outright inconsistency, such that the epistemically appropriate thing to do can sometimes be to believe p and not-p. As we will see in the next section, this brings us back to the issue of rebutting defeat.

To see the problem with Level-Splitting in the context of Suspected Hallucination, suppose you respond to the misleading evidence that you're merely hallucinating by entering an epistemically akratic doxastic state, as Level-Splitting recommends: you believe H but you also believe

that you're not enjoying a genuine perceptual experience[141] of your hand, that you're merely hallucinating, and thus that you lack evidence for H (since you know that your original evidential basis for believing H was putative perceptual evidence). The problem emerges if we consider more closely your perceptual beliefs. If you continue to believe H on the basis of your factive or relational perceptual experience, then you presumably also continue to believe (or at least would be justified in believing) something like *That is a hand*, where the demonstrative refers to your hand, to which you are now visually attending. But, equally, if you justifiably believe that your experience is hallucinatory, or that it is not a genuine perceptual experience, then, attending to your hand (which you can surely still do despite this belief about your experience), you are also permitted to believe something like *That is not a hand*, and perhaps something more positive like *That is merely a hallucinatory hand* or, at least on some accounts of hallucination, *That is a mere sense-datum*. After all, if you are merely hallucinating then it straightforwardly follows that whatever it is that you are now visually attending to must not be a hand or any other external-world object.[142] But such propositions are

[141] I use the phrase 'genuine perceptual experience' to refer to the kind of perceptual experience one has in a 'good' case: a factive or relational state or episode. Often this is referred to as a 'veridical experience', but this is potentially misleading, since hallucinations (and perhaps illusions) can be veridical.

[142] Admittedly, I am moving rather quickly here. The path from justified belief that you are merely hallucinating to justified belief that whatever you can now visually attend to is not, say, a hand may not be as straightforward as I am implying here. If it turns out that a subject who is justified in believing that they are hallucinating is *not* automatically justified in believing that *that is not a hand* in the kind of case I'm describing then Level-Splitting may turn out to be tenable after all in such cases. Of course, given that I am ultimately trying to defend infallibilism, that would good news.

straightforwardly inconsistent with *That is a hand*, on the plausible assumption that the demonstrative has the same reference across all these propositions. Therefore, the Level-Splitting view, if applied to Suspected Hallucination, seems committed to the absurd idea that it can be epistemically appropriate to believe straightforwardly inconsistent propositions.

The Level-Splitter might reply by denying that the demonstrative *that* has the same reference across the different propositions. But here we are supposing that you attend to your hand, perceptually presented by your factive or relational perceptual experience, which you are still enjoying whether you recognise it as such or not. When you do this, and use the demonstrative *that*, what else other than your hand could you possibly be referring to? There seem to be no obvious candidates. Instead, the Level-Splitter might say that *that* simply fails to refer when used in *That is not a hand*, so that this is not after all inconsistent with *That is a hand*. But it's difficult to see any reason to think this. It's true that if you actually were hallucinating, the demonstrative might well fail to refer (although this depends on what one's view of hallucinations is: if hallucinations involve sense-data then it may refer to a sense-datum). So it might be that if you believe you're hallucinating, you should believe that reference will fail. Still, it's difficult to see why this should make reference fail if you nevertheless go ahead and use the demonstrative to try to single out what is presented to you by your visual experience. And, indeed, it would be very strange if one could successfully think a thought with the content

That is a hand but be unable to think a thought that simply negates this content: *It's not the case that that is a hand*. Surely if one possesses the concept of negation and one can think a content *p* then it follows that one can think the content not-p. Perhaps this is not a general principle, since applied many times over it would yield propositions too long or too complex to think. But *It's not the case that that is a hand* is not an unthinkably long or complex proposition.

What about those Level-Splitters who are infallibilists about justified belief, and so who deny that we could ever be justified in falsely believing *I am merely hallucinating*? The position of these Level-Splitters, as applied to Suspected Hallucination, will say that you are justified in believing (since, according to them, you can come to know) H, while also being highly confident that you're merely hallucinating and so that your evidence doesn't support H. By the same argument as before, we can see that this leads to the idea that one should believe *That is a hand* and at the same time be highly confident that *That is not a hand*. Although this view doesn't involve licensing inconsistent beliefs, it is still surely implausible that one should have to have a high degree of confidence in the negation of a proposition that one should believe.

Can Level-Splitting ever be appropriately applied to cases of putative higher-order defeat of perceptual knowledge? Generally it cannot, since most (perhaps all) putative defeaters of perceptual knowledge (whether higher-order or of another sort) are arguably putative evidence that one

is not enjoying a genuine factive or relational perceptual experience, but is rather suffering from an illusion or hallucination.[143] Therefore, the argument just given against the application of Level-Splitting to Suspected Hallucination will apply generally to cases of putative defeat of perceptual knowledge. For in all such cases, if we apply the Level-Splitting approach, the subject receiving the putative defeater will be justified in believing that they are not genuinely having p perceptually presented to them, where p is the content of their putatively defeated perceptual knowledge. But if you're justified in believing that you're not enjoying a genuine, factive or relational perceptual presentation of p, then when you attend to your current visual experience, insofar as it seems to reveal p to be the case, you are justified in believing that aspects of your experience are not what they appear to be: *that is not a hand*, *that is not skin*, and so on. This is obvious in cases where your putative defeater justifies you in believing (or having high confidence) that you are merely hallucinating and so not having a genuine visual experience at all. It is perhaps less obvious in cases where your putative defeater justifies you in believing that you are suffering from an illusion, for if you are suffering from an illusion then you can still use demonstratives that refer to external-world objects and properties. But insofar as you have reason to think you're suffering from an illusion, you have reason to think that the objects and properties presented to you in your experience are not arranged in the way that they appear to be arranged, and so you can justifiably judge *Those are not thus*, where *those* refers to all the objects you

[143] See Silins 2014 for an argument that this is the case.

are now seeing and *thus* refers to their actual arrangement, which you can still demonstratively refer to since, in fact, you're not suffering from an illusion at all (we are, recall, only dealing here with misleading putative defeaters). Yet, on the Level-Splitting approach, you're also justified in believing the directly contradictory propositions *that is a hand*, *that is skin* and, in the illusory case, *those are thus*. So the argument against applying the Level-Splitting approach to putative perceptual defeaters seems quite general.

With that said, Williamson's discussion of inexact perceptual knowledge does provide an example of a case in which Level-Splitting may be appropriate within the context of perceptual knowledge.[144] It is worth seeing why my argument does not apply to the kind of case Williamson discusses. (I will stay agnostic about whether or not Level-Splitting is in fact appropriate in the case Williamson describes.) Some perceptual knowledge is inexact. Williamson describes an unmarked clock face, with a single pointer. We label the positions the pointer can be at 0, 1,…, n, where *n* is an extremely large number. Viewing the clock from a few metres away, you can reliably tell where the pointer is within a certain margin of error. Suppose the pointer is at position 10, you can reliably tell that it is within region R, where R is the region of the clock face consisting of only positions 9, 10 and 11, but you cannot reliably tell that the pointer is at position 10 because, were it at position 9 or 11 you might well still judge that it is at position 10. In other words, your

[144] Williamson 2011 and Williamson 2014.

judgement that the pointer is within R (call this proposition PWR) is safe enough to count as knowledge, but you cannot know that it is at position 10 because your judgement that this is so does not meet a safety condition on knowledge. R is the smallest region for which you can know that the pointer is in that region when it is pointing at 10. Your knowledge of PWR is inexact insofar as it does not exactly pinpoint the pointer's location, leaving open three different location possibilities. But now consider what you should think. You know PWR, but, assuming you know about your margin for error and the safety condition on knowledge, you can also work out that you only know PWR if the pointer is at position 10. For if the pointer were at position 9, say, then the smallest region you could locate the pointer within would not be R but rather R′, the region consisting of positions 8, 9 and 10. So, assuming that you cannot know that the pointer is at 10, you cannot know that you know PWR. Further, in a similar way, you can work out that you should have a low degree of confidence that you know that PWR. For, given that all that you know is PWR, your confidence that the pointer is at 10 should be 1/3. So your confidence that you know PWR should be 1/3. Williamson extends his argument to show that you can know a proposition p while it is extremely unlikely on your evidence that you know p.

If we identify your perceptual evidence in this case with the strongest proposition about the pointer's position that you can know, then your perceptual evidence is constituted by PWR. If we then identify your

perceptual evidence with a proposition presented by your perceptual experience, then you have a perceptual experience which presents PWR. But you also have evidence that supports thinking that you lack perceptual evidence supporting belief in PWR, and so that you are not having a perceptual experience which presents PWR. Suppose you consequently become epistemically akratic, believing PWR but also having high confidence that you are not having a perceptual experience which presents PWR. Can't you then reason in the following way?

> It's unlikely that I'm enjoying a genuine perceptual experience presenting the fact PWR. Thus, I should have low confidence that *That is in there* ('that' refers to the pointer; 'there' refers to R; both are perceptually attended to) and that *Those are thus* (see the above discussion about illusions). But, equally, the pointer *is* in R, so *That is in there* and, in general, *Those are thus*.

Again, it appears, Level-Splitting licenses obviously inconsistent beliefs when it is applied to a case of putative defeat of perceptual knowledge. However, this is not so. For in this case, your higher-order evidence is *not* evidence that you are undergoing a hallucination or an illusion, or more generally that you are not enjoying a perfectly genuine perceptual experience. Consequently, your high confidence, rationalised by your higher-order evidence, that you're not enjoying a genuine perceptual experience presenting the fact that the pointer is in R is not high confidence that you are hallucinating or suffering from illusion. To make

this clear, we could say that what you have high confidence in is the proposition that your genuine perceptual experience does not present the fact that PWR; this makes clear that you are not doubting (and have no reason to doubt) that you are undergoing a genuine perceptual experience. Instead, what you are doubting is what that perceptual experience is presenting to you as being the case. In particular, you rationally doubt that it is presenting to you PWR. But so long as you do not doubt that your experience is a genuine, non-illusory, non-hallucinatory experience, you have no more reason to move from this rational higher-order doubt to doubting that *That is there* or *Those are thus* than you do to move from this rational higher-order doubt to doubting PWR.

So the epistemically akratic position that Williamson's argument seems to recommend in this case does not lead to the kind of absurdities that epistemic akrasia leads to in the kinds of case we were considering earlier. We can sum up the difference between Williamson's case and the kinds of case we were considering earlier by saying that Williamson's case does not involve higher-order doubt about the kind of perceptual state you are in but rather about what the content of that state is, whereas cases like Suspected Hallucination involve higher-order doubt about what kind of perceptual state you are in.

Let's take stock. We are trying to respond to White's argument against the infallibilist conception of perceptual evidence or reasons. The

strategy we are pursuing is to deny that the kind of case White describes actually involves defeat. This would absolve the infallibilist from having to explain how defeat is possible. Since, from an infallibilist perspective, the putative defeater in Suspected Hallucination seems to be higher-order, an obvious suggestion for the infallibilist to pursue is that the case can be given a Level-Splitting treatment, according to which the subject should become epistemically akratic rather than losing her first-order perceptual knowledge. Unfortunately, however, this doesn't work in this kind of case because the levels cannot be isolated from each other, resulting in a recommendation that the subject should have outright contradictory beliefs. This doesn't mean that Level-Splitting can never be properly applied to cases in which one has higher-order evidence that one's perceptual belief is likely to be false. If the higher-order evidence in question is not direct or indirect evidence that one is hallucinating or suffering from an illusion, then Level-Splitting may well be appropriate. However, we will have to try a different approach to show that White's Suspected Hallucination case does not involve defeat.

4.4. Steadfastness and an error theory

In the previous section I mentioned that an infallibilist might be tempted to adopt the Steadfast strategy. This option is familiar from debates about peer disagreement, where those who advocate Steadfastness say

that when a knowledgeable or justified subject is confronted with disagreement from someone they take to be their epistemic peer, they should not give any ground but instead stick steadfastly to their knowledgeable or justified belief, which continues to be knowledgeable or justified.[145] Similarly, an infallibilist might simply say that in Suspected Hallucination you should continue to believe H, and that if you do this then you will retain your knowledge.[146]

If the Steadfast approach is not to simply amount to Level-Splitting, then the steadfast subject needs to be able to utilise her knowledge to, as it were, defeat the putative defeater. For even if she continues to know H, she will still have evidence that she is hallucinating. If she is to avoid being rationally required to believe or have high confidence that she's hallucinating, she needs to rebut or otherwise undermine this evidence. Fortunately, the discussion of Level-Splitting suggests that such a rebuttal will be possible. For although the truth of H is compatible with it being the case that one is hallucinating, one should also have perceptual knowledge that *That is a hand*, for example, which directly rebuts propositions which are entailed by the proposition that one is hallucinating, such as *That is not a hand*. One can therefore rely on one's

[145] See, for example, Kelly 2005.

[146] For this approach, see Lasonen-Aarnio 2010; Baker-Hytch and Benton 2015; Lasonen-Aarnio 2021. Cf. Lasonen-Aarnio 2014b. Lasonen-Aarnio holds that Level-Splitting is often the right policy in cases of higher-order putative defeat, whereas in cases of rebutting putative defeat she seems to suggest that one can get away with simply ignoring the putative defeater and thereby hanging on to one's knowledge. I am about to suggest that this leads to a view on which one can do more than ignore the putative defeater: one can rebut it with one's knowledge.

knowledge that *That is a hand* to infer that one is not hallucinating, thereby rebutting the putative defeater. Alternatively, as was suggested in the previous section, if one's perceptual experience puts one in a position to know that one sees that H, then one could rely on this to rebut the putative defeater.

The idea that one might legitimately respond to evidence (misleading or not) that one is suffering from a life-like hallucination by relying on one's perceptual experience to rebut this evidence is often written off as obviously and pre-theoretically absurd.[147] Indeed, it does seem that if the infallibilist is to make this approach plausible, they need to provide an error theory for our undoubted tendency to judge that in cases like Suspected Hallucination the rational thing to do is to cease to believe your eyes.

Lasonen-Aarnio has provided just this. She utilises this error theory in defence of both Level-Splitting and Steadfast approaches to putative defeat: the error theory aims to explain why, even in cases in which she thinks level-splitting or steadfastness is called for, we tend to judge that the subject would be irrational to level-split or to be steadfast.

Lasonen-Aarnio distinguishes between *rational*, *justified* and *knowledgeable* belief, on the one hand, and believing in a *reasonable* way on the other hand. Being epistemically reasonable involves managing your beliefs 'in a

[147] Besides White 2014, see Cohen 2002, for example.

way that makes sense given the goal of knowledge acquisition' and avoidance of false belief.[148] She claims that in cases of putative defeat of the kind we've been considering, one can continue to know p, and to be justified in believing p, but to continue to believing p is nevertheless unreasonable, given one's defeating evidence. Thus if one were to remain steadfast in this way then one would have 'unreasonable knowledge'.[149] One's knowledge would be unreasonable because it would manifest a disposition to ignore evidence that one has erred. This disposition will, more often than not, lead one astray, away from knowledge and truth. For a subject who is disposed to stubbornly continue to believe H even when she receives *misleading* evidence that she is hallucinating is also disposed to stubbornly retain her perceptual belief even when she receives *non-misleading* evidence that she is hallucinating. And evidence that one has erred is usually non-misleading. Therefore, given the general epistemic goal of knowledge or truth, this disposition is in a bad one to have: it leads one to manage one's beliefs in ways that don't make sense, given the goals of maximising knowledge and

[148] Lasonen-Aarnio 2010: 12.

[149] Ibid.

minimising ignorant or false belief.[150] It is much better to be disposed to revise belief in the face of evidence of ignorance.

Of course, it would be ideal to be disposed to revise belief when evidence of error is not misleading, and to be disposed to retain belief when evidence of error is misleading. But that is not 'humanly feasible', as Lasonen-Aarnio puts it.[151] Given our limited discriminatory capacities, we simply can't have dispositions that are sensitive in this way to whether or not evidence of error is misleading. Indeed, it is part of the description of cases in which we have made an error in coming to believe p that we don't notice that we have made an error. After all, if we noticed then we would stop believing p. So we need to manage our beliefs in a way that takes into account our fallibility: we make errors, and we need to take precautions to try to minimise our errors. This precautionary approach involves being disposed to revise beliefs when we get evidence that we have erred, and unfortunately this disposition can't be sensitive to whether or not this evidence is misleading. And of course, in practice we do tend to follow this approach to belief-revision. A

[150] I am assuming here that the fundamental norm of belief is knowledge, and that putatively defeating evidence is always evidence of ignorance. There is a large literature on the norm of belief, and it is becoming increasingly popular to endorse the knowledge norm. See Dutant and Littlejohn 2021 for an argument that defeaters are always evidence of ignorance. But Lasonen-Aarnio's approach is also compatible with holding that the fundamental norm of belief is truth, and with various different ways of characterising putatively defeating evidence. See in particular Lasonen-Aaarnio 2021: 109-10 on this compatibility.

