



THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

This work is protected by copyright and other intellectual property rights, which are retained by the thesis author, unless otherwise stated.

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author.

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

Epistemic Responsibility and Radical Scepticism

Cameron Boulton

PhD in Philosophy

Edinburgh University

2014

Dedicated to the memory of Jeff Boulton

Declaration

I have read and understood The University of Edinburgh guidelines on Plagiarism and declare that this written dissertation is all my own work except where I indicate otherwise by proper use of quotes and references.

Signature

Date & Place

Abstract

This thesis has two aims. One is to motivate the claim that challenging what I call a “sameness of evidence thesis” is a particularly promising approach to external world scepticism. The other is to sharpen an underexplored issue that arises when challenging the sameness of evidence thesis. The second aim is the primary aim of the thesis.

Pursuing the first aim, I start by examining a predominant formulation of external world scepticism known as the “closure argument” for knowledge. I examine three main strategies for responding to external world scepticism and highlight their major challenges (DeRose 1995; Dretske 1979; Nozick 1981; Sosa 1999). The goal is not to demonstrate that these challenges cannot be met, but rather to highlight a deeper issue that arises when responding to the closure problem for knowledge. In particular, I take the discussion to motivate looking at what I will call “scepticism about evidential justification” (Feldman 2000; Kornblith 2001; Pritchard *forthcoming*). The general argument in favour of a shift to scepticism about evidential justification is based on considerations about what an adequate response to external world scepticism should hope to achieve. I argue that one condition of adequacy is being able to account for radical forms of scepticism that challenge not only that our beliefs enjoy the epistemic status of knowledge (however that status is conceived) but also that our ordinary empirical beliefs are justified, or that we are reasonable in holding them. There are different varieties of scepticism about evidential justification. I focus in some detail on the anti-sceptical strategies of Pryor (2000; 2004) and Wright (2004) as examples of strategies that engage with scepticism about evidential justification. But I argue that one form of evidential scepticism known as the “underdetermination argument”—which Pryor and Wright do not directly engage with—is of particular importance.

The main assumption in the underdetermination argument I focus on is about the nature of evidence. More specifically, the underdetermination argument presupposes that one’s evidence is the same in so-called “bad” and “good” cases in which an agent forms an empirical belief. This is the “sameness of evidence thesis.” Pursuing the main aim of the thesis, I introduce two forms an anti-sceptical strategy that involves challenging the sameness of evidence thesis. The two forms I consider differ in their commitments concerning a condition of accessibility on our evidence. Pritchard (2006; 2007; 2012; *forthcoming*) maintains that one’s evidence is “reflectively accessible.” Williamson (2000; 2009) rejects this claim. The central issue I aim to

sharpen is that while accepting the condition of accessibility leads to serious challenges in rejecting the sameness of evidence thesis, rejecting it leads to counterintuitive consequences if we grant that there is a normative principle that requires us to proportion our beliefs to the evidence.

A central part of the thesis involves examining these counterintuitive consequences and showing what accounting for them requires. This is an underexplored project in the context of external world scepticism. I look at three different approaches to spelling out the counterintuitive consequences. My preferred account turns on a distinction between three different kinds of responsibility (Shoemaker 2011). I claim that there is a notion of responsibility – “attributability” – that is centrally connected to normative judgments. I argue for a “condition of accessibility” on attributability. Taken together, these two claims comprise an account of what is problematic about rejecting an access condition on our evidence. I then claim that there are two ways forward. One is to accept the condition of accessibility on our evidence that my account implies; the other is to challenge my claims about the connection between attributability and normative judgments, or the accessibility condition on attributability, or both.

Although I claim that the prospects look better for taking the second option when it comes to rejecting the sameness of evidence thesis, drawing on recent work from Gibbons (2006; 2013) and Daniel Greco (2013), I argue that the first option is still a live possibility. The main aim in this part of the thesis is not to decide what the best way of rejecting the sameness of evidence thesis is, but rather to examine the challenges that arise when we reject it in one way or another. The question of what sort of access we have to our normative requirements is the focus of an increasingly sophisticated discussion in contemporary epistemology. An important upshot of this thesis is that it brings the problem of external world scepticism directly within the scope of that debate.

Table of Contents

Abstract	i
Table of Contents	iii
Acknowledgements	v
Introduction	1
0.1 Introduction	1
0.2 Two Kinds of Internalism/Externalism	2
0.3 Normative Requirements	4
0.4 Evidence, Justification, and Knowledge	5
0.5 Chapter Outline	6
Chapter 1: Formulating External World Scepticism	10
1.1 Introduction	10
1.2 The Closure Argument	10
1.3 Dretske and the Denial of Closure	13
1.4 Nozick and the Denial of Closure	16
1.5 Premise Two	20
1.6 DeRose and Semantic Contextualism	23
1.7 Scepticism and Standards	27
1.8 Responding to Scepticism About Evidential Justification	31
1.9 Conclusion	34
Chapter 2: Entitlement and Immediate Justification	35
2.1 Introduction	35
2.2 Wright and Pryor on Scepticism	35
2.3 Entitlement of Cognitive Project	38
2.4 Relations of Support and the “Leaching Problem”	42
2.5 From Entitlement to Immediate Justification	48
2.6 Immediate Justification	51
2.7 Locating the Perceptually Basic and the Problem of Easy Knowledge	53

2.8	The Liberal/Conservative Cycle	57
2.9	Conclusion	58
Chapter 3: Sameness of Evidence		60
3.1	Introduction	60
3.2	McDowellian Neo-Mooreanism	60
3.3	Epistemological Disjunctivism	63
3.4	The Highest Common Factor Argument and the Explanatory Challenge	67
3.5	Williamson and Evidential Externalism	74
3.6	Same Evidence or Same Justification?	76
3.7	Justification and Rational Requirements	78
3.8	Conclusion	81
Chapter 4: The Natural Reaction		82
4.1	Introduction	82
4.2	Deliberation-Guidance	82
4.3	Normative Requirements and Blameworthiness	85
4.4	Normative Requirements and Responsibility	90
4.5	Attributability and Access	96
4.6	Further Reflections	101
4.7	Conclusion	103
Chapter 5: Considering the Options		105
5.1	Introduction	105
5.2	Anti-Luminosity	105
5.3	Access Externalism	108
5.4	Contextualism and Luminosity	116
5.5	Sameness of Evidence and the Closure Problem	120
5.6	Conclusion	128
Conclusion		129
Bibliography		133

Acknowledgements

I would like to thank the following people at Edinburgh University who have helped me along the way: Natalie Ashton, Tony Bolos, Evan Butts, Jamie Collin, Tim Kunke, Jie Gao, Jesper Kallestrup, Scott Kerzman, Shakira Lehmann, Alisa Mandrigin, Michela Massimi, Orestis Palermos, Fotini Panteledes, Chris Ranalli, Mike Ridge, Stephen Ryan, Kyle Scott, Nick Treanor, Lani Watson, Eusebio Waweru, Alan Wilson, Lee Whittington, and Sam Wilkinson. I would also like to thank Olle Blomberg, Dave Ward, and Shayne Ryan for weekly games of squash and the shared experience of venting our academic angst on the squash ball. Sebastian Köhler, Robin McKenna, Joey Pollock, and Ashley Taylor have been a particular source of inspiration and encouragement. I owe a special debt of gratitude to Sebastian—for conversations, advice, feedback, and simply being a good friend who also happens to be an excellent philosopher.

I would like to thank my supervisors Duncan Pritchard, Allan Hazlett, and Matthew Chrisman for all of their feedback and encouragement. They set an inspiring example and I am grateful to have had the opportunity to work closely with them. Duncan and Allan have provided me with extensive feedback on various parts of this thesis. Duncan has been a wonderful source of perspective and encouragement. The incisive and voluminous feedback that I have received from Allan over the years has improved my ideas as well as my appreciation of just how difficult philosophy really is. Matthew has also provided me with extremely helpful feedback and has generally been an inspiration. I am yet to meet another philosopher with quite the combined degree of charity and precision that characterizes Matthew's work and his approach to supervision.

Without the support of a PPLS Career Development Studentship, and a Global Overseas Research Scholarship from the College of Humanities and Social Sciences at Edinburgh University, my PhD studies would not have been possible.

I owe a huge debt of gratitude to my mother and step-father who have been there whenever I've needed them, and who have believed in me from day one. There are many other friends in Canada who have supported me from a distance, in particular Jessica Evans, Samantha Halas, Matt Johnston, Ryan Klatt, Sam Koulack, Beth Savickey, and especially Dave and Diane Johnston.

Epistemic Responsibility and Radical Scepticism: Introduction

0.1 Introduction

The primary aim of this thesis is to sharpen an issue that arises when responding to a central argument for external world scepticism. The central argument I focus on is called the “underdetermination argument”. There are many sceptical arguments about our knowledge of the external world. Part of this thesis examines the relationships between some predominant arguments and explains why some are more interesting than others. I will argue that the underdetermination argument is a particularly important and interesting argument. The main assumption in the underdetermination argument that this thesis focuses on is about the nature of evidence; more specifically, the underdetermination argument presupposes that one’s evidence is the same in so-called “bad” and “good” cases in which an agent forms an empirical belief. I call this the “sameness of evidence thesis”. I aim to sharpen an issue that arises when challenging the sameness of evidence thesis.

It seems true that people ought to proportion their beliefs to the evidence. For example, if you’ve got perfectly good evidence that there is an apple on the table, it’s natural to say that your beliefs about whether there is an apple on the table in some sense ought to reflect the fact that you’ve got such good evidence. It’s a complex matter how to spell out the details of such a claim, but in this thesis I’m interested in a very general idea. I’m interested in the idea that this claim reflects a normative principle about the appropriateness of forming beliefs under certain circumstances (and I explain more about what I mean by that below). I think that any adequate challenge to the sameness of evidence thesis must also account for the idea that people are normatively required to proportion their beliefs to the evidence.¹

There is a tension in rejecting the sameness of evidence thesis and accounting for this normative principle. The tension arises out of a strong temptation to be an externalist about evidence in rejecting the sameness of evidence thesis. Externalism about evidence has consequences for conditions of access that come with being normatively required to \emptyset . I claim that there is a “natural reaction” to these consequences (I borrow this phrase from Gibbons (2013)). I aim to articulate in some detail what this reaction amounts to. I also aim to explain

¹ That there is some such normative evidential principle is something that I assume in this thesis, along with the main parties to the debate that I’m engaging with.

why the externalist's consequences are so hard to avoid. A prominent strategy for rejecting the sameness of evidence thesis which would avoid the externalist's consequences is to endorse a condition of "reflective accessibility" on our evidence (Pritchard 2006; 2007; 2012; *forthcoming*). A key challenge in rejecting the sameness of evidence thesis while endorsing a condition of reflective accessibility on our evidence is being able to explain how our evidence could be different in phenomenologically indistinguishable good and bad cases. I call this the "explanatory challenge". It can easily look like the only way to meet the explanatory challenge is by endorsing evidential externalism, or at least by significantly weakening the condition of access on our evidence (perhaps so much so that the motivation for preserving whatever condition of access that remains is undermined). This is why it is important to explore the counterintuitive consequences of evidential externalism. My aim is to argue that an important desideratum for an anti-sceptical strategy that rejects the sameness of evidence thesis is to account for the natural reaction to the consequences of evidential externalism.

There are of course many ways of responding to scepticism and I explore a variety of proposals in this thesis. My discussion of the various anti-sceptical strategies I consider aims to motivate the idea that challenging the sameness of evidence thesis is a particularly interesting and promising approach. In any case, I would like to emphasise that my aim is not to solve the problem of external world scepticism. Rather this thesis is about sharpening what I think is a somewhat new, or at least underexplored issue that arises when we try to solve the problem. The question of what sort of access we have to our normative requirements is the focus of an increasingly sophisticated discussion in contemporary epistemology. An interesting upshot of this discussion is that the sceptical problem can be seen as falling within the scope of that debate. Perhaps if we can see the problem in this light, those epistemologists who know much more about this issue than I do will have some fruitful things to say about the problem. Before getting started, the remainder of this introduction addresses some preliminaries and then provides a chapter by chapter outline of how things will proceed.

0.2 Two Kinds of Internalism/Externalism

When I use the terms "internalism" and "externalism" I mean internalism and externalism about evidence and justification (unless I say otherwise). I endorse a distinction between two ways of spelling out the difference between internalism and externalism. One turns on considerations about "access". Until we get to Chapter 5, this will be the main thing that I

mean by “internalism”. This is the view that is commonly known as “access internalism” (Pryor 2001). Access internalism claims that one has a special kind of access to the facts about whether they are justified. One way of putting this is to say that, according to access internalism, if one is justified in believing that *p* then one is in a position to know that one is justified in believing that *p*. The idea of “being in a position to know” is admittedly vague. I will largely rely on an intuitive understanding of the idea. But in the context of his own discussion of accessibility, Williamson offers a helpful articulation of what it is to be in a position to know:

To be in a position to know, it is neither necessary to know *p* nor sufficient to be physically or psychologically capable of knowing *p*. No obstacle must block one’s path to knowing *p*. If one is in a position to know *p*, and one has done what one is in a position to do to decide whether *p* is true, then one does know *p*. The fact is open to view, unhidden, even if one does not yet see it (Williamson 2000, p.95).

An important distinction will arise in Chapter 3 between “strong” and “weak” ways of spelling this kind of special access out. Schematically, “strong” access says that one is in a position to know that condition *C* obtains when condition *C* obtains *and* that one is in a position to know that condition *C* does not obtain when it does not obtain. “Weak” access says merely that one is in a position to know that condition *C* obtains when condition *C* obtains. My default notion of the access condition in this thesis will be the strong version of access (this is because, as I discuss in Chapter 3, it’s not clear why we would endorse the asymmetry involved in denying “strong” but holding on to “weak”). But as I said, the distinction will come up in Chapter 3.

A different way of thinking about internalism and externalism—a way that will be important in Chapter 5—turns on considerations about the supervenience base of the determinants of justification. That is, internalism in this sense says that the determinants of justification supervene on facts that are in some sense internal to the agent—usually understood as reflectively accessible or introspectible facts. I’ll call internalism of this sort “supervenience internalism”. It’s important to note that while these two ways of thinking about internalism may look very similar, they are different. Supervenience internalism is plausibly *motivated* by access internalism in the sense that, if we are always in a position to know whether we are justified in believing that *p*, then it seems like the determinants of justification must supervene on facts that we are plausibly in a position to know about whenever they obtain. But one does not have to be an access internalist just because one is a supervenience internalist. One can

maintain that the supervenience base of justification supervenes on the internal, or the mental, or the reflectively accessible or introspectible, but deny that we are always in a position to know that we are justified in believing that p (Conee and Feldman 2004; Pryor 2001). That said, one might find it difficult to understand the motivation for such a view once the condition of special access to justification is denied. In any case, it's also possible to endorse a condition of special access but deny that the determinants of justification supervene on the internal, or the mental, or the reflectively accessible. In Chapter 5 I'll consider two approaches in the literature that maintain a special condition of access on justification, but deny an internal supervenience base.

0.3 Normative Requirements

The notion of “normative requirements” (and equally “normative judgments”) is difficult to clearly and informatively explain. The most important place in which I appeal to the idea of a normative requirement (and a normative judgment) is the discussion of evidential externalism (Chapters 3, 4, and 5). Because Williamson—my paradigm evidential externalist—himself appeals to the notion of normative requirements, and because I argue that according to everything that Williamson grants about this notion, his view about evidence is problematic, in one sense I think it is open to me to say that I mean what Williamson means when he talks about normativity. But of course it's important to say a bit more than that.

What I mean by normative requirements (and judgments) can be captured in deontological terms. When someone is normatively required to \emptyset , I take this to mean that they are obligated, or permitted, or have a duty or some similar such notion to \emptyset . This obligation or duty or permission is inescapable in Foot's sense (Foot 1972). It is not conditioned on one's desires or goals.² A normative judgment is just the judgment that an agent, or attitude, or action is subject to a normative requirement. In Chapter 4 I'll draw an important connection between normative requirements and a particular kind of responsibility that I call “attributability”. I think one way of spelling out what I mean when I say that a person is normatively required to

² It's plausible that there's more to the normative than inescapability. In Foot's sense, norms of etiquette are inescapable, but norms of etiquette are not plausibly normative. Some capture the feature of normativity I'm getting at by saying that normative judgments are categorical. The categorical is a sub-species of the inescapable. Inescapable judgements *apply* to you regardless of your desires or goals, but you may rationally believe that they apply and believe that they do not give you any reason to \emptyset . Categorical judgements, on the other hand, not only apply to you regardless of your desires or goals, but you cannot also rationally believe that they apply to you and believe that they do not give you a reason to \emptyset .

\emptyset is that the person is *responsible* for \emptyset -ing in a particular kind of way (which I'll explain). Judgments about responsibility are intuitive. We hold each other responsible all the time. That's one reason why I think it's helpful to draw this connection. But I'll be saying a lot more about this connection in Chapter 4, in particular in the context of a problem I think it creates for Williamson's view about evidence.

I should also note that when I use the term "norm" I follow Srinivasan's understanding of this term: "A norm is a universal generalisation about how one ought to conduct one's practical or doxastic affairs, involving a normative state N and a triggering condition C, of the form "N if/and/only if C" (Srinivasan *unpublished*, p.6). I understand "triggering condition" as meaning something like the condition under which a given normative requirement applies.

0.4 Evidence, Justification, and Knowledge

There are many different ways of thinking about the relations—or lack thereof—between evidence, justification, and knowledge. Whether or not evidence is necessary for justification and whether or not justification is necessary for knowledge, one thing I will argue in this thesis is that it looks very much like the sceptic has to assume both of these things to get her problem off the ground (although, as I explain in Chapter One, many think that the sceptic's claims about justification, as opposed to claims about knowledge, are in fact the most interesting sceptical claims). That the sceptic depends on these connections may not be immediately obvious. One task of Chapter 1 is to argue that this is the case. Insofar as the sceptical problem looks interesting, intuitive, and compelling, I think these are important claims. But this understanding of the relationship between evidence, justification and knowledge is also independently plausible.

In any case, the idea is not claim that this is the only or even the best way of understanding the relationship between evidence, justification, and knowledge. Rather, the idea is to grant the *prima facie* plausibility of these connections and engage with the sceptic on that basis. A guiding idea behind this thesis is that we learn lessons about the assumptions at work in the sceptical problem by looking at various responses to the problem, and seeing what seems to work and what does not. These assumptions are taken on board to the extent that the project aims to show what is required of an anti-sceptical strategy that engages as closely as possible with what motivates the sceptical problem.

0.5 Chapter Outline

The following is a detailed summary of how the thesis will proceed. In Chapter 1 I examine the best way to formulate of the problem of our knowledge of the external world. Much recent interest in the problem is due to a formulation known as the “closure problem” for knowledge. I examine some important responses to the closure problem and some of their challenges. My aim is not to demonstrate that these challenges cannot be met. Rather, my aim is to expose a deeper issue that I think arises when responding to the closure-based problem for knowledge. In particular, I take the discussion to motivate looking at what I will call “scepticism about evidential justification”. There are different varieties of scepticism about evidential justification and I will introduce as an example the form that I think is of particular importance, namely the underdetermination argument. But the general argument in favour of a shift to scepticism about evidential justification is based on considerations about what an adequate response to external world scepticism should hope to achieve. I argue that one condition of adequacy is being able to account for radical forms of scepticism which challenge not only that our beliefs enjoy the epistemic status of knowledge (however that status is conceived) but also that our ordinary empirical beliefs are justified or that we are reasonable in holding them.

In Chapter 2 I examine two predominant anti-sceptical strategies that engage with scepticism about evidential justification. In particular, these strategies engage with a sceptical argument that makes central use of a claim that we need antecedent justification to believe that some radical sceptical scenario does not obtain in order to have justification to believe ordinary empirical propositions. Wright’s “entitlement” strategy (2004) and Pryor’s “immediate justification” strategy (2000; 2004) make opposing claims about the antecedent epistemic support we require for certain presuppositions of our ordinary empirical beliefs in order to be justified in holding those beliefs. Wright claims that we require antecedent support for the relevant presuppositions, but that the sort of support we have is a kind of non-evidential warrant that he calls “entitlement”. He claims that this background of non-evidential warrant explains how we are justified in believing ordinary empirical propositions. Pryor denies that we require antecedent support for the relevant presuppositions of our ordinary empirical beliefs, but claims rather that we can have a kind of non-evidential “immediate justification” to believe certain ordinary empirical propositions. I argue that a closer examination of how Wright’s strategy avoids a problem to do with a tension he seems to introduce in our web of belief

makes his view look closer to Pryor's view than we might have originally thought. I claim that this motivates looking at Pryor's view. But I argue that the problems faced by Pryor's view—in particular the “problem of easy knowledge”—leads us back in the direction of Wright. In short I argue that the route of Wright in Pryor is a kind of cycle. I claim that in light of this, and in light of further considerations about the relationship between the sceptical argument they engage with and the underdetermination argument, we should turn our attention to the underdetermination argument, and more specifically the sameness of evidence thesis.

In Chapter 3 I turn to two broad possible ways of rejecting the sameness of evidence thesis, due to Pritchard (2006; 2007; 2012; *forthcoming*) and Williamson (2000; 2009). These strategies differ in their commitments to a condition of accessibility on our evidence. In particular, Pritchard maintains that our evidence is “reflectively accessible” and argues nevertheless that in the good case we have evidence in the form of *factive* mental states, while in the bad case we do not. Thus, our evidence in the good and bad cases is different (though the good and bad cases are phenomenologically indistinguishable), and so the sameness of evidence thesis is false. This motivates a rejection of a key premise in the underdetermination argument. I argue that Pritchard faces an explanatory challenge in rejecting the sameness of evidence thesis. This is the challenge of explaining how our evidence could be different in the good and bad cases while being phenomenologically indistinguishable, even though our evidence is reflectively accessible. I argue that a move that Pritchard may want to make in this context—endorsing merely a “weak” accessibility condition—undermines a central motivation for his view. Moreover, it seems to entail the very consequences that I will go on to discuss in the context of evidential externalism. A view that would have a comparatively easy time explaining how our evidence is different in the good and bad cases, even though the good and bad cases are phenomenologically indistinguishable, simply gives up on a condition of accessibility (both “weak” and “strong”). We can sometimes fail to be in a position to know what our evidence is. This is evidential externalism. I focus on Williamson as a paradigm form of evidential externalism. While this view has an easier time rejecting the sameness of evidence thesis, I argue that it has an important counterintuitive consequence. The consequence arises when we consider the normative principle that we ought to proportion our beliefs to the evidence. Williamson's view entails that we can be normatively required to \emptyset even though we are not in a position to know that we are normatively required to \emptyset . Although I claim there is an initial

intuitive reaction to this idea (which Williamson himself grants) I claim that it is not easy to spell out what this reaction amounts to.

In Chapter Four I look at three different approaches to spelling out what I call the “natural reaction” to Williamson’s view about evidence. I call the first approach the “guidance-deliberation” worry. I claim that this is not a good way of spelling out the natural reaction because it does not seem to entail that there is a problem with the sort of access to our evidence that Williamson allows we do have. The second approach turns on the idea that there is an essential connection between what Srinivasan (*unpublished*) calls the “deontological perspective”—the perspective from which we make judgments of whether an agent is obligated or required or has a duty to \emptyset —and the “hypological perspective”—the perspective from which we judge whether an agent is blameworthy for \emptyset -ing. According to this approach, judgments of blameworthiness entail a condition of accessibility on the conditions under which an agent is blameworthy, and since the hypological and deontological perspectives are essentially linked, Williamson’s view does not account for the kind of accessibility we ought to have to normative requirements concerning our evidence. I argue that this approach is not a promising way to go because it simply does not seem plausible that there is such an essential connection between these two perspectives. The final approach I consider is closely related but turns on a distinction between different kinds of responsibility (Shoemaker 2011). I note that there is a notion of responsibility—“attributability”—which comes apart from judgements of blameworthiness, but which may be centrally connected to normative judgments and imply a condition of access that conflicts with Williamson’s view about evidence. In particular, the upshot of the foregoing is the claim that we should always be in a position to know what our evidence is. The main aim of Chapter Four is to provide a better understanding of what the natural reaction amounts to and what accounting for it would involve.

In Chapter 5 I consider some options moving forward. The chapter begins by looking briefly at Williamson’s claim that an internalist access condition on our evidence is a “quaint relic of Cartesian epistemology” and rests on the mistaken assumption that there is such a thing as a “cognitive home”—a distinctive realm in which nothing is hidden and in which our judgments about what to do and what to think are always safe from the possibility of non-culpable error (Williamson 2000). The idea that we have a cognitive home is not a popular commitment in contemporary epistemology. Moreover, a commitment to the idea that we have a cognitive

home seems to amount to a commitment to accessibilist supervenience internalism. This introduces a potential in-principle reason to think that endorsing an internalist access condition and rejecting the sameness of evidence thesis is impossible. I explore the theoretical options for arguing that the idea that we are always in a position to know what our evidence is does not commit us to accessibilist supervenience internalism. I draw on work from Gibbons (2006; 2013) and D. Greco (2013). At this stage I note that I am still no further along in understanding what the best way to reject the sameness of evidence thesis is. In light of the difficulties raised so far for such a strategy, I conclude the chapter by taking a step back and considering whether the strategy is worth pursuing after all. I argue that it is. I conclude the thesis by taking stock and summarizing where we are when it comes to the problem of external world scepticism.

Chapter 1: Formulating External World Scepticism

1.1 Introduction

Much recent interest in external world scepticism is due to its formulation as a closure problem for knowledge. The problem exploits an intuitive relationship between our apparent lack of knowledge of the falsity of radical sceptical hypotheses, such as that one is a handless brain in a vat (BIV), and our knowledge of ordinary empirical propositions, such as that we have hands. In this chapter I examine some important responses to the closure problem and some of their challenges. My aim is not to demonstrate that these challenges cannot be met. Rather, my aim is to expose a deeper issue that I think arises when responding to the closure-based problem for knowledge. In particular, I take the discussion to motivate looking at what I'll call scepticism about evidential justification, a form of scepticism that embodies the issues left over once we've gone through the standard ways of approaching the closure-based problem for knowledge. I use the underdetermination argument as a paradigm example.

1.2 The Closure Argument

The closure-based sceptical argument goes as follows.

- 1) If S knows that he has hands then S knows he is not a handless BIV.
- 2) S doesn't know he is not a handless BIV.
- 3) Therefore, S doesn't know he has hands.

The argument is meant to establish a radical conclusion given that we can replace "S has hands" with just about any proposition about the external world, knowledge of which is supposedly incompatible with a lack of knowledge that one is not an envatted brain. The argument is often referred to as a paradox. It begins with apparently true premises and establishes, by an apparently valid rule of inference, an apparently false conclusion. Along with its highly implausible conclusion, the simplicity and seemingly mundane plausibility of the premises of this argument is precisely what is interesting about it. With closure-based scepticism, the notion that we can solve the sceptical problem simply by adjusting whatever conception of knowledge it depends on—think of scepticism as an obstacle merely to an

outmoded quest for Cartesian certainty—has given way to a more threatening concern. As DeRose puts it, the above sceptical argument (which he calls the “argument from ignorance” (AI)) “initially seems to threaten the truth of our ordinary claims—it threatens to boldly show that we’ve been wrong all along in thinking and saying that we know this and that. For it doesn’t seem as if it’s just in some ‘high’ or ‘philosophical’ sense that AI’s premises are true: they seem true in the ordinary sense of ‘know’” (DeRose 1995, p.5). If the premises of the sceptical argument appear true by ordinary standards, dealing with the closure-based problem becomes an exercise in understanding the nature of *our* conception of knowledge, not merely an uninteresting or optional “philosophical” conception of knowledge. To better see the supposedly paradoxical nature of the sceptical argument, let’s look a bit more carefully at its premises.

Premise (1) is said to trade on a highly intuitive and plausible epistemic principle, the closure principle. It’s not entirely clear how this principle should be formulated. For example, consider the following implausible principle:

(CK1) If S knows p, and p entails q, then S knows q.

Surely the sceptic is not demanding that we must know all propositions that are logically entailed by what we know. Indeed, we fail to *believe* most of the logical consequences of what we know (and we sometimes believe the negations of some of those consequences). Perhaps the sceptic merely requires that we know all propositions which we *know* to be entailed by what we know. A corresponding closure principle would go roughly like this:

(CK2) For all S, p, q, if S knows that p, and S knows that p entails q, then S knows that q.

There are countless examples from everyday life in which this closure principle (CK2) appears to be at play. For one, imagine a detective who reasons: “I know that the butler was in the garden at the time of the crime, and I know that *if* the butler was in the garden at that time then the butler was not at the scene of the crime (the kitchen); so I know that the butler was not at the scene of the crime”. The sceptic appears licensed to make her inference in precisely the same sort of way that the detective is licensed to make his. If I know that I have hands, and I know that this entails that I am not a handless BIV, then I must know that I am not a

handless BIV. If not, the sceptic wants to know what distinguishes this inference from the detective's.

However, it is a matter of some controversy whether (CK2) really is as plausible as I've suggested. For example, it doesn't account for the possibility that one could know p and know that p entails q , but know q on totally different grounds than their knowledge that p and that p entails q . It may be that an even weaker but closely related formulation of closure is the really intuitive one:

(CK3) If S knows p , and p entails q , and S competently deduces q from p on the basis of knowing that p entails q , then S knows q .

This principle seems to capture what's really intuitive about the relevant idea in the neighbourhood, which is that we can extend our knowledge by competent deductions.³ And we can understand (1) as intuitively motivated by this principle. When it comes to strategies that deny closure in an effort to undercut the motivation for (1) (which I turn to shortly), it is important that such strategies are capable of rejecting the principle that's really at play in the sceptical problem, as opposed to a closely related but ultimately irrelevant principle. That said, regardless of whether the strategies I consider below challenge the relevant principle, my discussion of these strategies aims to show that even by their own understanding of what the relevant principle is, they face serious issues.