[151] Lasonen-Aarnio 2021: 95.

subject in White's Suspected Hallucination case in practice is highly likely to give up her perceptual belief.

Nonetheless, says Lasonen-Aarnio, a stubborn subject can retain knowledge in such a case. For knowledge, '[l]ike most other things, […] can sometimes be achieved through bad strategies'.[152] And what she has said about the unreasonableness of knowledge in such cases suggests an error theory for our tendency to judge that a steadfast or stubborn subject lacks knowledge. Our judgement that knowledge is defeated in these cases is a product of our sensitivity to the bad disposition the subject is manifesting, or to the subject's unreasonableness. As Lasonen-Aarnio says, we mistakenly take the manifestation of epistemically bad dispositions to be incompatible with knowledge because in paradigm cases of knowledge manifesting good dispositions is part of what explains how one comes to know, and so checking whether good dispositions are manifested is 'a useful heuristic for evaluating whether [a belief] counts as knowledge.'[153] Using this heuristic in cases of putative defeat leads us to err, because such cases are not normal, in that they are cases in which the subject can know while manifesting an epistemically bad disposition. Lasonen-Aarnio thus provides an error theory for judgements about defeat, while at the same time doing justice to the judgement that there is a sense in which one ought to revise one's belief (whether or not it is

[152] Lasonen-Aarnio 2010: 17.

[153] Ibid., 16.

knowledgeable) when one receives putatively defeating evidence (misleading or not).

If Lasonen-Aarnio is right, or roughly right, then the infallibilist is in a much better position to be able to claim that, strictly speaking, there is no such thing as the defeat of knowledge. To be sure, knowledge can be lost by forgetting, by ceasing to believe, and so on. But putative defeat cases are cases where the subject is able to retain her belief, on the same basis as before, and yet even if the subject does this her belief will lose its positive epistemic status. This is what Lasonen-Aarnio provides a path to denying the reality of. This generality – the applicability of Lasonen-Aarnio's error-theory to all kinds of defeater – is helpful to the infallibilist. For while White's Suspected Hallucination case seems to involve a higher-order or undercutting defeater, similar cases can be construed with other kinds of defeater, such as rebutting or opposing defeaters. I find Lasonen-Aarnio's account persuasive. I suggest, then, that the infallibilist should respond to White's case by advocating Steadfastness, and invoking Lasonen-Aarnio's error theory to explain away the intuition that the subject's knowledge must be defeated in such a case. I turn now to defending this approach from a recent argument, due to Jessica Brown, according to which the reality of knowledge-defeat is not so easy to deny.

4.5. Brown on 'contributory' and 'revisionary' defeat

Brown argues that Lasonen-Aarnio's suggestion that knowledge-defeat is an illusion cannot be right. Brown's strategy involves distinguishing between 'contributory' and 'revisionary' notions of defeat. Contributory defeat takes place when you receive evidence which does not defeat a 'pre-existing' epistemic status of knowledge or doxastic justification, but rather pre-emptively prevents knowledge or doxastic justification from arising.[154] This happens when you receive the defeater before or at the same time as receiving the evidence which would, in the absence of the defeater, have been a basis for knowledge or doxastic justification. Revisionary defeat takes place when you receive evidence which defeats an epistemic status – propositional justification, doxastic justification or knowledge – which was already in place. Cases of revisionary defeat are what philosophers discussing defeat usually focus on; Brown argues that we can gain insight from considering contributory defeat. In particular, she argues that it is simply obvious that contributory defeat is an actual phenomenon. Then, 'given the plausible epistemic symmetry between cases of contributory and revisionary defeat'[155], if contributory defeat is actual then so is revisionary defeat. Thus revisionary defeat is actual, and Lasonen-Aarnio must be wrong.

[154] Brown 2018: 106–7.

[155] Ibid., 107.

Brown argues that it's obvious that contributory defeat takes place by asking us to consider 'Case 1':

> **Case 1**: You receive testimony from B that not-p and at the very same time you receive testimony from A that p. (Don't worry about whether this simultaneous receiving of testimony is actually possible!) A's testimony is reliable: A knows p and transmits the fact that p to you. B's testimony is faulty. However, as far as you know both testifiers are equally reliable, and you have no further evidence for or against p.[156]

Brown says that it's obvious that although you would have been able to come to know p had you only had A's testimony, the presence of B's testimony pre-empts this possibility of knowledge, defeating it. Similarly, you would have come to have propositional justification to believe p, had it not been for the defeater. To deny this, says Brown, would be 'to deny that justification is affected by one's total evidence, rather than just a part of that evidence', and that would be 'deeply problematic'.[157] So, she concludes, in Case 1 it's undeniable, on pain of absurdity, that contributory defeat takes place, so contributory defeat is a real phenomenon.

---

[156] I am paraphrasing ibid., 108.

[157] Ibid.

'Case 2' is like Case 1, except that now A's testimony comes first and is followed by B's testimony. You do not form any belief until after you have heard both pieces of testimony. Thus, to begin with, when you only have A's testimony and have not yet heard B, you have propositional justification for believing p that can be a basis for knowledge of p. But surely, says Brown, you no longer have all-things-considered propositional justification to believe p once you receive B's testimony: your prior propositional justification is thus defeated, and we have here a case of revisionary defeat, since you have lost the status of possessing all-things-considered propositional justification. After all, the ultimate evidential situation in Case 2 is just the same as in Case 1, and 'it's hard to see how the mere fact that the testimony from B arrives just a little bit later in Case 2 should affect whether the hearer ends up with [all-things-considered] justification to believe that p.'[158] So, if we accept that contributory defeat takes place in Case 1, we should accept that revisionary defeat takes place in Case 2.

'Case 3' is just like Case 2 except that in Case 3 you actually form the belief that p on the basis of A's testimony (before receiving B's testimony) and so come to know p. But again, says Brown, if one lacks all-things-considered propositional justification to believe p in Case 2 after receiving B's testimony, then surely one does in Case 3, too, since the evidential situation is the same in both cases. But one can't have knowledge if one lacks propositional justification; so one must cease to

[158] Ibid., 109.

know p when one receives B's testimony in Case 3. Thus one's knowledge is defeated, and we have here another case of revisionary defeat, this time of knowledge rather than merely of propositional justification. Brown's central point is that in cases 1–3 one ultimately ends up with the same evidence, and so if one's potential for knowing p is defeated in Case 1 then one's knowledge of p must be defeated in Case 3. As she puts it, 'once we accept the contributory notion of defeat, it is hard to deny the revisionary given the plausible epistemic symmetry between cases of contributory and revisionary defeat.'[159]

Brown notes that someone who identifies one's evidence with one's knowledge (E=K) will deny that one's evidential situation is the same in Case 3 as in cases 1 and 2, since in Case 3 one comes to know p whereas in cases 1 and 2 one does not. But she seems to treat this fact as a reductio of the E=K equation:

> [T]he current suggestion [that E=K] involves saying that the mere fact that the hearer forms the belief that p in Case 3 but not in Case 2 results in the hearer having better evidence for p in Case 3 than in Case 2 […]. But it seems grossly implausible to suppose that the mere fact that the hearer forms the belief that p in Case 3 but not Case 2 results in her having much better evidence for her belief that p.[160]

[159] Ibid., 107.

[160] Ibid., 110–11.

However, even if we accept that Brown is correct about this, thoroughgoing externalists about reasons and evidence, who hold that evidence and reasons are constituted by world-involving mental states, will object to an earlier stage of her argument. As we saw, she says of Case 1 that to deny that contributory defeat takes place here would be to deny, absurdly, that 'justification is affected by one's total evidence, rather than just a part of that evidence'.[161] But it's not clear why the thoroughgoing externalist should accept this. Brown seems to assume that the only way that we could deny that contributory defeat takes place is to deny that B's testimony has any evidential weight. But the thoroughgoing externalist needn't deny this: rather, they will simply deny that B's testimony has the same evidential weight as A's testimony. On the thoroughgoing externalist's approach, A's testimony makes manifest the fact that p, but this doesn't meant that B's testimony does not have any evidential force at all. Instead, just as Williamson says that the evidential force of a life-like hallucination as of p's being the case is to make manifest the fact that it appears to one as if p, so we can say that the evidential force of B's testimony is to make manifest that B reports that not-p.[162] This is good evidence against p; but given that one also has A's evidence, which makes manifest that p, one's *total* evidence supports believing p. So the thoroughgoing externalist can easily conclude that

[161] Ibid., 108.

[162] Williamson 2000: 198.

contributory defeat does not take place in Case 1, without denying that justification is affected by one's total evidence.

Brown's argument ignores this possibility. But once this possibility is recognised, it becomes much less obvious that contributory defeat takes place in Case 1, since we have a way of denying that contributory defeat takes place without denying that justification is affected by one's total evidence. But the claimed obviousness of contributory defeat in Case 1 provided Brown's grounds for arguing for revisionary defeat in cases 2 and 3. Consequently, Brown's argument for revisionary defeat fails at the first hurdle. Its only dialectical force is to make vivid the common intuition that defeat does take place by describing a putative kind of defeat – contributory defeat – which is not much discussed. But Lasonen-Aarnio has acknowledged that we do tend to judge that defeat takes place, and has endeavoured to explain why we make this error. Brown needs to show why this explanation of error fails before returning to appeals to cases about which we tend to judge that defeat takes place. Brown might try to show that Lasonen-Aarnio's error theory is inapplicable to our judgements about *contributory* defeat. But, on reflection, the error theory seems to apply nicely to these judgements, too. We judge that the subject's knowledge and justification is pre-empted in Case 1 because we are sensitive to the fact that someone with good epistemic dispositions will abstain from believing p, given the presence of evidence for not-p. Nevertheless, were a subject to manifest a bad disposition by ignoring B's evidence and believing A's evidence, she

can thereby come to have unreasonable knowledge of p, since her belief is based on good evidence for p.

4.6. Conclusion

In this chapter I have considered a challenge to the kind of infallibilist picture of evidence and reasons that I advocated in Chapter 3. The challenge is that infallibilism seems to struggle to accommodate defeat of knowledge, since defeaters apparently leave one's basis for belief intact while destroying one's knowledge, suggesting that one's basis is not knowledge-guaranteeing in the way that the infallibilist thinks. I suggested that the infallibilist should stand their ground, and deny the reality of knowledge-defeat. I first considered one popular strategy for denying defeat, Level-Splitting. This strategy tries to accommodate our judgement that subjects should react to putatively defeating evidence in some way by claiming that subjects should react by adopting conflicting beliefs at different levels. While level-splitting may sometimes be appropriate, I argued that it cannot be appropriate in cases like that described by White. The problem is that in cases where the putative defeater is evidence that one is hallucinating or suffering from an illusion, the higher-order belief that level-splitting recommends the subject adopt cannot be kept rationally isolated from the subject's first-order perceptual beliefs. I therefore suggested that a better infallibilist strategy to deal with standard cases of apparent defeat of perceptual knowledge is

Steadfastness. According to this approach, the subject can retain their knowledge by stubbornly sustaining their knowledgeable belief when they receive the putatively defeating evidence, and using this knowledgeable belief to rebut that evidence. Of course, it's highly counter-intuitive to think that this is the epistemically appropriate thing for the subject to do, and it's also quite counter-intuitive to think that the subject could retain their knowledge by adopting this policy. However, Lasonen-Aarnio has provided an error theory for our intuitions about such cases. This error theory vindicates our strong sense that there is something epistemically inappropriate about stubbornness in such cases, but it suggests that we are here being sensitive to the unreasonableness of the subject's belief-management policies, rather than to whether the subject retains her knowledge or not. In my view, Lasonen-Aarnio's error theory is compelling. For, at least for my part, alongside my intuition that there is something extremely epistemically inappropriate about stubbornly ignoring misleading evidence of error, I also have an intuition that a subject who does this remains connected to truth in a non-accidental way, via her appreciation of the rational significance of her original basis for belief.[163] Lasonen-Aarnio's error theory reconciles these apparently competing intuitions. There is much more to be said about all of these matters, but for now I conclude that the infallibilist should not be overly worried about apparent cases of knowledge defeat.

[163] Compare the intuition that I suggest we have about the views of Dretske in Chapter 5 (section 5.4).

Chapter 5

Easy knowledge

On the view that I have been developing in this thesis, perceptual knowledge is non-inferential but is rationally based on conclusive perceptual evidence, constituted by one's factive or relational perceptual experience or its worldly constituents. This weds me to Liberalism, according to which the rational support that perceptual experience offers perceptual beliefs is immediate; and it commits me to opposing Conservatism, according to which perceptual experiences offer only non-immediate support for perceptual beliefs. Here is Silins, describing an analogy that is often used to make vivid the two very different ways of thinking about perceptual knowledge encoded in Liberalism and Conservatism:

> According to the *Conservative*, […] [perceptual] experience is like a gas gauge: experience provides justification only in combination with information that it is not malfunctioning. […] According to the *Liberal*, the Conservative is entirely mistaken. The Liberal allows that experiences sometimes give us reasons for beliefs about the external world in a way that doesn't happen in virtue of our having independent reason to believe any claim. In other words, our experiences sometimes give us *immediate* justification for external world beliefs […]. At stake in the dispute is whether

> experiences provide us with a special kind of justification, different from the kind we get from gas gauges, thermometers, and the like.[164]

The Liberal-Conservative debate is different from but related to the debate about whether perceptual experience is a world-involving kind of mental state or a non-factive, non-relational mental state. In Chapter 3 I critically discussed the popular position of non-factivist Liberalism. In the present chapter I want to examine an apparent problem for Liberalism – in both its factive/relational and non-factive/non-relational varieties. For it is often suggested that Liberalism should be rejected since it is committed to holding that we can come to have knowledge in ways that seem to be too 'easy'. This is the problem of easy knowledge for Liberalism.[165]

The Liberal holds that one can form justified and knowledgeable perceptual beliefs without relying on antecedent justification to believe anything, and in particular without relying on antecedent justification to believe the negations of sceptical hypotheses or to believe that perception is a reliable guide to the truth. I'll talk in terms of perceptual evidence, rather than processes or methods, but the discussion could be reformulated in non-evidentialist terms. Now consider the following closure principle:

---

[164] Silins 2014: 505–506.

[165] Much of the recent discussion of this issue stems from Cohen 2002, which is where the label 'easy knowledge' comes from.

Deductive Knowledge Closure: If one knows P and competently deduces Q from P, thereby coming to believe Q, while retaining one's knowledge of P, one comes to know Q.[166]

Liberalism combines with Deductive Knowledge Closure to imply that perceptual evidence can provide us with knowledge that it has itself not misled us in various ways. For example, the Liberal holds that I can come to know that *that table is red* purely on the basis of my perceptual experience of the red table. By Deductive Knowledge Closure, I can infer from my *a priori* knowledge that *if that table is red then it's not white lit by tricky red lights* to knowledge that *that table is not white lit by tricky red lights*. That is, I can come to know that I am not suffering from an illusion caused by tricky lighting by making the following 'easy-knowledge argument' (henceforth EK-argument):

[166] Hawthorne 2013: 43; Williamson 2000: 116. Note that Cohen (2002) uses a closure principle with a slightly different formulation. The principle Cohen uses is weaker than this one, and some philosophers have argued that a Liberal can accept this weaker version of closure while holding that EK-arguments do not produce knowledge. E.g. Silins 2007. I think that Silins is wrong about this, and that Cohen's weaker principle still leads to the easy-knowledge problem. All the same, I will use Hawthorne and Williamson's principle for the moment, since I think this principle is correct and it can be claimed that it leads to the easy-knowledge problem without having to consider Silins' argument.

RED TABLE ARGUMENT

(EK1) The table looks red[167]

(EK2) The table is red

(EK3) If the table is red then it is not white with red lights shining on it

(EK4) The table is not white with red lights shining on it

Cohen (2002) says that you cannot come to know (EK4) like this: it is too 'easy'.[168] According to him, EK-arguments must be illegitimate and must not provide for genuine knowledge of their conclusions.

In EK-arguments one uses a piece of evidence or a method to come to know p, and then one infers q from p. The proposition q includes the information that a sceptical scenario does not obtain, where the sceptical scenario is such that, were it to obtain, p would be false but you would still believe p (and, for that matter, q, assuming that you make the

[167] I follow Cohen in presenting (1) in this way, but of course different theories of perceptual evidence will plug different things into (1). On some theories, such as Pryor's dogmatism, (1) is not a premise because it is not a proposition, but is rather a visual experience. And, of course, on my preferred view, (1) should be something that guarantees the truth of (2). The key point is that in an EK-argument the justification of (2) does not involve any justification to believe that one's perception is a reliable indicator of the truth, nor that it is accurate on this occasion, nor any justification to believe that the kinds of sceptical scenarios that figure in the conclusions of EK-arguments do not obtain. This is the Liberal's claim.
The problem also arises for non-evidentialist views, such as reliabilism. If we plug into (1) just the reliable perceptual process leading to belief in (2), then we get the same kind of EK-argument.

[168] For related a treatment of similar problems, see e.g. Vogel 2000. I focus on Cohen's presentation because it has been the focus of and inspiration for much recent discussion and it relates the problem in a perspicuous way to the kinds of view of perceptual evidence that I am concerned with throughout the dissertation.

inference to q in the actual case). The sceptical scenario may be but need not be radical, in the sense that it involves *all* of your beliefs being false. In the example just given the sceptical scenario just involves the table being lit by tricky lights.

Moorean arguments against sceptical hypothesis are often EK-arguments. For example:

NOT-BIV:

(1_{BIV}) [Visual experience of a hand].

(2_{BIV}) I have a hand.

(3_{BIV}) If I have a hand, then I am not a handless brain in a vat.

(4_{BIV}) Thus, I am not a handless brain in a vat.

Moore's original argument against the sceptic concluded with *There is an external world*, a positive claim that is incompatible with but is not the negation of a sceptical hypothesis, and so does not instantiate the EK-argument pattern.[169] However, Moore's original argument inspires a similar sense of being 'too easy', and our discussion of EK-arguments may be instructive when it comes to considering its merits.

---

[169] Wright's various discussions of Moorean arguments focus on Moore's original argument. E.g. Wright 2002, 2004a, 2008. But Silins 2007, for example, just talks about NOT-BIV. Pryor 2004 talks about both Moore's original argument and a variant of RED TABLE ARGUMENT; Pryor 2012 focuses on NOT-BIV. There is also a higher-order variation which, for example, Pritchard 2012 (Ch. 3) labels 'Mooreanism'. Wright also sometimes moves from discussing first-order Mooreanism to discussing higher-order Mooreansim. See Pryor's discussion of this at Pryor 2012: 293–4.

My main concern here, though, is not with vindicating Moore's response to the sceptic, but with defending Liberalism from the charge that it licenses EK-arguments when it should not. My defence will be that Liberalism is right to license EK-arguments, since they are legitimate routes to knowledge after all. To make this plausible, I need to supply an explanation of our strong sense that there is something wrong with EK-arguments; doing this is my main aim in this chapter.

First, we need to get clearer on why easy knowledge is a problem for Liberalism in particular, since it might seem that non-Liberals are just as committed to endorsing EK-arguments as Liberals. For suppose that the Conservative is right, so that perception does not justify beliefs about the world all by itself, since one's non-perceptual justification to believe that perception is reliable essentially contributes to perceptual justification of beliefs about the external world. Now suppose that a subject comes to perceptually know that the table is red. By Deductive Knowledge Closure, she can infer that the table is not white lit by red lights, and so come to know that this is so. Why isn't this too easy? The idea is that this doesn't strike us as being too easy, since in making the inference the subject is relying in part on her justification to believe that perception is reliable. It is supposed to seem quite reasonable to infer from the fact that you're having a perceptual experience as of a red table *and* the fact

that perception is reliable that the table is not white lit by red lights.[170] Easy knowledge is a problem for Liberal views because Liberalism implies that perception *all by itself* can generate justification to believe EK-conclusions.

There are two ways one might react to EK-arguments. The first is to agree with Cohen that we cannot come to know EK-conclusions, such as (EK4), via EK-arguments. One then needs to explain why this is, given that EK-arguments seem to be valid, and given that we generally believe their premisses. In any case, as the problem is presented in Cohen (2002), if one reacts to EK-arguments in this way then one cannot be a Liberal who endorses Deductive Knowledge Closure. Modifying Deductive Knowledge Closure, or going as far as giving up knowledge-closure principles altogether, is unattractive, for reasons brought out in Hawthorne (2013). As for giving up on Liberalism, the obvious difficulty is that, in order to avoid scepticism, one needs to explain how we can have justification to believe EK-conclusions antecedently of having perceptual knowledge. Suppose one says that there is an *a priori* source for this justification: label this source 'S'. Are we to be Liberals about S (in which case, why can't we be Liberals about perception as a source of justification?), or are we to hold that you need to have antecedent justification to believe that S is not misleading (in which case, aren't we on the way to a vicious regress?). This just scratches the surface of the

[170] Markie 2005 suggests that the easy-knowledge problem afflicts non-Liberal views as well as Liberal views. In this paragraph I have paraphrased Cohen 2005's response to Markie. For related discussion see White 2006: Section 8. 'Is Bootstrapping Everyone's Problem?'

difficulties. I won't here rehearse Hawthorne's arguments for endorsing Deductive Knowledge Closure, and I won't go into further detail about the difficulties facing non-Liberal accounts of perceptual justification. Instead, I will assume that it is preferable to avoid both of these options, if we can. My purpose is to try to show that, happily, we can avoid them, since there is a viable second way to react to EK-arguments.

The second way of reacting to EK-arguments is to deny that EK-arguments cannot be ways of coming to know EK-conclusions. Then we can continue to be Liberals who endorse Deductive Knowledge Closure. However, this way of responding must say something about why EK-arguments have the appearance of being epistemically illegitimate. Philosophers have tended to judge that EK-arguments are illegitimate: if we are going to insist that EK-arguments are ways of coming to know their conclusions, we need to pinpoint something about them that leads us to be uneasy about them. As I say, that task is the focus of this chapter.

Here is a plan of the chapter. In section 5.1 I consider more closely the arguments of those philosophers who deny that EK-arguments are routes to knowledge of their conclusions. We'll see there that such philosophers rely mostly on intuitive judgements about cases, rather than providing a principled diagnosis of what is wrong with EK-arguments. This leaves the way open for those of us who are committed to EK-arguments being routes to knowledge to try to explain away these intuitive judgements about cases. The sections which follow explore ways

in which we might do this. The focus of attention will be the idea that EK-arguments are in some way dialectically ineffective. In sections 5.2 and 5.3 I consider both Markie and Pryor's attempts to flesh out this idea, and find that Pryor's approach is the more promising because it purports to identify something distinctive about EK-arguments, namely that someone who doubts the conclusion of an EK-argument cannot use the argument to rationally dislodge their doubt in its conclusion. However, it turns out that Pryor's approach is incompatible with denying the reality of knowledge-defeat. In section 5.3 I therefore propose a variation of Pryor's approach, which extends Lasonen-Aarnio's notion of unreasonable-but-knowledgeable belief to articulate a sense in which it is epistemically inappropriate (because unreasonable) for someone who doubts the conclusion of an EK-argument to believe the argument's premises. After defending this approach (which I call the dispositional approach to dialectical ineffectiveness) from an objection in section 5.4, in section 5.5 I provide a further reason to prefer it over Pryor's approach. This turns out to have wider implications, providing a further reason to deny the reality of defeat and to endorse thoroughgoing externalism about evidence and reasons. In a nutshell, then, I argue for the dispositional approach to dialectical effectiveness as a way of explaining what is bad about EK-arguments while maintaining that they are genuine routes to knowledge, and I argue that this gives us a further reason to endorse the anti-defeat claim of the previous chapter and to endorse the thoroughgoing externalism that this thesis has been defending.

5.1. Against 'easy knowledge'

When philosophers claim that EK-arguments don't produce knowledge of EK-conclusions, they often try to elicit this judgement by telling a story involving an EK-argument. Cohen (2002) is the focus of much recent discussion of 'easy knowledge', and he tells the following story:

Red Table Conversation: Suppose my son wants to buy a red table for his room. We go in the store and I say, 'That table is red. I'll buy it for you.' Having inherited his father's obsessive personality, he worries, 'Daddy, what if it's white with red lights shining on it?' I reply, 'Don't worry – you see, it looks red, so it is red, so it's not white but illuminated by red lights.' Surely he should not be satisfied with this response.[171]

Cohen offers two reasons for thinking that there's something wrong with the father's argument. The first is just the judgement that were we in the son's position we would not be satisfied with it. The second is that, according to Cohen, if this kind of argument is legitimate then 'my inductive evidence against the possibility that there are red lights shining on the table turns out to be irrelevant to my knowing the table is not white with red lights shining on it', and this is 'strange'.[172]

---

[171] Cohen 2002: 314.

[172] Ibid., 313.

Cohen tells another story to bring out this strangeness:

How Do You Know?: Suppose my son is not worried about the possibility that the table is deceptively illuminated. He accepts that I know that the table is red and not white with red lights shining on it. He is just curious about *how* know it. I respond in just the same way. 'Oh, that's easy. It looks red, so it is red, so it is not white with red lights shining on it.'[173]

Again, surely the son will be unsatisfied by this response, says Cohen. Surely, in standard cases (where, for example, one hasn't asked the owner of the store, or made other special checks) one will rely on background evidence that tables in storerooms are not usually disguised by tricky lighting to know, via induction, that this table is not white lit by red lights.

It's not that clear what Cohen's point is here. Liberalism doesn't imply that the father lacks such inductive evidence, nor that he wouldn't rely on it in coming to know that the table is not white lit by red lights. Liberalism is consistent with one having available more than one source of evidence for such a proposition – perceptual, inductive, and potentially other sources too.

Perhaps the point is this. The son would rightly expect the father to refer to his inductive evidence, rather than to his perceptual evidence, in

[173] Cohen 2005: 420.

explaining how he knows that the table is not white with red lights shining on it. In fact, it would also be natural for the father to respond to the son by making special checks, such as looking up to check the lighting in the store, or by viewing the table outside under natural light before going ahead with the purchase. So either induction or special checks would be natural sources for the father to appeal to, whereas his original perceptual evidence would not be natural to appeal to. Noticing this makes more vivid Cohen's first point, which was that we instinctively judge the father's response to the son to be inadequate. This is made more vivid by noticing that there are alternative responses available to the father that we would judge to be adequate.

Cohen's stories also offer materials for generating an explanatory challenge to those of us who think EK-arguments are legitimate ways to know EK-conclusions. If they are legitimate ways of knowing EK-conclusions, then why do we tend to make use of inductive evidence and special checks to justify asserting EK-conclusions? After all, it might be thought, EK-arguments offer the prospect of *conclusive* evidence[174], whereas induction does not. And EK-arguments generally cost less in terms of time and energy than does making special checks. So why do we prefer induction and special checks to EK-arguments?

---

[174] I assume here that (EK2) is *known* by the subject making the EK-argument. Of course there are those who take (EK1) to be something that does not guarantee (EK2)'s truth. Still, if we accept that (EK2) is known, as any non-sceptical philosopher should, then this, together with (EK3), provides conclusive support for (EK4) if EK-arguments produce knowledge in EK-conclusions. In any case, on the view I defend in this dissertation, (EK1) is conclusive evidence for (EK2).

In summary, we can take the following points from Cohen's examples. First, we are left unsatisfied by the father's appeal to his original perceptual evidence in response to his son. This suggests that his original perceptual evidence is not up to the task of producing knowledge that the table is not white lit by red lights. Second, I suggested that the fact that in practice we use induction and special checks to come to know EK-conclusions calls out for explanation, since EK-arguments would offer conclusive evidence where induction would not, and EK-arguments generally cost less in terms of time and energy than making special checks. One possible explanation is that EK-arguments are not ways of coming to know EK-conclusions, whereas induction and special checks are. So far, then, in support of the claim that EK-arguments do not produce knowledge of EK-conclusions, we have intuitive judgements about cases and an explanatory challenge posed by our actual practices when it comes to justifying our beliefs in EK-conclusions.

However, if we assume for a moment that Cohen is right, and that EK-arguments are indeed not routes to knowing EK-conclusions, then we should want a diagnosis of what exactly is wrong with EK-arguments. For there remains a puzzle: viewed from one angle, EK-arguments look perfectly respectable. Their premisses seem to be unproblematically true, and the transition from (EK2) and (EK3) to (EK4) is deductively valid. So, to avoid ongoing puzzlement, we need to know why they don't produce knowledge. Furthermore, we might expect that the truth about

what it is that prevents EK-arguments from being routes to knowing EK-conclusions will guide us when it comes to choosing whether to give up Liberalism or Deductive Knowledge Closure, and what views to put in their place. If we were to accept Cohen's view, then, we would want to know what it is about EK-arguments that our negative judgements about them are sensitive to.

We might look for an answer to this question by considering the other problem of 'easy knowledge' that Cohen describes, the problem of 'bootstrapping'. Suppose one comes to know that the table is red on the basis of perceptual evidence or a reliable belief-forming method. One then knows that the table is red; further, one can introspectively learn that this very method or piece of evidence produced a belief that the table is red. One can then infer that this method or piece of evidence produced a true belief (since one knows that the table *is* red). If one repeats this process on many occasions of, for example, visual knowledge, then it seems one can thereby gain inductive support for believing that vision is a reliable source of evidence or a reliable method for coming to have true beliefs about the world. Cohen thinks this is too easy: one can't come to know that one's belief-forming method is reliable in this way.[175]

[175] Cf Vogel 2000; White 2006: 543ff. Something that Cohen does not refer to but which deserves to be noticed is that vulnerability to the bootstrapping version of the problem of easy knowledge also depends on one endorsing a closure principle, since otherwise one could simply deny that the inference (from *the table is red* and *method m produced in me a belief that the table is red* to *method m produced a true belief on this occasion*) produces knowledge. Black 2008 emphasises this point. One might also seek to block bootstrapping without denying closure or Liberalism by giving an account of inductive knowledge which provides a principled reason to think that one cannot

It is often claimed that bootstrapping is epistemically illicit because it involves using a method or piece of evidence to verify its own trustworthiness or reliability. Thus Bergman, for example, interprets critics of bootstrapping as appealing to what he calls the No-Self-Support-Principle:

NSS: One cannot obtain a justified or warranted belief that a belief source S is trustworthy by relying even in part on source S.[176]

Indeed, similar principles are sometimes stated as truisms, especially in the literature on easy knowledge. For example, according to one recent author, writing about EK-arguments rather than bootstrapping, 'Surely one cannot improve the epistemic status of a source merely by using it!'[177] Indeed, EK-arguments themselves seem to violate NSS; so if NSS is a true principle then it can explain what is wrong with both EK-arguments and bootstrapping.

Perhaps there are similar principles that are more promising, but they would need to be carefully worked out to avoid obvious counterexamples. For example, when a scientist uses visual perception to study the visual system of human beings, and thereby learns how it works, she can surely in this way gain evidence that human vision in

[176] Bergmann 2000: 168. Cited by Cohen 2002: 319.

[177] Bradley 2014: 387.

general, and so her own vision in particular, is reliable. But this seems to violate NSS. All the worse for NSS. White (2006) gives a similar example: surely I can learn that my vision is reliable by using my vision to read the optometrist's glowing report on my eye examination.[178] Here is a different kind of example. You look around you, and the coherence and vividness of your visual experience gives you reason (perhaps all by itself) to think you're not dreaming or hallucinating, and so that you're genuinely perceiving the world and so coming to know about it.[179]

White concludes that complaints about using a method or piece of evidence to gain justification for thinking that that method or piece of evidence is reliable need to be handled with care. Still, he thinks that there *is* something wrong with Bootstrapping and EK-arguments, such that they don't produce knowledge, and that it does have something to do with using evidence to verify its own credentials. But he offers us no improved formulations of any principle that might be being violated. Instead, he continues to rely on intuitive judgements about cases to distinguish problematic from unproblematic self-verification. Similarly, Cohen admits that 'we do not really have a precise definition of bootstrapping' – we have to make do with judgements about cases.[180] So considerations about self-supporting evidence have not taken us much beyond the intuitive judgements we began with. But they do alert to us

[178] White 2006: 530.

[179] For further critical discussion of NSS-like principles, see Markie 2005: 414–15.