Premise (2) of the closure-based argument is "S does not know that he is not a BIV". The sceptic of course needs this claim in order to deny, via closure, that S has all sorts of everyday knowledge. But what, if anything, makes this claim plausible? Interestingly, much of the literature seems to take premise (2) as primitively plausible (Brueckner 2010, p.367). Or at least, as DeRose has noted, explanations as to why we lack knowledge of the denials of sceptical hypotheses are given in a "circle of all-too-closely related terms of epistemic appraisal" (DeRose 1995, p.16). For example, Dretske claims that you don't know you're not a BIV

³ According to Schiffer, "[A] strand to our concept of knowledge deems one to know what one deduces with certainty from premises one knows" (Schiffer 1996, p.330). There are still a variety of further issues surrounding the proper formulation of the closure principle, but these would take us too far afield for my present purposes. See Hawthorne (2005, p.29) and Williamson (2000m, p.119). Hawthorne argues that we must add that the agent retains their knowledge that p over the competent deduction, since the agent could obtain counterevidence to p during the process. The point is to formulate the closure principle in such a way that this kind of case is not a counterexample.

because you can't "rule out" that you're not a BIV. To the contrary, DeRose suggests that this explanation is more informative in the *reverse*—namely, that the notion of being unable to rule out that p might be explained by the fact that you don't *know* that p. I will return to this question in more detail below. For now I'll simply note that one way of putting the basic idea is along the following lines. Sceptical scenarios are chosen for sceptical purposes precisely because they are scenarios we'd take ourselves *not* to be in, even if we were in them. It can easily seem like anything one might say or do to challenge the thought that they are in a sceptical scenario is something they could equally well do *even if the sceptical scenario obtained*.⁴

Regardless, whether (2) is true or not, many commentators seek to neutralize its ability to threaten our everyday knowledge. Indeed, many closure-denying anti-sceptical strategists are motivated by a conviction that (2) cannot be denied. I turn now to outline the basics of an anti-sceptical strategy that denies premise (1).

1.3 Dretske and the Denial of Closure

The seminal formulation of this strategy is Dretske's (1970) "Epistemic Operators". The basic idea behind Dretske's view is illustrated nicely by his famous zebra example. Imagine that you are at the zoo looking at some striped animals in a pen marked "zebras". You have plenty of reason to believe that the animals are zebras: you know what zebras look like and you can see the animals, not to mention a sign marked "zebras", clearly in front of you. It seems clear that you know the animals are zebras. But now imagine you are with a friend obsessed with conspiracy theories who insists that the animals are actually mules, cleverly disguised by the zoo keepers to look just like zebras. It would seem you *don't* know, just by looking, that the animals aren't cleverly disguised mules. After all, they've been disguised to look just like zebras and to expose such a hoax would seem to require special checks. "Have you checked with the zoo authorities? Did you examine the animals closely enough to detect such a fraud?" (Dretske 1970, p.1016). But you do know that the fact that the animals are *zebras entails* that the animals are not cleverly disguised mules.

⁴ For a helpful discussion on what makes for an effective sceptical scenario, see Neta (2003).

According to Dretske the example illustrates an everyday case in which the closure principle fails.⁵ “What I am suggesting is that we simply admit that we do *not* know that some of these contrasting ‘skeptical alternatives’ are *not* the case, but refuse to admit that we do not know what we originally said we knew” (Dretske 1970, p.1016). So, says Dretske, the closure principle is false. Perhaps, then, we can reject the sceptic’s appeal to closure in establishing the sceptical conclusion. The sceptic’s argument might be another instance of the failure of closure. Of course the sceptic may simply respond to such an example by claiming that, contrary to what we might have thought, in the zebra case, once the cleverly disguised mule possibility is raised, one doesn’t know that the animal is a zebra. And this is precisely because closure holds. The sceptic can argue that Dretske must assume the failure of closure from the outset in order to draw out the first intuition of the thought experiment. So what *argument* does Dretske offer to show that closure does not hold for knowledge? (Dretske 1970, p.1017).

Dretske aims to show that the epistemic operator ‘knows that’ “penetrates” only to relevant “contrast consequences” (what this means will become clear as we go along).⁶ Dretske argues by analogy, noting three other operators similar in key respects to ‘knows that’ and which do not penetrate to (known) contrast consequences.⁷ Because they are analogous in certain key respects, Dretske says, we should be happy to settle with the idea that epistemic operators don’t always penetrate to contrast consequences either. Dretske points out that even if his argument from analogy is a weak one, these analogous operators are useful for explaining why ‘knows that’ is not closed under known entailment.

Dretske’s explanation for why these operators don’t penetrate to certain contrast consequences involves the notion of “relevant alternatives”. In a nutshell, Dretske says that the operators

⁵ In order to argue that (CK3) fails, Dretske has to argue that this is a case in which S knows that the animal is a zebra, knows that the animal’s being a zebra entails that it is not a cleverly disguised mule, competently deduces that the animal is not a cleverly disguised mule on the basis of knowing this entailment, and yet fails to know that the animal is not a cleverly disguised mule.

⁶ The term “penetrates” refers to an operator’s ability to reach certain logical consequences of the proposition upon which it operates. Dretske notes different types of logical consequences, most important of which are “presuppositional” and “contrast” consequences. The latter are most important for our purposes since the sceptical scenarios are contrasting consequences of our ordinary knowledge claims. For example, “I am not a BIV” is a contrasting consequence of “I have hands”. To translate from Dretske’s terminology, when he claims that the epistemic operator ‘knows that’ does not penetrate to all contrasting consequences, he is saying that knowledge is not closed under known entailment.

⁷ These operators do not penetrate to certain other logical consequences, unlike a class of operators that Dretske refers to as “non-penetrating”—an example is ‘strange that’. It could be strange that there is a pink bird in the tree over there, you could know that this entails that there is a tree over there, and yet it is not at all strange that there is a tree over there. Non-penetrating operators also fail, as Dretske shows the epistemic operators do, to penetrate to presuppositional consequences.

penetrate only to those contrast consequences that are part of the relevant set of alternatives which served to define the content of the entailing proposition being operated on in the first place. As Dretske puts it:

To know that x is A is to know that x is A within a framework of relevant alternatives, B, C and D. This set of contrasts together with the fact that x is A, serve to define what it is that is known when one knows that x is A. One cannot change this set of contrasts without changing what a person is said to know when he is said to know that x is A. We have subtle ways of shifting these contrasts and hence, changing what a person is said to know *without changing the sentence that we use to express what he knows* (Dretske 1970, p.1022).

Consider the operator ‘explains that’ in the following example. When a waiter comes around with the dessert menu, Brenda says she’s not having any dessert. Dretske points out that, within a context of explanation (and likewise within a context of epistemic evaluation) “no fact is an island”. That is to say: “within the context of explanation and within the context of our other operators, the proposition on which we operate must be understood as embedded within a matrix of relevant alternatives” (Dretske 1970, p.1022). To continue the example, perhaps we could explain the fact about Brenda by noting that she’s on a diet. This explanation makes sense against the background of a set of relevant alternatives—for example, that Brenda didn’t *order dessert and eat it*. But consider the fact that we know that Brenda’s declining dessert entails, say, that she didn’t *order dessert and throw it at the waiter*. We can’t explain why she didn’t do *that* by bringing to light the fact that she’s on a diet. So our operator ‘explains that’ doesn’t penetrate to a known contrast consequence. But Dretske calls this contrast an irrelevant alternative: that is, it’s an alternative the introduction of which would change “what it is that is being explained and, therefore, what counts as an explanation, even though (as it were) the same fact is being explained” (Dretske 1970, p.1021). For this reason we shouldn’t *expect* the operator to penetrate to such a contrast consequence. Precisely the contrast consequences which do not serve to define the parameters of the original explanation are the “irrelevant” ones.

The crucial point is that epistemic operators, too, should be understood as operating upon propositions which are defined within a matrix of relevant alternatives. Introduce a contrast consequence which alters the matrix, and what S ‘knows that’ becomes something altogether different. Stine summarizes the point nicely: “If being a mule painted to look like a zebra became a relevant alternative, then one would literally mean something different in saying that

John knows that the animal is a zebra from what one meant originally and that something else may well be false” (Stine 1976, p.255). There seems to be no reason why we *ought* to know the denials of contrast consequences which would alter what it is that we can be said to know in the first place. According to Dretske, we can simply presuppose irrelevant alternatives. We do not have to rule them out in order to know that p.

Besides this sketchy illustration of his point, Dretske also suggests in a footnote that if one of the operators under discussion happens itself to be a necessary condition on knowledge then that operator’s failure to penetrate in certain cases will necessarily mean that ‘knows that’ fails to penetrate in the same cases.⁸ This would considerably strengthen the argument that Dretske has sketched. I turn, then, to Nozick’s argument that two necessary conditions on knowledge are not closed under known entailment. Interestingly, as we’ll see towards the end of the next section, it’s not clear that a refined version of Nozick’s approach supports a denial of closure. Indeed, it might be seen as paving the way for a denial of premise (2).

1.4 Nozick and The Denial of Closure

Nozick argues that two necessary conditions on knowledge are not closed under known entailment, and concludes that knowledge itself is not closed (Nozick 1981, p.206). According to Nozick’s necessary and jointly sufficient conditions, a true belief that p counts as knowledge when (and only when) S’s belief *tracks the truth* of p. What this means amounts to two further conditions: i) S wouldn’t believe p if p were false, and ii) S would believe p if p were true. Nozick claims that these conditions account for the intuitive idea that our beliefs that count as knowledge are counterfactually responsive to the facts (Nozick 1981, p.211).⁹ The main condition he focuses on is i), which is known as “counterfactual sensitivity.”¹⁰

We noted briefly above that the BIV scenario can be described precisely as one in which something like i) is not met. We said that even if you were a BIV you would believe you weren’t one. But, according to Nozick, condition i) holds for such mundane propositions as that I have hands. For if “I have hands” were false—if, say, I were to lose my hands in a car

⁸“I think that those who are inclined to give a causal account of knowledge should be particularly interested in the operator ‘R →...’ since, presumably, it will be involved in many instances of knowledge...”(Dretske 1970, p.1019 fn.4).

⁹ Nozick modifies his modal conditions in an important way in response to objections. I return to the modified forms shortly below.

¹⁰ Condition ii) is known as “adherence”. I’ll be setting that to one side since Nozick’s main focus is sensitivity.

accident—I wouldn't still believe that I have hands. And once again, I know that my having hands entails that I'm not a BIV. Nozick goes on to say, given that i) is a necessary condition of knowledge, it follows that I can know proposition p which I know to entail proposition q, and yet fail to know that q.¹¹

Nozick claims that the sceptic won't want to deny that condition i) is necessary for knowledge, given that he supports premise (2) precisely by invoking condition i) (according to Nozick). If this is true, then it seems to present the sceptic with a dilemma: either sensitivity is a necessary condition on knowledge, which problematizes the sceptic's appeal to premise (1); or sensitivity is not a necessary condition of knowledge, which problematizes the sceptic's appeal to (2). In other words, Nozick argues that the sceptic cannot have it both ways when comes to counterfactual sensitivity (Nozick 1981, p.209). Nozick takes himself to have established the right anti-sceptical conclusion. Although we don't know we're not BIVs, this has no impact on the fact that we know all sorts of everyday propositions which happily satisfy the conditions of his analysis of knowledge. But the proposal as it stands faces serious difficulties.

For one thing, the sensitivity condition arguably needs refining, as Nozick himself acknowledges. Consider the following case. A grandmother can reliably tell by looking at her grandson when he comes to visit whether he is healthy and well. If it were ever the case that the grandson was really not well, other members of the family would tell a white lie and assure the grandmother that the grandson is indeed healthy and well to avoid making the grandmother upset. It seems then, that when the grandson comes to visit and the grandmother forms the belief that her grandson is well, this belief is not sensitive. That is, were the grandson not well (in a close possible world) she would believe that he was. Yet it seems implausible to conclude that the grandmother does not know that her grandson is well. So the case is a challenge for the claim that sensitivity is necessary for knowledge (Nozick 1984, p.179).

Nozick deals with the case by modifying his conditions on knowledge. He modifies them as follows: S knows that p iff p is true, S believes, via method or way of coming to believe M, that p, and: i)* If p weren't true and S were to use M to arrive at a belief whether (or not) p, then S wouldn't believe, via M, that p; and ii)* If p were true and S were to use M to arrive at a belief whether (or not) p, then S would believe, via M, that p. The basic idea is that the modal profile

¹¹ Again, to challenge (CK3), Nozick must argue that one can know p, know that p entails q, competently deduce q from p on the basis of knowing that p entails q, yet fail to know q.

of S's belief that p must be evaluated in possible worlds in which the method by (or basis upon) which S forms the belief that p in the actual world is held fixed. The only possible worlds that are relevant to the assessment of whether S knows that p are worlds in which S forms the belief that p by the same method as in the actual world (Black 2002, p. 9). With these modifications, Nozick can get the right results in the grandmother case. After all, the case involves *different* methods in an important way. In the possible world in which she forms the false belief that her grandson is well, it's on the basis of testimony. However, the belief we evaluate for sensitivity is formed on the basis of visual perception. And it seems plausible that, holding this belief-forming method fixed across possible worlds, the grandmother's belief *is* sensitive. Thus, it remains open to Nozick to claim that she knows that her grandson is well.

Relativising to methods opens up an important worry for Nozick. It's questionable whether Nozick can motivate a denial of the closure principle given this modification of his modal principles. It is no longer clear that the closure principle fails when we hold belief-forming methods fixed. For example, we might insist that one's belief that they are not a BIV *is* sensitive according to the modified form of sensitivity. To be sure, in a BIV world one would form a false belief that one is not a BIV. But it seems one would do so by means of a different belief-forming method. In the actual world, it's plausible that one forms the belief that they are not a BIV on the basis of visually perceiving, say, that they have hands, and inferring that they are not a BIV. In the closest world in which one is a BIV, however, intuitively one doesn't form empirical beliefs on the basis of visual perception. Rather, one forms them via sensory experiences controlled by the scientists in charge. If S's belief that they are not a BIV is sensitive, then there is no need to deny closure. This is because both S's belief that she has hands, and S's belief that she is not a BIV meet the modified conditions on knowledge. This has a potentially important upshot for responding to scepticism beyond the implications for closure denial, and I will return to this point below.

Of course, it is open to Nozick to respond in the following way. He can insist that the belief-forming method does *not* differ between the actual world and the closest possible BIV worlds. Indeed, Nozick has a view about how to properly individuate belief-forming methods according to which BIVs and ordinary perceivers can be understood as forming beliefs with the same method. His view is that what's crucial for individuation of methods is the "final upshot" in experience that one has when using the method. When it comes to comparing one's

ordinary perception that they have hands to that of a BIV's, the final upshot in each case is the phenomenal experience as of p being the case. He claims that because of this these two methods count as the same (Nozick 1981, p.184).¹²

There are a number of things to say in response to this suggestion and I don't want to get caught up in the details here. I will note, however, that it seems odd for Nozick to individuate methods by this purely internal condition of phenomenal experience. First of all, it seems to be at odds with his externalist approach to knowledge (I discuss the externalist element of modal theories of knowledge below). As Pritchard points out:

If one is happy to endorse externalism, however, and note that this is not an uncontroversial move to make, then it is not obvious why one should accept that the defining mark of a method should be something that is reflectively accessible i.e., the nature of one's experiences. Why couldn't one's method be determined by facts which were not reflectively accessible to one as well? (Pritchard 2008, p.10).

Regardless of whether one has epistemic externalist leanings or not, it's plausible to assume that the belief forming method of visual perception involves such components as:

[...] the lenses of the eyes focussing ... light on the retinas, where a pattern of retinal cell stimulation occurs that sends electro-chemical impulses along the optic nerve to the visual cortex, where a pattern of brain cell stimulation occurs with the upshot that the subject has a visual experience (McLaughlin 1996, p.200).¹³

In any case, as Black notes, distinguishing what counts as visual perception in this way as opposed to Nozick's way allows us to avoid counterintuitive consequences, such as counting our dreams as instances of visual perception (imagine a dream in which the "final upshot" of the belief forming method is experientially indistinguishable from the final upshot of an experience had via one's visual perception in waking life), or cases in which a person who's been blind from birth has the experience as of something purple flashing before their eyes (Black 2002, p. 12). As I said, I don't want to get into the details of this complex issue here.¹⁴

¹²“Usually, a [belief-forming] method will have a final upshot in experience on which the belief is based, such as a visual experience, and then (a) no method without this upshot is the same method, and (b) any method experientially the same, the same ‘from the inside’, will count as the same method. Basing our beliefs on experiences, you and I and the person floating in the tank are using ...the same method” (Nozick 1981, p.184-185).

¹³ As quoted in Black (2002, p.13).

¹⁴ See Nozick (1984), p.233 for a challenge to externalist individuation of methods. For a subtle discussion of Nozick's arguments, see Williams (1991), p.336-346. See also Becker (2012) who argues both that methods should not be individuated externally, and that there is a way of interpreting Nozick's claims about methods which is consistent with his externalism.

My point is to note that it's far from clear that Nozick's view of knowledge licenses, and still less requires, a denial of closure. Indeed, we might think that these considerations suggest that there is hope for an anti-sceptical strategy that challenges premise (2).

1.5 Premise Two

A prominent way of responding to closure-based scepticism is known as "neo-Mooreanism". Neo-Mooreanism is inspired by G.E. Moore's "common sense" anti-sceptical strategy in "Proof of an External World" (1939). Moore's argument can be formulated in the following way:¹⁵

- 4) I know I have hands.
- 5) If I know I have hands, then I know I'm not a handless BIV.
- 6) Therefore, I am not a handless BIV.

Moore effectively *relies* on (as opposed to challenges) the closure principle. He argues that it's utterly clear and obvious that he knows that he has hands. Moreover, he knows that this knowledge entails that he is not a BIV. Therefore, he concludes, he knows that he is not a BIV. In other words, Moore turns the closure-based sceptical argument on its head. An important thing to note about this response is that Moore doesn't do much *else* in dealing with the sceptic (Coliva 2010; McGinn 1989; Stroud 1984). In a sense this strategy amounts to admonishing the idea that one does not know that they have hands. I don't think it's difficult to sympathise with Moore on this point. However, this "straight man" response seems pretty implausible to many, and for a variety of different reasons (McGinn 1989; Stroud 1984; Williams 1991; Wright 2002). But here is one general issue. We've seen that the sceptical problem can be understood as a triad of independently intuitive and plausible but jointly inconsistent claims. Simply employing commonsense in Moore's way at best seems to lead to an impasse with the sceptic. As it is often remarked in this context, one philosopher's *modus ponens* is another philosopher's *modus tollens*. And it's not clear what reason we have to go either way for all Moore has said. In any case, neo-Mooreans acknowledge that the worry about an impasse needs to be accounted for. In order to avoid an impasse, neo-Mooreans supplement their view in two crucial ways: a) they provide an explanation of *how* it is that we

¹⁵ This is a standard reconstruction in contemporary debate. I don't mean to suggest that Moore would agree with this reconstruction.

know the denials of sceptical hypotheses, and b) they provide a diagnosis of why it seems compelling that we do not have such knowledge.¹⁶

The reflections from last section suggest it might be possible to approach a) and b) with a basis-relative notion of sensitivity.¹⁷ But the details of that proposal are controversial—in particular, there are issues about how to individuate methods. So I'll look instead to a prominent neo-Moorean who approaches a) and b) in a closely related but different way. Sosa appeals to a “safety” condition on knowledge to explain how we can know the falsity of sceptical hypotheses (Sosa 1999). Safety is another modal condition.¹⁸ The basic idea behind safety is that our beliefs that count as knowledge could not have easily been false. One way of putting the idea is that we cannot have “lucky” knowledge. For example, consider the Gettier cases. Roughly, they present a problem for justified true belief accounts of knowledge by being cases of true justified belief that in some sense result from “luck”. These are familiar points, but Gettier cases can be used to illustrate the intuitive idea that knowledge excludes luck.¹⁹ Here is how Sosa puts safety:

Call a belief by S that p “safe” iff: S would believe that p only if it were so that p. (Alternatively, a belief by S that p is “safe” iff: S would not believe that p without it being the case that p; or, better, iff: as a matter of fact, though perhaps not as a matter of strict necessity, not easily would S believe that p without it being the case that p.) (Sosa 1999, p.142).

This principle is itself the subject of a good deal of controversy and stands in need of further refinement (for one thing, it will have to be relativised to a belief forming method to avoid counterexamples such as the grandmother case). But the basic idea for our purposes is that S's belief that p has the right modal profile for knowledge as long as it would be true within a certain non-remote range of possible worlds in which S believes that p (again, holding the method of belief formation in the actual world fixed). Crucially, according this condition, we

¹⁶ Actually, Pritchard (2007) identifies *six* shortcomings with Moore's argument. He thus identifies six corresponding desiderata for a compelling neo-Moorean position. I will only focus on three of these desiderata: the “no supporting epistemology objection”, the “no diagnosis objection”, and the “evidential scepticism objection”. This latter objection will take us into a fairly involved discussion about another sceptical argument.

¹⁷ Again, see Pritchard (2008), Black (2002), and Williams (1991).

¹⁸ Sosa claims that safety can be independently motivated as being preferable to sensitivity. He outlines three compelling counterexamples to sensitivity in Sosa (1999). Two examples are the problem of necessary truths and the garbage chute case. Further discussion would take us too far afield here.

¹⁹ There are complex issues that must be dealt with in explaining the relevant sort of “luck” that is supposed to be incompatible with knowledge (Pritchard 2005c, 2012b; Steglich-Petersen 2010). These details would take present discussion too far afield.

know both that we have hands and that we are not BIVs. So there does not seem to be any need to deny closure. Moreover, while it's true that this principle must also be basis relative, it is not merely in virtue of relativising to a basis that the idea that we have knowledge of the falsity of sceptical hypotheses is motivated. That is, we know we're not BIVs simply because the BIV scenario is putatively something that obtains in an outrageously far-off world.²⁰

With the safety condition independently motivated, Sosa can vindicate our sympathy with Moore's "straight man" response to the closure-based sceptical problem. He can explain in a theoretically motivated way how knowledge of the falsity of sceptical hypotheses is possible. Sosa also has a diagnosis of the felt-oddness of the idea that such knowledge is possible (recall b) from above). He notes that sensitivity and safety are easily confused. One may at first think that sensitivity is compelling, and come to believe that (2) is true on that basis. However, what one really finds compelling, Sosa says, is a safety condition: "Safety and sensitivity, being mutual contrapositives, are easily confused, so it is easy to confuse the correct requirement of safety (for knowledge and its correct attribution) with a requirement of sensitivity. It is easy to overlook that subjunctive conditionals do not contrapose" (Sosa 1999, p.148). Because safety does *not* support (2), people are mistaken—but understandably—when they think that (2) is true.

How plausible is Sosa's neo-Moorean strategy? First of all, I think this last line of thought is problematic. Although the two formulations of *conditions* on knowledge are easily confused, it's not at all clear that the intuitive thought behind each condition is easily confused. It's often said that what motivates both principles is the idea that our beliefs could not easily be false. But I think this glosses over an important difference in what's compelling about the principles. For instance, the intuitive thought behind the sensitivity condition is that our beliefs should be counterfactually responsive to the facts. It's intuitive that if I would believe *p* in the nearest possible world in which it isn't the case, then I don't know *p*. This is one way to say that a belief could easily have been false. Meanwhile, although safety captures the idea of counterfactual sensitivity as well, it goes beyond it. For example, when it comes to our knowledge that we are not BIVs, this is not due to counterfactual sensitivity. Rather, it's due to

²⁰ Meanwhile, when it comes to sensitivity, we do not get the result that we know we're not BIVs independently of relativisation to methods. This is because the sensitivity principle is formulated in such a way that possible worlds in which one is a BIV are taken to be relevant to our assessment of one's belief that they are not a BIV. In those worlds, one would continue to (falsely) believe that one is not a BIV. Indeed, this is precisely why the sensitivity principle is said to capture the idea that we don't know we're not BIVs.

the *de facto* modal distance of BIV worlds and the stubbornness of our beliefs that we're not BIVs (Pritchard 2008, p.13). Indeed, when it comes to motivating safety, it's plausible that what is at work is an intuition about *risk*, or something along those lines. Safety seems plausible because it's plausible that our beliefs that intuitively count as knowledge are ones that are not risky to have. We don't need to *worry* about what we'd believe in such far off worlds as BIV worlds because we're not at risk of being in those worlds. This is perhaps another way of understanding what it means to say that beliefs that count as knowledge wouldn't easily be false, but it's one that goes beyond the notion of counterfactual responsiveness. In addition to all this, it's not clear that sensitivity is the best explanation of the intuitiveness of (2) in the first place. But once again, I'll return to that point below.

The central issue with Sosa's neo-Moorean strategy, however, is that it seems so implausible to many to say that we can know the falsity of sceptical hypotheses. As we've seen, those who take the line of denying closure do so under the conviction that it's a desideratum on our theory of knowledge that we respect—what they take to be—the platitude that we could never have *this* kind of knowledge. Nozick is vocal on this: “The skeptic's possibilities make us uneasy because, as we deeply realize, we do not know they don't obtain; it is not surprising that attempts to show we do know these things leave us suspicious, strike us even as bad faith” (Nozick 1981, p. 201). I don't want to say anything one way or the other on this at the moment (I return to the issue below, as well as to the broader issue of responding to scepticism with an externalist theory of knowledge—which, as I'll explain, is what Sosa does here). I think there is an intuitive enough worry in the background at this point—whatever motivates it—to warrant pressing on to a final way of approaching the closure-based argument, one that does not require us to claim that we can support an outright denial of (2), nor that we need to reject closure.²¹

1.6 DeRose and Semantic Contextualism

DeRose offers an anti-sceptical strategy which, if successful, would eliminate the need to deny (1) or (2) outright. He borrows Nozick's notion of counterfactual sensitivity and puts it to

²¹ Perhaps one thing that motivates the idea that we could not possibly know the falsity of sceptical hypotheses in the way Sosa has explained we can, is the thought that our beliefs about not being BIVs are not counterfactually sensitive. The thought would be that it seems our knowledge of the falsity of sceptical hypotheses is due too much to lucky circumstances (compare Pritchard 2005).

work in a semantic contextualist theory of knowledge attributions. According to DeRose, the sentence “S knows that P”, strictly speaking, only expresses a complete proposition relative to standard N. This is because ‘knows that’ is a context sensitive term that only has a truth value relative to the standards at play in a given context.²² As Schiffer (1996) explains, according to the contextualist theory, the complete closure-based sceptical argument should look something like this (the term ‘Tough’ here refers to a high-standards context):

- 7) If S knows that he has hands relative to Tough then S knows he is not a handless BIV relative to Tough.
- 8) S doesn’t know he is not a handless BIV relative to Tough.
- 9) Therefore, S doesn’t know that he has hands relative to Tough.

According to Schiffer (and DeRose), this argument does not seem paradoxical.²³ The anti-sceptical thrust of the contextualist theory is that the sceptical “paradox” turns out to be a sound, but uninteresting argument about knowledge according to a particularly high-standards context of evaluation. The argument says nothing about what we know in ordinary contexts—although, we mistakenly think it does.

To fill in the theory, the contextualist needs to explain at least two crucial things: how do the standards of a given context get raised (and lowered)? And why didn’t we see the argument in this non-paradoxical way in the first place? To answer the first question, DeRose employs the notion of sensitivity along with the notion of “strength of epistemic position” to develop his “Rule of Sensitivity”, a mechanism which he says determines what standards are at play when

²² There are of course many contextual views in the literature and I can’t begin to address them all. I focus on DeRose because his view is a seminal semantic contextualist view, and his view has the paradigm features of contextualism that are relevant for present purposes (other views usually grouped together with DeRose-style contextualism include Cohen (2002) and Lewis (1996)). One particularly important distinction needs to be made between semantic contextualists like DeRose, and the contextualist views of Annis (1978), Neta (2003), and Williams (1991). There are of course a variety of differences between these last three contextualist views themselves, but one important difference they share concerning semantic contextualism is that they do not individuate contexts in terms of *higher* or *lower* epistemic standards, but rather in terms of standards that are different in kind. I return to an important point about the semantic contextualist’s appeal to high epistemic standards below.

²³ Note that this formulation would only so easily be attributed to a “scepticism-friendly” version of contextualism. DeRose notes that it is a complicated issue to settle whether, and when, the sceptic *does* successfully raise standards in conversation (DeRose 2006, p.333). Discussion of this point is not relevant to our present purposes.

we evaluate whether “S knows that p” is true. DeRose’s notion of “strength of epistemic position” goes like this:

[...] being in a strong epistemic position with respect to P is to have a belief as to whether P is true match the fact of the matter as to whether P is true, not only in the actual world, but also at the worlds sufficiently close to the actual world...The further away one gets from the actual world, while still having it be the case that one’s belief matches the fact at worlds that far away and closer, the stronger a position one is in with respect to P (DeRose 1995, p.34).

Take for example my belief, p, that I have hands. I am in a strong epistemic position with respect to p because in nearby as well as not-so nearby possible worlds my belief would match the facts. Indeed, it is difficult to imagine a scenario in which I would falsely believe that I have hands. Of course, one familiar scenario is that in which I am a handless BIV. In the possible world in which I am a BIV my belief that I have hands wouldn’t match the facts. But this is a *very* distant possible world; so, while my belief here wouldn’t match the facts, I am still in a strong epistemic position with respect to it. Consider also my belief, q, that I am not a BIV. This belief, too, is one that will match the facts in close possible worlds since it is only in very far off worlds that I would falsely believe it—i.e. worlds in which I *am* a BIV. So according to DeRose I am in a strong epistemic position with respect to this belief, too. Just *how* strong an epistemic position one must be in with respect to p in order to count as knowing p is a context-dependent matter. This is where DeRose’s “Rule of Sensitivity” comes into play:

When it is asserted that some subject S knows (or does not know) some proposition P, the standards for knowledge (the standards for how good an epistemic position one must be in to count as knowing) tend to be raised, if need be, to such a level as to require S’s belief in that particular P to be sensitive for it to count as knowledge (DeRose 1995, p.35).

DeRose notes that the distinctive feature of sceptical hypotheses is that you would need to be in an extremely strong epistemic position for your beliefs about them to be sensitive. The relevant possible worlds that come under consideration when assessing sensitivity are extremely far off. According to the Rule of Sensitivity, when an assertion or denial of a sceptical hypothesis is made, then, this tends to raise the standards for knowledge so that the strength of epistemic position one must be in with respect to a belief is such that your belief

would match the facts in a world in which you are a BIV—which is to say, your belief wouldn't match the facts.