[180] Cohen 2002: 321. Cohen describes some cases which might or might not count as problematic bootstrapping, saying 'I do not have a clear intuition here'.

the fact that EK-arguments (and bootstrapping) involve a piece of evidence or a method generating justification to believe that that very evidence or method is not misleading, and that our judgement about their epistemic illicitness may have something to do with this feature.[181]

Another suggestion that is sometimes made, at least with respect to EK-arguments, is that the arguments in question are problematic because they violate rules governing probabilistic reasoning. This is, for example, White (2006)'s diagnosis of EK-arguments. It's true that Liberalism, if combined with the view that perceptual evidence is non-conclusive, might well have to answer the charge that it licenses arguments that violate such rules, since some EK-arguments can seem to have this feature. For example, the probabilistic critique of arguments like RED TABLE ARGUMENT is that since if red lights were shining on a white table the table would look red, $Pr(EK1|\text{not-}EK4) = 1$, and so $Pr(\text{not-}EK4|EK1) > Pr(\text{not-}EK4)$ and $Pr(EK4|EK1) < Pr(EK4)$. But, still, $Pr(EK2|EK1) > Pr(EK2)$. That is, EK1 confirms EK2 while disconfirming EK4.[182] If justification follows probability relations then it seems that EK1 justifies believing EK2 but does not justify believing

[181] Sosa points out that there are two kinds of reliability or trustworthiness that EK-arguments and bootstrapping can lead to conclusions about: generic and specific. A source of evidence or method is generically reliable if it would lead us to form true beliefs over a wide variety of cases; it is specifically reliable if it would lead us to form a true belief in the present situation (i.e 'not easily would it in that instance lead us astray'). Sosa 2009: 225. In general, EK-arguments lead to conclusions that support the claim that the source of evidence in question is specifically reliable, whereas bootstrapping leads to the conclusion that the source of evidence in question is generically reliable.

[182] See White 2006 for this argument.

EK4 (in fact, it justifies lowering one's confidence in EK4). Yet by making RED TABLE ARGUMENT the subject ends up believing EK4 as a result of gaining EK1 as evidence.

Of course, the probabilistic critique of EK-arguments only goes through if we assume that perceptual evidence provides merely non-conclusive support for perceptual beliefs. However, if the probabilistic critique really did seem to capture what is wrong about EK-arguments then this might provide the grounds for an argument to the best explanation against the infallibilist view of perceptual evidence. The idea would be that, intuitively, there is something wrong with EK-arguments, and the best explanation for this intuition is that they violate rules governing probabilistic reasoning. Since this explanation involves the assumption that perceptual evidence provides less than conclusive support for perceptual beliefs, this would give us reason to think that perceptual evidence indeed does provide less than conclusive support.

However, the apparent violation of standard probabilistic assumptions can't be the fundamental problem with EK-arguments, because some of them don't violate any probabilistic rules (even assuming that perceptual evidence is non-conclusive) but still seem problematic. For example, we can't critique NOT-BIV in the same way because one's being a handless brain-in-a-vat does not entail that one will now be enjoying a perceptual experience as of hands, so it's not the case that $Pr(1_{BIV} | \text{not-}4_{BIV}) = 1$. Pryor makes the same point using an example of the following kind:

RED TABLE ARGUMENT 2

(A) The table looks red

(B) The table is red.

(C) The table is red and it's not white-and-illuminated-by-red-lights.[183]

(C) is logically equivalent to (B), so if (A) probabilistically confirms (B) to any extent then it confirms (C) to the same extent, so we don't here have a case in which one's perceptual evidence confirms a perceptual belief while disconfirming something which one deduces from that belief. But this argument is surely still an EK-argument of the problematic sort. So we won't find a general characterisation of what is wrong with EK-arguments by invoking probabilistic considerations.

We have now considered the main ways in which philosophers have developed the first way of reacting to EK-arguments, which is to hold that they cannot be routes to knowledge. We have found that this response is mainly motivated by judgements about cases and an inference to the best explanation.

The other way to react to EK-arguments is to hold that they are indeed routes to knowledge. This kind of response can only be made plausible if it is combined with an explanation of our discomfort with EK-arguments, thereby providing an error theory for our judgements about

[183] Pryor 2012: 282–3.

cases and answering the inference to the best explanation. Those of us who favour this kind of response need to identify something about EK-arguments which we are demonstrating sensitivity to in our negative appraisals of them and our refusal to rely on them in practice. In the literature so far, the main suggestion that has been made in response to this challenge is that EK-arguments have a peculiar kind of dialectical impotence, to which we are sensitive and which leads us our negative appraisals of them and our shunning of them in practice. I now turn to consideration of this approach. I argue that, while there is something to it, it cannot fully explain our negative appraisals of EK-arguments.

5.2. Dialectical ineffectiveness 1

Markie (2005) suggests that Cohen's use of a conversational context in RED TABLE CONVERSATION to elicit negative judgements about RED TABLE ARGUMENT shows that the problem these judgements are latching on to is no more than the dialectical misstep of begging the question. The father begs the question against the son, because the son doubts that the table's looking red is a reason to think that it is red, and yet the father then relies on that very reasoning in order to respond to the son's doubt that the table is not white with red lights shining on it. The father's reply to his son looks even worse because he has a non-question-begging reply available to him: he could utilise his inductive

evidence that tricksy red lights are very unusual or he could easy make special checks (by, for example, looking up at the ceiling to the lights) to persuade his son that the table is not white with red lights shining on it.

As Markie says, question-begging arguments can still be good arguments: someone who knows the premises can infer from them the conclusion and thereby come to know the conclusion. All arguments are question-begging when employed against someone who antecedently doubts a premiss or refuses to recognise the legitimacy of a transition in the argument. So if the sole sin of the father in RED TABLE CONVERSATION is that he begs the question against his son, then Cohen hasn't shown that there's anything wrong with RED TABLE ARGUMENT as an argument. It can still be a route to knowledge of its conclusion. It's just that, like any argument, one can't utilise it to rationally change the mind of someone who stubbornly doubts its premises or the support they offer for its conclusion. This is compatible with Liberalism and Deductive Knowledge Closure.

Markie is right that the father begs the question, and would be better to use inductive evidence or special checks in response to the son, but this is surely not the whole story about the badness of RED TABLE ARGUMENT. When we consider the argument on its own, not within any particular conversational context, it still looks bad. Cohen makes this point by redescribing the case in such a way as to 'eliminate entirely any dialectical context'. Suppose, he says, that the son accepts that the table

is red and that it is not white with red lights shining on it, but wants to know *how* the father knows this. If the father responds by saying 'Oh, that's easy. It looks red, so it is red, so it is not white with red lights shining on it', then, Cohen insists, this still looks like a bad answer.[184] But our judgement here can't be latching onto dialectical ineffectiveness, since there is no argumentative dialectic in play; instead, there is mere curiosity about the source of knowledge.

A more promising attempt to argue that the real problem with RED TABLE ARGUMENT and arguments like it has to do with dialectical ineffectiveness is that of Pryor (2004; 2012). Markie's suggestion implied that there is nothing special about EK-arguments that we might be latching onto when we judge that RED TABLE ARGUMENT's use is illegitimate in RED TABLE CONVERSATION, since any kind of argument can be deployed in the same sort of question-begging way. Pryor's suggestion differs in that he professes to find a way in which EK-arguments are dialectically ineffective that marks them out as unlike other good arguments. According to Pryor, EK-arguments seem fishy because they can never rationally persuade anyone who antecedently disbelieves or suspends judgement in their conclusions (where suspension of judgement is conceived of as an attitude towards the proposition, rather than the absence of any attitude[185]).

---

[184] Cohen 2005: 420.

[185] See Friedman 2013.

This suggestion has the advantage of providing a formula for classifying arguments as EK-arguments: they are such that if one antecedently doubts (i.e. disbelieves or suspends judgement in) their conclusions then one cannot be rationally persuaded by them to believe their conclusions. Markie's claim was just that since in RED TABLE CONVERSATION the son happened to doubt the premises or the reasoning, those very premises or that very reasoning couldn't be appealed to without begging the question. This didn't provide any way of singling out EK-arguments as distinctive. Pryor's suggestion is that anyone who doubts the conclusion of an EK-argument will be unable to rationally accept that argument's reasoning. This makes EK-arguments dialectically hamstrung in a specific way by their very nature: they can only be used by someone who antecedently has no attitude towards or who already believes their conclusions. This doesn't mean that they can't generate knowledge, but it does mark them out as having a special kind of fragility that other arguments don't have. In short, Markie's suggestion was that EK-arguments, like all arguments, can never be used to rationally persuade someone who doubts their premises or their reasoning; Pryor's suggestion is that EK-arguments, unlike other arguments, can never be used to rationally persuade someone who antecedently doubts their conclusions. This makes them unhelpful in many contexts in which one might want to use them.

Pryor defends his suggestion by appealing to the notion of defeat of epistemic statuses like knowledge and doxastic justification. Assuming the

reality of defeat, the negations of EK-conclusions are all propositions evidence in favour of which would count as defeating evidence for your knowledge of or justified belief in one of the premisses of the argument. For example, evidence that the table is white with red lights shining on it defeats your perceptually based knowledge that it is red. Likewise, evidence that you are a handless brain-in-a-vat defeats your perceptually formed knowledge that you have a hand. Given that the negations of EK-conclusions are potential defeaters for EK-premises, 'it follows', says Pryor, 'that if the subject *just happens* to doubt the conclusion, without having any reason to do so, her doubt will rationally obstruct her from believing the argument's premises'.[186]

Pryor calls the notion of rationality in play here 'hypothetical rationality'. Let's suppose that standard rationality consists in responding properly to one's reasons and evidence. Pryor calls this 'categorical rationality'. Pryor's hypothetical rationality is related to categorical rationality in the following way:

> [I]f some states A would, *when you had justification for them*, ground or justify other, inferentially 'downstream' attitudes B, then even when you lack justification, having A still puts you in *some* kind of interesting normative relation to B. But not the relation of simply justifying B. […] When you're justified in some beliefs A that obviously support a conclusion B, I'll say you have categorical justification for both A and B. When the

[186] Pryor 2004: 365–6.

justification for A is removed but you retain your (not unjustified) beliefs in A, then I'll say that they merely hypothetically justify or support having belief B.[187]

Hypothetical rational requirements are to be understood as genuine normative requirements. They are requirements to believe only in ways that would be categorically justified if all of your current beliefs were themselves justified. These hypothetical 'rational requirements' are what, in Chapter 3, we saw Kolodny categorise under the '"ought" of rationality' (contrasted with the '"ought" of reasons').[188]

Crucially, as Pryor seems to be thinking of things, following relevant hypothetical requirements is a necessary condition for having categorically rational beliefs. Thus, so long as you have an attitude – say, a belief or a suspension of judgement – which makes it hypothetically irrational to believe p then even if your total evidence supports believing p you cannot have a doxastically justified (or categorically rational) belief in p. The hypothetical requirement not to believe p 'rationally obstructs' believing p in a way that is categorically doxastically justified. Thus, a subject who has categorically irrational doubt in an EK-conclusion might have all-things-considered propositional justification to believe the EK-argument's premises, but she will be unable to form a doxastically

[187] Pryor 2018: 126. Pryor's emphasis. As Pryor notes, hypothetical rationality has been discussed under several names in the literature: 'structural' or 'attitudinal' rationality, or 'coherence requirements'. Cf. for example Broome 1999.

[188] Kolodny 2005.

justified belief in these premisses while doubting the conclusion, since in doubting the conclusion she is hypothetically rationally committed to believing *as though* her total evidence does not support believing the argument's premisses. She will therefore be unable to come to know the conclusion of the EK-argument, and so to rationally assuage her doubt, because she is unable to come to know the argument's premisses. Doubt about the conclusion rationally pre-empts the EK-argument.

Sceptics, of course, do doubt the conclusions of anti-sceptical EK-arguments, or Moorean arguments. Consequently, Moorean arguments cannot rationally dissuade a sceptic, and so are rationally impotent against them. Similarly, the son in RED TABLE CONVERSATION might be interpreted as doubting that the table is not white-lit-by-red-lights. If that's right, he will, according to Pryor's suggestion, be unable to rationally accept the father's reasoning, which will therefore be rationally impotent in this context.

5.3. Dialectical ineffectiveness 2: the dispositional approach

In the previous chapter I endorsed the view that knowledge can't be defeated. 'Defeat' involves one's knowledge that p being destroyed merely by one coming to have misleading evidence that one does not know p. If I was right to endorse this view then we are deprived of Pryor's defence of his suggestion in terms of defeat. For if defeat of

knowledge is not possible, then doubt about an EK-argument's conclusion can't defeat knowledge of its premisses. Thus, even if a subject doubts the conclusion of an EK-argument, she should be able to use the argument to assuage her doubt and to come to know the conclusion. My denial of the reality of defeat therefore seems to force me to reject Pryor's suggestion about what is wrong with EK-arguments.

However, I also defended Lasonen-Aarnio's claim that there is a sense in which it is unreasonable for a subject who knows p to ignore misleading evidence that they don't know p. To stubbornly continue believing p in the face of such evidence may be a way of retaining knowledge, but it involves manifesting a disposition to ignore evidence of ignorance. Usually, evidence of ignorance is not misleading and alerts us to actual ignorance. Therefore, given that our epistemic goals are to maximise knowledgeable beliefs and to minimise unknowledgeable beliefs, this is an epistemically bad disposition to have.

As we saw in the previous chapter, noticing this fact about dispositions gives the infallibilist a way of responding to the criticism that they cannot accommodate knowledge-defeat. The problem for the infallibilist was that the conditions that infallibilism adverts to as conferring on a belief the status of knowledge (such as: meeting a safety condition, being based on knowledge-guaranteeing evidence, etc.) can continue to obtain, unaffected by the subject coming to possess misleading evidence of error or ignorance. Of course, in practice, in most such cases the subject stops

believing p as a result of receiving the misleading evidence, and for that reason stops knowing p. The apparently implausible infallibilist commitment is the idea that the subject is still in a position to know p, and that a stubborn subject who maintains their belief in p would continue to know p, and in addition that the subject's total evidence still supports believing p so that there is a sense in which the subject *ought* to continue believing p. All of these can seem like implausible claims that infallibilists of various kinds are committed to.

Lasonen-Aarnio suggested a response. A stubborn subject continues to know p, but manifests an epistemically bad disposition in doing so, and in this sense is unreasonable. We hesitate to attribute knowledge to this subject because we tend to latch onto the dispositions that the subject is manifesting by believing as she does, treating these dispositions as guides to whether they know or not. I argued that this is a compelling error theory for our judgements about defeat, which allows us to retain our

judgement that there is a sense in which one ought to give up one's knowledgeable belief when one receives misleading evidence of error.[189]

I suggest now that we can apply Lasonen-Aarnio's notion of reasonableness in a way that parallels Pryor's application of the notion of defeat. Just as, on Pryor's approach, having misleading evidence for the negation of an EK-argument's conclusion defeats knowledge of an EK-premise, so on Lasonen-Aarnio's approach having this misleading evidence makes it reasonable to avoid believing the relevant EK-premise. Pryor's next step is to claim that doubting the EK-conclusion for no good reason makes it hypothetically rational for you to avoid believing the

[189] A different way that an infallibilist might try to use the resources provided by Lasonen-Aarnio is by claiming that not manifesting epistemically bad dispositions is a necessary condition on knowledge. On this approach, knowledge-defeat *is* real, but the infallibilist can accommodate it after all. Since manifesting bad dispositions is compatible with safe belief and, arguably, with believing on the basis of one's factive or relational evidence, not manifesting bad dispositions would have to be a genuine, non-redundant extra condition on knowledge. Manifesting bad dispositions is compatible with safe belief despite the fact that bad dispositions are bad precisely because their manifestation more often than not produces or sustains non-knowledgeable or false beliefs. For, as Lasonen-Aarnio puts it, 'circumstances can conspire so as to make applications of generally unsafe or unreliable methods yield safe beliefs' (Lasonen-Aarnio 2010: 7). A belief formed or sustained in part because of an epistemically bad disposition might be safe from error *in the circumstances*. Putative defeat cases exemplify this. If you perceptually know that the wall is red, and then get misleading evidence that it's white with red lights shining on it, but you stubbornly maintain your perceptual belief that the wall is red, this belief remains safe so long as there is no nearby possible world in which the wall is white with red lights shining on it. Virtue epistemologists do include as a necessary condition on knowledge something like the manifestation of epistemically good dispositions, and this is certainly another option for infallibilists. Lasonen-Aarnio recognises this, while offering some criticisms of the virtue-theoretic approach. I will put aside the virtue theoretic approach, too; what I say in the rest of this section of the chapter is, though, compatible with developing Lasonen-Aarnio's suggestions in a virtue-theoretic way.

premise. On the approach I am suggesting – which we can call the 'dispositional approach to the complaint of dialectical ineffectiveness', or 'the dispositional approach' for short – the notion of hypothetical rationality plays no role. Instead, we can simply say that if you doubt the EK-conclusion, whether with good reason or not, then if you believe the relevant (EK2)-like premise on the basis of (EK1)-like evidence, then you are manifesting an epistemically bad disposition.