The interesting thing here is that sensitivity and strength of epistemic position can come apart. In some contexts, a belief may count as known on strength of epistemic position, even if it isn't a sensitive belief. By avoiding the claim that sensitivity is necessary for knowledge, DeRose can claim that most of the time we *do* know that we're not BIVs. He does so by appeal to the strength of our epistemic position with respect to such beliefs. Only in contexts in which the BIV or other sceptical scenarios have been raised do we need to be in such an extremely strong epistemic position that beliefs, even in ordinary propositions, would no longer match the facts. Thus, DeRose thinks he can preserve our intuitions about sensitivity, and thus, he claims, the plausibility of (2) in the closure-based sceptical argument; yet he also thinks he can preserve our intuitions about (1) along with our intuitions about the denial of the sceptic's conclusion.

In particular, DeRose can preserve closure by noting that the modal distance at which my belief that I have hands would fail to match the facts is also the modal distance at which I would falsely believe that I am not a BIV. In all those worlds in which I'm not a BIV, I *wouldn't* falsely believe that I am not one. Not until we consider the very distant possible world in which I *am* a BIV would I falsely believe that I am not one. As DeRose puts it, my epistemic position with respect to my belief that I'm not a BIV is at least as strong as my belief that I have hands. In BIV worlds, my belief that I have hands would also fail to match the facts. DeRose can thus maintain that closure is true, regardless of the standards at play.

What about the second task involved in filling in a contextualist anti-sceptical strategy mentioned above? The contextualist needs to explain why we didn't notice that the sceptical paradox isn't a paradox after all. The fastest way to do this is to claim that it's easy to miss the context-sensitivity of "S knows that p". As Schiffer puts it, closure-based scepticism "strikes us as presenting a profound paradox merely because we're ignorant of what it's really saying, and this because we don't appreciate the indexical nature of knowledge sentences" (Schiffer 1996, p.325). DeRose doesn't seem to offer a positive theory that would explain such semantic blindness. Instead, he objects to the fact that Schiffer is concerned with the need for a

contextualist in particular to provide such an error theory. According to DeRose, *any* theory of knowledge, whether invariantist or contextualist, is going to have to attribute semantic blindness to ordinary speakers. DeRose cites as an empirical fact that about half of ordinary speakers say “yes” and half say “no” to the question of whether the sceptic’s denial of S’s knowledge that he has hands contradicts with an “ordinary” assertion that S knows he has hands. If an invariantist theory is true, then all of those epistemic agents claiming “no” (presumably because they at least implicitly take ‘knows that’ to be contextually determined) are semantically blind to the invariant nature of ‘knows’. DeRose maintains that there *just is* a good deal of semantic blindness when it comes to the hard questions about knowledge that contextualism raises (DeRose 2006, p.322). Critics of contextualism on this score are numerous (Hofweber 1999; Rysiew 2001; Chrisman 2007).²⁴ The details of these criticisms would take us too far afield.

Perhaps the most prominent objection to semantic contextualism is that it mistakenly assumes that external world scepticism requires an appeal to especially strong or high epistemic standards. According to the semantic contextualist, the sceptical context is explicitly stated as being one that demands high standards for correct knowledge-attributions. However, the objection goes, we need not think of scepticism as invoking exceptionally high epistemic standards. It is worth examining this objection in some detail.

1.7 Scepticism and Standards

Opponents of contextualism often point out that the really interesting form of external world scepticism denies that there are important epistemological distinctions to be made concerning our empirical beliefs, regardless of whether we are in “high-standards” or “low-standards” contexts for knowledge-attributions.²⁵ Consider Kornblith:

²⁴ “...contextualism requires the posit of a deep and hard to explain semantic blindness to anyone [who] resists contextualism about ‘knows’ or is perplexed by [the closure-based argument]. This makes the contextualist response to skeptical paradoxes such as [the closure-based argument] seem unacceptably ad hoc, so call this the *ad hoc semantics problem*” (Chrisman 2007, p.233).

²⁵ Kornblith (2000) makes the point forcefully. The point is also made in Williams (2001). See also Feldman (1999 and Neta (2003). There are a variety of arguments we might appeal to, and as we’ll see, my preferred approach is the underdetermination argument. Feldman (1999), on the other hand, claims that there is an argument for “justification-scepticism” which relies on the plausible idea that ordinary “low” standards for justification preclude reliance on circular arguments (p.22). One thing to note about Feldman’s argument is that regardless of whether contextualism speaks to this sort of argument, it’s not clear that we should endorse the idea that justification

The reason I don't know anything about the external world, according to the skeptic, is not that I have a small degree of justification for my beliefs when knowledge requires a larger degree of justification. Rather, the skeptic claims that I have no degree of justification whatever for my claims about the external world. None. Let us call this view *Full-Blooded Scepticism* (Kornblith 2000, p.25).

Kornblith's "Full-Blooded Sceptic" denies that we have *any* justification for ordinary empirical beliefs such as that we have hands. While DeRose has argued that in high-standards contexts knowledge attributions are false, he has not said anything about whether we have justification to believe that we have hands, or whether we have evidence that we have hands in those contexts. Moreover, the closure problem for knowledge gives us no reason to think not. Kornblith argues that DeRose simply grants that we can have justification for ordinary empirical beliefs in both high- and low-standards contexts. He claims that DeRose's response to the sceptical problem is therefore uninteresting (from an epistemological perspective):

Since [DeRose's] skeptic agrees that we can make meaningful and important distinctions about how well justified we are in various claims, and agrees with us about which claims we should believe and act upon, nothing much turns on it. It is like dealing with the Vermonter who insists that he won't say that it is cold outside unless it is at least 25 degrees below zero Fahrenheit. If he recognizes that there are important distinctions to be made in temperatures above minus 25, and that these distinctions have a bearing on how one should interact with the world, then the only difference between him and us is a bit of charming linguistic eccentricity (Kornblith 2000, p.26).

So what kind of sceptical argument does Kornblith have in mind? One example is the underdetermination argument:

- 1*) If my evidence does not favour that I have hands over that I am a BIV (being stimulated to believe that I have hands), then my evidence does not justify that I have hands.
- 2*) My evidence does not favour that I have hands over that I am a BIV (being stimulated to believe that I have hands).
- 3*) My evidence does not justify that I have hands.
- 4*) Therefore, I do not have justification to believe that I have hands.

requires being able to defend one's belief with non-circular arguments. Indeed, there's a plausible distinction to be made between what Williamson calls "inferential" justification (justification in the form of providing arguments in support of one's belief) and "normative" justification, which has to do simply with whether it's "OK" to believe that p (Williamson 2009). Williamson, and others (Pryor 2005) would argue that inferential justification is neither necessary nor sufficient for normative justification.

5*) Therefore, I do not know that I have hands.

Insofar as we can run this argument equally well for just about any ordinary empirical proposition, the radical conclusion of this argument is that we don't have any justification to believe any ordinary empirical propositions. Why is this a particularly interesting form of Full-Blooded Scepticism? Like the closure argument, the argument is said to be paradoxical. That is, it consists of a series of highly intuitive premises and establishes by way of a valid rule of inference a highly counterintuitive conclusion. The underdetermination argument works by appealing to commitments that are reflected in our ordinary epistemic practices. It then employs particular instances of these commitments to derive the conclusion that we do not have any justification to believe that we have hands. Let's consider the motivation for the premises.

Premise (1*) is motivated by an underdetermination principle (UP). That principle can be stated:

(UP) For all S, p, q, if S's evidence for believing p does not favour p over some hypothesis q which S knows to be incompatible with p, then S's evidence does not justify S in believing p

The idea is that in order for some body of evidence, E, to provide justification for a proposition, p, it is plausible to maintain that E should favour p over alternatives that are known to be incompatible with p. What does "favouring" mean? Briesen (paraphrasing Brueckner (1994)) understands "favouring" in the following way: "If my evidence favours p over q, then p has some epistemic credit which q lacks. In other words: If my evidence favours p over q, then it is more reasonable for me to believe p than q" (Briesen 2010, p.4). Pritchard prefers simply to speak of favouring in terms of a body of evidence making one proposition *more likely to be true* for an agent than another proposition. I think there's a family resemblance in these basic ways of spelling out the notion of favouring. And I think the notion is intuitive enough that we can rely on this rough understanding in what follows. But to make things more concrete, consider the following example:

If...my evidence for believing that I am currently in the town's Odeon cinema does not favour this proposition over the competing scenario (which I know is incompatible) that I am in the town's other the cinema, the Multiplex, perhaps

because I have just woken up from a drunken sleep with only a dazed recollection of the chain of events that led to my being in a cinema, then it is hard to see how that evidence could possibly support a justification for my belief that I am at present in the Odeon cinema (Pritchard 2005, p.40).

Premise (2*) appeals to the idea that one's evidence would be the same whether one was a BIV or not. We can call this the "sameness of evidence thesis". The very notion of a BIV scenario is designed to draw out the intuitiveness of this thesis. If we find the possibility of a BIV scenario epistemically threatening, a good explanation of this is that it's because it seems obvious that in such a scenario we would have the same evidence and thus that our evidence doesn't favour the proposition that one is a BIV over the proposition that one is not. As we'll see in Chapter 3, this idea is often associated with a "highest common factor" conception of evidence, on which our evidence always and everywhere consists in what is common to good and bad cases. There are deep and interesting questions about what motivates such a conception, and it will be helpful to examine these more carefully in the context of Pritchard's view, where he motivates an alternative view precisely by challenging the idea that this sort of conception is all that intuitive after all.

With (UP) and the favouring claim motivated in this way, it's a short step to premise 3*), the claim that one's evidence does not justify that one has hands. With the further assumption that if one's evidence does not justify that one has hands then one does not have justification to believe that they have hands, the sceptic arrives at the conclusion that one does not have justification to believe that they have hands (and thus, as noted, that one does not have justification to believe just about any ordinary empirical proposition). From there, insofar as we think that justification is necessary for knowledge, the Full-Blooded Sceptic concludes that we do not know we have hands.

The Full-Blooded Sceptic's assumptions about the connections between evidence and justification, and justification and knowledge are of course controversial. Indeed, as Kornblith points out, DeRose has his own epistemological story about our knowledge of the external world. This involves a central role for the sensitivity condition and the notion of "strength of epistemic position".²⁶ It is open to DeRose to insist that our ordinary beliefs are often in good epistemic standing because they have the right modal profile. He can respond to the underdetermination argument by claiming that evidential justification is not necessary for

²⁶ Recall, contextualism comes into play with the "Rule of Sensitivity" (DeRose 1995, p.35).

knowledge. However, as Kornblith notes, this feature of the view alone is not a contextualist feature. It is an *externalist* feature. As such, whether or not it is correct is irrelevant to the question of whether semantic contextualism is correct (Kornblith 2000, p.27). Indeed, any one of the anti-sceptical strategies considered above can appeal to the sensitivity or safety condition on knowledge and note that meeting these conditions does not imply that one has evidential justification for one's belief that p.²⁷ To be sure, while safety or sensitivity are endorsed by many as *necessary* conditions on knowledge, a further question remains about what, if anything, must be added to these conditions for knowledge. But Sosa's neo-Moorean response to the sceptic, for example, explains how we can know the falsity of sceptical hypotheses *merely* in terms of the safety condition. Whether or not he thinks safety is sufficient for knowledge, he appears to claim that safety is sufficient for our knowledge of the falsity of sceptical hypotheses.²⁸ All of this said, it's important to consider exactly how far denying the evidential sceptic's views about evidential justification and knowledge gets us in an anti-sceptical strategy. As I'll explain, I don't think it gets us very far.

1.8 Responding to Scepticism about Evidential Justification

Regardless of whether evidential justification is necessary for knowledge, it's doubtful that even the externalist would deny that we sometimes have evidential justification for our empirical beliefs (and have knowledge on this basis). It seems, then, that when it comes to an argument like the underdetermination argument, the externalist has two options. Either she can continue to talk about knowledge while remaining silent on the issue of knowledge that comes from (is based on) evidential justification (thus remaining silent on the argument in question); or she can maintain that evidential justification is never relevant to our knowledge of ordinary empirical propositions. Both of these options would be good to avoid in responding to such a sceptical problem if possible. This is because, taking the first option, we'd be talking past an interesting and important sceptical issue. Taking the second option, we'd be endorsing a radical revision of our concept of knowledge (surely we at least sometimes know things in virtue of evidential justification). Pritchard frames this idea nicely:

²⁷ As Bergmann (2008) notes, externalist responses to scepticism may involve any combination of appeals to reliabilism (Goldman 1979, 1986; Alston 1985), cognitive virtues (Sosa 1991; Greco 1999, 2000), modal conditions (Sosa 1999; Nozick 1981), or proper function (Plantinga 1993; Bergmann 2004, 2006: ch. 5).

²⁸ In more recent work, Sosa (2007) "bifurcates" knowledge. He divides knowledge into "animal knowledge" and "reflective knowledge". To put it very roughly, only reflective knowledge requires that the agent have some kind of perspective on the conditions she meets as a knower. Discussion of Sosa's more recent epistemology would take us too far afield.

In any case, while the general shape of such a response to radical scepticism is appealing, the devil (as so often in philosophy) lies in the detail. For although the epistemic externalist is keen to break the logical link between knowledge and rational support, such that one can have the former without the latter, they surely do not wish to disengage our everyday knowledge from rational support altogether. Or, at least, the sceptic can force a dilemma here. On the first horn of the dilemma is the charge that the epistemic externalist is ultimately offering no response at all to the sceptical problem. On the second horn is the charge that the epistemic externalist is presenting us with an epistemological proposal which is so revisionist, so discontinuous with our ordinary epistemic practices, that not even the most ardent proponent of epistemic externalism would find it palatable (Pritchard *forthcoming*, Ch.1, p.12).

Notice that when it comes to the closure argument for knowledge, it may be open to someone like Sosa to allow that we have evidential justification to believe that we have hands, but deny that this is a necessary condition on knowledge, and therefore support our knowledge of the falsity of sceptical hypotheses with a safety condition. Such a move is not available with the underdetermination argument. And that is because the underdetermination argument directly threatens our evidential justification to believe that we have hands by claiming that this evidence must favour that we have hands over that we are not BIVs (if it is to provide epistemic support at all). Thus, if we want to hold on to evidential justification, we must either show that we have this favouring evidence or resist the idea that we require it. This is precisely a matter of responding to the underdetermination argument.²⁹

But it's important to note that the underdetermination argument is just one form of Full-Blooded Scepticism. Pritchard himself makes his point about the dilemma the externalist faces in responding to evidential scepticism by reformulating the closure argument directly in terms of "rationally supported knowledge" (Pritchard *forthcoming*). But the idea is essentially the same. Even if the externalist anti-sceptic is willing to allow that we have rationally supported knowledge when it comes to ordinary empirical propositions, as soon as this rationally supported knowledge is connected in the relevant way to a requirement of rationally supported

²⁹ In my Boulton (2013), I defend Brueckner's claim that a refutation of the underdetermination argument would entail a refutation of the closure argument, but not vice versa. The main claim I defend is that the underdetermination sceptic's appeal to an underdetermination principle is needed in order to motivate premise (2) of the closure argument. My defence consisted in responding to an objection due to Cohen (1998). I take Cohen to present a particularly elegant alternative way of motivating premise (2) of the closure argument which appeals to an explanatory principle. I argued that Cohen's explanatory principle does not provide a natural motivation for premise (2) of the closure argument. One upshot of that discussion is that, in addition to being a form of Full-Blooded Scepticism and taking priority over the closure argument for knowledge in that sense, we might see the underdetermination argument as taking priority in another sense. Namely, if Brueckner is right, a refutation of the underdetermination argument would entail a refutation of the closure argument but not vice versa. I won't defend or explore this possibility any further here. I return to a related issue in Chapter 5.

knowledge of the falsity of sceptical hypotheses (in Pritchard's case via a closure principle for rationally supported knowledge), the externalist anti-sceptic must either show how we have that rationally supported knowledge (a matter of dealing with what I'm calling evidential scepticism), or face Pritchard's dilemma.

I'd like to note that these points about externalism and scepticism are not the same as the familiar point that externalist responses to scepticism are in some sense "philosophically unsatisfying". Bergmann notes there are a variety of different ways to spell out the idea that an externalist response to scepticism is philosophically unsatisfying.³⁰ A popular claim is that an externalist response is a mere *conditional* response. Here is Stroud:

[the externalist epistemologist]...is at best in the position of someone who has good reason to believe his theory if that theory is in fact true, but has no such reason to believe it if some other theory is true instead. He can see what he *would* have good reason to believe if the theory he believes were true, but he cannot see or understand himself as knowing or having good reason to believe what his theory says (Stroud 1989, p.46).

Here we might read Stroud as equating "philosophical satisfactoriness" with a kind of philosophical *understanding*. He is claiming that the so-called conditional nature of externalist epistemology (such as an epistemology that treats true belief meeting the safety condition as sufficient for knowledge) is inadequate to the task of a first-personal quest for understanding how it is that our knowledge of the world around us is possible (Stroud 1989).

I don't intend to endorse Stroud's line of thought, nor any of Bergmann's other ways of spelling out the notion of "philosophical satisfactoriness".³¹ Rather my claim is that denying the evidential sceptic's conditions on knowledge doesn't get us very far. Insofar as it's highly intuitive that we at least sometimes have knowledge that is supported by evidential justification, simply going externalist forces a dilemma. Again, either we must radically revise our concept of knowledge, or talk past an interesting and important sceptical argument. If all

³⁰ More specifically, he says that there are four ways (Bergmann 2008, p.14). These are: i) externalism is a mere conditional response; ii) externalism permits an epistemically circular response to the sceptic; iii) the externalist is herself uncomfortable with going externalist on higher-order beliefs, and this is suspicious; iv) anything goes if the externalist is right – for example, what stops a believer in the Great Pumpkin from legitimately appealing to externalist considerations? Getting into the details of Bergmann's four ways would take us too far afield. It's important to note that Bergmann claims each of these considerations is equally a problem for internalists, and also that there are satisfactory externalist replies to each of these considerations (Bergmann 2008).

³¹ For penetrating discussions of this issue, see also Fumerton (1995) and Williams (1991).

of this is right, we should remain focused on Full-Blooded Scepticism. We should remain focused on sceptical arguments that challenge our evidential justification to believe ordinary empirical propositions.

1.9 Conclusion

The chapter started by examining three main ways of responding to the closure problem for knowledge. I looked at a variety of issues for each strategy. As I said, the aim has not been to demonstrate that these issues cannot be overcome. Rather, the aim has been to highlight a deeper issue. While the closure argument for knowledge is certainly interesting in the sense that it at least seems to appeal to our ordinary epistemic commitments in drawing out paradoxical results, it's not clear that it's interesting in the sense of being entirely radical (although of course the idea that we don't know we have hands is certainly radical enough for many). I've argued in favour of a shift in focus to scepticism about evidential justification. This is because a central desideratum in responding to the external world sceptic is to understand how we can rightly take ourselves to at least some of the time have evidential justification in support of our ordinary empirical beliefs.

Chapter 2: Entitlement and Immediate Justification

2.1 Introduction

This chapter examines two anti-sceptical strategies that engage with evidential scepticism. In particular, these strategies engage with a sceptical argument that makes central use of a claim that we need antecedent justification to believe that we're not BIVs in order to have justification to believe ordinary empirical propositions. Wright (2004) concedes that we lack evidence for the falsity of sceptical hypotheses, but that we have a different kind of warrant—"entitlement"—to trust that they are false and this supports our evidential justification for ordinary empirical beliefs. Pryor (2000; 2004) claims that we simply don't need antecedent support for the denials of sceptical hypotheses in order to have justification to believe ordinary empirical propositions. We have "immediate justification" for certain ordinary empirical beliefs (and moreover, this can in turn give us justification to believe the falsity of sceptical hypotheses). I will argue that the route of Wright and Pryor looks like a kind of cycle. I don't claim that there is no way out of the cycle. But I claim that once we combine these difficulties, along with considerations about the relationship between the sceptical argument Wright and Pryor engage with and the underdetermination argument from last chapter, we should turn our attention to the underdetermination argument, and more specifically the sameness of evidence thesis.

2.2 Wright and Pryor on Scepticism

I begin with an important point about Wright *vis-à-vis* what I've called the basic externalist response to scepticism. This will help us further understand how Wright sees himself as engaging with scepticism in a deeper way than any of the strategies considered in Chapter One. In a discussion of the externalism/internalism distinction and the satisfactoriness of externalist responses to scepticism,³² Wright analyses the shortcomings of externalism in the following way:

³² Wright understands the internalism/externalism distinction in terms of access. He's talking about *access* externalism in the sense I've identified in the introduction.

Fully to address this reservation would need a complex and extensive discussion. But one immediate observation is that what is put in doubt by sceptical argument is—of course—not our *possession* of any knowledge or justified belief—not if knowledgeability, or justification, are conceived as constituted in aspects of the external situation in which we come to a belief. (How indeed could armchair ruminations show anything about that?). What is put in doubt is rather our right to *claim* knowledge and justified belief. It is this which the project of making out entitlements tries to address and which, on what seems to me to be a correct assumption, externalism is impotent to address (2004, p.210).

The “correct assumption” Wright mentions here is the idea that there is a bifurcation of “epistemic values”, analogous to the bifurcation of “objective” and “subjective” evaluations in the moral domain. What Wright has in mind is the familiar division of evaluations of an attitude or action in terms of its consequences, or actual outcomes, and evaluating an attitude or action in terms of its provenance —i.e. the question of whether it is the result of conscientious or responsible deliberation, for example. Wright’s point seems to be that while externalism tells a compelling story about the rules governing the first sort of evaluation, when it comes to scepticism, what we are really interested in is whether we “respect the canons of intellectual integrity”, i.e. the rules governing the second sort of evaluation.

Wright claims that the sceptical problem is not merely about whether we, as a matter of fact, enjoy a certain kind of cognitive contact with reality (indeed, he thinks *that* is not a question of interest to philosophy). It is a problem about whether we can legitimately *say* that we do, with respect to the most ordinary of our beliefs. And according to Wright, that seems to be a matter of whether, on careful reflection, we are aware of, or have a grasp of, or in some sense have access to whatever it is that supports or justifies these beliefs. At any rate, he claims that meeting this challenge is the most interesting sceptical problem. Here is Wright connecting the desire for the right to claim to know to “intellectual integrity”:

The right to claim knowledge, as challenged by scepticism, is something to be understood in terms of—and to be settled by—canons of intellectual integrity. The paradoxes of scepticism are paradoxes for the attempt at a systematic respect of those canons. They cannot be addressed by a position which allows that in the end thoroughgoing intellectual integrity is unobtainable, that all we can hope for is fortunate cognitive situation (2004, p.211).

Thus, we can see why Wright aims to respond to evidential scepticism: he sees the most interesting form of sceptical paradox as one that challenges our right to claim knowledge; and the right to claim knowledge is tied to having reflectively available support for one's beliefs. But what form of evidential scepticism is Wright interested in? Wright claims that the best and most interesting sceptical paradoxes proceed by identifying a "cornerstone" proposition and claiming that we have no warrant for it.³³ He uses the term "cornerstone" to refer to propositions that we must have warrant to believe if we are to have warrant to believe large classes of certain everyday propositions (Wright 2004, p. 167). That we are not presently BIVs being deceived to believe that we have hands is what Wright calls a "cornerstone of cognitive competence". The sceptic proceeds by claiming that since we don't have warrant for this cornerstone we don't have any warrant for everyday empirical propositions.

Pryor agrees with this set up. Pryor essentially defines the sceptical problem in terms of an appeal to a "sceptical principle of justification" (SPJ), which says:

(SPJ) If you're to have justification for believing p on the basis of certain experiences or grounds E , then for every q which is 'bad' relative to E and p , you have to have antecedent justification for believing q to be false—justification which doesn't rest on or presuppose any E -based justification you have for believing p .

(Pryor 2000, p.531)

It's clear that a lot hangs on what "bad" means, and I'll discuss that below. But we can see that this principle captures the basic idea in Wright's way of understanding the problem. That is, it captures the idea that the sceptic proceeds by claiming that if we are to have justification to believe ordinary empirical propositions, we must already have justification to believe certain presuppositions—namely "cornerstones." So we've seen how both Wright and Pryor understand scepticism. Let's examine how they respond.

³³ Wright's notion of "warrant" covers both evidential justification and "entitlement." Presumably, given Wright's comments about externalism, warrant is the property a belief has when it meets the "canons of intellectual integrity," whether this is via evidential justification or the notion we are about to discuss—i.e. entitlement.

2.3 Entitlement of Cognitive Project

Wright concedes that we don't have evidential justification for cornerstones. But he argues that it does not follow from this that we lack warrant in general concerning cornerstones. He argues that certain cornerstone propositions enjoy a kind of non-evidentially based warrant that he calls entitlement. As we will see below (and as the section above would suggest), Wright is not making a familiar move to some form or another of externalist epistemology. Given that he is not an externalist about warrant, what does he have in mind?

Wright's most compelling defence of entitlement is his argument for "entitlement of cognitive project" (Wright 2004, pp.188-197).³⁴ Entitlement of cognitive project is a kind of non-evidential warrant for a particular propositional attitude towards certain *presuppositions* of our "cognitive projects" (more on the type of propositional attitude in a minute). By "cognitive project", Wright means something like inquiry, or perhaps more specifically the process of acquiring justification for believing a particular proposition by the appropriate exercise of certain cognitive capacities, such as perception, introspection, memory, or reasoning (he does not define "cognitive project"). Here is Wright's definition of a presupposition:

P is a presupposition of a particular cognitive project if to doubt P (in advance) would rationally commit one to doubting the significance or competence of the project (Wright 2004, p.191).

For example, that the world did not spring into existence five minutes ago replete with evidence of a distant past is a presupposition of my historical inquiry about whether Julius Caesar was really assassinated. This is because, if I were to doubt the former, this would rationally commit me to doubting the significance or competence of my historical inquiry. For reasons we'll see in a moment, even though Wright concedes we could not possibly support

³⁴Wright also sketches arguments in favour of "strategic entitlement", "entitlement of rational deliberation", and "entitlements of substance". Strategic entitlement bears a close resemblance to the notion of a "dominant strategy" in game theory (the paradigm example of an argument for this sort of entitlement is Reichenbach's defence of induction (Wright 2004, p.178; Reichenbach 1938)). Entitlement of rational deliberation derives from the constitutive requirements of rational *action* as opposed to inquiry. As Wright puts it, "The generic thought is that since rational agency is nothing we can opt out of, we are entitled to place trust in whatever (we have no evidence against and which) needs to be true if rational decision-making is to be feasible and effective" (Wright 2004, p.198). Entitlements of substance differ from each of the above insofar as the above are all *methodological* entitlements, whereas entitlements of substance are *ontological*. They concern our entitlement to accept the *existence* of, for example, the external world. Wright only offers a highly guarded provisional defence of a very limited form of such an entitlement (Wright 2004, p.200-203).

belief in the denial of this proposition with evidence, we are entitled to take a particular kind of propositional attitude towards it.

The key move behind Wright's strategy is his approach to the idea that *all* inquiry comes with presuppositions. He claims that this is a necessary truth. In undertaking a given inquiry, one *cannot* but take certain things for granted on pain of the given inquiry ceasing altogether. While this may seem like a sceptical thought, Wright flips the seemingly sceptical import on its head:

If there is no such thing as a process of warrant acquisition for each of whose specific presuppositions warrant has already been earned, it should not be reckoned to be part of the proper concept of an acquired warrant that it somehow aspire to this—incoherent—ideal.” (p.24)

The idea is that if it's simply impossible to acquire a presuppositionless warrant, then failure to acquire a presuppositionless warrant is not a shortcoming. The ideal is *incoherent* in the first place. The lesson to draw, Wright says, is that warranted belief must be understood as possible within an epistemological picture on which presuppositions are ineliminable. Wright makes the point as follows:

...[W]arrant is acquired whenever investigation is undertaken in a fully responsible manner, and what the paradox [i.e. his point above that, necessarily, there is no presuppositionless inquiry] shows is that full epistemic responsibility cannot, *per impossibile*, involve an investigation of every presupposition whose falsity would defeat the claim to have acquired a warrant (2004, p.191).

It is important to emphasise that the kind of support Wright has in mind is what he calls *rational* support. By this I take Wright to indicate that he is not making an externalist move. He makes a point of articulating his view in terms of responsible believing (in line with the points discussed in the previous section), where what he means by this is tied to access internalism.³⁵ Wright's key claim in making sense of this is that some of our presuppositions enjoy a kind of “unearned,” or non-evidential support, but it is support one can access by means of a priori reflection. This is entitlement of cognitive project.

³⁵ ...[E]ntitlements, it appears, in contrast with any broadly externalist conception of warrant, are essentially recognisable by means of traditionally internalist resources—a priori reflection and self-knowledge—and are generally independent of the character of our actual cognitive situation in the wider world—indeed, are designed to be so (2004, p.210).

Not just any presupposition is a candidate for entitlement. Wright claims that a presupposition P is an entitlement of cognitive project for S if and only if it meets two criteria:

i) S has no sufficient reason to doubt P.³⁶

ii) Further investigation into P would only lead to more presuppositions in turn of no more secure a prior standing; the attempt to justify P would only lead to a regress of evidential justification. In short, rational investigation into P would merely lead to *cognitive paralysis* as opposed to further epistemic rigour.³⁷

Cornerstones fit precisely this bill. For example, that my eyesight is functioning properly enough to locate my keys is an entitlement of cognitive project as I search for my car keys this morning. Image that there is nothing unusual in the way things appear to me this morning. I have no reason to doubt that my eyesight is functioning well enough to locate my keys. Moreover, halting my search for my keys and investigating whether or not my eyesight is functioning properly enough to locate the keys would itself only lead to further presuppositions which would themselves be of no more secure epistemic standing. For instance, I might stop my search, walk to the doctor's office, and get her to test my eyes. But,

³⁶ Imagine that I undertake to make predictions about today's results at the horse races by reading tea leaves. The effectiveness of this method is a presupposition of my cognitive project. However, clearly we do not want to say that I am entitled to the presupposition that this method is effective. This is precisely the rationale behind condition i). The thought is that I have every reason to doubt that my method is effective. Moreover, the presupposition does not satisfy condition ii). I can engage in further inquiry in discovering reasons to doubt that my method is not effective, and it is not the case that the presuppositions of this further inquiry are of less secure standing than the presupposition that my method is effective. Indeed, Wright notes that the fact that I have reason to doubt its effectiveness entails that ii) will not be satisfied (Wright 2004, p.196). See next footnote for a related issue.