Lasonen-Aarnio claimed only that it is unreasonable to stubbornly retain a knowledgeable belief when one receives evidence of error or ignorance. Thus, Lasonen-Aarnio would claim that if you receive evidence that you are a handless brain-in-a-vat then you should cease believing that you have hands on the basis of your perceptual experience. I am here claiming, further, that if you *merely believe* or *suspend judgement with respect to* the proposition that you are a handless brain-in-a-vat, then it is unreasonable for you to believe that you have hands on the basis of your perceptual experience. It doesn't matter whether this attitude (belief or suspension) towards the proposition that you are a handless brain-in-a-vat is itself rational and evidence-based or not. Is this extension of Lasonen-Aarnio's position plausible?

I suggest that it is. Consider a case where you doubt, without reason, the conclusion of an EK-argument. By the description of the case, you have failed to notice that you lack reason for this doubt. For had you noticed this, you wouldn't doubt. The attitude you have – whether outright belief

in the negation of the EK-argument's conclusion or suspension – is one that seems to you to be rational. Of course, there is a sense in which this case differs 'from the inside', as it were, from one in which you doubt the EK-conclusion for good reason, since in the latter case you possess a reason for doubt which in the former you lack. But as Lasonen-Aarnio says, 'even if the good and bad cases differ in some respects "from the inside"'[190] (where a good case is a knowledgeable case with misleading evidence of error, and a bad case is an error-case with non-misleading evidence of error), our dispositions may not be able to pick up on the differences. After all, as we noted already, error cases in this context just are cases in which one has failed to be sensitive to an error, at least before receiving evidence of error. Consequently, you can't have a disposition to react in different ways to the good and the bad case. And the same point holds with respect to doubting an EK-conclusion for good reason and doubting an EK-conclusion for no good reason: we humans can't have dispositions that are sensitive to the differences between these two cases. Thus, any disposition to avoid believing the (EK-2)-like-premise on the basis of the (EK-1)-like-evidence when you have good reason to doubt the relevant (EK)-conclusion will also be a disposition to avoid believing the (EK-2)-like-premise on the basis of the (EK1)-like-evidence when you merely doubt (for no good reason) the relevant (EK)-conclusion. This claim is parallel to Lasonen-Aarnio's claim that any disposition to revise a false belief in p on the basis of non-misleading evidence that you've erred in believing p will also be a disposition to revise a knowledgeable

[190] Lasonen-Aaarnio 2021: 102.

belief in p on the basis of misleading evidence that you've erred in believing p. And since reasonable subjects *do* have a disposition to avoid believing (EK2)-like premisses on the basis of (EK-1)-like evidence when they have reason to doubt the relevant (EK)-conclusions, reasonable subjects are disposed to avoid believing (EK2)-like premisses on the basis of (EK-1)-like evidence when they merely doubt the relevant (EK)-conclusions without reason. So the extension of Lasonen-Aarnio's view that I am suggesting is as plausible as her original view, which I defended in the previous chapter.[191]

Whereas Pryor's approach involves the claim that a subject who doubts an EK-conclusion cannot know (because they cannot categorically rationally believe) the relevant EK-premise, the dispositional approach involves the claim that a subject doubting an EK-conclusion cannot believe the relevant EK-premise without being unreasonable. Of course, on Pryor's approach, the doubting subject could give up their irrational doubt in the EK-conclusion, and then come to know the premisses of the EK-argument and the conclusion too. His point is that the EK-argument by itself cannot rationally persuade the doubting subject to do this, since its rational force is inaccessible to them while they doubt the EK-conclusion. On the dispositional approach, the doubting subject *could* come to know the premisses of the EK-argument without first ceasing to doubt the EK-conclusion, but in doing this they would be manifesting an

[191] Of course, a subject who doubts an EK-conclusion for no good reason should cease doubting the EK-conclusion, but they should do this because they lack reason for doubting it rather than in response to being presented with an EK-argument.

epistemically bad disposition. There is thus, on the dispositional approach, no way for a subject who antecedently doubts an EK-conclusion to be persuaded to *reasonably* (in Lasonen-Aarnio's sense) believe an EK-conclusion by an EK-argument.

This is a less severe indictment of EK-arguments than that of outright dialectical impotence in the face of a doubting interlocutor, as on Pryor's approach. But since most humans don't have the bad disposition in question, the dispositional treatment of EK-arguments does imply that in practice they will be unable to persuade interlocutors in exactly the same situations as on Pryor's suggestion.

In summary, then, the dispositional approach to dialectical ineffectiveness allows us to retain the claim that EK-arguments display a peculiar kind of dialectical ineffectiveness, without needing to defend this claim by appealing to the notion of defeat. Instead, we can defend the claim by appealing to the idea that subjects who ignore misleading evidence that their knowledgeable belief is false, unsupported or in some other way epistemically inappropriate are thereby manifesting an epistemically bad disposition. EK-arguments are peculiar in that someone who antecedently doubts their conclusions (whether for good reason or not) cannot, without being unreasonable, use them to overcome their doubt.

5.4. Evaluating the dialectical ineffectiveness diagnosis

Pryor's suggestion was, like Markie's, that EK-arguments arguments are dialectically ineffective in a certain way. I have suggested a variation of Pryor's suggestion, which I called the dispositional approach to dialectical ineffectiveness. But, as we saw, Cohen responded to Markie's suggestion by pointing out that even if we strip away the dialectical context of RED TABLE CONVERSATION, the RED TABLE ARGUMENT still seems problematic. It might seem that surely the same response can be given to both Pryor's suggestion and to the dispositional variation on his suggestion that I have defended.

However, Pryor's diagnosis of EK-arguments arguments is arguably more protected from this kind of response than Markie's. As we saw, Pryor's suggestion, unlike Markie's, provides a way of categorising EK-arguments as a distinctive kind. Pryor presents us with a distinctive feature that all EK-arguments have, which other kinds of argument lack. Thus, although this feature has to do with their role and power in dialectical settings, it's plausible that we might latch on to this feature of EK-arguments even when we are presented with them outside of any dialectical context. For this feature is a result of the relationship between the conclusion and premises of EK-arguments, namely that evidence for the negation of the conclusion is evidence that suggests that (EK1) does not support (EK2). This relationship is still obvious even when the

argument is displayed outside of a dialectical context; and it's plausible that it makes us feel uneasy about the credentials of the argument. So Cohen's response to Markie is less suited to undermining Pryor's suggestion. Of course, the dispositional approach to dialectical ineffectiveness inherits this advantage, too. It cannot be denied, though, that Cohen's complaint retains some force even against Pryor and the dispositional approach. I come back to this below, in the conclusion.

All the same, I think that the dispositional approach to dialectical effectiveness tracks a definite feature of EK-arguments which undoubtedly influences how we respond to them. EK-arguments are, as a result of the relationship between their conclusions and premisses, dialectically ineffective against anyone who doubts their conclusion, in that such a doubter would be epistemically unreasonable in Lasonen-Aarnio's sense – manifesting a bad disposition – if they allowed themselves to be persuaded by the relevant EK-argument.

5.5. An argument for thoroughgoing externalism

Should we prefer the dispositional approach to dialectical effectiveness over Pryor's approach? The main reason suggested so far for preferring the dispositional approach is simply the conclusion of the previous chapter, that knowledge-defeat is not a reality. Since Pryor's approach is only defensible given the reality of defeat, the dispositional approach,

which makes no reference to defeat, is preferable. However, there is a further reason to prefer the dispositional approach to Pryor's approach, and this further reason has a wider significance, lending support to the claim that knowledge-defeat is non-actual and to the thoroughgoing externalist claim that one's evidence in the 'good' case differs from one's evidence in the corresponding 'bad' case.

Dretske, like many philosophers, held that EK-arguments are not ways of coming to know EK-conclusions.[192] But, further, he held that we humans have no means of coming to know the negations of radical sceptical hypotheses. For example, according to Dretske, we cannot come to know that we are not handless brains in vats. We *can* can come to know that the table is not white-lit-by-red-lights, since we can use induction or special checks to know such things, but these resources are unavailable to us when it comes to inquiring about radical sceptical hypotheses. Dretske's stance thus differs from that of philosophers like Wright, who agrees with Dretske that EK-arguments are not ways of coming to know EK-conclusions, but who holds that we do have an alternative source of entitlement to accept the negations of radical sceptical hypotheses.[193] From Dretske's perspective we should suspend judgement with respect to the negations of sceptical hypotheses, and thus with respect to the conclusions of anti-sceptical EK-arguments. Yet, at the same time, and also unlike Wright, Dretske holds that transitions

[192] See Dretske 1970; Dreske 1971.

[193] See Wright 2004a; Wright 2007.

from (EK1)-like evidence (taken on its own) to belief in (EK2)-like propositions is epistemologically legitimate, and is a way of coming to know (EK2)-like propositions.[194] So Dretske is someone who suspends belief with respect to the conclusions of anti-sceptical EK-arguments, yet who is also happy to base his belief in (EK2)-like propositions on (EK1)-like evidence. For example, he will suspend judgement about whether he is a handless brain-in-a-vat while believing on the basis of his perceptual experience that he has hands.

However, if Pryor's suggestion is right then Dretske, by suspending judgement with respect to propositions like *it's not the case that I am a handless brain-in-a-vat*, has actually ceased to have any perceptual knowledge at all. For, according to Pryor, if one suspends judgement with respect to such propositions then one cannot come to know that, for example, one has hands on the basis of one's perceptual experiences. Although Dretske has presumably not considered all possible radical sceptical scenarios, and so has not formed an attitude of suspension with respect to all sceptical hypotheses, the fact that he has considered the proposition that he is not a handless brain-in-a-vat and suspended judgement with respect to it is enough to deprive him of all his perceptual knowledge, assuming Pryor's suggestion is right, since evidence for that proposition would (according to Pryor) be defeating evidence for all of his perceptual knowledge.

---

[194] To be clear, Wright thinks that such transitions are legitimate, but denies that the (EK1) evidence can do the work *all by itself*.

This result about Dretske is reason to doubt Pryor's suggestion. I don't think it's plausible that Dretske has, through his work in epistemology, inadvertently rendered himself without any perceptual knowledge. Dretske continues to navigate the world using his perceptual faculties; he continues to be as reliable as ever in his perceptual beliefs; and they are formed in a way that is amenable to the standard rationalising explanations we give for perceptual knowers.

On the other hand, the dispositional approach to dialectical ineffectiveness does not imply that Dretske lacks perceptual knowledge. Instead, on the dispositional approach, Dretske retains all his perceptual knowledge, and is merely manifesting an epistemologically bad disposition by retaining his perceptual beliefs at the same time as doubting anti-sceptical EK-conclusions. As we discussed in the previous chapter, one can be manifesting an epistemologically bad disposition by believing p and yet one's belief in p can still be knowledge. So the dispositional approach is preferable to Pryor's approach to dialectical ineffectiveness. It captures a sense in which there is something odd about Dretske's position, while not depriving him of all of his perceptual knowledge. It is more conceivable that philosophy might drive someone to a position which involves managing one's beliefs in a way that is not knowledge-conducive, than it is that it might deprive someone of all their perceptual knowledge.

Dretske is not the only philosopher who denies that we can know the negations of sceptical hypotheses. Nozick is another who has, on Pryor's account, deprived himself of any perceptual knowledge.[195] However, this denial that we can know the negations of sceptical hypotheses is, in the wake of criticism of Dretske and Nozick, now often accompanied by a contextualist treatment of knowledge attributions.[196] On some versions of contextualism, merely raising a sceptical hypothesis in conversation raises the standards for knowledge so high that one counts as knowing neither that the sceptical hypothesis is false nor that one has hands (or whatever other everyday proposition is at issue). The result is that in everyday situations, when the standards for knowing are low enough for one to perceptually know that one has hands, one is in a position to know that one is not a handless brain-in-a-vat. However, if coming to actually know that one is not a handless brain-in-a-vat involves entertaining that proposition, then one cannot come to know it, since by entertaining the proposition one will raise the standards for knowledge to a level that one cannot satisfy. It is unclear what standing, non-occurrent attitude we should have towards sceptical hypotheses, according to contextualists. If the standing attitude should be one of suspension, then they too, according to Pryor's suggestion, will have inadvertently allowed misguided philosophical doctrine to deprive them of any perceptual

---

[195] Nozick 1981.

[196] See, e.g., Cohen 1988; DeRose 1995. Given their contextualism, it is unclear whether or not Cohen and DeRose advocate suspending judgement with respect to the negations of sceptical hypotheses.

knowledge. Again, I regard this as implausible: contextualism may be false, but contextualists don't lack perceptual knowledge.[197, 198]

Could Pryor alter his view to avoid this implausible consequence? Could he hold that it is only outright *disbelief* in an EK-conclusion that makes it impossible to know an (EK2)-like premise on the basis of (EK1)-like evidence, so that merely *suspending* with respect to an EK-conclusion doesn't render one unable to have everyday perceptual knowledge? After all, Dretske doesn't outright disbelieve EK-conclusions; he merely suspends. Not even sceptics disbelieve EK-conclusions; they just claim that EK-conclusions cannot be known and so, presumably, the appropriate attitude to take to them when they are raised is suspension. However, altering Pryor's suggestion in this way is not tenable, for two reasons. First, if the dialectical ineffectiveness of EK-arguments merely amounts to their inability to rationally persuade someone who

---

[197] Contextualism was in part developed to provide a diagnosis of EK-arguments. There is no space here for a thorough critique of contextualism; suffice it to say that I find persuasive both Pryor's specific criticism of contextualist treatments of EK-arguments and the independent considerations that have been brought to bear against contextualist theories of knowledge in general. For Pryor's criticism see Pryor 2004: 349–50. There is a vast literature on contextualism. The most developed recent defence is to be found in DeRose 2009 and DeRose 2017. See also Unger 1984; Cohen 1988; Lewis 1996; Cohen 1999 and DeRose 1995. Among the many critiques of contextualism, see Brueckner 1994; Feldman 1999; Kornblith 2000; Sosa 2000; Klein 2000; Feldman 2001; Stanley 2004; Dretkse 2013.

[198] To be clear, the important point is not that Dretske, Nozick and their contextualist colleagues are committed by their own epistemological theories to suspending judgement with respect to sceptical hypotheses. The important point is rather that it is implausible to hold that someone who was persuaded by philosophical theory to suspend judgement with respect to sceptical hypotheses would thereby be deprived of all perceptual knowledge. But that is what Pryor's suggestion implies.

antecedently outright disbelieves their conclusions, this dialectical deficiency would almost never play any role in practice, since in practice no one (not even the sceptic) disbelieves the conclusions of anti-sceptical EK-arguments, or EK-arguments like RED TABLE ARGUMENT. Since the deficiency would therefore be irrelevant in most contexts, it's puzzling why we should display our sensitivity to this feature of EK-arguments by appraising them negatively, and even holding that they cannot be ways of knowing their conclusions. So dialectical ineffectiveness must amount to more than this, as it does with the inclusion of the idea that if you antecedently suspend with respect to the negations of sceptical hypotheses (as the sceptic does) then you can't rationally believe EK-premisses.

Second, altering Pryor's suggestion in this way is not tenable because, intuitively, there *is* something wrong with suspending judgement with respect to whether you are a handless brain-in-a-vat while at the same time believing, on the basis of perception, that you have hands – and the same goes for other EK-conclusions and premisses. But if Pryor limits his suggestion in the way under consideration, he will lack any account of what is wrong with doing this. So Pryor should maintain that both disbelieving and suspending judgement with respect to EK-conclusions rationally obstructs believing (EK2)-like propositions on the basis of (EK1)-like evidence. But, of course, this leads him into the problem I have been describing.