³⁷ The issue surrounding the first condition prompts a bigger one which ultimately has to do with the second condition, and which I will only very roughly outline here. Wright asks: "Suppose I postulate a tract of reality—it might be the realm of non-actual possible worlds as conceived by Lewis—which is spatio-temporally insulated from the domain of our usual empirical knowledge, and a special faculty—as it may be, our non-inferential 'modal intuition'—whose operation is supposed to allow us to gather knowledge about it. Do I have an entitlement of cognitive project to trust the (alleged) faculty on any particular occasion? If not, why not?" (Wright 2004, p.196). The basic worry here is that conditions i) and ii) indeed seem to be met by the presupposition this faculty functions reliably, and as such his conditions on entitlement may seem too weak insofar as we are inclined to resist the idea that Lewis's views about the nature of modality could turn out to be a matter of entitlement (it's not clear to me whether Wright's implying that Lewis's views about the tract of reality, or the faculty by which we come to know it are at issue here). But Wright makes an important distinction in response to this. He grants that he may be forced to accept that we are entitled to something like the reliable functioning of a faculty of modal intuition (p.197). The bullet he does not want to bite, nor does he think he must bite, however, is that his conditions allow entitlement to the specific propositions involved in Lewis' particular view about the nature of modality. The details are complex and would take us too far afield (see Wright 2004, p.196-197).

given that I have no special reason to doubt that my eyesight is functioning properly enough to locate keys, getting information via such a procedure would involve various presuppositions of no more secure an epistemic standing than the presupposition of my present cognitive project. I'm about as sure that my eyes are functioning properly enough to see my car keys right now (in good lighting) as I would be concerning the question of whether the doctor's instruments are functioning accurately, or that the information she gives me is not distorted by pressure to sell me a new pair of glasses, etc.³⁸

The basic anti-sceptical thrust of all this is that, since the sceptic (according to Wright) proceeds by identifying cornerstones and arguing that we don't have warrant to believe them, the sceptic can be defeated. Although we do not have *evidential* justification for cornerstones, we do have entitlement—when it comes to our cognitive projects, cornerstones are presuppositions that meet the criteria for presuppositions to which we are entitled. The general reliability of one's senses is something to which one is entitled when undertaking everyday empirical cognitive projects. This is a cornerstone of cognitive competence. Since the reliable functioning of one's senses excludes precisely the sorts of radical scenarios we've been discussing, that I am not now a BIV is something to which I am entitled. And, again, this is true *despite* the fact that I have no evidence in support of this proposition.

Wright's aim is to secure evidential justification for all of the ordinary sorts of things we take ourselves to have evidential justification to believe—such as that there is a table and two chairs in one's office. The idea is that a background of entitlement to rationally trust cornerstones of cognitive competence can secure evidential justification to believe precisely those things that radical scepticism threatened. I will return to the details of how this is supposed to work below. Wright puts it in terms of the following structure.

- I) It appears to me as though I have hands.
- II) I have hands.
- III) I am not a BIV being deceived to believe that I have hands.

³⁸ Compare a quote from Wittgenstein: "My having two hands is, in normal circumstances, as certain as anything I can produce in evidence for it. That is why I am not in a position to take the sight of my hand as evidence for it" (Wittgenstein 1969, remark 250) Wright acknowledges that his entitlement view is inspired by Wittgenstein.

The thought is that while some kind of support is necessary for III in order for I to provide any kind of evidential support for II, the sceptic is wrong to assume that *evidential* support is needed for III. Rather, so long as we are rationally entitled to III, I provides evidential support for II.

The non-evidential nature of entitlement of cognitive project leads Wright to an important caveat (one that I've been skirting around in the discussion so far). This concerns the sort of propositional attitude that we are entitled to have when it comes to cornerstones of cognitive competence. Wright claims that we are entitled to *rationally trust* cornerstones. He contrasts this with *belief*—which he concedes, for the sake of argument, is essentially evidentially constrained—and with other attitudes of acceptance—such as *taking for granted*—which he claims are compatible with doubt or agnosticism, and which, as such, would introduce an unacceptable level of tension in our web of doxastic attitudes when it comes to connections between propositions that we are entitled to and those for which we have evidential justification. In the next section I discuss these details and an issue that Wright faces in light of them, as well as other concerns for the proposal.

2.4 Relations of Support and the “Leaching” Problem

A well-known objection to this strategy is that Wright has done nothing more than provide *pragmatic* justification for our beliefs. The thought is that an appeal to the avoidance of cognitive paralysis in vindicating our epistemic position with regard to cornerstones is no *epistemic* vindication—it's a practical reason, akin to a “dominant strategy” in decision theory. However, this would be too quick, since Wright clearly states that the sort of paralysis that would ensue were we to doubt cornerstones is *cognitive* paralysis. Thus, rather than giving us a practical reason to trust cornerstones, it seems Wright's view licenses the idea that we have an *instrumental* epistemic reason to trust cornerstones. There is of course a large debate over whether, in the first instance, there can be practical reasons for belief, but also whether epistemic reasons are instrumental reasons. I will not attempt to add anything to this sort of concern for Wright's view here (Pritchard 2005b; Jenkins 2007; Pedersen 2009).³⁹ Another

³⁹ Indeed, it seems to me that the debate about practical reasons for belief, as well as the debate about epistemic instrumentalism is one of the most interesting and active areas in epistemology (Kelly 2003; Berker 2013; Foley 1987; Leite 2007; Asbjorn Steglich-Petersen 2011). As such, criticizing Wright successfully on this score is bound to require a highly sophisticated view about these matters. At the very least this would take us much too far afield. But also, it strikes me that simply pointing out that Wright's view at best licenses instrumental reasons for belief

prominent objection is that since Wright claims we cannot *believe* cornerstones, it's not clear that we can *know* them (insofar as belief is necessary for knowledge) (Pritchard *forthcoming* Ch.3, p.16). And if that's the case, his view has unsettling implications for closure. In particular, it seems that even though he allows we can know ordinary propositions, and that we know that these entail the truth of cornerstone propositions, he must deny that we know the relevant cornerstone propositions. Pritchard's thought is that this undermines the view in virtue of the counterintuitiveness of such a consequence and insofar as Wright himself endorses closure (Wright 2004, p.208). But I'd like to focus on a (related but different) issue that I think pushes us in the direction of a more general worry. This issue surrounds the possibility of the relations of support that Wright is committed to when it comes to our evidential justification to believe ordinary empirical propositions and our entitlement to cornerstones. I turn now to a brief critique of Wright on this score.

According to Wright, belief in everyday propositions can be saved given a background of rational trust in cornerstones. It will be helpful to start with a concern raised by Davies about this. Davies considers the question of whether, once we have secured evidential justification for an everyday belief—say, that I have hands—we cannot “perform a kind of epistemic alchemy” on our warrant for a corresponding cornerstone (Davies 2004, p.222). That is, he notes that the proposition that I have hands *entails* that there is indeed a material world. In light of this entailment, he wonders, what prevents my evidential justification for the ordinary empirical proposition from becoming evidential justification for a belief in the cornerstone proposition?⁴⁰ Davies asks, “Does the I-II-III argument serve to transform the lead of rational trust into the gold of justified belief?” (Davies 2004, p.220).⁴¹ Since Wright has secured evidential support for one's ordinary belief that they have hands, why is it not available to one to competently deduce that since one has evidential support for the belief that they have hands, and since one's having hands *entails* that one is not a BIV, that one therefore has evidential justification to believe that one is not a BIV?

can be taken to beg the question against Wright—we can understand his point about the avoidance of cognitive paralysis as precisely a sort of positing of the notion of an instrumental reason to trust cornerstones. That said, Jenkins' critique engages with the issue in detail and does not argue against Wright simply on such a basis (see Jenkins 2007, p.36).

⁴⁰ Thus, Davies' concern is closely related to Pritchard's just noted above. But I go on to pursue this line of thought in a different direction than Pritchard's.

⁴¹ By “I-II-III argument” Davies just means an argument for a particular view about the inferential/justificatory structure of claims I, II, and III above. The thought is of course similar to the basic structure on the sort of neo-Moorean strategy considered in Chapter One.

The answer lies in Wright’s well-known claims about “transmission of warrant”. Wright discusses this notion in various other papers (Wright 2000; 2002). He calls attention to it in the context of the present discussion: “Type-III propositions cannot be warranted by transmission of evidence provided by type-I propositions for type-II propositions across a type-II to type-III entailment—rather it’s only if one already has warrant for the type-III proposition that any type-II propositions can be justified in the first place” (Wright 2004, p.172). The thought is that because belief in a type-II antecedently depends upon the warrant one has for a type-III, the justification for II cannot work to give new justification for III. In fact, Wright suggests that this condition on transmission of warrant underlies our intuitive dissatisfaction with, for instance, Moore’s proof: it is a case of begging the question, according to Wright, to claim that because one is justified in believing one has hands that one is justified in believing there is an external world (Wright 2002).⁴² It is only against the background, antecedent, belief—or more precisely, rational trust—that there is an external world that one’s belief that one has hands is justified.

Davies suggests, however, that Wright’s condition on transmission of warrant might raise a general concern about the ability of entitlement to save justified belief in type-II’s. Davies notes, “Someone might query whether something less than an antecedent warrant to *believe* the type-III proposition can really secure this favourable outcome for the type-II proposition” (Davies 2004, p.220). In other words, once the point about antecedent support is exhibited explicitly in the considerations about warrant transmission, it seems odd that a belief in a proposition that antecedently depends upon mere rational trust in another can achieve the status of evidentially justified belief in the first place (Davies 2004, p.222). Davies claims that the limitation to mere rational trust in a type-III, “taken together with the idea that we do indeed have a warrant⁴³ to believe the type-II proposition, seems to impose some strain on our

⁴²It is important to note the difference between the failure of *transmission* and the failure of *closure*. Wright endorses the possibility of the former but not the latter. The closure principle, as I finally formulated it in Chapter One, says: If S knows p, and p entails q, and S competently deduces q from p on the basis of knowing that p entails q, then S knows q. A related but different “transmission principle” can be stated as follows: If S knows p, and p entails q, and S competently deduces q from p on the basis of knowing that p entails q, then S knows q *on the very same basis* as the basis on which S knows that p. The difference, of course, is the additional point about the very same basis *transmitting* across the known entailment. And this is what Wright claims can fail. In particular (and this is what I’m calling Wright’s “condition of transmission of warrant”), if q is a proposition for which I require an *antecedent* warrant in order to have warrant for p, then I can know p (or have warrant for p), competently deduce that q from p on the basis of knowing that p entails q, and yet fail to know that q on the very same basis as that upon which I know that p. We can maintain closure, however, by allowing that one knows that q on a different basis (perhaps entitlement).

⁴³ It seems to me that Davies must mean “evidential justification” here.

ordinary thinking about the proper management of our web of belief”(Davies 2004, p.222). Although Davies doesn’t consider this “strain” necessarily a cause for rejecting Wright’s view, he does take it seriously enough to consider a different approach to the problem of our epistemic relationship to type-III’s, or cornerstones. In particular, Davies thinks it might be fruitful to consider whether, rather than a positive propositional attitude towards cornerstones, all that is required to save evidential justification for everyday beliefs is that we *not doubt* cornerstones. He develops a notion of “negative” entitlement which draws on Pryor’s “liberalism” about justification (Davies 2004, p.243; Pryor 2004). As I’ve said, I will discuss Pryor below. For now, note that the basic idea is to challenge the initial assumption which Wright grants to the sceptic—namely, that we need *antecedent* warrant of some kind for cornerstones in order to have justification for everyday propositions (Davies 2004, p.243). However, before considering whether this is a good route to go, I want to examine Wright’s treatment of Davies’ concern.

Wright refers to this concern as the “leaching” problem (Wright 2004, p. 207). He acknowledges that it will be important for his view to show what prevents mere rational trust in cornerstones, upon which (he claims) justified belief in type-II’s antecedently depends, from supplanting genuine evidential justification for those type-II’s—or as Wright puts it, from “leaching upwards from the foundations, as it were like rising damp, to contaminate the products of genuine cognitive investigation” (Wright 2004, p.207). He considers the example of one’s perceptual justification that they have hands: “If [such a] cognitively earned warrant...is achieved subject to mere entitled acceptance that there is a material world at all, then why am I likewise not merely entitled to accept that there is a hand in front of my face, rather than knowing or fully justifiably believing that there is?” (Wright 2004, p.207). Wright’s answer is that there *is* leaching “but it is at one remove and can be lived with” (Wright 2004, p.207). We need to get clear on what he has in mind.

Wright maintains that entitlement to trust cornerstones is all that is needed to secure *first order* evidential justification for everyday empirical beliefs. He seems to take this as an intuitive point, and asserts that if I am entitled to trust that there is a material world, or, say, that my eyes are working properly, then my perceptual experience that I have hands counts as evidence that there are indeed hands in front of me—I am “entitled to claim that my vision is right now a source of reliable information about the local perceptible environment and is hence at the

service of the gathering of perceptual knowledge”(Wright 2004, p.207). Short of arguing for this, Wright appears instead to try to redirect the feeling that there is something odd about it: he notes that what I *don't* have, in virtue of the above entitlements to trust, is *second order* evidential justification to believe the relevant type-II proposition. That is, I do not have evidential justification to believe that I have such evidential justification. “In general, the effect of conceding that we have mere entitlements for cornerstones is not uniformly to supplant evidential cognitive achievements—knowledge and justified belief—with mere entitlements right across the board but to qualify our claims to *higher order* cognitive achievement” (Wright 2004, p.208). According to Wright, the problem of leaching amounts to a limitation on second order evidential justification.

Against the background of an entitlement to trust, I cannot claim that I know that I have evidential justification to believe that I have hands. This is presumably to be understood in light of Wright’s commitments about the connections between knowledge claims and the “cannons of intellectual integrity”. Claims to know invite requests for reasons, they invite requests for evidential justification. But reflection on the presupposition upon which one’s evidential justification for their ordinary empirical beliefs is ultimately based should indicate that I can indeed only *rationally trust* that I have this evidential justification. Importantly, Wright thinks this is enough to go on in a response to scepticism. Recall again from earlier that Wright’s condition of adequacy on a response to scepticism is that we are able to claim the knowledge which scepticism calls into question. The basic idea here is that, while Wright’s strategy does not enable us to claim that we know that we have evidential justification for beliefs such as that we have hands, it does enable us to claim that we know we have hands. And concerning whether or not we should see that as a shortcoming, Wright thinks not. He says, despite our lack of second-order justification for type-II’s, the sorts of knowledge claims we can justifiably make “will still have every point if enough has been done to ensure that all that remains to put the knowledge claim at risk is the possible failure of conditions in which everybody, speaker and audience, (rationally) trusts”(Wright 2004, p.209). That is to say, a lack of second order justification should not seem threatening if all there is to threaten first order justification is the sort of sceptical worry that Wright has shown we have an entitlement to dismiss. I turn now to suggest that these comments of Wright’s about leaching may give rise to a potential worry about the notion of entitlement.

Wright claims that second order evidential justification for type-II propositions would require reflection on what kind of warrant one has for the relevant type-III. As Wright puts it, having second order evidential justification for a type-II requires that one “be able to know the presuppositions of its truth, some of which—we are taking it—sceptical argument has put beyond evidence” (Wright 2004, p.208). In contrast, according to Wright, first-order evidential justification for type-II propositions does not require such reflection. A crucial difference, then, between first- and second-order evidential justification for a belief about a type-II is precisely that, in the first-order case, one is not required to reflect on the kind of warrant one has for the type-III—the presupposition, as Wright has just put it—even though one’s evidential justification in some sense epistemically depends on it. But notice that rather than not *requiring* reflection on one’s warrant to trust cornerstones, Wright’s strategy for limiting the leaching problem seems to imply that agents *must not* reflect on their warrant to trust cornerstones, on pain of the leaching of mere entitlement to trust everyday empirical propositions. If this is so, perhaps entitlement to trust cornerstones is doing less work than Wright would have us believe.

For one thing, it’s not clear that this doesn’t undermine Wright’s commitment to providing an internalist response to the sceptic. After all, entitlements are contrasted with externalist warrant in the sense that they are a priori accessible. But even if Wright can maintain that his view is consistent with internalism, it’s not clear that his internalist approach is as different as he would like to have us think from competitors. Recall Davies’ suggestion that we consider a different approach to cornerstones altogether—namely, what he calls “negative entitlement”. Again, Davies’ hypothesis is that instead of requiring some positive propositional attitude towards cornerstones, all that is required to secure our justification for everyday propositions is that we *not doubt* cornerstones. If the only evidential justification that the putative rational trust in cornerstones secures is that which we have so long as we do not actively reflect upon this rational trust, it is difficult to see the difference between the respective roles that Wright’s rational trust and Davies’ lack of doubt are meant to play in our “web of belief”.

If we want to avoid leaching, we must not reflect on our warrant for cornerstones. Fortunately, our evidential justification to believe ordinary empirical propositions does not require such reflection. But this is the sense in which Wright’s view ends up looking more like his opponent’s than he might want to admit. Wright’s strategy in avoiding the leaching problem

seems to amount precisely to advocating that evidential justification for everyday beliefs is saved so long as we don't reflect upon the ultimate source of that justification—namely, rational trust in a relevant cornerstone. To be sure, the view would dramatically overintellectualize what it takes to have evidential justification if it required otherwise.⁴⁴ But the point is that the foregoing should make us wonder what is so important about having this antecedent warrant in the first place. Why not consider the possibility of a simpler view according to which we must merely not doubt cornerstones in order to have evidential justification to believe ordinary empirical propositions? At this stage in the dialectic, it is worth considering whether a positive trust in cornerstones, as opposed to an absence of doubt in cornerstones, is what is needed.

2.5 From Entitlement to Immediate Justification

Pryor's anti-sceptical strategy contrasts with Wright's in a helpful way. He argues that we do not require justification (or "warrant") to believe the denials of radical sceptical hypotheses antecedently to having justification to believe that we have hands, in order to have justification to believe that we have hands. Thus he rejects the need for entitlement. Pryor's view is that we have *immediate justification* to believe certain "perceptually basic" propositions. This justification can support our knowledge of (and justification to believe) ordinary empirical propositions, as well as the falsity of radical sceptical hypotheses.

As we've seen above, Pryor essentially defines the sceptical problem in terms of an appeal to a "sceptical principle of justification." It is worth restating this principle:

(SPJ) If you're to have justification for believing p on the basis of certain experiences or grounds E , then for every q which is 'bad' relative to E and p , you have to have antecedent justification for believing q to be false—justification which doesn't rest on or presuppose any E -based justification you have for believing p .

⁴⁴Wright does not require of people that they undergo his philosophical reflections in order to have warrant for their ordinary beliefs.

We can see the plausibility of the principle by considering the following example. Imagine that you intend to walk to parking lot 22 and drive home. So you conclude that you will soon be driving your car out of parking lot 22. The basic idea behind (SPJ) is that it seems that in order for this belief to be justified, you need antecedent justification to believe that your car will still be in parking lot 22 when you get there. If you did not *already have* justification to believe that your car will be in lot 22 when you get there (for example, imagine you know that there have been very few car thefts in the neighbourhood in recent years), then your mere intention to drive home does not support a justification for the belief that you will soon be driving home.

However, as I noted earlier, a lot hangs on what “bad” means. Here’s what Pryor says:

So we might want to count a hypothesis as “bad” for the purpose of a skeptical argument just in case it is (and is recognized to be) incompatible with what you purport to know, but it is nonetheless “allowed” by your grounds E, in the sense I described (Pryor 2000, p.527).

And here’s what Pryor means about a hypotheses being “allowed” by grounds E (i.e. “in the sense he described”):

Say that some grounds E you have “allow” a possibility *q* iff the following counterfactual is true: if *q* obtained, you would still possess the same grounds E. (Pryor 2000, p.527).

Pryor says that sceptical scenarios are “bad” in the sense that they are incompatible with ordinary beliefs about the external world, and that even if the scenario obtained one would possess the same grounds for believing what one does about the external world. According to Pryor, this is what *explains* the plausibility of the sceptic’s appeal to (SPJ). That is, it’s because sceptical hypotheses are bad in this way that we must have antecedent justification to believe that they are false in order to have justification to believe ordinary empirical propositions. The way Pryor puts it, both Wright and the sceptic accept the idea that in order to have justification to believe that you have hands, you must have antecedent justification to believe that a “non-perceiving hypothesis” does not obtain. These are such hypotheses as that you are not dreaming, or not a BIV being deceived to think you presently have hands (Wright’s “cornerstones of cognitive competence”).

Pryor rejects (SPJ). In cases of “perceptually basic” propositions, one can have immediate justification to believe them (more on what “perceptually basic” means in a moment). This is a

kind of justification that requires no antecedent justification to believe that a “non-perceiving hypothesis” does not obtain. An example of a perceptually basic proposition is that there is a red patch on the wall before you. Pryor thinks the belief that “here is a hand” is perceptually basic. So, Pryor claims, we can have immediate justification to believe that we have hands. As Pryor puts it, whether you are “liberal” or “conservative” about justification depends on which hypotheses you require an agent to have antecedent justification to believe. If you do not require antecedent justification to believe that a non-perceiving hypothesis does not obtain (“I am not a BIV”, “I am not dreaming”), then you are a liberal about perceptual justification. Liberals do not claim that all presuppositions can be safely ignored—Pryor accepts the parking lot example, for example (and he discusses others (Pryor 2004, p.362)). But the claim about non-perceiving hypotheses distinguishes his strategy from Wright’s in the relevant way.

To be sure, Pryor claims that if one *doubts* whether or not one is a BIV, then they will lack “doxastic justification” to believe that they have hands. The distinction between propositional and doxastic justification is a familiar one in epistemology. The difference is between *having justification* to believe a given proposition (propositional) and coming to believe, or *basing* your belief on that justification in the right sort of way (doxastic). Part of what *defines* liberalism is that there can be relations of epistemic dependence such that were the agent to *doubt* q, they would lack doxastic justification to believe that p. The thought is that even though this sort of dependence between p and q may hold, this does not imply that the agent needs antecedent justification to believe q in order to have justification to believe p. Thus, another crucial distinction for Pryor comes into play here. Doubting whether one is a BIV is a matter of (a problem with) the beliefs one happens to have. Such doubt does not affect what one *has justification* to believe.⁴⁵ Pryor discusses the way in which a doubt about the truth of radical sceptical hypotheses can *rationaly obstruct* you from believing that you have hands (and thus obstruct you from having doxastic justification). But it does not affect the justification that you have to believe that you have hands.

⁴⁵ Pryor’s view is about what he calls “*prima facie* justification”. He allows that if one has defeating evidence against a cornerstone proposition, one may lack all things considered propositional justification for an ordinary empirical belief. But so long as one lacks defeating evidence—or so long as one’s defeating evidence is itself defeated—one’s *prima facie* justification constitutes all things considered justification. By defeating evidence, Pryor stipulates that he means “ordinary evidence of the sort employed by the man in the street and by the working scientist” (Pryor 2000, p.534), and maintains that a priori sceptical arguments do not introduce defeating evidence according to his usage.

Pryor claims, then, for example, that Moore’s argument from his hands exhibits a good “justificatory structure”.⁴⁶ What he means is that Moore’s argument is *not* a case of transmission failure. You can have justification to believe that you have hands, competently deduce that if you have hands then you are not a BIV, and come to have justification to believe on that very basis that you are not a BIV. The whole point of his anti-sceptical strategy is to argue that an argument like Moore’s has this good justificatory structure, but to defend this claim by drawing crucial distinctions—for example between liberalism and conservatism and other “targets of evaluation” when it comes to assessing arguments—to convince his readers that they can enjoy and, if need be, make use of this justificatory structure.⁴⁷ His story is intended to remove an agent’s rational obstruction to having justification to believe the falsity of sceptical hypotheses. And if one is not rationally obstructed, one does not need to hear Pryor’s philosophical story in order to have justification to believe the falsity of sceptical hypotheses.

2.6 Immediate Justification

Given the way Pryor defines the sceptical problem, it’s no surprise that liberalism has anti-sceptical consequences. Liberalism entails the denial of (SPJ) (on any sceptical reading of it). But Pryor obviously needs some kind of story to tell about where our justification to believe ordinary propositions and the denials of sceptical hypotheses comes from. How does he defend the idea that we have *immediate* justification to believe certain “perceptually basic” propositions? And what is it in virtue of which immediate justification makes people justified?

Pryor is careful to note what immediate justification is *not*. Immediate justification is not infallible justification—there is nothing about this kind of justification that suggests it is not defeasible, according to Pryor (Pryor 2000, p.532). Immediate justification does not imply *autonomy* of belief: it is compatible with the notion of immediate justification that to hold an

⁴⁶ Recall the formulation of Moore’s argument from Chapter 1.

⁴⁷ Pryor claims there are at least five targets of evaluation when assessing an argument like Moore’s. We can assess its validity as a “proof” (i.e. a sequence of propositions and derivation-rules), its justificatory structure, the reasoning Moore engaged in, the dialectical power of the argument, and its effectiveness as a philosophical response to the sceptic. The most important distinction, I think, which cuts across this five-way distinction in different ways, is the distinction between rational obstruction and having justification to believe. The reason why it is important is that it allows Pryor to recognize an important sense in which Moore’s argument is “defective” (i.e. it can’t help you rationally overcome your doubt in the conclusion, if indeed you already doubt the conclusion), but which is not a defect that undermines its anti-sceptical potential (since antecedently doubting the conclusion can only be due to a fault in *your* beliefs, not to the—supposed lack of—justification you have to believe that radical sceptical hypotheses are false).

immediately justified belief requires that one has *other* beliefs; the point is that other beliefs need not be involved in the justification itself. They can be background beliefs, in the way that background beliefs can still play a role in *a priori* knowledge (Pryor 2000, p.534). Crucially, Pryor also notes that it may be difficult to assess when a proposition is perceptually basic, as opposed to when it is the result of inference. *Only* perceptually basic propositions are immediately justified (Pryor 2000, p.539). He claims that although this is a tricky issue, it does not detract from the idea of immediate justification. It seems clear, says Pryor, that at least some things can immediately justify (Pryor 2000, p.539).⁴⁸ He casts his view as a sort of “sensible philosophical conservatism” (Pryor 2000, p.538). That is to say, he thinks of his position as commonsensical and phenomenologically accurate. As such, he sees the main task in defending it as showing why various objections are only apparently good.

Pryor says that there are a variety of positive stories to tell that are consistent with his theory about *why* perceptually basic propositions have the epistemic standings that they do.⁴⁹ The story he favours, however, is that the phenomenology of its “seeming to you that you ascertain that p” is something that has the relevant epistemic significance.

I think there’s a distinctive phenomenology to 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 (Pryor 2004, p.357).

According to Pryor, this distinctive phenomenology is what *makes* you justified in cases of immediate justification to believe perceptually basic propositions. It’s important to note that Pryor is not a reliabilist. He does *not* claim that this distinctive phenomenology makes you immediately justified because it is a reliable indicator of the truth. He endorses “simple internalism” (Pryor 2001) (for our purposes, this can be understood as the same thing as supervenience internalism). That is, he claims, for example, that if your recently envatted twin were to come to believe that they have hands because of this distinctive phenomenological experience (and were they to have been otherwise epistemically responsible), they would,

⁴⁸ This may be more of an issue for Pryor than he makes it out to be. I return to this below.

⁴⁹ For example, we might think that perceptually basic experiences are irresistible, and that justification is a matter of believing what you ought to believe, and because an ought implies can principle is true, it can’t be the case that you ought not believe something that it’s irresistible for you to believe. Or, we might think that there is something special about the concepts we employ in our perceptual beliefs that makes those beliefs epistemically appropriate responses to the experiences. Pryor says that he isn’t sympathetic to either of these approaches (Pryor 2000, p.356).

according to Pryor, count as justified. It does not affect the agent's justification that they formed their belief in a completely unreliable way.

2.7 Locating the Perceptually Basic and the Problem of Easy Knowledge

I want to raise one potential problem for Pryor before raising a deeper worry. It has to do with the notion of perceptually basic propositions. This notion stands in need of a good deal of clarification. And it's not clear where Pryor is going to find the resources for clarification. Recall that Pryor casts his view as a kind of "modest conservatism", which is to say that he thinks the main business for systematic epistemology in arguing for his view is to defend the view against objections. The reason for this is that he thinks the view is phenomenologically accurate and occupies a sort of privileged place in commonsense. Here is a quote in which Pryor explains his defence of immediate justification in Pryor (2000): "I emphasized there how simple and intuitively appealing the view is; I said the main work for systematic epistemology should be to defend the view against challenges" (Pryor 2004, p.356). He adds in (2004) that systematic epistemology can do more, but he doesn't elaborate on the notion of perceptually basic propositions (Pryor 2004, p.537). Insofar as he needs to appeal to commonsense and ordinary epistemic practice to do so, it looks like he is in trouble. This is because everyday claims about what propositions people "seem to ascertain" to be true outstrip the sorts of propositions Pryor will be comfortable defending as ones for which we can have immediate justification. Pryor himself makes the point:

It's not easy to discern what propositions we 'seem to ascertain,' and what propositions we merely unreflectively infer. Our perceptual reports don't track the difference very closely. When you have the right kinds of background evidence, you'll unhesitatingly say things like 'It looks as if the police have arrived' and 'I see that the Smiths have already left for Australia'. But I think there is a real difference. And my theory is that what we seem to ascertain, we thereby have immediate *prima facie* justification to believe (2004, p.357).

But we need more of a story about an in-principle line to draw between what we seem to ascertain and what propositions we unreflectively infer. There may be a serious tension between Pryor's ability to draw a principled line around the perceptually basic and his ability to defend his view as a kind of modest conservatism. Do the folk share his intuitions about what counts as perceptually basic? The way we talk doesn't obviously suggest so. The tension is serious insofar as much of Pryor's defence of immediate justification *depends* on casting it as a

kind of modest conservatism. That is to say, he is unable to provide much by way of positive support for the idea of immediate justification (Pryor 2000, p.532-539; 2004, p.537).

A more serious and well-known objection to liberalism about justification generally is the “problem of easy knowledge” (Cohen 2002; Vogel 2008; Wright 2002; Kallestrup 2012).⁵⁰ Consider the following case. You are driving along in your car and you read that the gas gauge says “full”. So you form the belief that the tank is full. You continue driving and the gauge reads “half full”. So you form the belief that the tank is half-full. Repeat this process any number of times. After a time, imagine that you reason as follows: at time t1 the gauge read full, and the tank was full, at time t2 the gauge read half-full and the tank was half full, and so on. You conclude that the gauge is working properly on the basis of this excellent track record. Thus, you’ve come to form a belief that the gauge is reliable. But this seems too easy: one cannot form a justified belief about the reliability of one’s gas gauge on the basis of the readings of the gauge alone. It seems like that’s the sort of thing you would need to do some kind of independent check in order to verify (such as using a dipstick). If the liberal about justification allows that you do not need antecedent justification to believe that the gauge is working properly in order to form a justified belief that the gas tank is full on the basis of what the gauge says, then it’s not clear how the liberal about justification can avoid this too easy sort of knowledge (I return to the question whether the liberal would indeed claim you do not need such antecedent justification). At each stage in this “track record argument”, your belief about what is in the tank is justified independently of any antecedent support for believing that the gauge is working properly. Your belief about what the gauge says is also of course justified. So what’s preventing you from reasoning in this way to the general conclusion about the reliability of your gas gauge?