This simple argument in favour of the dispositional approach to dialectical ineffectiveness has wider significance, especially in relation to thoroughgoing externalist conceptions of evidence and reasons. Endorsing the dispositional approach to dialectical ineffectiveness instead of Pryor's approach is only sustainable given a conception of perceptual evidence as factive or relational. For if perceptual evidence is not world-involving in this way, then the idea that perceptual knowledge can be defeated by misleading evidence is irresistible. We saw this in the previous chapter. In section 4.5 I considered Brown's argument that to deny the reality of knowledge-defeat is to deny, absurdly, that 'justification is affected by one's total evidence, rather than just a part of that evidence'.[199] I agreed with Brown that one's total evidence is what determines what one should believe, but I pointed out that on a thoroughgoing externalist conception of evidence, according to which one's evidence in 'good' cases differs from one's evidence in corresponding 'bad' cases, we can accommodate this fact while denying that knowledge-defeat is actual. This response to Brown is not available to someone who is not a thoroughgoing externalist about evidence and reasons. Thus, someone who is not a thoroughgoing externalist is liable to find Brown's argument, according to which the defeat of knowledge (including perceptual knowledge) is a genuine phenomenon, irresistible. And if perceptual knowledge can be defeated, then evidence that you are a handless brain-in-a-vat is surely a paradigm case of defeat of your knowledge that you have hands. If so, then Pryor's approach to the

[199] Brown 2018: 108.

dialectical ineffectiveness of EK-arguments is, I think, in turn irresistible, and so we end up having to deny that Dretske has any perceptual knowledge. So not only should we favour a dispositional approach to dialectical ineffectiveness, but we now have a further consideration in favour of the conclusion of the previous chapter, that knowledge defeat is not a genuine phenomenon. For only if we deny the reality of knowledge defeat can we avoid having to claim that Dretske lacks any perceptual knowledge. Furthermore, since the only plausible way to resist Brown's argument that knowledge-defeat is actual is by being a thoroughgoing externalist, we have here a – perhaps surprising – reason to endorse thoroughgoing externalism.

Recall from the previous chapter White's argument against the idea that perceptual evidence is conclusive. He presents a case, Suspected Hallucination, which he says is clearly a case of knowledge-defeat. Since the subject's basis remains in place, he says, that basis cannot have been conclusive. The argument I have now given for a thoroughgoing externalist conception of evidence rests on a case that is a mirror-image of White's Suspected Hallucination. I have presented a case in which, intuitively, the subject (Dretske) retains his knowledge, but about which a non-factive, non-relational account of perceptual evidence predicts that his knowledge is defeated (assuming here the cogency of Pryor's argumentation given the assumption that defeat is a real phenomenon). Of course the non-factive, non-relational theorist could respond by saying that Dretske only appears to retain his knowledge, and that the

intuition about his situation that I have been pushing is mistaken. So this is not a conclusive argument. But it does show, at the very least, that it is not only infallibilists who are forced into awkward manoeuvres by their stance with respect to defeat.

5.6. Conclusion

In this chapter I have embraced the idea that EK-arguments are hamstrung by a special sort of dialectical ineffectiveness. The main claim has been that this is what explains our negative reactions towards EK-arguments, and our tendency in practice not to employ them. But this dialectical ineffectiveness is compatible with EK-arguments being genuine routes to knowledge; we therefore need not give up either Liberalism or Deductive Knowledge Closure.

I have also argued that consideration of Pryor's account of the dialectical ineffectiveness of EK-arguments leads to an argument for thoroughgoing externalism and the unreality of knowledge-defeat. For Pryor's account is extremely convincing given the reality of defeat, and the reality of defeat is extremely convincing given the falsity of thoroughgoing externalism. But Pryor's account of the dialectical ineffectiveness of EK-arguments carries the unfortunate commitment of having to claim that Dretske and Nozick deprived themselves of perceptual knowledge (and, indeed, other kinds of knowledge) by

suspending judgement with respect to radical sceptical hypotheses. So to deny thoroughgoing externalism leads, via defeat and Pryor's account of what's wrong with EK-arguments, to a quite unattractive commitment.

I should note, however, that I do not think that the dialectical ineffectiveness diagnosis of EK-arguments is the full story about what makes us uneasy about them. For one thing, there is probably no unified account of why various philosophers have been uneasy about EK-arguments: a variety of different assumptions might be at work. Furthermore, it remains true that if we completely strip EK-arguments of dialectical contexts, they still seem somehow inadequate. Of course, one might claim that this inadequacy simply is their dialectical ineffectiveness, which we can pick up on even when we contemplate them outside of any dialectical context. However, what I think it hard to explain by appeal to dialectical ineffectiveness is that, as seems likely to me, someone who has no attitude towards a particular EK-conclusion – perhaps because they've never entertained it before – might, on encountering an EK-argument for that conclusion, be led, just by considering that argument, to newly doubt the EK-conclusion. For the EK-argument might simply seem obviously inadequate for producing knowledge of its conclusion, and one might then worry that one has no other route to knowledge of its conclusion. But if they didn't antecedently doubt the EK-conclusion, why should they be so resistant to the EK-argument?

The dialectical ineffectiveness suggestion will also presumably be unable to persuade those philosophers like Dretske and Nozick who, as we saw, deny that we have *any* way of knowing some EK-conclusions, but who nevertheless think that it's perfectly legitimate to believe (EK2)-premisses on the basis of (EK1)-evidence. The dialectical ineffectiveness suggestion is precisely that if you doubt an EK-conclusion then it's *not* legitimate to believe (EK2)-premisses on the basis of (EK1)-evidence. So Dretske and Nozick's claim that EK-arguments are not ways of knowing their conclusions is not best understood as driven by their unconscious sensitivity to the dialectical ineffectiveness of EK-arguments. Any such sensitivity would have to be very deeply repressed, and completely inoperative in their doxastic dispositions. What should we say to such philosophers? I do not have answers to these questions.

Bibliography

Antognazza, Maria Rosa (2015), 'The Benefit to Philosophy of the Study of Its History', *British Journal for the History of Philosophy*, Vol. 23, No. 1: 161–184. DOI: 10.1080/09608788.2014.974020

Alston, William (1988), 'An Internalist Externalism', *Synthese*, 74, No. 3: 265–83. DOI: 10.1007/BF00869630.

Audi, Robert (1986), 'Acting for Reasons', *The Philosophical Review*, Vol. 95, No. 4: 511–546. DOI: 10.2307/2185049.

Austin, J. L. (1962), *Sense and Sensibilia*, Oxford: Clarendon Press.

Baker, G. P. and P. M. S. Hacker (1984), *Language, Sense and Nonsense*, Oxford: Blackwell.

Baker-Hytch, Max and Matthew Benton (2015), 'Defeatism Defeated', *Philosophical Perspectives*, Vol. 29, No. 1: 40–66. DOI: 10.1111/phpe.12056

Bergmann, Michael (2000), 'Externalism and Skepticism', *The Philosophical Review*, Vol. 109, No. 2: 159–194. DOI: 10.1215/00318108-109-2-159

Bergmann, Michael (2006), *Justification Without Awareness*, Oxford: Oxford University Press.

Black, Tim (2008), 'Solving the Problem of Easy Knowledge', *Philosophical Quarterly*, Vol. 58, No. 233.

Block, Ned (1995), 'On a confusion about a function of consciousness', *Brain and Behavioural Sciences*, Vol. 18, No. 2: 227–47.

Boghossian, Paul (2008), 'Epistemic Rules' in *Content and Justification: Philosophical Papers*, Oxford: Oxford University Press.

Boghossian, Paul (2014), 'What Is Inference?', *Philosophical Studies*, Vol. 169, No. 1: 1–18. DOI: DOI: 10.1007/s11098-012-9903-x

BonJour, Laurence (1976), 'The coherence theory of empirical knowledge', *Philosophical Studies*, Vol. 30, No. 5: 281–312. DOI: 10.1007/BF00357928

BonJour, Laurence (1978), 'Can Empirical Knowledge Have a Foundation?' *American Philosophical Quarterly*, Vol. 15: 1–14.

BonJour, Laurence (1985), *The Structure of Empirical Knowledge*, Cambridge, MA: Harvard University Press.

Boyle, Matthew (2009), 'Active Belief', *Canadian Journal of Philosophy*, Vol. 39: 119–147. DOI: 10.1080/00455091.2009.10717646

Boyle, Matthew (2011), '"Making up Your Mind" and the Activity of Reason', *Philosophers' Imprint*, Vol. 11, No. 17.

Bradley, Darren (2014), 'A Relevant Alternatives Solution to the Bootstrapping and Self-Knowledge Problems', *The Journal of Philosophy*, Vol. 111, No. 7: 379–393. DOI: https://doi.org/10.5840/jphil2014111726

Breckenridge, Wylie (2007), 'Against one reason for thinking that visual experiences have representational content', *Philosophical Perspectives*, Vol. 21, No. 1: 117–123. DOI: 10.1111/j.1520-8583.2007.00122.x

Brewer, Bill (1999), *Perception and Reason*, Oxford: Oxford University Press.

Brewer, Bill (2004), 'Realism and the Nature of Perceptual Experience', *Philosophical Issues*, Vol. 14: 61–77.

Brewer, Bill (2005), 'Perceptual experience has conceptual content', in Ernest Sosa and Matthias Setup (eds), *Contemporary Debates in Epistemology*, Oxford: Blackwell.

Brewer, Bill (2006), 'Perception and Content', *European Journal of Philosophy*, Vol. 14, No. 2: 165–181. DOI: 10.1111/j.1468-0378.2006.00220.x

Brewer, Bill (2011), *Perception and Its Objects*, Oxford: Oxford University Press.

Brewer, Bill (2018), 'Perceptual Experience and Empirical Reason', *Analytic Philosophy*, Vol. 59, No. 1: 1–18. DOI: 10.1111/phib.12123

Brewer, Bill (2019), 'Empirical Reason: Answers to Gupta, McDowell, and Siegel', *Philosophical Issues*, Vol. 29, No. 1: 366–377. DOI: 10.1111/phis.12160

Brewer, Bill, David de Bruijn, Chris Hill, Adam Pautz, Raja Rosenhagen, Miloš Vuletić, Wayne Wu (2018), 'Discussion of Bill Brewer's "Perceptual Experience and Empirical Reason"', *Analytic Philosophy*, Vol. 59, No. 1: 19–32.

Broome, John (1999), 'Normative Requirements', *Ratio*, Vol. 12, No. 4: 398–419. DOI: 10.1111/1467-9329.00101

Broome, John (2003), *Rationality Through Reasoning*, Chichester: Wiley-Blackwell.

Brown, Jessica (2018), *Fallibilism: Evidence and Knowledge*, Oxford; Oxford University Press.

Brueckner, A. (1994), 'The Shifting Content of Knowledge Attributions', *Philosophy and Phenomenological Research*, Vol. 54, No. 1: 123–126. DOI: 10.2307/2108358

Brueckner, A. (2009), 'E=K and Perceptual Knowledge', in Patrick Greenough and Duncan Pritchard (eds), *Williamson on Knowledge*, Oxford: Oxford University Press.

Burge, Tyler (1979), 'Individualism and the Mental', *Midwest Studies in Philosophy*, Vol. 4, No. 1: 73–121. DOI: 10.1111/j.1475-4975.1979.tb00374.x

Burge, Tyler (2003), 'Perceptual Entitlement', *Philosophy and Phenomenological Research*, Vol. LXVII, No. 3: 503–548. DOI: 10.1111/j.1933-1592.2003.tb00307.x

Byrne, Alex (1996), 'Spin Control: Comment on John McDowell's *Mind and World*', *Philosophical Issues*, Vol. 7: 261–273.

Byrne, Alex (2009), 'Experience and Content', *The Philosophical Quarterly*, Vol. 59, No. 236: 429–451.

Campbell, John (2000), 'Wittgenstein on Attention', *Philosophical Topics*, Vol. 28, No. 2: 35–48. DOI: 10.5840/philtopics200028213

Campbell, John (2002), *Reference and Consciousness*, Oxford: Oxford University Press.

Carroll, Lewis (1895), 'What the Tortoise Said to Achilles', *Mind*, Vol. 4, 14: 278–80.

Carruthers, Peter (2011), *The Opacity of Mind: An Integrative Theory of Self-Knowledge*, Oxford: Oxford University Press.

Carter, J. Adam and Duncan Pritchard (2016), 'Perceptual knowledge and relevant alternatives', *Philosophical Studies*, Vol. 173: 969–990. DOI: 10.1007/s11098-015-0533-y

Chudnoff, Elijah (2011), 'The nature of intuitive justification', *Philosophical Studies*, Vol. 153, No. 2: 313–333. DOI: 10.1007/s11098-010-9495-2

Chudnoff, Elijah (2011b), 'What Should a Theory of Knowledge Do?', *Dialectica*, Vol. 65, No. 4: 561–579. DOI: 10.1111/j.1746-8361.2011.01285.x

Chudnoff, Elijah (2012), 'Presentational Phenomenology', in Sofia Miguens and Gerhard Preyer (eds), *Consciousness and Subjectivity*, De Gruyter.

Cohen, Stewart (1984), 'Justification and Truth', *Philosophical Studies*, Vol. 46, No. 3: 279–95. DOI: 10.1007/BF00372907

Cohen, Stewart (1988), 'How to be a Fallibilist', *Philosophical Perspectives*, 2: 91–123. DOI: 10.2307/2214070

Cohen, Stewart (1999), 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives*, 13: 57–90. DOI: 10.1111/0029-4624.33.s13.3

Cohen, Stewart (2000), 'Contextualism and Skepticism', *Noûs*, Vol. 34: 94–107. DOI: 10.1111/0029-4624.34.s1.12

Cohen, Stewart (2002), 'Basic Knowledge and the Problem of Easy Knowledge', *Philosophy and Phenomenological Research*, Vol. 65, No. 2: 309-329.

Cohen, Stewart (2005), 'Why Basic Knowledge is Easy Knowledge', *Philosophy and Phenomenological Research*, Vol. 70, No. 2: 417–430. DOI: 10.1111/j.1933-1592.2005.tb00536.x

Conee, Earl and Richard Feldman (2004), *Evidentialism: Essays in Epistemology*, Oxford: Oxford University Press.

Cook Wilson, John (1926), *Statement and Inference*, Oxford: Oxford University Press.

Cummins, Robert (1998), 'Reflection on Reflective Equilibrium', in M. DePaul and W. Ramsey (eds), *Rethinking Intuition*, Lanham, MD: Rowman & Littlefield.

Cunningham, J. J. (2018), 'Are Perceptual Reasons the Objects of Perception?', in Johan Gersel, Rasmus Thybo Jensen, Morten S. Thaning, and Søren Overgaard (eds), *In the Light of Experience: New Essays on Perception and Reasons*, Oxford: Oxford University Press.

Dancy, Jonathan (2000), *Practical Reality*, Oxford: Oxford University Press.

Dancy, Jonathan (2014), 'On Knowing One's Reason', in Clayton Littlejohn and John Turri (eds), *Epistemic Norms: New Essays on Action, Belief, and Assertion*, Oxford: Oxford University Press.

Davidson, Donald (1980), 'Actions, Reasons, and Causes', in *Essays on Actions and Events*, Oxford: Oxford University Press.

Davidson, Donald (1983/2001), 'A Coherence Theory of Truth and Knowledge', in *Subjective, Intersubjective, Objective: Philosophical Essays Volume 3*, Oxford: Oxford University Press, 2001.

Davidson, Donald (1999), 'Reply to John McDowell', in *The Philosophy of Donald Davidson*, L. E. Hahn (ed.), Chicago, IL: Open Court.

Davies, Alex (2018), 'Infallibilism and Easy Counter-Examples', *Grazer Philosophische Studien*, Vol. 95, No. 4: 475–499. DOI: 10.1163/18756735-000054

Dennett, Daniel (2009), 'Intentional Systems Theory', in Ansgar Beckermann, Brian P. McLaughlin, and Sven Walter (eds), *The Oxford Handbook of the Philosophy of Mind*, Oxford: Oxford University Press.

DeRose, Keith (1995), 'Solving the Skeptical Problem', *Philosophical Review*, Vol. 104: 1–52.

DeRose, Keith (2002), Review of *Knowledge and Its Limits*, *British Journal for the Philosophy of Science*, Vol. 53, No. 4: 573–577.

DeRose, Keith (2009), *The Case for Contextualism: Knowledge, Skepticism, and Context*. Vol. 1. Oxford: Clarendon Press.

DeRose, Keith (2017), *The Appearance of Ignorance*: *Knowledge, Skepticism, and Context*. Vol. 2. Oxford: Clarendon Press.

DeRose, Keith (2017b), 'Contextualism and Fallibilism', in Jonathan Jenkins Ichikawa (ed.), *The Routledge Handbook of Epistemic Contextualism*, Oxford: Oxford University Press.

Dogramaci, Sinan (2013), 'Intuitions for Inferences', *Philosophical Studies*, Vol. 165: 371–399. DOI: 10.1007/s11098-012-9955-y

Drayson, Zoe (2014), 'The Personal/Subpersonal Distinction', *The Philosophy Compass*, Vol. 9, No. 5: 338–346.

Dretske, Fred (1969), *Seeing and Knowing*, Chicago: Chicago University Press.