Notice that there is a disanalogy between the gas-gauge case and someone reasoning from their justification to believe that they have hands to the conclusion that they are therefore not a BIV, as Pryor explicitly allows. Cohen (2002) claims that the latter may look like a plausible enough thing to allow, but for the following reason: “I think it may look plausible only because

⁵⁰ The problem is also raised in the context of the internalism/externalism debate (i.e. as a problem for reliabilism). But I think the problem is more clearly relevant to the liberalism/conservatism debate. This is because the problem plausibly arises for the reliabilist precisely because of the reliabilist’s liberalism about justification. Moreover, as we’ve seen, Pryor is himself an internalist (a supervenience internalist) – so if he faces this problem, it looks like the problem of easy knowledge is a problem for internalists as well as reliabilists, and shouldn’t be seen as motivation for internalism *per se* (compare Vogel 2008).

it is obscure in general how we know global skeptical alternatives do not obtain, e.g., how we know we're not brains-in-a-vat. And so insofar as we *are* inclined to say we do know such things, this can seem like a reasonable hypothesis about how we know” (Cohen 2002, p.313). Meanwhile, the idea of coming to know about the reliability of one's gas gauge via the sort of reasoning just discussed does not look plausible because it is not obscure in general how we know about the reliability of our gas gauges. That is to say, such knowledge is intuitively gained by independent checks, such as comparing what the gauge says with a dipstick (see Vogel 2008, p. 518). Thus, the case brings to light an implausible implication of immediate justification in the sense that it demonstrates how the view licenses knowledge via methods that are at odds with intuitive judgments about what would normally be required for that knowledge. So, while the problem of easy knowledge may not have been immediately apparent when discussing Pryor's view about immediate justification qua anti-sceptical strategy, the implications it has for beliefs, say, about gas gauge reliability, are more intuitively worrying.

There are a number of things that the liberal can say in response. For one thing, they would do well to note that this issue afflicts a broad range of views in epistemology, not just a view like Pryor's (Comesaña 2013). If this is a problem it is a problem for good deal of contemporary epistemology, so it's unfair to burden Pryor's view about immediate justification in particular with the objection. Moreover, a liberal may want to bite the bullet and claim that we can indeed form justified beliefs about the reliability of our gas gauges in this way—or at least, they can claim that this sort of reasoning does provide *some* justification to believe that one's gauge is working properly (Pryor seems to take this approach). Another response might be to allow for liberalism while resisting the track record argument in some other way. For example, one might argue that this is a case of “bootstrapping”, and while there is a problem with bootstrapping, this is not a problem with liberalism *per se* (Kallestrup 2012).

Perhaps most importantly, we should take note of the fact that Pryor is not a liberal about *all* presuppositions.⁵¹ We've seen how he is a liberal about perceptual justification, where this means that he thinks we do not need antecedent justification to believe that “non-perceiving hypotheses” obtain in order to believe such things as that we have hands. Again, these are hypotheses such as that one is now dreaming, or that one is a BIV. It is open to Pryor to respond to the gas gauge case by claiming that he is not a liberal about the relevant

⁵¹ See Pryor 2004, p. 362, and below.

presuppositions in that case. He may indeed want to require that we have antecedent justification to believe that the gauge is working properly in order to have justification to believe what it says is in the tank. But this raises a question about Pryor's demarcation of the sorts of hypotheses it is legitimate to be a liberal about. Why exactly would one be liberal about non-perceiving hypothesis but not hypotheses about the reliable functioning of our gas gauges? After all, it's not as though we go around asking each other whether we know that our gas gauges are reliable when someone says they need fuel. And this is precisely the sort of motivation driving the idea that we don't need antecedent justification to believe that we are not BIVs when we make claims about our immediate environment. Consider that Pryor is not a liberal about the parking lot case discussed above. Again, the reasoning was as follows: I intend to walk to lot 22 and drive home; so I *will* walk to lot 22 and drive home; so my car will be in lot 22 when I get there. Pryor explains: "This argument sounds bad to us because we think your intention is not enough, by itself, to justify you in believing you'll succeed in driving your car out of the lot. We think you also need antecedent justification to believe your car is still in the lot" (Pryor 2004, p.360). Notice that this explanation is very similar to my diagnosis of the gas gauge case. That is, we need antecedent justification for the relevant hypothesis because a route to justification is proposed that conflicts with intuitive judgments about what would normally be required in order to have such justification.

Although Pryor may deny that he is a liberal about the latter sorts of cases, he does not deny that he is a liberal about an additional sort of case. Imagine that you are looking at what appears to be a red wall. Necessarily, if the wall is red then it is not a white wall being illuminated by tricky red lighting. So, it seems open to you to competently deduce that you are not looking at a white wall being illuminated by tricky red lighting (Pryor 2004, p,362).⁵² What Pryor requires in this case is that one not already doubt that the wall is being trickily lit (that would rationally obstruct you from having justification to believe that the wall is not being lit by tricky red lighting). This is important because I think precisely the same diagnosis as to why this is counterintuitive applies in this case as in other cases in which Pryor is not a liberal. That is to say, just as in the gas gauge case it's intuitive that in order to know that one's gas gauge is functioning reliably one must make some kind of check independently of reading the dial, it's

⁵² "Here too I think your visual justification to believe the premise makes the conclusion more credible for you. Your justification to believe the wall is red contributes to the credibility of the claim that the wall isn't white but lit by tricky red lights" (Pryor 2004, p.362).

intuitive that in order to come to know whether a wall is lit by tricky lighting one must make some kind of check independently of what the present appearance of redness indicates (such as going in for a close inspection of the room). After all, the appearance of a red wall is precisely what the tricky lighting is designed to imitate. So I think, again, this is pressure for more of an in-principle demarcation of the sorts of cases it's legitimate to be a liberal about.

2.8 The Liberal/Conservative Cycle

All of these cases are precisely the sort of case that Wright would characterize as exhibiting transmission failure. Wright would claim that you cannot come to have justification to believe that you are not looking at a white wall being illuminated by tricky red lighting in this way. And this is because—he says—you require antecedent justification to believe that the latter hypothesis does not obtain in order for the appearance as of a red wall to provide evidence that you are looking at a red wall. Indeed, Wright's entitlement strategy is designed to avoid the problem of easy knowledge. He claims that we do require antecedent support for propositions such as that the gas gauge is working properly in order to have justification to believe what's in the tank on the basis of what the gauge says. Thus the present sorts of cases are appealed to by Wright and other conservatives precisely to motivate conservatism. The trouble with conservatism when it comes to sceptical scenarios is that it's just not clear where the antecedent support comes from. This is what forces Wright to concede that we do not have *evidential* support for cornerstones of cognitive competence. He maintains that we have entitlement to rationally trust such cornerstones. However, I've argued that given Wright's own constraints on an adequate anti-sceptical strategy, it's difficult to see what work entitlement actually does in the justificatory structure of our everyday empirical beliefs. I've argued that Wright's condition of non-reflection on entitlement to trust cornerstones makes it difficult, on closer inspection, to see the difference between his conservatism and Pryor's liberalism. Perhaps it's not clear that entitlement can help with the problem of easy knowledge after all. Indeed, we seem to be stuck in a cycle rather than getting any closer to a satisfactory way of thinking about external world scepticism.

I don't want to claim that there is no way to break out of the cycle. But I take the foregoing to suggest that we should return to Pryor's point about what explains the plausibility of the sceptic's appeal to (SPJ). Pryor himself says that sceptical scenarios are "bad" in the sense that they are incompatible with our ordinary beliefs about the external world, but that one's

evidence would be the same even if the sceptical scenario obtained. According to Pryor, this is what explains the plausibility of (SPJ). It's because sceptical hypotheses are bad in this way that we must have antecedent justification to believe that they are false in order to have justification to believe ordinary empirical propositions.

Notice that one potential way of rejecting the sceptic's appeal to (SJP) would be by challenging the sameness of evidence thesis that's also at play in the underdetermination argument. That is, we might try to undermine the plausibility of the sceptic's appeal to (SPJ) by claiming that it's *false* that one's evidence (or "grounds") would be the same if one were a BIV (that is, that one would still possess the same grounds E if q obtained). In effect, we'd be denying that our grounds E always "allow" the obtaining of a radical sceptical scenario in Pryor's sense above, and thus that a condition on the hypothesis as counting as "bad" is not met. In other words, perhaps a successful response to the underdetermination argument would provide resources for a response to Wright and Pryor's sceptical problem as well.⁵³

In any case, it seems clear that considerations about underdetermination of evidence are at work in Wright and Pryor's formulation of the sceptical problem. Moreover, it's not clear whether either Wright or Pryor has resources to deal with the underdetermination problem. At the very least, they do not explicitly consider that problem. When it comes to Pryor's view, Pritchard puts his concern in the following way: "Let us state the matter baldly: why should its merely seeming to one as if p give one better rational support for believing that p than for believing that one is the victim of a known to be incompatible radical sceptical hypothesis?" (Pritchard *forthcoming*, p.14). Regardless of whether Pryor's view has consequences for the underdetermination problem, Pritchard's point is intended to motivate his own approach to the sceptical problem given that he has a sophisticated story to tell about how one's evidence can favour the proposition that they have hands over that they are a BIV.

2.9 Conclusion

We've looked at two strategies for responding to a kind of scepticism about evidential justification. I've argued that they move in a cycle. We don't seem to be getting any further along with external world scepticism by engaging with (SPJ). Moreover, I've suggested that a crucial feature of what makes the sceptic's appeal to (SPJ) plausible is a commitment to

⁵³ I won't be considering this point in any detail.

precisely the claims that are at work in the underdetermination argument. In particular, the sceptic's appeal to (SPJ) seems plausible at least in part because the sameness of evidence thesis seems plausible. Important issues may remain when it comes to the particular form of scepticism Wright and Pryor are interested in, but the discussion of this chapter indicates that we should turn our attention to the sameness of evidence thesis.

Chapter 3: Sameness of Evidence

3.1 Introduction

In this chapter I introduce two ways of challenging the sameness of evidence thesis. These strategies engage directly with the underdetermination sceptic's claim that one's evidence does not favour that one has hands over that one is a BIV being stimulated to believe that one has hands. I begin with an approach known as "McDowellian neo-Mooreanism" (Pritchard 2006; 2007; 2012; *forthcoming*). McDowellian neo-Mooreanism is motivated in part by its ability to accommodate internalist and externalist intuitions in a particularly effective way. I'll argue that this strategy faces a difficulty in meeting a particular kind of explanatory challenge. This is the challenge of explaining how one's evidence could be different in good and bad cases that are phenomenologically indistinguishable. Pritchard's difficulties motivate turning to Williamson's evidential externalist strategy (2000; 2009). If plausible, this approach to evidence would have a relatively easy time rejecting the sameness of evidence thesis. This is because it abandons the condition of "reflective access" that Pritchard aims to preserve. However, I will argue that this view raises an important issue about epistemic normativity.

3.2 McDowellian neo-Mooreanism

Pritchard claims that McDowellian neo-Mooreanism satisfies core intuitions motivating both internalist *and* externalist epistemology. The strategy aims to vindicate knowledge of the denials of sceptical hypotheses. It argues that we have this knowledge in virtue of our evidence.⁵⁴ The evidence we have comes in the form of *factive* mental states. These are mental states that entail the truth of some proposition, such as that one can *see* that things are thus and so. Despite its appeal to *factive* mental states, Pritchard claims that he is not an evidential externalist (i.e. someone that claims that we do not have a special kind of access to facts about what evidence we have). Rather, McDowellian neo-Mooreanism maintains a condition of reflective

⁵⁴ I should note that in recent work Pritchard switches from talking about evidence to talking about "rationally supported knowledge" (Pritchard *forthcoming*). This move may allow Pritchard to cast his anti-sceptical strategy in terms that would be acceptable to anti-evidentialists. But that all depends on what Pritchard means by "rationally supported knowledge". I think the best way to understand what he has in mind is in terms of evidence. Rationally supported knowledge is evidentially supported knowledge. In any case, I do not think this switch is important for my purposes. Moreover, I think that given the way I've set up the debate, thinking of McDowellian neo-Mooreanism in these terms is a fruitful strategy. That said, Pritchard has also switched to casting underdetermination-based scepticism as a problem about rationally supported knowledge. So he would obviously be uninterested in the claim that thinking in terms of evidence is fruitful.

accessibility on one's evidence. We'll need to get clear on why one might think that reflective access to our evidence is an interesting and important thing to have. I'll take up the bulk of that question in the next chapter. For now, I'll note that one reason why Pritchard thinks so is because he thinks that this is the ordinary or commonsensical way of thinking about what's required for justification. He takes it as an important constraint on epistemological theorizing to avoid revisionism whenever possible.

For example, comparing McDowellian neo-Mooreanism with the strategy of Sosa from Chapter One, Pritchard offers the following remarks about why McDowellian neo-Mooreanism is a preferred form of neo-Mooreanism:

[O]n the face of it the McDowellian view is plausibly in a better position to be counted as the true heir to the Moorean tradition on account of how it retains core internalist intuitions. After all, part of the desiderata of neo-Moorean positions is that they are able, where possible, to accommodate our pre-theoretical intuitions, and internalist intuitions are surely highly embedded within folk epistemology. Moreover, since both views appear to endorse some form of disjunctivism, it seems that one could similarly argue that what the McDowellian picture highlights is how disjunctivism, properly understood, enables one to evade the sceptical problematic *without* having to resort to the revisionism of classical epistemic externalism (Pritchard 2007, p.96).

And here is Pritchard commenting on the advantage of McDowellian neo-Mooreanism over externalist strategies generally:

Taking the externalist route out of the sceptical problem is not without its own difficulties, however, since there clearly is a strong intuitive pull towards epistemic internalism. Accordingly, the worry one might have is that by opting for an externalist version of neo-Mooreanism one fails to stay true to the commonsense credentials of Mooreanism. If such revisionism is unavoidable, then this price might be worth paying to get a way out of the sceptical predicament. My interest here, however, is in whether there is a version of neo-Mooreanism available which is true to our internalist intuitions but which avoids the problems facing classical internalist renderings of the neo-Moorean thesis (Pritchard 2006, p.10).

The key theme is that Pritchard seems to understand internalism as a pre-theoretical, or natural epistemological position in line with common sense epistemology. At least part of his motivation for arguing in favour of the reflective accessibility of our evidence stems from a commitment to avoiding revision. I will leave to one side the issue of whether Pritchard is correct about the claim that internalism enjoys the title of a pre-theoretical or commonsense epistemological outlook. Indeed, one can easily imagine externalists objecting to this by

pointing out their own preferred range of cases, such as cases of small children, animals, and chicken-sexers. But another motivation for Pritchard's commitment to an internalist condition of reflective accessibility turns on considerations about epistemic responsibility. Here is Pritchard:

For if the facts in virtue of which one's beliefs enjoy a good epistemic standing are not reflectively available to one, then in what sense is one even able to take epistemic responsibility for that epistemic standing? (Pritchard 2012, p.2)

Pritchard seems to accept the idea that an agent counts as in some sense epistemically responsible just in case the facts in virtue of which one is justified meet an internalist access condition. I will return to this point below. The point here is simply to get a handle on why Pritchard is interested in reflective accessibility.

Of course we haven't explained in any detail yet what "reflective access" is according to Pritchard. Indeed, the burden is on the McDowellian neo-Moorean to provide an adequate account of reflective access, since, as we've just seen, this is a key notion when it comes to the McDowellian neo-Moorean's putative advantage over other forms of neo-Mooreanism. Here is what Pritchard says about access in connection with his understanding of internalism:

By epistemic internalism, I mean *access* internalism such that what makes an epistemic condition (i.e., a condition which, perhaps in conjunction with other epistemic conditions, can turn true belief into knowledge) an internal epistemic condition is that the agent concerned is able to know by reflection alone those facts which determine that this condition has been met. Meeting the justification condition, for example, at least as it is standardly conceived, involves the possession of grounds in support of the target belief, where these grounds—and the fact that they are supporting grounds—is reflectively accessible to the subject (Pritchard 2006, p.6).⁵⁵

It's not entirely clear how we are to understand the notion of "reflection". But I think we can understand it in terms of "introspection", where the main point is that this is not an ordinary empirical kind of access. And presumably the notion of "access" is to be spelled out in epistemic terms. Having access to one's evidence is being in some sort of position to *know* that one has that evidence. The idea that this is reflective or introspective access indicates that, according to Pritchard, the internalist is interested in a special sort of epistemic position. As I

⁵⁵ Note that this is not an articulation of mere supervenience internalism. Pritchard claims that, according to the internalism he's interested in, the *fact that* one's supporting grounds are supporting grounds is accessible to one. This is an essential condition of what I've called access internalism. As we'll see below, Pritchard claims that the McDowellian neo-Moorean brand of internalism is a "non-classical" form of access internalism, however, in that it allows that such supporting grounds can be *factive*.

said in the introduction, internalists are understood as requiring that agents are in an epistemic position with respect to their evidence such that they could not be non-culpably mistaken about what their evidence is, or that they are always in a position to know what they're evidence is (plausibly two very closely related ways of putting the same idea). In other words, it seems fair to assume that Pritchard is interested in what I have called the "default" notion of accessibility.⁵⁶ However, we will return to this point.

Insofar as it claims that we have reflective access to factive mental states, Pritchard claims that McDowellian neo-Mooreanism sits somewhere between "classical" internalist and externalist views (Pritchard 2006; 2007; 2012). It putatively satisfies the accessibility condition of internalism, yet it simultaneously satisfies the truth conduciveness condition that externalist views are so well-positioned to satisfy (after all, factive reasons to believe that *p* entail that *p* is true). Because one's evidence in the good case is factive, it "reaches all the way out to the world". It is precisely because of this position between classical internalism and externalism that Pritchard's strategy looks like an interesting approach to supporting a rejection of the sameness of evidence thesis. According to Pritchard, in any case, it is in virtue of an ability to remain non-revisionary and to respect intuitions about epistemic responsibility that his strategy is a good way of rejecting the sameness of evidence thesis.

3.3 Epistemological Disjunctivism

The root of McDowellian neo-Mooreanism is a kind of *disjunctivism*. We can draw a distinction between two broad forms of disjunctivism. Disjunctivism about *perceptual content* (also known as *metaphysical* disjunctivism) says that veridical and non-veridical perceptual experiences have two metaphysically different sorts of content.⁵⁷ Perhaps counter-intuitively, veridical and non-veridical experiences do not share a "highest common factor", despite being phenomenologically indistinguishable. A highest common factor is simply something in common between cases of veridical and non-veridical perception. We might think that such a view has immediate epistemological consequences:

⁵⁶ That is, a notion of accessibility understood in terms of a "strong" access condition. As I said in the introduction, "strong" access says that one is in a position to know that condition *C* obtains when condition *C* obtains *and* that one is in a position to know that condition *C* does not obtain when it does not obtain. All I claim here is that given Pritchard's interest in preserving access internalist intuitions, it is fair to understand Pritchard's access condition in this way. However, we'll see below that things are not so clear.

⁵⁷ For recent discussions of metaphysical disjunctivism, see Haddock & Macpherson (2008*b*), Byrne & Logue (2009*b*), Fish (2009), Brogaard (2010) and Dorsch (2008).

One would naturally expect [content] disjunctivism to have ramifications for one's epistemology, since on most standard views of perceptual knowledge the nature of one's perceptual experience forms part of one's perceptual evidence, and thus if one's experiences are different in the two cases just considered, then so too must be one's evidence (Pritchard 2006, p.11).

However, Pritchard has argued elsewhere that the kind of disjunctivism he is interested in does not entail content disjunctivism (Pritchard 2012). The basic thought seems to be that there is no immediately obvious reason why one's experiences in different cases could not be metaphysically different while nevertheless giving one the same evidential support for a given belief in each case. In any case, this is why Pritchard goes on to say:

Rather than consider the general issue of whether disjunctivism about perceptual experience has this consequence for perceptual evidence, I want to focus instead on the specific McDowellian version of disjunctivism which directly expresses the disjunctivist point in epistemic terms (Pritchard 2006, p.11).

Pritchard is interested in a McDowellian form of disjunctivism, to be distinguished from content disjunctivism. He casts this form of disjunctivism directly in epistemological terms. It says that there is a difference in our *evidence* for believing that p in the good case and in the bad case. Again, this is true despite the phenomenological indistinguishability of the good and bad cases. In the good case, our evidence is *factive*—one's evidence is that one can *see* that p. One's evidence reaches all the way out to the world in the sense that it entails that such and such is the case. In the bad case, one's evidence is a mere non-factive mental state (it does not entail the truth of p)—for example, one's evidence is that it merely *seems* to one that p. Crucially, again, according to Pritchard, that one *has* this evidence is reflectively accessible to one. One has reflective access to the fact that they can *see* that p.

Before getting into further details, let's take a closer look at the precise way in which McDowellian neo-Mooreanism responds to radical scepticism. I've said that the central move is to challenge the sceptic's sameness of evidence thesis. As we've seen (Chapter One), that thesis motivates the sceptic's favouring claim:

S's evidence does not favour that they have hands over that they are a BIV being stimulated to believe that they have hands.

By hypothesis, the sceptic claims that our experiences would be phenomenologically the same in each case and this is why one's evidence doesn't favour that they have hands over that one is a BIV. The move from phenomenological sameness of experience to evidential sameness is precisely what needs further investigation. But this claim works in combination with the underdetermination principle in order to draw the sceptical conclusion. Recall that (UP) goes as follows:

(UP) For all S, p, q, if S's evidence for believing p does not favour p over some hypothesis q which S knows to be incompatible with p, then S's evidence does not justify S in believing p.

The sceptic claims that because our evidence does not favour the proposition that one has hands over the proposition that one is a BIV being stimulated to believe that one has hands, one's evidence does not provide justification for one's belief that they have hands (and insofar as we need evidential justification for knowledge, we do not have knowledge that we have hands).⁵⁸ The McDowellian neo-Moorean challenges the favouring claim in this line of reasoning. He claims that, in the good case, our evidence *does* favour the proposition that we have hands over the proposition that we are BIVs. After all, in the good case, our evidence consists in the fact that we can *see* that we have hands—which entails that one has hands. Thus in a very strong sense our evidence favours that we have hands over that we are BIVs (in the good case). Moreover, since having hands in turn entails that one is not a (handless) BIV, by performing the relevant deduction, one can know that they are not a BIV. In other words, like all neo-Moorean anti-sceptical strategies, a distinctive feature of the view is that we can have knowledge of the falsity of radical sceptical hypotheses.⁵⁹ Moreover, we can have evidentially-justified beliefs about the falsity of radical hypotheses. As such, this sort of view, if correct, may provide a way of responding to Wright's and Pryor's (SPJ) sceptic as well.

⁵⁸Pritchard motivates McDowellian neo-Mooreanism in part by its ability to grant this conception of the connection between evidence and justification and knowledge (see Pritchard 2007, p.94).

⁵⁹McDowellian neo-Mooreanism is cast as being able to respond both to the closure problem and the underdetermination problem (see Pritchard 2007; Brueckner 2010; Cohen 1998). I will return to this question in Chapter 5. There are a variety of features of McDowellian neo-Mooreanism that I am suppressing in this discussion in the interest of staying focused on my central issue. For example, it is claimed that the view has resources for accounting for the conversational oddness of explicit claims to know the falsity of sceptical hypotheses (Pritchard 2006; 2012).

A crucial question for McDowellian neo-Mooreanism concerns the motivation for epistemological disjunctivism. After all, appealing to epistemological disjunctivism is precisely the way in which the view is supposed to make headway on other forms of neo-Mooreanism—such as Sosa’s—by allowing for a response to the sceptical problem that speaks to the question of the evidential basis of our knowledge of the falsity of sceptical hypotheses as well as the idea that we have reflective access to our evidence. Pritchard notes that we can say three sorts of things to support epistemological disjunctivism: a) it has strong anti-sceptical potential; b) it fits with an ordinary way of talking about evidence; and c) the best argument against it, the “highest-common factor” argument (which I will discuss shortly), is a bad argument.⁶⁰ Note that right away, it’s plausible that Pritchard’s main argument must derive from c). This is because, first, there are many views which, if true, would have strong anti-sceptical potential, and second, Pritchard himself undermines the dialectical support deriving from b) by claiming that support from everyday language in this area is highly defeasible (Pritchard 2006, p.14).⁶¹ In any case, c) will be my focus in what follows.

⁶⁰ In fact, Pritchard considers and rejects a number of arguments against epistemological disjunctivism. Pritchard does a convincing job of explaining why his form of epistemological disjunctivism does not face a variety of challenges we may initially think would rule out the availability of the view. The most important of these in the sceptical context is the highest common factor argument. But I’ll note one other, namely the so-called “McKinsey problem” for epistemological disjunctivism (see Neta and Pritchard 2007; Silins 2005). The problem is analogous to the famous objection to content externalism in light of intuitive claims about first-person authority concerning our mental states. That is to say, because the McDowellian neo-Moorean claims that we have reflective access to *factive* mental states, and because it’s plausible that one can know by reflection alone that *if* one can *see* that p, then p, it seems to follow that one can come to know some specific empirical fact about the world by reflection alone. And this would intuitively be a reductio of the view.

Pritchard responds by claiming that the problem only arises if his view licenses the *acquisition* of non-empirical knowledge of empirical facts in this way. But, he claims, it does not. This can be seen once we note the following. First, according to Pritchard, having *factive* evidence that p, and basing one’s belief on that evidence, is sufficient for empirical knowledge that p. Second, having reflective access to the fact that one has *factive* evidence that p *entails* that one has *factive* evidence that p (this follows so long as we understand having “access” *factively*). And finally, the thought continues, although one may simply fail to form the belief that p, it’s clear that one cannot run through the chain of reasoning proposed in the McKinsey problem without *thereby* gaining empirical knowledge that p. This is because in doing so they meet a sufficient condition on empirical knowledge that p—according to the view. Pritchard says: “given the further trivial claim that if one has empirical knowledge of p then one cannot also have reflective (i.e., non-empirical) knowledge of p, it follows that there is no prospect of acquiring non-empirical knowledge in this case, and thus McDowell’s McKinsey-style difficulty disappears” (Pritchard 2006, p.20). If the way in which one acquires reflective knowledge that one can see that p necessarily involves gaining empirical knowledge that p, the McKinsey problem does not arise: it’s simply not true that on the McDowellian neo-Moorean picture we can come to know by reflection *alone* that some specific empirical fact about the world obtains, simply because we have reflective access to *factive* mental states.

⁶¹ A common way to provide support from ordinary language for the idea that *factive* mental states are a kind of evidence is with the following sort of example. Imagine that someone calls you at work and asks if Bob is in today. As it happens, Bob is sitting in the next cubicle. You tell your friend that Bob is in, but (for some reason) your friend asks, “Well, but how do you know he’s in?” A natural reply is: “Because I can *see* him, he’s right there!” and not “Because it appears to me as though Bob is right there!” However, as Pritchard notes, it’s possible to interpret these sorts of claims as implicitly bracketing error possibilities which would undermine the *factivity* of

3.4 The Highest Common Factor Argument and the Explanatory Challenge

The highest common factor argument can be put as follows:

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

(P2) The 'good' and 'bad' cases are phenomenologically indistinguishable. (Premise).

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

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

(Pritchard 2006, p.15)

The point of this conclusion is that, even in the good case, the best we have as evidence in support of our perceptual beliefs are facts about the way the world *seems* to us. Thus, the epistemological disjunctivist's claim that there is such thing as evidence in the form of reflectively accessible *factive* mental states is challenged. I discuss the point in more detail below, but I will note briefly here that abandoning an accessibility requirement makes responding to this argument much easier. Namely, it enables us to more easily resist the move from P2 to C1 by arguing that just because the good and bad cases are phenomenologically indistinguishable, it doesn't follow that the evidence one has in the good case is just as good as the evidence one has in the bad case. After all, it will be open to us to insist that one has evidence in the good case that one does not have in the bad case, it's just that one is not in a position to know that one has it.

According to McDowellian neo-Mooreanism the move from P2 to C1 is no good. There is a reflectively accessible difference in one's evidence in the good and bad cases, even though this difference is not *phenomenologically* detectible. In other words, the McDowellian neo-Moorean

the evidence cited. In other words, we can read claims such as "Because I can see him" as elliptical for "Given that certain error possibilities do not obtain (which are shared presuppositions of the conversational context), then my reason for believing that Bob is over there is that I can see him." Although this sort of reading is perhaps equally consistent with the data, it provides no support for epistemological disjunctivism. This is obviously a very complex issue and further discussion would take us too far afield. See Pritchard (2003) for further discussion.

holds that a mere lack of a phenomenological difference between the good and bad cases does not settle the question of whether the evidence one has (and to which one has reflective access) is the same in each case. Thus, Pritchard urges, the proponent of the highest common factor argument owes us a justification for the move from P2 to C1. More precisely, he says:

The proponent of the highest common factor argument thus owes us an explanation of what is wrong with such a view before they can help themselves to the key move from (P2) to (C1) (Pritchard 2006, p.17).

A proponent of the highest common factor argument may indeed have something to say here (which I'll discuss momentarily). But note, too, that it seems available to them to respond to Pritchard in precisely the same way—that is, to demand, on the contrary, that the move from P2 to C1 is quite intuitive indeed, and that the McDowellian neo-Moorean therefore owes *her* an explanation of what's *right* about the disjunctivist view such that she *cannot* help herself to the key move. In other words, it is not difficult to imagine a standoff between the proponent of the highest common factor argument and the McDowellian neo-Moorean when it comes to the move from P2 to C1. Whether this is a genuine stand-off concerning who counts as holding the “default position”—and thus who owes a justification to whom—is an important question.

I think the proponent of the highest common factor argument *does* have something to say here. They can respond by claiming that the best way to understand the support for the move from (P2) to (C1) is in terms of an *explanatory* challenge. The proponent of the highest common factor argument will claim that the best explanation of the phenomenological indistinguishability of our evidence in the good and bad cases is that our evidence *is* the same. And they'll challenge an opponent to come up with a better explanation. I think we can highlight this thought by considering a slightly different formulation of the highest common factor argument, due to Macarthur:

(P1*) Any experience of an object may be phenomenally indistinguishable from dreaming or hallucinating that object.

(P2*) In dreaming or hallucinating an object one is not aware of a real or existent object.

(P3*) In order to explain such phenomenal indistinguishability we must suppose that one is perceptually aware of something in common to perceiving an object and dreaming or hallucinating the same thing. Call this a sensory ‘idea’ of an object.

(C*) Therefore, in experiencing an object one is ‘directly’ aware of a sensory ‘idea’ of that object, which forms a defeasible basis for judging that there is an object there.

(Macarthur 2003, p.177)

Again, the idea is that the best explanation of the phenomenological indistinguishability of one’s evidence in the good and bad cases is that one’s evidence is indeed the same. But this of course prompts a question. As Macarthur puts it: “No doubt the existence of phenomenally indistinguishable cases of deceptive and non-deceptive experience may tempt one to posit a highest common factor in order to explain this fact. But *must* we explain the fact in this way?” (Macarthur 2003., p.178). This is where a central feature of Pritchard’s interpretation of McDowell is of crucial importance. We’ve seen that Pritchard casts McDowell’s view as one according to which our evidence must meet a condition of reflective accessibility. But this seems to block an alternative explanation that might be offered to meet the explanatory challenge. In particular, the alternative explanation of how our evidence could be phenomenologically indistinguishable yet different in the good and bad cases could be that are not always in a position to know what our evidence is. Indeed, it seems that McDowell himself appeals to a claim like this in offering such an explanation.⁶² Here is McDowell:

Of course we are fallible in our judgments as to the shape of the space of reasons as we find it, or—what comes to the same thing—as to the shape of the world as we find it. That is to say that we are vulnerable to the world’s playing us false; and when the world does not play us false we are indebted to it. But that is something we must simply learn to live with, rather than recoiling into the fantasy of a sphere in which our control is total (McDowell 1995, p.887).

⁶²It’s crucial to note that Pritchard is not interested in doing “McDowell Studies”. He’s merely interested in taking inspiration from some of McDowell’s works and seeing if he can develop a plausible view that hasn’t been adequately explored in the literature. My interest in what McDowell’s own view may or may not be is not in whether Pritchard accurately accounts for it; rather, I’m interested in an alternative interpretation of McDowell’s view only insofar as I’m interested in the potential range of responses to the explanatory challenge under discussion. To be sure, interpreting McDowell as some kind of externalist in the way I am here is controversial. For interpretations of McDowell qua externalist see Greco (2004) and Macarthur (2003).

It seems that McDowell's own way with the highest common factor argument is to *deny* an internalist condition of reflective accessibility on our evidence. Moreover, one of his main diagnostic moves in the process is to claim that the internalist condition of accessibility represents a kind of fantasy view, albeit one with an "intelligible motivation". Here is McDowell:

I have described that conception in a way that equips it with an intelligible motivation. The aim is to picture reason as having a proper province in which it can be immune to the effects of luck; not in the sense of sheer chance, but in the sense of factors that reason cannot control, or control for. The idea is that reason can ensure that we have only acceptable standings in the space of reasons, without being indebted to the world for favors received; if we exercise reason properly, we cannot arrive at defective standings in the space of reasons, in a way that could only be explained in terms of the world's unkindness (1995,p.885).

As I've said, I will return to the issue of what motivates an accessibility condition next chapter.⁶³ For now what's important to note is that McDowell does *not* endorse the "fantasy view" – he merely says it's an intelligible, though ultimately mistaken way of thinking. Once we abandon that way of thinking, an explanation of the phenomenological indistinguishability of our evidence in the good and bad cases which does not appeal to the idea that our evidence *is* the same in the good and bad cases opens up. We are simply not in always in a position to know what our evidence is. But this does not preclude the idea that, in the good case, the evidence we have, and therefore what we *ought* to believe, can be determined by our factive mental states. The fact that one *sees* that p can be a *reason* for believing that p. As McDowell repeatedly says: believing as one ought to believe is always in part due to a "kindness of the world" (McDowell 1995). Interestingly, a key move underpinning McDowell's ability to endorse this idea is something that looks for all the world like a kind of externalism. Importantly, it's not the basic externalism of a neo-Moorean view like Sosa (1999). Rather, it construes knowledge as (always and everywhere) a "standing in the space of reasons."⁶⁴ However, the reasons one has, and whether one accords with them, can ineliminably depend

⁶³I consider something very close to, if not the same as, McDowell's diagnosis of the "fantasy view" in the next chapter.

⁶⁴I think McDowell's notion that knowledge is a standing in the space of reasons means that knowledge is a normative status. It is normatively constrained and normatively constraining (Wright 2004, p.183). What I mean by normatively constrained is that a necessary condition on knowing that p is that one has formed one's belief in response to reasons (where this is understood in normative terms). Moreover, what I mean by the idea that knowledge is normatively constraining is that when one knows that p, certain things are normatively implied, for example, perhaps it is OK for S to assert that p or to act on p (I don't know what McDowell's views on assertion and action are).

on the way things really are out there in the world (whether one is in a position to know how things are or not). As McDowell continually emphasises, this just means that we must embrace the idea that enjoying a proper standing in the space of reasons is always in part due to the kindness of the world.

Macarthur articulates this approach to the explanatory challenge in a helpful way: “Only in the context of the view that the mind’s contents are wholly and infallibly accessible to itself does the failure to detect a difference between two appearances imply that they are the very same—an admission that forces us to accept a ‘veil of ideas’ theory of perception and its sceptical consequences” (Macarthur 2003, p.178).⁶⁵ We can meet the explanatory challenge by embracing the idea that we are not always in a position to know what our evidence is. Meanwhile, Pritchard’s commitment to a reflective access condition seems to render this second way of meeting the explanatory challenge unavailable. In other words, it seems that Pritchard faces the following dilemma in explaining how our evidence could be phenomenologically indistinguishable yet different in the good and bad cases: on the one hand, he can appeal to sameness of evidence. But since denying sameness of evidence is the whole point of disjunctivism, this move is obviously unavailable; on the other hand, he can deny that we are always in a position to know what our evidence is. But because of his distinctive commitment to the “reflective accessibility” of our evidence, it’s not clear that this move is available either.⁶⁶

However, at this point Pritchard might respond in the following way. He can distinguish between the “weak” and “strong” access conditions that I noted in the introduction (Smithies 2013). A weak access condition on our evidence says that one is always in a position to know that one has evidence e when evidence e obtains. Strong access says that one is always in a position to know that one has evidence e when it obtains *and* that one is always in a position to

⁶⁵Macarthur’s appeal to “infallibilism” is a red-herring here. What’s relevant is the idea that we are always in a position to know that p – mistakes about whether p are still possible even when one is in a position to know that p . As Williamson notes: “To deny that something is hidden is not to assert that we are infallible about it. Mistakes are always possible. There is no limit to the conclusions into which we can be lured by fallacious reasoning and wishful thinking, charismatic gurus and cheap paperbacks. The point is that, in our cognitive home, such mistakes are always rectifiable” (Williamson 2000, p.94). I return to Williamson’s metaphor of a “cognitive home” in Chapter 5.

⁶⁶This is because of the connection I’ve drawn above between endorsing a reflective access condition and the kind of epistemic position this seems to imply we ought to be in with respect to the facts about what justification we have. Once again, although Pritchard explicitly claims that he is offering a non-standard form of internalism, what he means by this is simply that he allows that we can have reflective access to our *factive* mental states (a view that he notes classical internalism has traditionally eschewed).

know that one does *not* have evidence *e* when it does not obtain. In Pritchard (2012), Pritchard indeed claims that epistemological disjunctivism is only committed to the claim that one has reflective access to one's evidence in the good case, not the bad case. In the bad case, one is not in a position to know that they do not have a factive reason to believe that *p*. And this is OK, because he does not endorse the claim that one is always in a position to know that one does not have evidence *e* when *e* does not obtain. The response to the explanatory challenge I have in mind, then, is that we can explain how it is that we can have different evidence in phenomenologically indistinguishable good and bad cases in terms of the fact that we *can* fail to be in a position to know what our evidence is when we're in the bad case.

This may be a way of responding to the explanatory challenge. But I think it raises an important question about Pritchard's commitment to internalism in general. We've seen that he motivates the commitment by appeal to avoiding revisionism where possible. But we've also seen that he explains the motivation for internalism in terms of the notion of epistemic responsibility (Pritchard 2012, p.2). Again, Pritchard seems to accept the idea that an agent counts as in some sense epistemically responsible just in case the facts in virtue of which one is justified meet an internalist access condition. But if one's evidence is not accessible to one in the bad case, it seems that Pritchard is committed to claiming that one cannot be held responsible in the bad case. As Smithies points out, however—and as I've indicated in the introduction—it's something of a mystery why we would posit this asymmetry (Smithies 2013). Why would we find it any more compelling to claim that an agent in the good case is epistemically responsible than an agent who believes in accord with how things appear to them (and behaves otherwise identically in their manner of inquiry) but who happens to be in a bad case? Usually, the thought taken to motivate internalism is precisely that, even in the bad case we are at least epistemically responsible (i.e. do not make culpable mistakes about our evidence). But moreover, a central motivation for internalism is that the agent in the bad case can be justified in believing that *p*.

Indeed, it's crucial to note that Pritchard explicitly denies what is known as the "New Evil Genius" thesis. Roughly, this is the intuition that *S*'s recently envatted counterpart is equally justified or reasonable in believing things about the world around them as *S* is under ordinary circumstances. Pritchard argues by way of rejecting the highest common factor conception of evidence that his internalist view about evidence does not entail the New Evil Genius thesis.

He considers the following argument that his internalism does entail the New Evil Genius thesis:

- i) *Accessibilism*: S's internalist epistemic support for believing that p is constituted solely by facts that S can know by reflection alone.
- ii) *The Highest Common Factor Thesis*: the only facts that S can know by reflection alone are facts that S's recently envatted physical duplicate can also know by reflection alone.
- iii) *The New Evil Genius Thesis*: S's internalist epistemic support for believing that p is constituted solely by properties that S has in common with her recently envatted physical duplicate. [from (1), (2)].

(Pritchard 2012, p.39)

Pritchard resists (iii) by rejecting (ii). He obviously resists that thesis given his view that in the good case one can have reflective access to *factive* reasons. This may be a legitimate way of separating his internalism from the New Evil Genius thesis. But the point I want to make is that doing so calls into question a central motivation for epistemological disjunctivism—namely its putative ability to reconcile externalist and *internalist* intuitions.⁶⁷ In any case, it is not clear how Pritchard can meet the explanatory challenge without raising important questions about epistemic responsibility. As I just said, I will explore such questions further in the context of a card-carrying externalist view about evidence. For now it's important to take stock. What we've seen is that Pritchard's commitment to the reflective accessibility of our evidence is important when it comes to respecting the internalist intuitions he aims to respect. However, precisely how we understand the kind of access this implies really matters. Pritchard is getting pushed in two different directions. That is, if we endorse the strong access condition (which seems important for respecting internalist intuitions), we face the explanatory challenge. It's tough to see a possible route for explaining the phenomenological indistinguishability of cases when one is always in a position to know what one's evidence is—both when one has it and when one does not—other than by maintaining that one's evidence is indeed the same in good

⁶⁷Pritchard aims to account for internalist intuitions in this respect by claiming that agents in the bad case believe unjustifiedly but *blamelessly*. For my purposes, I think it is safe to interpret this move as similar enough to Williamson's point about *excuses* (which I'll discuss briefly below, and next chapter) to reserve discussion of the basic idea until next chapter.

and bad cases. Endorsing a weak access condition perhaps allows for a response to the explanatory challenge. But we're getting further away from one of the central motivations for the view in the first place. Moreover, by rejecting the New Evil Genius intuition, it seems that Pritchard's view will generate the same consequences that I will go on to discuss about evidential externalism in Chapter Four. I take these issues about the explanatory challenge and access to motivate looking at a more thoroughgoing evidential externalism.

3.5 Williamson and Evidential Externalism

Perhaps the most well-known proponent of evidential externalism is Williamson (2000). Williamson's particular form of evidential externalism is motivated by his thesis that one's total evidence is equal to one's total *knowledge*, or "E = K". Williamson is an externalist about knowledge. Thus, his externalism about knowledge leads to externalism about evidence. According to Williamson, we are not always in a position to know what our evidence is.⁶⁸ We do not have this sort of access to our evidence because we are not always in a position to know what we know. Williamson's thesis about evidence partly motivates, and is partly motivated by, his broader "knowledge first" programme in epistemology, the details of which will largely be put to one side for our purposes. In this section I briefly examine the motivation for E=K independently of considerations about knowledge first epistemology.

Why should someone who is not already convinced by the knowledge first programme think that our total evidence is equal to our total knowledge? Williamson argues from the *function* of evidence. Evidence plays a role in inference to the best explanation, probabilistic reasoning, and the exclusion of hypotheses. He claims that in all three of these cognitive activities, only E=K can make sense of the role that evidence plays.⁶⁹

For instance:

When we prefer an hypothesis *b* to an hypothesis *b** because *b* explains our evidence better than *b** does, we are standardly assuming *e* to be known: if we do not know *e* why should *b*'s capacity to explain *e* confirm *b* for us? (Williamson 2002, p.200).

⁶⁸We can take Williamson to be denying both the weak *and* strong access conditions.

⁶⁹These arguments constitute part of a "theory of evidential status" as opposed to a "theory of evidential relevance" (these terms are due to Joyce (2004)). Williamson is careful to point out that there are distinct questions about what it is for evidence to *favour* belief in a proposition *p* (he cashes this out in terms of probability relations), and about what sort of epistemic standing putative evidence must have in order to *be* evidence. The theory being offered here is that it must be *knowledge*. He provides a theory of relevance in Williamson (2000), Ch. 10.

Moreover:

It is likewise hard to see why the probability of b on e should regulate our degree of belief in b unless we know e . Again, an incompatibility between b and e does not rule out b unless e is known (p.200).

And finally, consider an example:

Suppose that balls are drawn from a bag, with replacement. In order to avoid issues about the present truth-values of statements about the future, assume that someone else has already made the draws; I watch them on film. For a suitable number n , the following situation can arise. I have seen draws 1 to n ; each was red (produced a red ball). I have not yet seen draw $n+1$. I reason probabilistically, and form a justified belief that draw $n+1$ was red. Consider two false hypotheses:

h : Draws 1 to n were red; draw $n+1$ was black.

h^* : Draw 1 was black; draws 2 to $n+1$ were red.

It is natural to say that h is consistent with my evidence and that h^* is not. In particular, it is consistent with my evidence that draw $n+1$ was black; it is not consistent with my evidence that draw 1 was black. Thus my evidence does not include the proposition that draw $n+1$ was red. Why not? After all, by hypothesis I have a justified true belief that it was red. The obvious answer is that I do not *know* that draw $n+1$ was red (p.201).

The moral of the example is that only $E=K$ enables us to discriminate between hypotheses in the way we want to. We might wonder whether only knowledge can make sense of the function of evidence in these instances (Goldman 2009, Joyce 2004, Fumerton 2000, Owens 2004). Alvin Goldman (2009), in particular, has challenged Williamson's claims that nothing other than knowledge can make sense of the role of evidence. He offers an alternative conception of evidence which says, roughly, that a proposition p is evidence for S iff p is non-inferentially justified for S . Goldman claims that evidence so construed can do the job just as well as $E=K$. I will not get into the question of whether Williamson, or Goldman, is correct about this (Littlejohn 2011, 2011b; Silins 2005). Instead, I'd like to consider some important consequences of $E=K$.

The consequence immediately relevant for our purposes is that when it comes to the underdetermination argument for scepticism, the sameness of evidence thesis is readily challenged. If Williamson is right about evidence, there *is* a fundamental difference in the evidence one has in each case. Very roughly, in a BIV case one does not have knowledge of

the world around them. As a BIV, you do not know that you have hands. After all, you are a BIV being stimulated to believe that you have hands. Thus, the proposition “I have hands” is not part of your evidence. However, even when you are in the bad case you do know that it appears to you that you have hands. Thus, the proposition “It appears to me that I have hands” remains part of your evidence. Meanwhile, in the ordinary case, you know you have hands. The proposition “I have hands” *is* part of your evidence. Thus, in the bad case one would *not* have the same evidence that one has in the ordinary case in which one knows lots of things about the world around them. Although it needs some careful spelling out, the view promises anti-sceptical results.

It’s worth emphasising here that once we abandon the idea that our evidence is always accessible, denying the sameness of evidence thesis seems relatively straightforward. This is arguably *why* it is so straightforward. In particular, this is exactly the sort of move that makes things so much easier when dealing with the proponent of the highest common factor argument. Indeed, we’ve just seen that it seems as though Pritchard’s adherence to reflective accessibility was what causes him problems in that respect. Of course, Williamson’s approach to scepticism is not without difficulties.

3.6 Same Evidence or Same Justification?

There are a variety of objections to Williamson’s anti-sceptical strategy in the literature. Getting into these here would take us too far afield. However, an important recurring theme is that Williamson’s strategy for dealing with the sameness of evidence thesis is compelling, but that it does not do the relevant work in challenging the evidential sceptic (Fumerton 2000; Owens 2004; Schiffer 2009).⁷⁰ The basic idea is that there is a difference between a subject having the *same* evidence in the good and bad cases and the subject being *equally rational* or *justified* in believing what they do on their evidence in the good and bad cases.⁷¹ Many epistemologists are inclined to think that a subject can be equally rational in both the good and bad cases (this is just the New Evil Genius intuition I mentioned above). Owens calls this the “equal rationality thesis”. He goes on to claim:

⁷⁰ Presumably we might take the following objections to apply to disjunctivist strategies generally. I return to a related concern in Chapter 5.

⁷¹ There are of course complex issues surrounding all of the different things people mean by “rational” and “justified”, especially in the present context. What is of central interest to me is the idea of an agent \emptyset -ing as they are normatively required to \emptyset . Though of course there are equally complex issues surrounding what that means. However, I’ve stipulated what I mean by this in the introduction, and I discuss some further details next chapter.

[I]f the equal rationality thesis is true, the evidence the subject in the bad case has that he is in the good case must be just as strong as the evidence that the subject in the good case has that he is in the good case. The subject in the bad case is no less reasonable in supposing that he is in the good case than he would be were he actually in the good case, so the evidence supporting his belief in the bad case must be just as strong as it would be were he actually in the good case (Owens 2004, p.3).

Owens thus draws a distinction between *strength* of evidence and *sameness* of evidence. While he thinks that the subject in the good and bad case has equally strong evidence he does not think this has a bearing on the question of whether the subject has the same evidence. And so, he claims, whether or not we can successfully reject the sameness of evidence thesis does not necessarily have a bearing on the equal rationality thesis. As Owens puts it: “to establish that two subjects have equally strong evidence for p is not to establish that they have exactly the same evidence for p” (Owens 2004, p.3). And he goes on to ask: “But why need the sceptic take any further interest in sameness of evidence?” (2004, p.3). The question is rhetorical. Owens (and others) claims that all the sceptic needs is the equal rationality thesis to undermine our justification to believe ordinary propositions such as that we have hands.⁷² The argument can be put as follows:

- i) You think you have justification to believe that you are in the good case, but if you were in the bad case you would have the same justification to believe you were in the good case.
- ii) Therefore, you don’t really have any more justification to believe that you are in the good case than that you are in the bad case.
- iii) Therefore, you do not have justification to believe that you are in the good case (i.e. that you have hands).

(Owens 2004, p.4)

⁷² Schiffer (2009) argues that Williamson’s response to scepticism does not generalize to a formulation of the sceptical argument in terms of justification. He argues that the sceptic can plausibly argue that one has the same justification in the good and bad cases and from this derive a sceptical conclusion. He claims that Williamson’s theory that E=K fails to support a claim of asymmetry of justification in the good and bad cases (the details are complex and would take us too far afield here). Schiffer’s argument prompts Williamson to claim that we don’t have the “full normative status” required for justification in the bad case. We need to look more carefully at the implications of Williamson’s view about the relationship between epistemic normativity and justification and rationality, which I do immediately below.

It seems that the obvious question to ask about this argument is *why* we would accept premise i).⁷³ The answer that comes immediately to mind is the assumption that your evidence is the same in the good and bad cases, and in both cases you respect your evidence equally well. Owens might resist and say that what's really doing the work for i) is just the idea that you respect what you take to be your evidence equally well in both cases, regardless of what your evidence actually *is*. I don't think it's worth getting hung up on the difference between these ways of spelling out the motivation for i).⁷⁴ Indeed, I'm going to set to one side the question of whether this is a good argument. This is because I think it's clear that Williamson takes his claims about sameness of evidence (or lack thereof) in the good and bad cases to directly imply a claim about sameness of justification and rationality. So he would reject the first premise of Owens' argument. I think this highlights a particularly important consequence of his evidential externalism which will be the subject of the next chapter. But let me explain how we get there.

3.7 Justification and Rational Requirements

Williamson talks about being “justified” in normative terms. He says that a belief that p is justified just in case it is “OK” to believe that p.⁷⁵ He goes on to explain what determines whether it is OK to have a belief that p:

On the account in *Knowledge and its Limits*, justification is to be explained in terms of knowledge rather than *vice versa*. Belief that falls short of knowledge is to that extent defective. If we conceive belief as the mental analogue of assertion, then just as the rule for assertion is ‘Assert p only if you know p’, so the norm for belief is ‘Believe p only if you know p’. In this sense, a belief is fully normatively justified if and only if it constitutes knowledge (Williamson 2009, p.359).

Thus, according to Williamson, because the agent in a bad case who believes that p does not know that p, the agent violates the rule “believe that p only if you know p”. The agent thereby fails to have the “full normative status” for justification.⁷⁶ To be sure, Williamson is sensitive

⁷³It's not immediately obvious to me what the rationale behind the move from i) to ii) is either. Perhaps it's something like an analogue of the underdetermination principle.

⁷⁴For example, it seems open to the disjunctivist to insist that what it is to respect your evidence well just is to proportion your beliefs to what your evidence *really* is. Indeed, I think this is a natural thing to say, and it's how I understand Williamson to proceed (next section).

⁷⁵He distinguishes this sense of being justified from another “inferential” sense, according to which a belief is justified just in case one can support the belief with further argument. He claims this is neither necessary nor sufficient for a belief to count as being OK in the normative sense.

⁷⁶Precisely what Williamson means by “normative” is obviously an important question. For now I'll take him to use the term in the same way I've spelled it out in the introduction. As I'll discuss in the next chapter, there is

to the idea that a belief formed in the bad case can have *some* good normative standing. For example, believing that you have hands as a BIV when it appears to you just as though you have hands is better than believing under the same circumstances that you have hooks as opposed to hands.

It is important to note that Williamson does not limit himself to talk of justification. He also talks about “requirements of rationality”. According to Williamson, the requirements of rationality tell you to proportion your beliefs to the evidence (Williamson 2000, Ch.8). “If rationality requires one to respect one's evidence, then it is irrational not to respect one's evidence” (Williamson 2000, Ch.8). This is a normative requirement. Although the details are complex, the basic idea is that your degree of confidence that *p* ought to be proportioned to the probability that *p* given your evidence. If the probability that *p* given your evidence is 1, then you ought, for example, simply outright believe that *p* (believe to a degree of confidence of 1). If the probability that *p* given your evidence is less than 1, then you ought to have something less than outright belief that *p*. Perhaps you ought to believe that *p* is probable to *n* degrees, or perhaps you ought to have some distinct propositional attitude that is in some sense less committal about the fact that *p* – the details are unimportant here.

Given this normative principle about what you ought to believe under certain conditions, what should we say about what *S* ought to believe in good and bad cases, respectively? Let's start with the bad case. The evidence (what *S* knows) is that it *appears* to them that they have hands. After all, they are in the bad case. Since they do not have hands, and since knowledge is factive, they cannot know that they have hands—thus, *S*'s evidence cannot be that they have hands. However, *S*'s evidence (that it appears to *S* just as though *S* has hands), does make it highly probable that they have hands. So, proportioning their belief to the evidence, *S* ought to believe with a high degree of confidence that they have hands. In the good case, *S*'s evidence includes the proposition that they have hands. *S* *knows* that they have hands. According to Williamson, if a subject knows that *p*, then the probability of *p* on her evidence is 1.⁷⁷

Proportioning one's belief to the evidence, one ought to believe outright that they have hands.

As Williamson puts it:

another way of understanding what he means, but then the upshot of his view concerning rational requirements will not be as controversial as he himself and many others seem to acknowledge that it is.

⁷⁷ As Williamson notes, if we are to avoid scepticism about knowledge, this implies that a proposition can be evidence for itself. Brown (2013) argues that this is an unpalatable upshot of the view based on considerations about the infelicity of citing *p* as evidence for itself. For a response, see Littlejohn (*unpublished*).

[P]robabilities on one's evidence are what one's degrees of confidence must match if one is to satisfy the norm of proportioning one's belief to the evidence. In Bad, if one knows *e* and falsely believes [that they have hands], one deviates from that norm with respect to [the proposition that they have hands] by only a little, even if that result is partly the outcome of one's blunder in misconstruing the extent of one's evidence (Williamson 2009, p.362).

Of course we are familiar with the idea that one cannot *tell* whether one is in the good or bad case.⁷⁸ This is built into the description of the bad case. The point is that although different degrees of confidence are normatively required of you, depending on which case you are in, you are not in a position to know which case you are in. You are thus not in a position to know what you ought to do. And yet, because you ought to do one thing or the other, failure to do so entails that you are at some sort of normative fault.⁷⁹ Williamson's spells this out in terms of a lack of justification, and in terms of failing to meet the requirements of rationality. In the bad case, when you believe outright that you have hands, you are not justified in believing that you have hands. As he puts it, your belief fails to have the "full normative status" of justification (Williamson 2009). Indeed you are not even fully rational.

To be sure, Williamson allows that while it is true that you violate a normative requirement (and thus are not justified or even fully rational in believing that *p*) in such a case, you nevertheless have a "cast iron excuse" for believing that *p* (Williamson 2009, p.360). This is presumably the sense in which your belief is still better than had you believed, as in the example above, that you have hooks rather than hands. I will return to this point about excuses in the next chapter. For now I simply want to claim that, given the foregoing considerations, it's becoming clear that Williamson's view about evidence—specifically, the externalist feature of it—has some important consequences. In particular, it leads to claims about the conditions under which we are normatively required to \emptyset that are highly counterintuitive in at least some kind of way (as many epistemologists in the literature, including some who I've cited above would agree). It turns out, however, that it's not easy to spell out what this claim amounts to.

⁷⁸ This does not mean merely that some subject can't tell but that there is no difference *for* the subject. Nothing is available to the subject with which they could tell.

⁷⁹ For an interesting argument that a consequence of Williamson's view is that we ought to have lower credences in the *good* case than in the bad case, see Silins (2005).

3.8 Conclusion

We can challenge the sameness of evidence thesis in at least two broad ways: by retaining an internalist access condition on our evidence or denying that an internalist access condition is necessary. Pritchard motivates his McDowellian neo-Mooreanism by saying that he can preserve internalist and externalist intuitions in a particularly satisfying way. He thinks this is important because it allows us to preserve a commonsensical way of thinking about our epistemic commitments as well as intuitions about epistemic responsibility. We saw that he runs into trouble when it comes to explaining how one's evidence could differ between the phenomenologically indistinguishable cases. Of course, at that point he can appeal to the "weak access" condition. But then, I argued, his view starts to lose its internalist-friendly appeal. I then turned to the card-carrying evidential externalism of Williamson. Denying an internalist access condition makes challenging the sameness of evidence thesis a lot easier. But I've started to draw out an unsettling implication that comes with this approach. I turn now to consider in more detail the costs of going evidential externalist.

Chapter Four: The Natural Reaction

4.1 Introduction

Gibbons claims that there is a “natural reaction” to the general idea that one can be normatively required to \emptyset when that requirement is in some sense outside of one’s first person perspective or inaccessible to one, and thus (according to Gibbons) incapable of playing a guiding role (Gibbons 2013).⁸⁰ Because I think the case with Williamson’s view about evidence is an instance of this more general issue, I will call a basic concern we might have with Williamson’s view the “natural reaction” to Williamson’s view about evidence.⁸¹ Gibbons uses this term because he wagers that many people have it when they reflect on the idea of what’s involved in being normatively required to \emptyset . But more needs to be said about how to spell out the natural reaction. We’ll see that it’s not easy to articulate exactly what it amounts to. Before engaging in this project, I’d like to note that the following is not intended as a comprehensive and watertight account of what this concern with Williamson’s view amounts to. Rather, the point of the label “natural reaction” is to indicate that I take it that (at least part of) the burden is on Williamson (or evidential externalism generally) to explain why this intuitive worry is misguided or mistaken. I’ll consider an important part of Williamson’s diagnosis in Chapter 5. But I think Williamson goes against the grain of our intuitions about the relationship between agents’ capacities to accord with normative requirements and the conditions under which those normative requirements obtain, whatever the comprehensive theory of this relationship is. In this chapter I offer a brief account of some features of this relationship. The goal is to provide enough of a story about the natural reaction to make accounting for it look like an important desideratum for any anti-sceptical strategy that seeks to reject the sameness of evidence thesis.

4.2 Deliberation-Guidance

One way of spelling out the natural reaction is as follows. We’ve seen that Williamson argues for his view about evidence by considerations about the function of evidence in conducting our cognitive lives in various ways. Another crucial function of evidence, missing from the list last chapter, is that we use evidence in our deliberations about what to believe, and thus indirectly in our deliberations about what to do. For example, when thinking about whether

⁸⁰ Gibbons’ main concern is with what he calls “the ethicists’ notion of *objective* reasons” (Gibbons 2013, Ch.1).

⁸¹ Insofar as I’m using Williamson’s view as a paradigm case of evidential externalism, I will also refer to this as the natural reaction to evidential externalism generally.

you stole my bike or not, I consider what evidence I have in support of thinking this. This is a good example because a lot hangs on something like accusing someone of stealing my bike. And it's often the case that a lot hangs on what one believes given their evidence. We might wonder, then, whether the sort of opacity of our evidence on Williamson's view has counterintuitive consequences in light of this further important function of evidence. Williamson recognizes this: "How can one follow the rule 'proportion your belief in p to your evidence for p ' when one doesn't know exactly what one's evidence is?" (Williamson 2000, p.192). I will call this the "deliberation-guidance" worry.