Dretske, Fred (1970), 'Epistemic Operators', *Journal of Philosophy*, Vol. 67, No. 24: 1007–1023. DOI: 10.2307/2024710

Dretske, Fred (1971), 'Conclusive Reasons', *Australasian Journal of Philosophy*, Vol. 49, No. 1: 1–22. DOI: 10.1080/00048407112341001

Dretske, Fred (1981), *Knowledge and the Flow of Information*, Cambridge, MA: MIT Press.

Dretske, Fred (2006), 'Information and Closure', *Erkenntnis*, Vol. 64, No. 3: 409–413.

Dretske, Fred (2013), 'The Case Against Closure', in E. Sosa and M. Steup (eds), *Contemporary Debates in Epistemology*, Second Edition, Oxford: Blackwell. (First Edition: 2005)

Dretske, Fred (2017), 'Golden Gettier: What we (should have) learned', in R. Borges, C. De Almeida and P. Klein (eds), *Explaining Knowledge: New Essays on the Gettier Problem*, Oxford: Oxford University Press.

Dutant, Julian (2015), 'The legend of the justified true belief analysis', *Philosophical Perspectives*, Vol. 29, No. 1: 95–145. DOI: 10.1111/phpe.12061

Dutant, Julien and F. Dorsch (forthcoming), *The New Evil Demon: New Essays on Knowledge, Rationality and Justification*, Oxford: Oxford University Press.

Dutant, Julian and Clayton Littlejohn (2021), 'Defeaters as Indicators of Ignorance', in Jessica Brown and Mona Simion (eds), *Reasons, Justification, and Defeat*, Oxford: Oxford University Press.

Dutant, Julian and Fabian Dorsch (eds) (forthcoming), *The New Evil Demon*, Oxford: Oxford University Press.

Eilan, Naomi (1998), 'Perceptual Intentionality, Attention and Consciousness', *Royal Institute of Philosophy Supplements*, Vol. 43: 181–202.

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

Feeney, Megan and Susanna Schellenberg (2020), 'Bayesian Liberalism', in Katja Maria Vogt and Justin Vlasits (eds), *Epistemlogy After Sextus Empiricus*, Oxford: Oxford University Press.

Feldman, Richard (1995), 'Authoritarian Epistemology', *Philosophical Topics*, Vol. 23, No. 1: 147–169. DOI: 10.5840/philtopics199523119

Feldman, R. (1999), 'Contextualism and Skepticism', *Philosophical Perspectives*, Vol. 13: 91–114. DOI: 10.1111/0029-4624.33.s13.4

Feldman, R. (2001), 'Skeptical Problems, Contextualist Solutions', *Philosophical Studies*, Vol. 103: 61–85. DOI: 10.1023/A:1010393022562

French, Craig (2016), 'The Formulation of Epistemological Disjunctivism', *Philosophy and Phenomenological Research*, Vol. XCII, No. 1: 86–104.

French, Craig (2019), 'Epistemological Disjunctivism and its Representational Commitments', in Duncan Pritchard, Casey Doyle and Joe Milburn (eds), *New Issues in Epistemological Disjunctivism*, London: Routledge.

Friedman, Jane (2013), 'Suspended Judgement', *Philosophical Studies*, Vol. 162: 165–181.

Fumerton, R. (1985), *Metaepistemology and Skepticism*, Lanham, MD: Rowman & Littlefield.

Gersel, Johan, Rasmus Thybo Jensen, Morten S. Thaning, and Søren Overgaard (eds) (2018), *In the Light of Experience: New Essays on Perception and Reasons*, Oxford: Oxford University Press.

Ginsborg, Hannah (2006), 'Reasons for Belief', *Philosophy and Phenomenological Research*, Vol. 72, No. 2: 286–318. DOI: 10.1111/j.1933-1592.2006.tb00562.x

Ginsborg (2006b), 'Kant and the Problem of Experience', *Philosophical Topics*, Vol. 34, No. 1/2: 59–106.

Ginsborg, Hannah (2011), 'Perception, generality, and reasons', in A. Reisner and A. Steglich-Petersen (eds), *Reasons for Belief*, Cambridge: Cambridge University Press.

Ginsborg, Hannah (2018), 'Empiricism and Normative Constraint', in Johan Gersel, Rasmus Thybo Jensen, Morten S. Thaning, and Søren Overgaard (eds), *In the Light of Experience: New Essays on Perception and Reasons*, Oxford: Oxford University Press.

Glüer, Kathrin (2009), 'In Defence of a Doxastic Account of Experience', *Mind and Language*, Vol. 24, No. 3: 297–327. DOI: 10.1111/j.1468-0017.2009.01364.x

Gomes, Anil (2011), 'McDowell's Disjunctivism and Other Minds', *Inquiry*, Vol. 54, No. 3: 277–292. DOI: 10.1080/0020174X.2011.575001

Gopnik, Alison (1983), 'How We Can Know Our Minds: The Illusion of First-Person Knowledge of Intentionality', *Brain and Behavioural Science*, Vol. 16: 1–14.

Harman, Gilbert (1973), *Thought*, Princeton: Princeton University Press.

Harman, Gilbert (1986), *Change in View*, Cambridge, MA: MIT Press.

Harman, Gilbert (1990), 'The Intrinsic Quality of Experience', *Philosophical Perspectives*, Vol. 4: 31–52. DOI: 10.2307/2214186

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

Hawthorne, John (2013), 'The Case for Closure', in E. Sosa and M. Steup (eds), *Contemporary Debates in Epistemology*, Second Edition, Oxford: Blackwell. (First Edition: 2005)

Hawthorne, John and Ofra Magidor (2018), 'Reflections on the Ideology of Reasons', in Daniel Star (ed.), *The Oxford Handbook of Reasons and Normativity*, Oxford: Oxford University Press.

Heck, Richard Kimberly (2000), 'Nonconceptual Content and the "Space of Reasons"', *The Philosophical Review*, Vol. 109, No. 4: 483–523. Originally published under the name 'Richard G. Heck, Jr'.

Hinton, J. M. (1973), *Experiences*, Oxford: Clarendon Press.

Hornsby, Jennifer (2008), 'A Disjunctive Conception of Acting for Reasons', in Adrian Haddock and Fiona Macpherson (eds), *Disjunctvism: Perception, Action, Knowledge*, Oxford: Oxford University Press.

Huemer, Michael (2007), 'Compassionate Phenomenal Conservatism', *Philosophy and Phenomenological Research*, Vol. 74, No. 1: 30–54. DOI: 10.1111/j.1933-1592.2007.00002.x

Hutto, Daniel D. And Erik Myin (2012), *Radicalizing Enactivism: Basic Minds Without Content*, Cambridge, MA: MIT Press.

Hyman, John (1999), 'How Knowledge Works', *Philosophical Quarterly*, Vol. 49, No. 197: 433–451. DOI: 10.1111/1467-9213.00152

Johnson-Laird, P. N. (2006), *How We Reason*, Oxford: Oxford University Press.

Johnson-Laird, P. N. (2012), 'Inference with Mental Models', in Keith J. Holyoke and Robert G. Morrison (eds), *The Oxford Handbook of Thinking and Reasoning*, Oxford: Oxford University Press.

Johnston, Mark (2006), 'Better than Mere Knowledge? The Function of Sensory Awareness', in Tamar Szabo Gender and John Hawthorne (eds), *Perceptual Experience*, Oxford: Oxford University Press.

Kalderon, Mark Eli (2011), 'Before the Law', *Philosophical Issues*, Vol. 21: 219–244.

Kelly, Thomas (2005), 'The Epistemic Significance of Disagreement', in T. Gendler and J. Hawthorne (eds), *Oxford Studies in Epistemology*, Vol. 1, Oxford: Oxford University Press.

Kripke, Saul (2011), 'On Two Paradoxes of Knowledge', in *Philosophical Troubles: Collected Papers*, Volume 1, New York: Oxford University Press.

Kolodny, Niko (2005), 'Why Be Rational?', *Mind*, Vol. 114, No. 455: 509–563.

Kornblith, H. (2000), 'The Contextualist Evasion of Epistemology', *Philosophical Issues*, Vol. 10, No. 1: 24–32. DOI: 10.1111/j.1758-2237.2000.tb00004.x

Kotzen, M. (2019), 'A Formal Account of Epistemic Defeat', in Branden Fitelson, Rodrigo Borges and Cherie Braden (eds), *Themes from Klein: Knowledge, Scepticism and Justification*, Synthese Library, Vol. 404, Cham, Switzerland: Springer.

Lasonen-Aarnio, Maria (2010), 'Unreasonable Knowledge', *Philosophical Perspectives*, Vol. 24, No. 1: 1–21. DOI: 10.1111/j.1520-8583.2010.00183.x

Lasonen-Aarnio, Maria (2014a), 'Higher-order evidence and the limits of defeat', *Philosophy and Phenomenological Research*, Vol. 88, No. 2: 314–345.

Lasonen-Aarnio, Maria (2014b), 'The Dogmatism Puzzle', *Australasian Journal of Philosophy*, Vol. 92, No. 3: 417–432. DOI: 10.1080/00048402.2013.834949

Lasonen-Aarnio, Maria (2015), 'New rational reflection and internalism about rationality', *Oxford Studies in Epistemology*, Vol. 5: 145–171.

Lasonen-Aarnio, Maria (2020), 'Enkrasia or Evidentialism? Learning to Love Mismatch', *Philosophical Studies*, Vol. 177: 597–632. DOI: https://doi.org/10.1007/s11098-018-1196-2

Lasonen-Aarnio, Maria (2021), 'Dispositional Evaluations and Defeat', in Jessica Brown and Mona Simion (eds), *Reasons, Justification, and Defeat*, Oxford: Oxford University Press.

Le Morvan, Pierre (2017), 'Knowledge Before Gettier', *British Journal for the History of Philosophy*, Vol. 25, No. 6: 1216–1238. DOI: 10.1080/09608788.2017.1320968

Leplin, Jarrett (2009), *A Theory of Epistemic Justification*, *Philosophical Studies Series 112*, Dordrecht: Springer Netherlands. DOI: 10.1007/978-1-4020-9567-2

Lewis, David (1996), 'Elusive Knowledge', *Australasian Journal of Philosophy*, Vol.74, No. 4: 549–567. DOI: 10.1080/00048409612347521

Littlejohn, Clayton (2012), *Justification and the Truth-Connection*, Cambridge: Cambridge University Press.

Littlejohn, Clayton (2013), 'The Russellian Retreat', *Proceedings of the Aristotelian Society*, New Series, Vol. 113: 293–320. DOI: 10.1111/j.1467-9264.2013.00356.x

Littlejohn, Clayton (2017), 'How and Why Knowledge is First', in J. Adam Carter, Emma C. Gordon, and Benjamin Jarvis, *Knowledge First: Approaches in Epistemology and Mind*, Oxford: Oxford University Press.

Littlejohn, Clayton (2018), 'Evidence and its Limits', in Conor McHugh, Jonathan Way and Daniel Whiting (eds), *Normativity: Epistemic and Practical*, Oxford: Oxford University Press.

Littlejohn, Clayton (2020), 'Do you see what I know? On reasons, perceptual evidence, and epistemic status', *Philosophical Issues*, Vol. 30, No. 1: 205–220. DOI: 10.1111/phis.12181

Littlejohn, Clayton (2021), 'On what we should believe (and when (and why) we should believe what we know we should not believe)', in Kevin McCain and Scott Stapleford (eds), *Epistemic Duties: New Arguments, New Angles*, New York: Routledge.

Logue, Heather (2018), 'World in Mind: Extending Phenomenal Character and Resisting Skepticism', in Johan Gersel, Rasmus Thybo Jensen, Morten S. Thaning, and Søren Overgaard (eds), *In the Light of Experience: New Essays on Perception and Reasons*, Oxford: Oxford University Press.

Lycan, William (2001), 'Moore against the new skeptics', *Philosophical Studies*, 103: 35–53.

Lycan, William (2006), 'On the Gettier Problem problem', in S. Hetherington (ed.), *Epistemology Futures*, Oxford: Oxford University Press.

Markie, Peter (2005), 'Easy Knowledge', *Philosophy and Phenomenological Research*, Vol. 70.

Martin, M. G. F. (1997), 'The Reality of Appearances', in M. Sainsbury (ed.), *Thought and Ontology*, Milan: FrancoAngeli. Reprinted in Alex Byrne and Heather Logue (eds), *Disjunctivism: Contemporary Readings*, Cambridge, MA: MIT Press, 2008.

Martin, M. G. F. (2000), 'Beyond Dispute: Sense-Data, Intentionality and the Mind-Body Problem', in Tim Crane and Sarah Patterson (eds), *History of the Mind-Body Problem*, London: Routledge.

Martin, M. G. F. (2002a), 'The Transparency of Experience', *Mind & Language*, Vol. 17, No. 4: 376–425. DOI: 10.1111/1468-0017.00205

Martin, M. G. F. (2002b), 'Particular Thoughts and Singular Thought', in Anthony O'Hear (ed.), *Logic, Thought and Language*, Cambridge: Cambridge University Press.

Martin, M. G. F. (2004), 'The Limits of Self-Awareness', *Philosophical Studies*, Vol. 120, No. 1: 37–89.

Martin, M. G. F. (2006), 'On Being Alienated', in Tamar Szabo Gendler and John Hawthorne (eds), *Perceptual Experience*, Oxford: Oxford University Press.

Martin, M. G. F. (2010), 'What's In a Look?' in Bence Nanay (ed.), *Perceiving the World*, Oxford: Oxford University Press. Pp. 160–225.

Martin, M. G. F. (2020), 'Variation and Change in Appearances', in Katja Maria Vogt and Justin Vlasits (eds), *Epistemlogy After Sextus Empiricus*, Oxford: Oxford University Press.

McCulloch, G. (1993), 'The very idea of the phenomenological', *Proceedings of the Aristotelian Society*, Vol. 93: 39–58. DOI: 10.1093/aristotelian/93.1.39

McCulloch, Gregory (2002), 'Phenomenological Externalism', in Nicholas H. Smith (ed.), *Reading McDowell: On Mind and World*, London: Routledge.

McDowell, John (1978), 'Are Moral Requirements Hypothetical Imperatives?', *Proceedings of the Aristotelian Society, Supplementary Volume*, Vol. 52: 13–29.

McDowell, John (1982), 'Criteria, Defeasibility, and Knowledge', *Proceedings of the British Academy*, Vol. 68: 455–479. Reprinted in McDowell, *Meaning, Knowledge, and Reality*, Cambridge, MA: Harvard University Press, 1998.

McDowell, John (1986), 'Singular Thought and the Extent of Inner Space', in Philip Pettit and John McDowell (eds), *Subject, Thought, and Context*, Oxford: Clarendon Press. Reprinted in McDowell, *Meaning, Knowledge, and Reality*, Cambridge, MA: Harvard University Press, 1998.

McDowell, John (1993/1998), 'Knowledge by Hearsay', in McDowell, *Meaning, Knowledge, and Reality*, Cambridge, MA: Harvard University Press. Originally published in B. K. Matilal and A. Chakrabarti (eds), *Knowing from Words: Western and Indian Philosophical Analysis of Understanding and Testimony*, Kluwer, Dordrecht: 1993. All my page references are from the 1998 reprint.

McDowell, John (1995), 'Knowledge and the Internal', *Philosophy and Phenomenological Research*, Vol. 55, No. 4: 877–893. DOI: 10.2307/2108338

McDowell, John (1996), *Mind and World: With a New Introduction*, Cambridge, MA: Harvard University Press.

McDowell, John (1996b), 'Reply to Gibson, Byrne, and Brandom', *Philosophical Issues*, Vol. 7, Perception: 283–300. DOI: 10.2307/1522913

McDowell, John (2002a), 'Responses', in Nicholas H. Smith (ed.), *Reading McDowell on Mind and World*, London: Routledge.

McDowell, John (2002), 'Knowledge and the Internal Revisited', *Philosophy and Phenomenological Research*, Vol. 64, No. 1: 97–105.

McDowell, John (2003), 'Subjective, Intersubjective, Objective', *Philosophy and Phenomenological Research*, Vol. 67, No. 3: 675–681.

McDowell, John (2004), 'Reply to Åsa Wikforss', *Theoria*, Vol. 70: 296–297. DOI: https://doi.org/10.1111/j.1755-2567.2004.tb00997.x

McDowell, John (2006), 'Reply to Dancy', in Cynthia Macdonald and Graham Macdonald (eds), *McDowell and His Critics*, Malden, MA: Blackwell.

McDowell, John (2008), 'The Disjunctive Conception of Experience as Material for a Transcendental Argument' in Haddock and Macpherson (eds), *Disjunctivism: Perception, Action, Knowledge*, Oxford: Oxford University Press.