The deliberation-guidance worry can be sharpened in the following way. Call a norm "operationalizable" just in case, "if one knows the norm, one is always able to use the norm as a premise in a competent deduction, together with one's knowledge of whether the norm's triggering condition obtains, to come to know whether one is acting in conformity with the norm" (Srinivasan *unpublished*, p.11).⁸² For example, the norm "One ought to proportion one's beliefs to the evidence" is operationalizable just in case one can use it in a competent deduction, together with their knowledge that they now have evidence E and degree of credence C on E , to come to know whether they believe in conformity with the norm. The idea is that such knowledge is necessary for being "guided" by a norm. One must be able to know whether one is acting in accord with the norm or not in order to properly be counted as being guided by the norm. Without such knowledge, the deliberation-guidance worry goes, one's accordance with a norm is at best a sort of accident—in which case it's questionable that one deserves a normative status for believing in accordance with the norm. So the problem with Williamson's view, according to this worry, is that it renders the evidential norm "unoperationalizable."

How serious is the deliberation-guidance worry? On one hand, Williamson argues that the opacity of our evidence, on his view, is not extensive enough to give this concern any bite. He says: "Although we have no infallible procedure for determining whether we know p , in practice we are often in a position to know whether we know p " (Williamson 2000, p.193). The point here is that the fact that we are "often" in a position to know whether we know p is

⁸² As I said in the introduction, we can think of Srinivasan's helpful phrase "triggering condition" as meaning something like the condition under which a given normative requirement applies. And again, here's Srinivasan's explanation of her own terminology: "A norm is a universal generalisation about how one ought to conduct one's practical or doxastic affairs, involving a normative state N and a triggering condition C , of the form " N if/and/only if C " (Srinivasan *unpublished*, p.6).

enough to dismiss the deliberation-guidance worry. While it might be true that we could not use our evidence to guide deliberation (and indirectly action) if we were *never* in a position to know what our evidence is, this doesn't entail that we cannot use our evidence to guide deliberation given that we are often (though not *always*) in a position to know what our evidence is. Compare the following example from Srinivasan:

SEDER NORM: When setting the table for Passover, one ought to set as many places as there will be Seder guests plus one.

Srinivasan notes that, while it's true that one is not always in a position to know how many guests one will have at one's Passover Seder (there may be unexpected arrivals, or someone might get sick), we can clearly employ such a norm in our reasoning about what to do at Passover. Norms can guide us even if we're not always in a position to know what they tell us to do.⁸³

This seems fair enough.⁸⁴ But we might press the basic spirit of the deliberation-guidance worry in the present context. We might insist on something that Williamson himself notes: "None of this would be much consolation if our beliefs about our knowledge were hopelessly unreliable. Sceptics say that those beliefs have no rational basis..." (Williamson 2000, p.193). The thought is something like this. If scepticism is true, then it is simply *false* that we are often (even if not always) in a position to know what our evidence is (given E=K). Of course, Williamson is not worried about what the sceptic says. He reminds us: "We have found [the sceptic's] reasons for saying so inadequate" (p.193). And what he means (very roughly) is that we can reject the sceptic's appeal to the sameness of evidence thesis. Perhaps at this stage in the dialectic of my thesis (i.e. considering whether a knowledge-first account of evidence can support a rejection of the sameness of evidence thesis and thus provide a response to

⁸³ Williamson offers an example in this context, too. He compares the evidential norm to the rule, "Proportion your voice to the size of the room". He claims we can perfectly well be guided by such a rule even though we are fallible about the size of the room we are in. "To have applied the rule 'Proportion your voice to the size of the room', one needs beliefs about the size of the room, but they need not have been true—although if they were false, one's application was faulty. Similarly, to have applied the rule 'Proportion your belief in *p* to the support that *p* receives from your knowledge', one must have had beliefs about how much support *p* received from one's knowledge, and therefore about one's knowledge, but those beliefs need not have been true—although if they were false, one's application was faulty (2000, p.192)." For a critique of Williamson's analogy, see Owens (2004).

⁸⁴Although one might want to press the question, "*how* often is often enough?"

scepticism), endorsing a move like this is inappropriate.⁸⁵ But rather than pursuing that question, I think a deeper issue is as follows.⁸⁶

4.3 Normative Requirements and Blameworthiness

It's plausible that what's behind the deliberation-guidance worry is a deeper worry about the idea that one could *ever* be normatively required to \emptyset when one is not in a position to know that one is required to \emptyset .⁸⁷ Srinivasan suggests that what lies behind such a worry is a desire to align the "deontological" and "hypological" perspectives (Srinivasan, p.14). The deontological perspective is just the perspective from which we make judgments about whether an agent conforms to normative requirements – i.e. their obligations, duties, or what they ought to do or think, etc. The hypological perspective is the perspective from which we make judgments about whether an agent is blameworthy or praiseworthy for their actions or attitudes. One reason we might think that these perspectives should be in alignment is if we think that judgments of blameworthiness and praiseworthiness are in some sense central to the very idea of normative requirements. Indeed, the thought goes, we might think that they are so central that it's difficult to understand the possibility of a mismatch between the conditions under which someone is blameworthy and the conditions under which they count as violating a normative requirement. As I will explain further below, whether one is blameworthy intuitively depends on whether they are in a position to know that the conditions under which they would be blameworthy obtain (Williamson, for example, seems to agree, as we can see from his use of *excuses* in his theory of evidence and justification – more on this below). Thus, allowing that one might *ever* violate a norm when one is not in a position to know that its triggering conditions obtain effectively introduces a mismatch between conditions of blameworthiness and conditions of norm-violation. And so, on the present assumption, it is to

⁸⁵ What I mean is the following. If the guidance-deliberation worry is to be avoided, scepticism must be false (so we're assuming here, at any rate). I wouldn't want to assume that scepticism is false in considering whether Williamson's view can support a rejection of the sameness of evidence thesis and successfully respond to scepticism. But note that I don't intend the point here to apply to Williamson's own appeal to this anti-sceptical remark.

⁸⁶ Srinivasan points out a second objection to the deliberation-guidance worry. This is that the deliberation-guidance worry seems to presuppose that our practical and cognitive lives, when they go well, are a rule-governed affair. That is, it seems to presuppose that well-lived practical and cognitive lives involve actively applying general rules to specific cases. She notes that there are many reasons to think that is an implausible (Srinivasan *unpublished*, p.13).

⁸⁷ As opposed to a worry about whether we are *often enough* in the relevant position to know.

misperceive something central about normativity. We can sharpen the idea that blameworthiness is “central” to normativity in the following way. Call it the “hypological principle” (H):

(H): S ought to \emptyset only if, if S doesn't \emptyset , then S is blameworthy for not \emptyset -ing.

The principle says that it's only true that S is normatively required to \emptyset if, were they not to \emptyset , they would be blameworthy for not \emptyset -ing. The basic idea behind the present way of spelling out the natural reaction is that, once we grant (H), a condition of access on our normative requirements follows. This condition is incompatible with Williamson's view about evidence. One reason we might think that judgments of blameworthiness give rise to these considerations about access has to do with the intuitive incompatibility of judgments of blameworthiness with what Williams (1976, 1981) and Nagel (1979) call “moral luck” or “resultant luck” (respectively). Imagine the following case:

Suppose Kate, a generally thoughtful person, brings some soup to her friend Mary, who is down with the flu. And suppose further that Kate isn't in a position to know that the soup contains an allergen that, rather than making Mary feel better, will make her more ill. In trying her best to help her friend, Kate inadvertently makes her friend worse off (Srinivasan, p.15).

This is an example of “bad resultant luck.”⁸⁸ Intuitively Kate is *not* blameworthy for what she's done to Mary. And an issue about access seems like precisely the relevant feature of the case. That is, the best explanation for why she is not to blame is that she was not in a position to know that the soup contains the allergen (imagine, too, that she was very cautious in checking all the allergens and that she asked Mary about her allergies). Thus the idea of resultant luck brings out the intuitive plausibility of accessibility conditions on blameworthiness.⁸⁹ As

⁸⁸ As Srinivasan notes, when Williams and Nagel speak of resultant luck, they are speaking about both the presence and the absence of the phenomenon of luck exhibited in the Kate and Mary case. We can contrast the Kate case above with a case in which the soup does *not* contain allergens that will make Mary ill. According to Williams and Nagel such a case is an example of the absence of bad resultant luck. Kate would be said to benefit from the absence of bad resultant luck. In this chapter, I will only have in mind the phenomenon of the presence of bad resultant luck when I speak of resultant luck.

⁸⁹ It's important to differentiate resultant luck from other kinds of luck that are intuitively *compatible* with blame. For example, behavioural dispositions that are not the result of an agent's free will, but rather the result of natural events and laws are regarded by many to be compatible with judgments of blame. Moreover, the kind of luck involved growing up in a community of slave owners may plausibly be compatible with blame. Also, the kind of luck involved in finding oneself in a situation such that one performs an action that one would otherwise have had no inclination to perform—for example the actions of a concentration camp prison guard—seems to many to be compatible with blame (Srinivasan *unpublished*, p.15). To the extent that we have intuitions that agents subject to these kinds of luck are *not* the proper subjects of blame, it is again plausibly because of the agents' lack of epistemic access to the relevant conditions.

Srinivasan puts it, “this form of luck is present when one’s ignorance of the relevant circumstantial facts prohibits one from controlling the outcome or result of one’s actions” (Srinivasan unpublished, p.15).

Notice that Kate violates the normative requirements of at least some possible normative theories. For example, to put it roughly, if we assume that Kate only had the choice of either giving Mary the soup or not, and that Kate’s giving Mary the soup reduces the overall good, then Kate violates a consequentialist norm of action that says: “One ought to \emptyset iff \emptyset -ing would maximise the good”. Taken together, the consequentialist norm and the intuitive judgement about Kate’s blameworthiness suggest that blameworthiness and norm violation indeed come apart. However, if we think that blameworthiness is central to normativity in the way specified above, we might be inclined to say that the consequentialist norm—in virtue of having opaque triggering conditions—misses, or fails to account for or respect something central about normativity. In other words, it conflicts with the (H) principle. Consider the following claim:

Kate ought to maximize the good only if, if Kate doesn’t maximize the good, then Kate is blameworthy for not maximizing the good.

Because Kate is intuitively not to blame for not maximizing the good here, (H) tells us that it’s false that she really *ought* to have done so—i.e. it’s false that she violated a normative requirement in bringing Mary the soup. Likewise then, when it comes to Williamson’s theory of evidence, the evidential norm’s triggering conditions are opaque (in the relevant sense). It’s possible to violate the norm due to resultant luck. As such, it will conflict with the (H) principle. Consider the following claim:

S ought to proportion her beliefs to what she knows only if, if S doesn’t proportion her beliefs to what she knows, then S is blameworthy for not proportioning her beliefs to what she knows.

Because there are many cases in which one is intuitively not to blame for failing to proportion one’s belief to what they know (as Williamson of course agrees), (H) tells us that it’s false that they really *ought* to have done so when in fact they did not. This is one explanation of the desire to make the triggering conditions of our normative requirements accessible in the sense that

we are always in a position to know about them. That is, doing so enables us to eliminate the possibility of resultant luck in our conformity to norms (as in the example with Kate and the consequentialist norm). Once a norm's triggering conditions are accessible in the right way, violation of the norm goes hand in hand with the appropriateness of judgments of blameworthiness. A case like the Kate and Mary case will never count as the violation of a normative requirement. Our norms won't conflict with the (H) principle. More schematically, this interpretation of the natural reaction can be put as follows:

- i) Blameworthiness is central to normativity (S ought to \emptyset only if, if S doesn't \emptyset , then S is blameworthy for not \emptyset -ing). ((H) principle)
- ii) Blameworthiness is incompatible with resultant luck. (premise)
- iii) The right kind of access condition eliminates resultant luck. (premise)
- iv) Thus, the connection between blameworthiness and normativity generates an access condition on our normative requirements (i.e. it motivates the natural reaction). (from i, ii, and iii)

So far we only have a sort of structural understanding of why we might think that connecting blameworthiness and normativity leads to the idea that we should always be in a position to know what our normative requirements are. What's missing in all of this, of course, is any reason to think that blameworthiness is central to normativity, or in other words that (H) is true. It's interesting that Srinivasan, the person to whom this diagnosis of the natural reaction is due, doesn't really provide much of an explanation for why we might think this in the first place. Rather, her suggestion is accompanied by assertions such as: "The hypological perspective has a strong grip on contemporary normative philosophy, and perhaps contemporary culture more generally. Perhaps this is a result both of overconfidence in our theoretical powers and naïveté about the place of luck in human life" (Srinivasan unpublished, p. 34). Srinivasan's aim in introducing this diagnosis is to challenge the natural reaction. But I think her comments on the plausibility of this explanation of the natural reaction suggest that it is probably not quite the right place to be looking in the first place. Here are the comments I am referring to:

Consider a political leader who, doing his best to act on his evidence and follow his conscience, leads his country into an illegal, immoral and materially devastating war. That he is in some sense excused for so doing is and should be of little interest to those who are not his biographer. Indeed, a first personal fixation on one's own culpability seems to indulge a narcissism deeply at odds with the normative. It is fine and important to tell children that all that matters is that they try their best. But when this thought becomes a complete consolation for failure – when I, harming you, care only that I am trying my best not to – I have abandoned the deontological perspective. Williams launches a similar criticism against the Christian outlook when he says that it “associates morality simultaneously with benevolence, self-denial, and inner directedness or guilt (shame before God or oneself)” [(Williams 1981, p. 244)]. Seeking to avoid blame, moral or epistemic, is a form of cowardice, a turning inward of the normative perspective (Srinivasan unpublished, p.34).

Ultimately, I think I agree with the spirit of these remarks. But I also think this suggests that the foregoing is not a good way of spelling out the natural reaction. For example, the case of the political leader seems like a counterexample to (H). In other words, the claim that

S ought not to have led his country into an illegal, immoral and materially devastating war only if, if S did so, then S is blameworthy for leading his country into an illegal, immoral and materially devastating war

seems false. Srinivasan's remark that the political leader is in some sense excused suggests we can imagine the case such that he's not to blame for leading his country into an illegal, immoral and materially devastating war. But her point is well taken that it seems true that normatively he ought not to have done so—or that something could be said to have gone horribly morally wrong for the leader. The details are of course complex, and I wouldn't want to rest much on this case.

Moreover, it's important to keep some things in mind about blameworthiness before assessing (H). It's important to note, first, that the idea should not be taken to imply that violation of normative requirements goes hand in hand with the appropriateness of *blame*. There are crucial differences between the conditions of the appropriateness of blaming someone, and whether that person is *blameworthy*. For one thing, blaming is highly *positional* in nature—for example, being implicated in the wrongdoing of some other agent can place constraints on the appropriateness of one blaming that agent. Meanwhile, while *blameworthiness* is also considered by some to be to a certain degree positional in nature, it is arguably less so (Kelly *forthcoming*). And we might think that since violating normative requirements is not a positional matter, the

idea that there is some essential connection between blameworthiness and normativity is at least more plausible than the idea that there is some essential connection between the appropriateness of actual blame responses and normativity. Moreover, it's also important to note that blameworthiness and the sorts of responses that this may elicit from people can be wide-ranging and subtle. Kelly notes that:

These responses are commonly associated with condemning and punishing. But blaming responses might also be understood, more broadly, to include less retributive sentiments, such as disappointment, sadness, or resignation. Although they are not oriented punitively, these attitudes respond to an agent's demonstrated capacity for wrongdoing and take it seriously. They may be accompanied by behavioral responses that are not retributive in nature, such as calling for an explanation or apology, lowering or otherwise changing expectations regarding future interactions, breaking off a relationship, trying to make the wrongdoer feel guilty or make amends, and so on. I submit that any of these responses might qualify as blame when they accompany or express moral criticism of a person (Kelly *forthcoming*, p.2).

We might think that taking this into consideration serves to weaken, and thus render more plausible, the claim that (H) is true (i.e. in the sense that one can easily imagine cases of norm violation that do not go hand in hand with, say, *punitive* measures – for example, violations of the evidential norm).

However, even granting all of this, it's still not clear to me that the foregoing is the best way of spelling out the natural reaction. I suggest that the way to investigate the matter further is by looking more closely at the notion of *responsibility*. Blameworthiness and praiseworthiness are obviously closely connected to the notion of responsibility. So we might think that the view that blameworthiness is central to normativity just is the view that responsibility is central to normativity. But there are at least three different notions of responsibility relevant to our purposes. Examining the differences between these notions will allow us to see why (H) is implausible. However, I'll argue that a picture emerges that provides a more plausible, though closely related, way of spelling out the natural reaction.

4.4 Normative Requirements and Responsibility

There is a standard distinction in the literature on responsibility between accountability and attributability (Shoemaker 2011; Watson 2004, 2011).⁹⁰ Accountability is a notion of

⁹⁰ For a dissenting view, at least when it comes to *moral* responsibility, see Smith (2012). Smith argues directly against Shoemaker's view in favour of a unified account of moral responsibility as "answerability". It's not clear to

responsibility that is often spelled out *in terms of* the appropriateness of praise and blame or other similar reactive attitudes (Strawson 1962). Shoemaker articulates accountability in the following way. He says that we are responsible in the sense of being accountable when we flout “relationship-defining demands”. He adds that a certain kind of blame is fitting when these demands are flouted. That is, to hold someone to account “is precisely to sanction that person, whether it be via the expression of a reactive attitude, public shaming, or something more psychologically or physically damaging” (Shoemaker 2011, p. 623). Contrast this with attributability. There are different ways of spelling out the notion of attributability. Scanlon (1998) and Smith (2008; 2012), for example, treat attributability as coextensive with S being *answerable* for an action or attitude \emptyset . Being answerable amounts to the idea that S can be called upon to give reasons in order to justify her action or attitude. For example, Smith claims that the actions I perform “reflect my assessment of reasons, and therefore I can, in principle, be called upon to defend them and am open to rational (and in some cases moral) criticism if an adequate defence cannot be provided” (Smith 2008, p.385).

Shoemaker argues that this way of spelling out attributability is too strong. An action or attitude can also simply be *attributable* to an agent where this picks out a genuine kind of responsibility for the action or attitude which comes apart from having the capacity to give reasons in justification of one’s action or attitude. Shoemaker points to *aretaic* assessments to fill this out.⁹¹ ⁹²We often make judgments about an agent’s character and the actions and attitudes that manifest that character without requiring that an agent is able to justify the relevant attitude or action. Shoemaker gives the example of Huck Finn:

me how to assess the consequences of all of the different views in the literature about moral responsibility for epistemic responsibility. Smith herself takes the view she defends to have direct consequences for our understanding of responsibility for *belief*. However, what she has in mind is *moral* responsibility for belief: “I have argued that this condition of moral responsibility implies that we are morally responsible not only for our intentional actions but also for the majority of our attitudes (our beliefs, desires, emotions, etc.)...” (Smith 2012, p.578). In any case, responding to Smith’s conception of responsibility would take us too far afield in the present context.

⁹¹ See Watson (2004). According to Smith (2012), the notion of aretaic appraisal was introduced by Watson. According to Watson, aretaic appraisal is an evaluation of an agent’s “excellences and faults—or virtues and vices—as manifested in thought and action” (Watson 2004, p.266).

⁹² Shoemaker also argues from cases of irrationality. He argues that a consequence of the answerability view is that agents are not responsible for their irrationality in certain cases where this is implausible (Shoemaker 2011, p.607-609). Smith replies in Smith (2012), p.579.

[W]e rightly assess Huck Finn as courageous, loyal, and kind, regardless of what he thought of himself and independently of what he took his reasons to be when he failed to seize the opportunity to turn in the slave Jim. Furthermore, these are assessments of his actual character, of his authentic agential nature, and as such they have real depth. Answerability demands to Huck, however, would be senseless. After all, he lacks access to any reasons there are that ground his evaluative commitment (Shoemaker 2011, p.614).

In the case of Huck Finn, according to Shoemaker, what is essential is that the action or attitude is something that discloses Huck's "authentic agential nature", not his ability to justify the action or attitude.⁹³ The crucial thing about attributability on both views is that the action or attitude for which the agent is responsible is one that discloses the agent's evaluative commitments and as such the agent's authentic agential nature.⁹⁴ The disagreement between Scanlon and Shoemaker is about what is required in order for this latter condition to be met. I tend to favour Shoemaker's position. For one thing, I find the Huck Finn case plausible, and I find it plausible that the ability to justify one's action or attitude \emptyset comes apart from whether \emptyset is in some relevant sense a manifestation of one's agential self.

In any case, one point I want to make is that distinguishing between accountability and attributability allows us to see why the (H) principle above is implausible. I should acknowledge before continuing, however, that we seem to be getting continuously further away from the central issue of this chapter —namely, spelling out the natural reaction to Williamson's view about *evidence*. We seem to be a long distance from doing epistemology. However, I think the distinctions that I am drawing, and the various (non-epistemic) cases I'm discussing are helpful for elucidating a broader point about normativity in general, one that I

⁹³ Shoemaker appeals to cases of arational emotional commitments for a more general defence of the claim that there are cases of responsibility as attributability that come apart from the appropriateness of a demand for justification (Shoemaker 2011, pp.611-612). Shoemaker claims of himself, for example, that his deepest emotional commitments "reflect on me, on my deep self, and in particular on who I am as an agent in the world, but they are not grounded in any evaluative reasons (at least of the sort I take to justify my attitudes or actions)" (pp. 611–12). Smith responds to this idea in Smith (2012), p.582. Smith, responding to Shoemaker's cases of aretaic appraisal, invokes cases of OCD, depression, and Tourette's syndrome as counterexamples. She says: "I think a more careful look at [Shoemaker's] cases will reveal that we are willing to apply these aretaic predicates [i.e certain judgments of character] to agents only in cases where we regard the agent as answerable for the attitudes and conduct on which these appraisals are based. And I fear that if we deny this presupposition, then we will have to regard agents as morally responsible for many psychological conditions for which, intuitively, they do not appear to be responsible in any sense."(Smith 2012, p.586). I'm inclined to bite the bullet here in favour of Shoemaker. That is, I'm inclined to say that an action or attitude of a person with OCD is attributable to the person (and that they are in that sense responsible for it) even when it's in some sense the manifestation of their condition. What I don't think is plausible is that such an action or attitude is one that the person is *answerable* for.

⁹⁴ Note that I don't intend to suggest that there are not complex issues surrounding how to spell out what counts as attributable or what attributability is.

take to apply in the epistemic domain (i.e. insofar as our concern has to do with a normative principle about evidence). So before bringing this all back to epistemology, let me continue with the distinction between accountability and attributability and the implications it seems to have for the (H) principle.

It will be helpful to look at another case to spell this out. Shoemaker has a compelling account of responsibility-judgments concerning psychopaths.⁹⁵ Shoemaker argues that psychopaths can be attributable but not accountable in at least some central cases of actions or attitudes. The reason why is that accountability, by virtue of its essentially social structure—recall, it is spelled out in terms of a sensitivity to relationship-defining demands—implies certain capacities which the psychopath lacks. These are capacities to recognize the *reasons* implied by the relationship-defining demands at play in their relationships with others. As Watson puts it (in a discussion of the case of psychopaths that shares significant similarities with Shoemaker's):

[P]sychopathy (as I read the evidence) precisely involves an incapacity to recognize the interests of others as making any valid claims on them. Consequently, they are disabled from standing in the reciprocal relations or (to use another idiom) from engaging in the mutual recognition that lies at the core of moral life. In John Rawls's sense, they lack the features of moral personality: a sense of justice and a conception of the good (Watson 2011, p.307).

By appeal to the “ought implies can” principle, Shoemaker argues that the psychopath cannot be held accountable for her actions and attitudes (importantly, Shoemaker means *morally* accountable – he claims that there is no reason why we cannot stand in other kinds of *non-moral* relationships with psychopaths, such as an “aesthetic relation” or “fellow citizenry relation,” and thereby hold the psychopath accountable in these non-moral arenas).⁹⁶ This is because the psychopath lacks the relevant capacities required in order to do so. Nevertheless, we can still understand the psychopath as responsible in the sense of attributable. The psychopath's

⁹⁵ I rely on Watson's (2011) characterization of the psychopathic profile: “Psychopaths (as such) suffer no psychoses or notable neuroses, or general deficits of intelligence. Their condition is characterized by callous interpersonal relations, lack of significant attachment to other people or institutions (shallow and fleeting “friendships,” amorous relations, and loyalties), lack of a sense of shame and guilt, lack of self-criticism, refusal to take responsibility for the troubles caused to others or oneself, and lack of sincere commitment to long-term goals. Psychopathy is present from childhood and, by all indications, endures a lifetime. As this description implies, this condition has proven unamenable to psychopharmacological or psychological therapies and treatments” (Watson 2011, p. 308).

⁹⁶ Shoemaker's key point is that psychopaths are incapable of recognizing *moral* relationship-defining standards, recognition of which is necessary (according to Shoemaker) for an agent to be held morally accountable.

actions and attitudes disclose the agent's evaluative commitments, they are expressions of his self qua agent.

Imagine a psychopath who cheats an elderly woman out of her life savings. If Shoemaker is right, the psychopath is not *accountable* for this action. The psychopath is insensitive to the key moral relationship-defining demand that he be sensitive to the elderly woman's interests.⁹⁷ Thus, Shoemaker claims that "The kind of moral address involved in holding someone accountable is thus pointless with respect to the psychopath" (Shoemaker 2011, p.629). This explains why there is a good case to be made that the psychopath is not an appropriate target of blame.⁹⁸ This is not to say that blame wouldn't be *understandable* in such a case (i.e. it's understandable that, for example, the elderly woman would blame the psychopath). The point is that blame would be *unreasonable*.⁹⁹ Nevertheless, it's intuitive to treat the psychopath as morally responsible for this action. And the notion of responsibility Shoemaker thinks is appropriate for this case is attributability, as I've just explained. Notice, however, that it's plausible to regard the psychopath's action as violating some moral requirement. The psychopath does something wrong in cheating the elderly woman out of her life savings. This is a case, then, in which blameworthiness and a violation of normative requirements plausibly come apart. In

⁹⁷The point is not that the psychopath cannot understand that some moral relationship-defining demand applies to him. The point is that the psychopath is incapable of seeing that the demand is reason-giving. "Let me be clear: the psychopath may well understand that there is a demand being made of him; what he cannot understand is that the demand is reason-giving. There are facts being presented to him that he judges irrelevant—"Yes, I understand that when I do this you will be in pain and you won't like it," he says, "but so what?"—so they simply fail to constitute any sort of constraint on his deliberations or motivations (Shoemaker 2011, p.629).

⁹⁸ Although of course it depends on what we mean by "blame". Shoemaker distinguishes between "Scanlonian blame" and "accountability blame", and allows that the psychopath is "Scanlonian" blameworthy (Shoemaker 2011, p.630). The details would take us too far afield here. For now I'm depending on the intuitiveness of the idea that if an agent is incapable of recognizing that a key moral relationship-defining demand (being sensitive to others' interests) is reason giving, and thus that these demands are incapable of constraining his deliberations or motivations in practical reasoning, then the sort of reactive attitude that is involved in holding someone accountable – i.e. blaming them – is simply lost on that agent. This supports the claim that such an agent is not the appropriate target of a blame response. Shoemaker provides a helpful analogy for the case of psychopaths. It involves a community of aliens who are morally just like the members of one's own community, except that they believe that grass has a kind of moral status. He argues that while one may be able to stand in a variety of moral relations with these aliens, when it comes to the grass, one is simply incapable of seeing what the aliens think is genuinely reason giving when it comes to grass (i.e. its being crushed or breaking) as genuinely reason-giving. And so, one is incapable of recognizing certain relationship-defining demands in one's moral relationship with the aliens –i.e. "respect the moral status of the grass". For example, when you walk across some grass to get a better look at a particular species of bird in the forest, an alien spots you and shouts at you with indignation. Shoemaker asks whether you are the appropriate target of this blame-response. He of course claims that you are not. And the point is that the sort of insensitivity to a certain relationship-defining demand –"respect the moral status of the grass" – is analogous to the psychopath's insensitivity to moral relationship-defining demands. See Talbert (2008) for opposing view on the issue of whether psychopaths are blameworthy.

⁹⁹ Talbert (2008) makes this distinction.

other words, perhaps it's another case like Srinivasan's case of the political leader above. It seems like a counterexample to the following claim:

S ought not to have cheated the elderly woman out of her life savings only if, if S did cheat the elderly woman out of her life savings, then S is blameworthy for cheating the elderly woman out of her life savings.

There are many other examples we could appeal to. However, what's important about this case is that it is still a case of attributability. And I think this motivates a different but closely related way of spelling out the natural reaction. We could put what I have in mind in the form of the following principle. Call it the "attributability principle" (A):

(A): S ought to \emptyset only if, if S doesn't \emptyset , then S is attributable for not \emptyset -ing

The idea is that attributability, as opposed to blameworthiness, is central to the notion of a normative requirement. What (A) says is that it's only true that S is normatively required to \emptyset if, were they not to \emptyset , they would be attributable for not \emptyset -ing. This way of spelling out the natural reaction avoids the implausible implications of the blameworthiness view. For example, the psychopath case does not seem like a counterexample to the following:

S ought not to have cheated the elderly woman out of her life savings only if, if S did cheat the elderly woman out of her life savings, then S is attributable for cheating the elderly woman out of her life savings.

But more positively, I think the plausibility of this idea is derived from the plausibility of the simple idea that we make normative demands on the actions and attitudes of *persons* or *agents*. We don't make normative demands, for example, on inanimate objects, or the proceedings of nature. The domain of normative requirements does not plausibly extend beyond the actions and attitudes of persons (and perhaps *persons* themselves (i.e. one's character)). That said, here's a possible objection against this way of positively motivating a connection between attributability and normative requirements. The objection is that we do attribute things to inanimate objects and to the proceedings of nature. And we use the language of responsibility to do so. For example, it's perfectly intelligible to say that the avalanche was responsible for the destruction of the ski lodge. But we would not plausibly say that the *avalanche* violated some normative requirement by destroying the ski lodge. In response, I don't think it's plausible to understand

this in terms of attributability as I've outline that notion. Rather, I think what's plausible is that the avalanche is causally responsible for the destruction of the ski lodge. That said, it's not difficult to imagine someone putting pressure on me to explain exactly what the difference is. But my first response would be to say that the avalanche does not disclose any evaluative commitments, or manifest its agential self in destroying the ski lodge. It doesn't have any evaluative commitments or an agential self.¹⁰⁰ Getting into the details of this enormously complex issue would take us to far afield.