McDowell, John (2011), *Perception as a Capacity for Knowledge*, Milwaukee: Marquette University Press.

McDowell, John (2013), 'Perceptual Experience: Both Relational and Contentful', *European Journal of Philosophy*, Vol. 21, No. 1: 144–157. DOI: 10.1111/ejop.12005

McDowell, John (2013b), 'Acting in the Light of a Fact', in David Bakhurt, Brad Hooker, and Margaret Olivia Little (eds), *Thinking About Reasons: Themes From the Philosophy of Jonathan Dancy*, Oxford: Oxford University Press.

McDowell, John (2018), 'Perceptual Experience and Empirical Rationality', *Analytic Philosophy*, Vol. 59, No. 1: 89–98. DOI: 10.1111/phib.12120

McDowell, John (2019), 'Comments on Brewer, Gupta, and Siegel', *Philosophical Issues*, Vol. 29: 338–347. DOI: 10.1111/phis.12158

McGinn, Marie (2009), 'McDowell's Minimal Empiricism', *Philosophical Topics*, Vol. 37, No. 1: 77–94.

McGinn, Marie (2012), 'Non-Inferential Knowledge', *Proceedings of the Aristotelian Society*, Vol. 112, No. 1: 1–28. DOI: 10.1111/j.1467-9264.2012.00323.x

Miracchi, Lisa (2015), 'Competence to Know', *Philosophical Studies*, Vol. 172, No. 1: 29–56. DOI: 10.1007/s11098-014-0325-9

Moore, G. E. (1939), 'Proof of an External World', *Proceedings of the British Academy*, Vol. 25: 273–300. Reprinted in Moore (1959), *Philosophical Papers*, London: Allen and Unwin.

Moretti, Luca (2015), 'In defence of dogmatism', *Philosophical Studies*, Vol. 172: 261–282. DOI: 10.1007/s11098-014-0293-0

Nagel, Jennifer (2013), 'Knowledge as a Mental State', in Tamar Szabó Gendler and John Hawthorne (eds), *Oxford Studies in Epistemology*, Vol. 4, Oxford: Oxford University Press.

Nagel, Jennifer (2017), 'Factive and Nonfactive Mental State Attribution', *Mind & Language*, Vol. 32: 525–544.

Nagel, Jennifer (2021), 'Losing Knowledge by Thinking about Thinking', in Jessica Brown and Mona Simion (eds), *Reasons, Justification, and Defeat*, Oxford: Oxford University Press.

Nozick, Robert (1981), *Philosophical Explanations*, Cambridge, MA: Harvard University Press.

Pink, Thomas (2016), *Self-Determination: The Ethics of Action, Volume 1*, Oxford: Oxford University Press.

Pollock, John (1974), *Knowledge and Justification*, Princeton: Princeton University Press.

Pollock, John (2006), *Thinking about Acting: Logical Foundations for Rational Decision Making*, Oxford: Oxford University Press.

Pollock, John and Joseph Cruz (1999), *Contemporary Theories of Knowledge*, Totowa, NJ: Rowman and Littlefield.

Prior, A. N. (1960), 'The Runabout Inference Ticket', *Analysis*, Vol. 21, No. 2: 38–39. DOI: 10.1093/analys/21.2.38

Pritchard, Duncan (2003), 'Virtue Epistemology and Epistemic Luck', *Metaphilosophy*, Vol. 34, Nos 1-2: 106–30.

Pritchard, Duncan (2012), *Epistemological Disjunctivism*, Oxford: Oxford University Press.

Pritchard, Duncan (2012b), 'Anti-luck Virtue Epistemology', *Journal of* Philosophy, Vol. 109, No. 3: 247–279.

Pryor, James (2000), 'The Skeptic and the Dogmatist', *Noûs*, Vol. 34, No. 4: 517–549. DOI: 10.1111/0029-4624.00277

Pryor, James (2001), 'Highlights of Recent Epistemology', *The British Journal for the Philosophy of Science*, Vol. 52, No. 1: 95–124.

Pryor, James (2004), 'What's Wrong with Moore's Argument?' *Philosohpical Issues*, Vol. 14, No. 1: 349–378. DOI: 10.1111/ j.1533-6077.2004.00034.x

Pryor, James (2012), 'When Warrant Transmits', in Annalisa Coliva (ed.), *Mind, Meaning, and Knowledge: Themes from the Philosophy of Crispin Wright*, Oxford: Oxford University Press.

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

Pryor, James (2013b), 'There Is Immediate Justification', in Matthias Steup, John Turri and Ernest Sosa (eds), *Contemporary Debates in Epistemology*, 2nd Edition, Oxford: Blackwell.

Pryor, James (2018), 'The Merits of Incoherence', *Analytic Philosophy*, Vol. 59, No. 1: 112–141. DOI: 10.1111/phib.12118

Pryor, James (ms), 'Uncertainty and Undermining'. Accessed at http://www.jimpryor.net/research/papers/Uncertainty.pdf (3 November 2021). The manuscript is dated 2007.

Putnam, Hilary (1975), 'The Meaning of "Meaning"', in Keith Gunderson (ed.), *Language, Mind and Knowledge*, Minneapolis: University of Minnesota Press. Reprinted in Hilary Putnam, *Philosophical Papers*, *Volume 2: Mind, Language and Reality*, Cambridge: Cambridge University Press.

Putnam, Hilary (1981), *Reason, Truth and History*, Cambridge: Cambridge University Press.

Putnam, Hilary (1992), *Renewing Philosophy*, Cambridge, MA: Harvard University Press.

Quilty-Dunn, Jake (2015), 'Believing in Perceiving: Known Illusions and the Classical Dual-Component Theory', *Pacific Philosophical Quarterly*, Vol. 96, No. 4: 550–575. DOI: 10.1111/papq.12115

Quilty-Dunn, Jake and Eric Mandelbaum (2017), 'Inferential Transitions', *Australasian Journal of Philosophy*, Vol. 96, No. 3: 532–547. DOI: 10.1080/00048402.2017.1358754

Raz, Joseph (2011), *From Normativity to Responsibility*, Oxford: Oxford University Press.

Roessler, Johannes (2009), 'Perceptual Experience and Perceptual Knowledge', *Mind*, Vol. 118, No. 472: 1013–1041. DOI: 10.1093/mind/fzp131

Roessler, Johannes (2014), 'Reason explanation and the second-person perspective', *Philosophical Explorations*, 17:3, 346–357.

Schellenberg, Susanna (2010), 'The particularity and phenomenology of perceptual experience', *Philosophical Studies*, Vol. 149: 19–48. DOI: 10.1007/s11098-010-9540-1

Schwitzgebel, Eric (2011), *Perplexities of Consciousness*, Cambridge, MA: MIT Press.

Silins, Nicholas (2005a), 'Deception and Evidence', *Philosophical Perspectives*, Vol. 19, *Epistemology*: 375–404. DOI: 10.1111/j.1520-8583.2005.00066.x

Silins, Nicholas (2005b), 'Transmission Failure Failure', *Philosophical Studies*, Vol. 126, No. 1: 71–102.

Silins, Nicholas (2007), 'Basic Justification and the Moorean Response to the Skeptic', in John Hawthorne and Tamar Szabo Gendler (eds), *Oxford Studies in Epistemology*, Vol. 2, Oxford: Clarendon Press.

Silins, Nicholas (2014), 'The Agony of Defeat', *Philosophy and Phenomenological Research*, Vol. 88, No. 3: 505–532. DOI: 10.1111/phpr.12087

Smithies, Declan (2014), 'Can Foundationalism Solve the Regress Problem?', in Ram Neta (ed.), *Current Controversies in Epistemology*, Routledge.

Smithies, Declan (2019), *The Epistemic Role of Consciousness*, Oxford: Oxford University Press.

Snowdon, P. F. (1981), 'Perception, Vision and Causation', *Proceedings of the Aristotelian Society*, New Series, Vol. 81: 175–192.

Snowdon, P. F. (1990), 'Perception and its Objects', *Proceedings of the Aristotelian Society*, *Supplementary Volume*.

Snowdon, P. F. (2009), 'McDowell on Skepticism, Disjunctvism, and Transcendental Arguments', *Philosophical Topics*, Vol. 37, No. 1.

Sosa, Ernest, (1999), 'How to Defeat Opposition to Moore', *Philosophical Perspectives*, Vol. 13: 141–153.

Sosa, E. (2000), 'Skepticism and Contextualism', *Philosophical Issues*, Vol. 10: 1–18. DOI: 10.1111/j.1758-2237.2000.tb00002.x

Sosa, Ernest (2007), *Apt Belief and Reflective Knowledge, Volume 1: A Virtue Epistemology*, Oxford: Clarendon Press. DOI: 10.1093/acprof:oso/9780199297023.001.0001

Sosa, Ernest (2009), *Apt Belief and Reflective Knowledge, Volume 2: Reflective Knowledge*, Oxford: Oxford University Press.

Soteriou, Matthew (2000), 'The Particularity of Visual Perception', *European Journal of Philosophy*, Vol. 8, No. 2: 173–189. DOI: 10.1111/1468-0378.00107

Srinivasan, Amia (2015), 'Are We Luminous?', *Philosophy and Phenomenological Research*, Vol. XC, No. 2: 294–319. DOI: 10.1111/phpr.12067

Stampe, Dennis (1987), 'The Authority of Desire', *The Philosophical Review*, Vol. 96: 335–381.

Stanley, J. (2004), 'On the Linguistic Basis for Contextualism: Context-Sensitive Expressions', *Philosophical Studies*, Vol. 119: 119–146. DOI: 10.1023/B:PHIL.0000029353.14987.34

Stroud, Barry (1984), *The Significance of Philosophical Scepticism*, Oxford: Oxford University Press.

Stroud, Barry (1999), 'Radical Interpretation and Philosophical Scepticism', in L. E. Hahn (ed.), *The Philosophy of Donald Davidson*, Chicago: Open Court.

Stroud, Barry (2002), 'Sense-experience and the grounding of thought', in Nicholas H. Smith (ed.), *Reading McDowell: On Mind and World*, London, Routledge.

Stroud, Barry (2018), 'The Epistemological Promise of Externalism', in *Seeing, Knowing, Understanding: Philosophical Essays*, Oxford: Oxford University Press.

Stroud, Barry (2018), *Seeing, Knowing, Understanding: Philosophical Essays*, Oxford: Oxford University Press.

Sutton, Jonathan (2007), *Without Justification*, Cambridge, MA: MIT Press.

Sylvan, Kurt (2016), 'Epistemic Reasons I: Normativity', *Philosophy Compass*, Vol. 11, No. 7: 364–376.

Thune, Michael (2010), '"Partial Defeaters" and the Epistemology of Disagreement', *The Philosophical Quarterly*, Vol. 60, No. 239: 355–372. DOI: 10.1111/j.1467-9213.2009.611.x

Travis, Charles (2004), 'The Silence of the Senses', *Mind*, Vol. 113, No. 449: 57–94. DOI: 10.1093/mind/113.449.57

Travis, Charles (2005), 'A Sense of Occasion', *The Philosophical Quarterly*, Vol. 55, No. 219: 286–314. DOI: 10.1111/j.0031-8094.2005.00399.x

Travis, Charles (2007), 'Reason's Reach', *European Journal of Philosophy*, Vo. 15, No. 2: 225–248. DOI: https://doi.org/10.1111/j.1468-0378.2007.00257.x

Travis, Charles (2013), *Perception: Essays After Frege*, Oxford: Oxford University Press.

Unger, Peter (1975), *Ignorance: A Case for Scepticism*, Oxford: Clarendon Press.

Unger, Peter (1984), *Philosophical Relativity*, Minneapolis: University of Minnesota Press.

Velleman, J. David (2000/2015), *The Possibility of Practical Reason*, Ann Arbor, MI: Maize Books. Open Access DOI: http://dx.doi.org/10.3998/maize.13240734.0001.001

Vendler, Z. (1972), *Res Cogitans*, Ithaca, NY: Cornell University Press.

Vogel, J. (1990), 'Cartesian Skepticism and Inference to the Best Explanation', *Journal of Philosophy*, Vol. 87, No. 11: 658–666. DOI: 10.5840/jphil1990871123

Vogel, J. (2000), 'Reliabilism Levelled', *Journal of Philosophy*, Vol. 97, No. 11: 602–623. DOI: 10.2307/2678454

Wedgwood, Ralph (2012), 'Justified inference', *Synthese*, Vol. 189, No. 2: 273–295. DOI: 10.1007/s11229-011-0012-8

Weir, Alan (2004), 'An ultrarealist theory of perception', *International Journal of Philosophical Studies*, Vol. 12, No. 2: 105–128.

White, Roger (2006), 'Problems for Dogmatism', *Philosophical Studies*, Vol. 131: 525–557. DOI: 10.1007/s11098-004-7487-9

White, Roger (2014), 'What Is My Evidence that Here is a Hand?', in Dodd and Zardini (eds), *Scepticism and Perceptual Justification*, Oxford: Oxford University Press.

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

Williamson, Timothy (2007), 'On Being Justified in One's Head', in Mark Timmons, John Greco and Alfred Mele (eds), *Rationality and the Good: Critical Essays on the Ethic and Epistemology of Robert Audi*, Oxford: Oxford University Press.

Williamson, Timothy (2008), 'Why Epistemology Cannot be Operationalized', in Quentin Smith (ed.), *Epistemology: New Essays*, Oxford: Oxford University Press.

Williamson, Timothy (2009), 'Replies to Critics' in Patrick Greenough and Duncan Pritchard (eds), *Williamson on Knowledge*, Oxford: Oxford University Press.

Williamson, Timothy (2011), 'Improbable Knowing', in Trent Dougherty (ed.), *Evidentialism and Its Discontents*, Oxford: Oxford University Press.

Williamson, Timothy (2014), 'Very Improbable Knowing', *Erkenntnis*, Vol. 79, No. 5: 971–999. DOI: 10.1007/s10670-013-9590-9

Williamson, Timothy (2017), 'Acting on Knowledge', in J. Adam Carter, Emma Gordon, Benjamin Jarvis (eds), *Knowledge First: Approaches in Epistemology and Mind*, Oxford: Oxford University Press.

Williamson, Timothy (forthcoming), 'Justifications, Excuses, and Sceptical Scenarios' in Julien Dutant and Fabian Dorsch (eds), *The New Evil Demon*, Oxford: Oxford University Press.

Wittgenstein, Ludwig (1969), *On Certainty*, Ed. G. E. M. Anscombe and G. H. Von Wright, trans. D. Paul and G. E. M. Anscombe, Oxford: Blackwell.

Wright, Crispin (1985), 'Facts and Certainty', *Proceedings of the British Academy*, Vol. 71: 429–472.

Wright, Crispin (1991), 'Scepticism and Dreaming: Imploding the Demon', *Mind*, Vol. 100, No. 397: 87–115.

Wright, Crispin (1992), *Truth and Objectivity*, Cambridge, MA: Harvard University Press.

Wright, Crispin (2002), '(Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell', *Philosophy and Phenomenological Research*, Vol. LXV, No. 2: 330–348.

Wright, Crispin (2004a), 'On Epistemic Entitlement. I. Warrant for Nothing (And Foundations for Free)?', *Proceedings of the Aristotelian Society, Supplementary Volumes*, Vol. 78: 167–245.

Wright, Crispin (2004b), 'Scepticism, Certainty, Moore and Wittgenstein', in Max Kölbel and Bernhard Weiss (eds), *Wittgenstein's Lasting Significance*, Taylor and Francis.

Wright, Crispin (2007), 'The Perils of Dogmatism', in Nuccetelli and Seay (eds), *Themes from G. E. Moore: New Essays in Epistemology*, Oxford: Oxford University Press.

Wright, Crispin (2008), 'Comment on John McDowell's "The Disjunctive Conception of Experience as Material for a Transcendental Argument"', in Haddock and Macpherson (eds), *Disjunctivism: Perception, Action, Knowledge*, Oxford: Oxford University Press.

Wright, Crispin (2012), 'Replies Part IV: Warrant Transmission and Entitlement', in Annalisa Coliva (ed.), *Mind, Meaning, and Knowledge: Themes from the Philosophy of Crispin Wright*, Oxford: Oxford University Press.

Wright, Crispin (2014), 'On Epistemic Entitlement (II): Welfare State Epistemology', in Dodd and Zardini (eds), *Scepticism and Perceptual Justification*, Oxford: Oxford University Press.

Wright, Crispin (2014b), 'Comment on Paul Boghossian, "What is Inference"', *Philosophical Studies*, Vol. 169, No. 1: 27–37.