4.5 Attributability and Access

One way of putting what the proponent of the blameworthiness view has in mind is that the presence of an excusing condition serves to undermine the judgment that a normative requirement has been violated. On this view, having an excuse *just means* that one's action or attitude fails to violate the relevant norm(s). One way of putting the present proposal is to say that, not just any sort of excusing condition, but rather a particular kind of excusing condition serves to undermine the judgment that normative requirements are violated. In particular, when an excusing condition is such that the agent no longer counts as attributability-responsible for some action or attitude, this undermines the judgment that an action or attitude violates a normative requirement. Under such conditions, the relevant action or attitude no longer counts as an expression of the agent's self *qua* agent.

Compare two cases. Imagine that you absentmindedly run someone over after an extremely long-winded drive through a blizzard. Because you were so tired you lost your focus just at the moment someone suddenly appeared in front of the car. In such a case, the action is attributable to you in the sense that you were being *absent minded*. You weren't being careful enough. And the thing that's important is that *you* weren't being careful enough. It remains plausible to insist that you did something *wrong* here, even though our inclination to blame might be mitigated by certain features of the case. But now imagine that you are driving in an obstacle course that's full of life-like dummies, and your objective is to run all of the dummies over. You've been led to believe that this is just a game and there are no real people on the obstacle course. Unbeknownst to you, however, a real person has snuck onto the course and in the process of doing your best to succeed in the game you run the person over. Intuitively, the

¹⁰⁰ See Gibbons (2013, Ch.6) for an interesting discussion of some of the difficulties in this area.

fact that you run this person over does not *per se* disclose anything about your evaluative commitments, or manifest your agential self. The action is not attributable to you (though you are of course *causally* responsible). Indeed, I think this is where a central point for my purposes needs to be made.

I want to claim that attributability implies an accessibility condition. We can call it an “accessibility condition on attributability”. It’s not clear exactly how this should be spelled out. For a start, we might say:

Necessarily, an agent is attributable for Ø-ing only if they are in a position to know that they are Ø-ing.

Note that this condition says that the agent must be *in a position to know* and not that the agent must *know* that they are Ø-ing. There is a big difference. For example, one can fail to know that they are Ø-ing where it would be appropriate to say, “but they *should have known*”. However, if one is not even in a *position* to know, it’s far less clear there’s any sense to the claim that they should have known. After all, knowing is not something the agent is in a position to do. That said, we may also need to add that an agent can of course fail to be in a position to know when they ought to have been in a position to know. For example, doctors are obligated to read about the latest important findings in their area. Now imagine a GP who doesn’t bother, and so ends up making an error that kills a patient. The doctor was not in a position to know that what she was doing was an error (according to the latest findings), but she ought to have been in a position to know. So perhaps the access condition should go:

Necessarily, an agent is attributable for Ø-ing only if they are in a position to know that they are Ø-ing (and they have not culpably failed to put themselves in a position to know that they are Ø-ing).

There are no doubt further refinements we could make to such a condition. But the basic idea behind this condition of accessibility on attributability is that if S is not in a position to know that they are Ø-ing (and they have not culpably failed to put themselves in a position to know), it’s not at all clear that their Ø-ing discloses their evaluative commitments or manifests their authentic agential nature. For ease of exposition, from now on I will take the further condition that “they have not culpably failed to put themselves in a position to know” as implicit when I

talk about being in a position to know. Taking the access condition on attributability together with (A), we get the result that when an agent is not in a position to know that they are \emptyset -ing, this undermines the judgment that in \emptyset -ing the agent violates some normative requirement. One way of putting the thought here seems simply to be that attributability, like blameworthiness, is incompatible with the notion of resultant luck discussed above. Recall, as Srinivasan put it, “this form of luck is present when one’s ignorance of the relevant circumstantial facts prohibits one from controlling the outcome or result of one’s actions” (Srinivasan *unpublished*, p.15).

To illustrate we can bring this back to the BIV case. The BIV’s failure to proportion their beliefs to what they know is plausibly not something that is attributable to them. The BIV simply has no way of knowing that they are not proportioning their beliefs to what they know. And because the BIV’s failure to proportion their beliefs to what they know is not attributable to them, this undermines the judgment that they violate a normative requirement in so doing. So this conflicts with Williamson’s view about evidence. Moreover, I don’t think Williamson’s appeal to excusing conditions helps here. When he claims that the BIV violates the evidential norm but has an excuse, the kind of “excuse” he has in mind is precisely that the BIV is not attributable (though Williamson of course doesn’t put it this way) for failing to proportion his beliefs to what he knows. Again, this is because the BIV is not in a position to know what he knows.

When it comes to (A), it’s important to avert a potential confusion. The confusion might arise, for example, if we return to the case of Kate and the soup. Someone might want to point out that of course Kate’s action of bringing the soup to Mary is attributable to Kate. But the intuition in the case was supposed to be that it’s implausible that Kate violated some moral requirement. To avoid a counterexample to (A), we need to distinguish the following: while bringing the soup is attributable to Kate, what’s not plausible is that bringing a soup with allergens that will make Mary ill (*qua* such a soup) is an action that is attributable to Kate.¹⁰¹ And that’s just because she was in no position to know the soup would have the bad effects. Bringing such a soup to Mary does not disclose Kate’s evaluative commitments, it is not an

¹⁰¹ I recognize that I’ve already been in tricky territory in the philosophy of action, and I’m now getting into even trickier territory. But again, the point here is to spell out the beginnings of a general way forward in understanding what the intuition that there’s *something* odd about Williamson’s view about evidence (vis-à-vis the evidential norm) could possibly amount to.

expression of her self qua agent. And the same thing needs to be said in cases in which Williamson's evidential norm is putatively violated, where the person is not in position to know that its triggering conditions do not obtain. Take the BIV: while it's of course true that the agent's belief that *p* is attributable to the agent (insofar as the agent, say, is responsive to the way things appear and checks for relevant defeaters, the belief reflects the agent's evaluative commitments). What's not plausible, in an analogous way to the Mary case, is that the agent's belief *qua* belief that's based on something less than knowledge, is attributable to the agent. After all, the agent is in no position to know that she does not have the knowledge she thinks she has.

Here's one way to put the point: in both the BIV case and the Mary case, *had the agent known* that the conditions under which their actions/attitudes come out less than ideal obtained, they would not have performed the action/formed the attitude. If, say, I had known that the appearance as of a hand in front of me was indeed a *mere* appearance, I wouldn't have formed the belief that there is a hand in front of me. Williamson would no doubt claim that all this explains is why I have an excuse. But what I'm suggesting is that the excusing condition is one such that the faulty belief (*qua* faulty belief) is not attributable to me. Thus, given (A), the excusing condition undermines the idea that I violate a normative requirement.

Another important confusion to avert is the following. We have to distinguish between an agent's being in a position to know that the conditions under which they count as \emptyset -ing obtain, and an agent's being in a position to know what their evaluative commitments are. I certainly do not mean to imply by my accessibility condition on attributability that agents must be in a position to know what their evaluative commitments are in order to be attributable for \emptyset -ing (perhaps something like that would be required if I was talking about *answerability*, but I'm not sure). It seems to me perfectly plausible that an agent's evaluative commitments could be rather opaque to them (an agent who is in denial or who represses some part of themselves could be an example, but there are probably many ways in which one's evaluative commitments are opaque to oneself).¹⁰² Rather, I take the accessibility condition on attributability to be that the agent must be in a position to know that the conditions under which they count as \emptyset -ing obtain, in order for \emptyset -ing to count as disclosing the agent's evaluative commitments, or manifesting their agential self.

¹⁰² This distinction also seems important when it comes to the Huck Finn case, for example.

If the basic story I'm telling is correct, then we need some other understanding of the evidential norm's triggering conditions in order to make sense of the evidential norm as a genuine normative requirement. We cannot simultaneously endorse E=K and the idea that the evidential norm is a genuine normative requirement. A BIV case is one particularly compelling way of bringing this out. But more schematically, this interpretation of the natural reaction can be put as follows:

- i*) Attributability is central to normativity (S ought to \emptyset only if, if S doesn't \emptyset , then S is attributable for not \emptyset -ing). ((A) principle)
- ii*) Attributability for \emptyset -ing implies a certain kind of accessibility to the fact that one is \emptyset -ing. (accessibility condition on attributability)
- iii*) Thus, the connection between attributability and normativity generates an access condition on our normative requirements (i.e. it motivates the natural reaction).
(from i*, and ii*)

The key difference between this argument and the argument from the (H) principle is that the (A) principle is more plausible.¹⁰³ We've located something that is more plausibly connected to normative requirements than blameworthiness, but which nevertheless implies the same sort of accessibility conditions. Thus, when it comes to questions about evidence in light of the normative principle about proportioning our beliefs to the evidence, accounting for the natural reaction would amount to endorsing a view about evidence that accounts for (A). It would either amount to endorsing a view about evidence such that we are always in a position to know what our evidence is, insofar as that is what is required for us to be attributable for proportioning our beliefs to the evidence on any given occasion; or it will involve rejecting (A), or the accessibility condition (or both), and explaining away their plausibility.

¹⁰³ That said, I've defended the claim about access in a different way as well. Although I think it might be possible to spell out the plausibility of the access condition in terms of issues about resultant luck (as I noted above), I've tried to avoid making that a central claim in the argument.

Concerning the first option, Williamson of course thinks that there can be no such view of evidence, and I'll return to his reason why in Chapter 5.¹⁰⁴ Thus, he would of course take the second route. So I'm suggesting that part of the burden on Williamson, in explaining why there's something highly counterintuitive about his view about evidence, is to explain what's wrong with at least one of the premises above. All of this needs a lot more development and defence in order to make a compelling case. But I turn now to offer a few more general points that I think lend support to the foregoing.

4.6 Further Reflections

It is important to keep in mind that while we can speak of the “normative” broadly in contrast with the “descriptive”, we can also speak of the “normative” more narrowly in contrast with the “evaluative”. For example, many moral philosophers distinguish between judgments about what one is *required*, or *obligated*, or has a *duty*, or *ought* to do on the one hand, and judgments about what would be *good* or *best* on the other. I think this is important when assessing claims—such as Williamson's, for example—that a recently envatted BIV who behaves with perfect epistemic responsibility is unjustified but nevertheless has a cast iron excuse for believing outright that they have hands. In particular, I think there is pressure on those who would claim that the BIV is *excused* but not *justified* for believing responsibly to demonstrate that the contrast they intend to latch onto is that between what the agent is *excused* for doing versus what they are *normatively required* to do, as opposed to a contrast between what the agent is *normatively required* to do versus what would be *good* or *best* for the agent to do (Broome 2013; Gibbons 2013). After all, it is easy to interpret the case in both ways (although I wager it is much easier to interpret in the second way). It is easy to imagine the BIV case as one in which the BIV does what she normatively ought to do (say, she forms her perceptual belief on the basis of careful enough attention to the look of things) while nevertheless does not do what would be best (for example, refrain from believing she has hands).

¹⁰⁴ As we'll see, Williamson's diagnosis of the natural reaction suggests that the *only* possible option for accounting for it could be rejecting (A) and/or the accessibility condition, since, as he puts it, there are no (non-trivial) luminous conditions. But moreover, as I will discuss, setting the putative impossibility of (non-trivial) luminous conditions to one side, Williamson's diagnosis might be taken to imply that taking the first option in accounting for the natural reaction would make it in-principle impossible to support a rejection of the sameness of evidence thesis. And this is because of the claim that the idea that we are always in a position to know what our evidence is entails accessibilist supervenience internalism (and that view, in turn, is plausibly tantamount to the idea that our evidence is the same in the good and bad cases).

One might want to support the idea that the BIV violates certain epistemic normative requirements by pointing out that it's perfectly natural to say that the BIV fails to believe what she *ought* to believe. After all, she believes something completely false and continues to do so systematically. However, it is worth noting that pointing to intuitive 'ought'-claims concerning the BIV's doxastic states won't settle anything in this matter. This is because the various "flavours" of ought in ordinary language vary in their semantic contribution to a sentence which is plausibly determined by the context of utterance.¹⁰⁵ For instance, it is easy to imagine a context for which it would be natural to read the following as an 'ought' that states an obligation or duty or requirement (and thus something paradigmatically "normative"):

- a) You ought to stop eating animals.

But not all uses of 'ought' state an obligation or duty or requirement. For example, people often use 'ought' in the following sorts of ways:

- b) They ought to be here by now.
- c) Milton you ought to be living at this hour (as England needs you!).
- d) If you don't want to blow the fuse, you ought to turn off one of your appliances.

None of these are naturally read as stating an obligation, duty, or requirement. For example, b) is referred to as the "epistemic ought" or "predictive ought". It has to do with something like what is likely or probable given our evidence. c) is sometimes referred to as the "evaluative ought". It evaluates some state of affairs relative to a given standard. In other words, it is one way of saying that it would be *good* or perhaps even *best* if Milton were alive (say relative to standards of what would be morally and politically good for England).¹⁰⁶ d) is an example of a "teleological ought". What this sentence tells you to do is entirely contingent on your goals (namely, say, not blowing the fuse, or more generally, keeping the power on, etc.). If a characteristic feature of the "normative ought" is its detachability, then this is not a normative ought.¹⁰⁷ So there are a variety of seemingly quite different uses of 'ought' in ordinary English.

¹⁰⁵ See Chrisman *forthcoming*.

¹⁰⁶ The example is borrowed from Chrisman (*forthcoming*). But it comes from Wordsworth's sonnet "London, 1802".

¹⁰⁷ An aside: many think it's possible to account for all of these different "flavours" of ought in a unified semantics by thinking of 'ought' as a unary modal operator. As Chrisman notes, "the orthodox view in theoretical semantics is that 'ought' is monosemous operator (akin to deontic necessity) but semantically underdetermined, such that its uses require significant contextual augmentation to determine a definite sense" (Chrisman 2012, p.3). The inspiration for this project is inspired primarily by Kratzer (1977, 1981, 1991). There are interesting issues

This complicates our understanding of any linguistic data used to motivate claims about normativity.

In light of this, I think the metanormative theorist who would challenge the centrality of attributability to normative requirements faces a second related challenge. This is the challenge of explaining what's distinctive of the normative, as opposed to, say, the evaluative (a notion that seems readily available in explaining away, for example, Williamson's interpretation of the BIV case) without appealing to judgments about responsibility. I won't try to do this on anyone's behalf here.¹⁰⁸

4.7 Conclusion

Taking stock, I've examined three potential ways of spelling out what the natural reaction to Williamson's view about evidence amounts to. The first way was the deliberation-guidance worry. I claimed that this worry did not entail that we ought to have any sort of access to our evidence that would be inconsistent with Williamsons' view. We can use the evidential norm to guide deliberation because, as Williamson claims, in practice we are often in a position to know what we know. I said that perhaps a deeper worry behind the deliberation-guidance worry concerns the idea that one could *ever* be normatively required to \emptyset when one is not in a position to know that that they are required to \emptyset . The first way of spelling out why this mightn't be possible turned on considerations about a deep connection between the hypological and deontological perspectives. I put this in terms of the (H) principle. I then turned to an examination of three different kinds of responsibility and suggested that the (H) principle is not compelling. Rather, a closely related (A) principle says something more compelling about normative requirements, but nevertheless implies the same sort of accessibility conditions as the (H) principle. This accessibility condition conflicts with Williamson's view about evidence. And this, in a nutshell, is what I've suggested the natural reaction amounts to. I've also suggested that accounting for the natural reaction (when it

that arise in understanding the 'ought' of obligation in terms of a unary modal operator (Chrisman 2012; *forthcoming*; Schroeder 2011). Chrisman provides a compelling account of how this might work in Chrisman (*forthcoming*).

¹⁰⁸To be sure, explaining what's distinctive of the normative might just be an impossible task (i.e. if the notion of normativity is an explanatory primitive) (see Broome 2013). I myself have only given a very rough outline of what I mean by "normative", with no pretention to have explained what's distinctive about it. Still, the objection can be pressed in a slightly different way: the challenge is to explain how we should understand the claim that the BIV 'ought' not believe that she has hands in terms that are not merely evaluative, but also not essentially connected to attributability.

comes to questions about evidence in light of the normative principle about proportioning our beliefs to the evidence) means doing one of two things. Either we endorse a view of evidence according to which we are always in a position to know what our evidence is, or we reject (A), or the accessibility condition on attributability (or both), and explain away their intuitive plausibility.

Chapter 5: Considering the Options

5.1 Introduction

In this chapter I consider some options moving forward. I start by examining the possibility that when it comes to accounting for the natural reaction and supporting a rejection of the sameness of evidence thesis, only the second option is available—that is, the option of rejecting (A) or the accessibility condition on attributability (or both) and explaining away their intuitive plausibility. I don't suggest that this would necessarily be a bad thing, but it's worth examining whether it's the case. The basic thought is that endorsing an internalist condition of access on our evidence makes rejecting the sameness of evidence thesis impossible. I appeal to recent work from Gibbons (2006; 2013) and Daniel Greco (2013) to show that there are theoretical resources available for claiming that it does not. The upshot of the discussion is that an internalist approach to accounting for the natural reaction and rejecting the sameness of evidence thesis is still an option. However, we are no further along in understanding how such a view would go, or what the best way to reject the sameness of evidence thesis is. In light of the difficulties raised so far for disjunctivist strategies, I conclude the chapter by taking a step back and considering whether the strategy is worth pursuing after all. In particular, I examine whether a strategy that rejects the sameness of evidence thesis has adequate scope for responding to scepticism about evidential justification generally. I argue that it does.

5.2 Anti-Luminosity

I've said that there are two ways of accounting for the natural reaction to evidential externalism. Either we endorse a view of evidence according to which we are always in a position to know what our evidence is (thus rejecting evidential externalism), or we reject (A), or the accessibility condition on attributability (or both), and explain away their intuitive plausibility. Williamson would of course take the second option. This is necessary insofar as he wants to endorse his particular form of evidential externalism. But he also has powerful resources for arguing that this is the only way to account for the natural reaction. In particular, Williamson claims that a commitment to the idea that we are always in a position to know what our evidence is boils down to an untenable commitment to the idea that we have a "cognitive home". This is a distinctive realm in which nothing is hidden, so long as we pay enough attention—our judgements about what to do and think can safely be grounded upon

“luminous” conditions in our cognitive home. Here is Williamson on the notion of a cognitive home:

There is a constant temptation in philosophy to postulate a realm of phenomena in which nothing is hidden from us. Descartes thought that one's own mind is such a realm. Wittgenstein enlarged the realm to everything that is of interest to philosophy. That they explained this special feature in very different ways hardly needs to be said; what is remarkable is their agreement on our possession of a cognitive home in which everything lies open to our view. Much of our thinking—for example, in the physical sciences—must operate outside this home, in alien circumstances. The claim is that not all our thinking could be like that (Williamson 2000, p.94).

Williamson claims that the notion of a cognitive home is nothing more than a “quaint relic of Cartesian epistemology”. He spells out luminosity as follows:

For every case x , if in x condition C obtains, then in x one is in a position to know that C obtains.

I've certainly spelled out the natural reaction in terms of something like the idea that our normative requirements, or their triggering conditions, must be luminous. This is because of the connections I've drawn between being in a position to know and attributability, and attributability and normative requirements via (A). But Williamson wants to know, what sorts of conditions could possibly be luminous? The condition of having my computer in front of me is intuitively not a luminous condition. It is not true that whenever this condition obtains, necessarily, no obstacle blocks my path to knowing that it obtains. For example, my computer might be in front of me and I could be asleep. Williamson talks about pain, meaning things with one's words, and statements of the form “it appears to S that A ”, as characteristic examples of luminous conditions. So a candidate example might be: necessarily, when it *appears* to me that my computer is in front of me, I am in a position to know that it appears to me that my computer is in front of me. Traditionally, luminous conditions obtain in the realm of the reflectively accessible, or introspectible.

A crucial point to draw from this, then, is that if the natural reaction to Williamson's view amounts to the idea that normative requirements are luminous, it's starting to look very much like accounting for it in the first way (i.e. by endorsing the idea that we are always in a position to know what our evidence is – thus rejecting evidential externalism) will simply entail a

commitment to accessibilist supervenience internalism about evidence and justification.¹⁰⁹ The reason for this is implicit in the *prima facie* plausibility of Williamson's examples of luminous conditions. If it is plausible that anything is a luminous condition, it is an internal condition—a condition that obtains in the realm of the reflectively accessible, or introspectible. On the face of it, then, this looks like bad news for the project of rejecting the sameness of evidence thesis. And that is because, if our evidence supervenes, for example, on our non-factive mental states (say, on the way things appear to us), it simply follows that one's evidence is the same in the good and bad cases. After all, for example, by stipulation the BIV case is one in which everything appears to you to be exactly the same as it does now.¹¹⁰ So we might think that, in addition to making it difficult to meet the explanatory challenge, maintaining an internalist condition of access on our evidence simply implies that it is impossible to reject the sameness of evidence thesis. This is not necessarily a surprising or even negative result. However, it is worth examining whether it in fact follows from endorsing an internalist condition of accessibility.

Importantly, Williamson argues that there are no non-trivial luminous conditions. He has an “anti-luminosity” argument.¹¹¹ Thus, he's not interested in examining the upshots of an internalist condition of accessibility on our evidence in the context of scepticism. He simply thinks such a position is unavailable. I don't want to get into the issue of whether the anti-luminosity argument is successful. An enormous literature is devoted to this issue (Cohen 2010; Greco *forthcoming*; Owens 2004; Smithies 2012). What I want to do is agree with Williamson that we do not have a cognitive home, but challenge his assumption that the natural reaction can be successfully diagnosed in terms of a commitment to the idea that there must be a distinctive internal realm in which luminous conditions obtain—that is, that it can be successfully diagnosed as entailing a commitment to the notion of a cognitive home. I agree with Williamson that the internal world—the realm of reflective and introspective access—is

¹⁰⁹ Of course, as I note shortly below, Williamson denies that there is any reason to think conditions obtaining in the internal world are more likely to be luminous than conditions obtaining in the external world. However, as I note presently, independently of Williamson's arguments against luminosity, there is more *prima facie* plausibility to the idea that internal conditions are luminous (at least according to many epistemologists). This is, to my mind, the best explanation of the internalists' interest in reflective or introspective accessibility.

¹¹⁰ Note that we can construe Pritchard's epistemological disjunctivism as a kind of supervenience internalism, provided that we understand supervenience internalism as the view that our evidence and justification supervenes on our *mental states* and that Pritchard's view can be understood as claiming that our evidence (in the good case) consists of *factive* mental states.

¹¹¹ See Williamson (2000), Ch.4.

not as important for justification as his quaint Cartesians take it to be.¹¹² However, I'm not convinced that the natural reaction (even understood in terms of normative luminosity) must be understood as implying that it is. I turn now to the project of showing how one can legitimately think both of these things at the same time.

5.3 Access Externalism

The first view I want to look at in order to spell this out is due to Gibbons (2006; 2013). Gibbons argues for a view that he calls “access externalism”. Access externalism is a view about justification that says that we have “privileged access” to the facts about what we have justification to believe (as we'll see below, this entails that we are always in a position to know the facts about what we have justification to believe), but that the supervenience base for facts about justification is not necessarily internal.¹¹³ Here is an instructive quote:

One strategy for explaining privileged access to our genuine requirements is to identify the reasons metaphysically, as things of a certain sort. Maybe they're all nonfactive mental states, or maybe it's the nonfactive mental states plus the necessary truths...The other strategy is to make your requirements, in a particular situation, out of whatever you have plain old regular access to in that situation (Gibbons 2013, p.184).

This talk of a “particular situation” is important for Gibbons and I'll return to it briefly below. But for now, note that according to Gibbons' usage, we can have a kind of “privileged access” to conditions that we are in a position to know through *empirical* investigation—or as Gibbons puts it, “plain old regular access”. In other words, privileged access is not the same thing as *reflective* or introspective access. This might seem odd. But his use of the notion of privileged access has nothing to do with a special source of knowledge, like introspection. As I understand Gibbons, the way to think about privileged access starts with intuitive judgments about responsibility.¹¹⁴ An agent has privileged access to a condition just in case judgments of

¹¹² See Smithies for a defence of what Williamson would call quaint Cartesianism (Smithies 2012). Smithies argues in favour of a sophisticated version of a luminosity condition on justification within a supervenience internalist framework.

¹¹³ I should note right away that although access externalism is a view about justification, all of the crucial claims Gibbons makes for our purposes can be framed in terms of *propositional* justification. If we assume that having evidence that p is necessary for having propositional justification that p, then we can translate Gibbons' claims about justification into claims about evidence. So, even if Gibbons does not endorse that assumption, I will interpret his discussion in this way when I want to talk about evidence.

¹¹⁴Gibbons doesn't explain what he means by “responsibility”, although he does explain what he means when he talks about *epistemic* responsibility: “Epistemic responsibility as I understand it is not a responsibility for doing something. It is a responsibility for knowing something. Not all obligations, even moral obligations, are

the agent's responsibility concerning their knowledge that the condition obtains would *always* be warranted. It doesn't matter how the agent is able to know about the condition (whether it's through perception, or introspection). The important thing is that, having privileged access to a condition amounts to the idea that, when the condition obtains, this is something they are responsible for knowing (and thus the agent can be said to violate some normative requirement if they fail to proportion their beliefs accordingly).

So how does Gibbons argue for access externalism? He starts with the following case:

The other morning, I went downstairs to make a mushroom, jalapeno, and cream cheese omelette. I had checked the night before to make sure we had all of the ingredients. Since Sunny rarely eats breakfast, it was reasonable for me to believe that the ingredients were still there. I went to the refrigerator and pulled out the eggs and mushrooms. While chopping, I firmly believed that I would soon have a mushroom, jalapeno, and cream cheese omelette. Unfortunately, in plain sight on the door of the refrigerator, there was a note. 'We're out of cream cheese.' I didn't notice the note, but I should have. After all, this is where we leave notes in our house (Gibbons 2006, p.22).

Gibbons firmly believed that he was about to have a cream cheese omelette for breakfast, but he should have known better. An important principle that Gibbons endorses is that if one should have known that not-p, then one is not justified in believing that p.¹¹⁵ Granting this principle, because Gibbons claims that he should have known that he was not going to have a

obligations to act. You ought to care about the well-being of others, even if caring is not an action. To say that you are epistemically responsible for the fact that *p* is to say, roughly, that whether you are aware of the fact or not, you ought to know that *p*" (Gibbons 2006, p.23). It's interesting to note that he likely doesn't have *accountability* in mind, at least insofar as accountability is linked to blameworthiness: "While I will rely on your intuitions about when people ought to know and believe things, nothing I say will rely in any way on anyone's intuitions about when blaming someone is appropriate. The appropriateness of the action does not completely determine the appropriateness of blaming. If I stop blaming you because I forgive you, something that happens every day, this is a situation in which I think you did something wrong, but I do not think blame is appropriate. Forget about blame, and think about what ought to be" (Gibbons 2006, p.22). That said, in this quote Gibbons is talking about *blame* responses as opposed to blameworthiness. In any case, the text is consistent with interpreting Gibbons as thinking about responsibility in terms of attributability.

¹¹⁵ This principle is motivated by two further principles. The first is a principle that says that if you ought to know that not-p, then you oughtn't believe that p. As long as we're understanding 'ought' as an epistemic (or doxastic) 'ought', this seems relatively straightforward. Consider an example: Jim ought to know that he's not going on vacation tomorrow, since recently learned that all the flights have been cancelled. But Jim believes that he's going vacation (say, he's got a bad case of wishful thinking). Jim violates the principle (although it's true that it seems he violates more than one principle here). The second principle is essentially an expression of a deontological view of justification (which is consistent with the way I've been talking about justification all along in this thesis). It says that if you ought not believe that p (again, epistemically speaking), then you are not justified in believing that p (Gibbons 2006, p.24). Gibbons notes that it's controversial to say that meeting one's epistemic deontological requirements (if there are any) is sufficient for justification, but he claims (and I agree) that it's necessary. Putting these two principles together, we get the principle above: if one should have known that not-p, then one is not justified in believing that p.

cream cheese omelette for breakfast (because of the note), he is not justified in believing that he is about to have a cream cheese omelette.¹¹⁶ Gibbons uses this case as a counterexample to supervenience internalism. To see why, we have to look at his next case:

In a nearby possible world, things are the same inside me, but slightly different inside and outside the refrigerator. In that possible world, there is cream cheese in the fridge and no note on the door. My belief that I would have cream cheese for breakfast was based on the same reasonable grounds. Unlike the first case, there was no evidence (or potential evidence) available to me to override these reasonable grounds. So in this case I am justified in my belief about my breakfast (Gibbons 2006, p.22).

By stipulation there is no difference in the internal facts in each case. However, in the first case, there *is* a note on the fridge, and in the second case there is not. According to Gibbons, the note in the first case is something Gibbons is responsible for knowing about. Since it's in the place where they normally keep such notes, he should have known about it. As Gibbons says, if he should have known that not-p, then he's not justified in believing that p. However, in the second case he is justified. So because the cases involve internal duplicates with different justificatory statuses, justification does not supervene on the internal world.

What is the motivation for the idea that Gibbons is responsible for knowing about the note? It seems there are at least two things that Gibbons cannot say about this (if he wants to maintain that justification does not supervene on the internal world). We said that by stipulation the cases involve internal duplicates. So the motivation for the idea that Gibbons is responsible for knowing about the note cannot be that he is in some sense aware of the note but has not yet noticed it. If that's why we think it's plausible that he ought to know about it, then it's not true that the cases involve internal duplicates. Gibbons' awareness of the note in the first case would imply that there is a difference in his internal mental life, even though he hasn't noticed the note. If so, the case does not function as a counterexample to supervenience internalism.

Alternatively, we might find it plausible that Gibbons is responsible for knowing about the note in the sense that he has background beliefs about where notes of this sort are usually placed and that he has some responsibility to check that spot under such circumstances. Gibbons seems to draw on precisely this thought with the point that he knows that this is the place where notes are usually kept. However, if this is the motivation for the idea that he is

¹¹⁶ Although this is a claim about doxastic justification, the reason Gibbons claims he is not justified is because he *has* justification to believe not-p, and he fails to believe accordingly. So I think we can also interpret this latter point as meaning that he has *evidence* that not-p.

responsible for knowing about the note, it is doubtful that the case functions as a counterexample to supervenience internalism. This is because the lack of justification will be due to a lack of responsible inquiry on Gibbons' part. And that is compatible with the idea that the determinants of one's justification are internal to one. All the case serves to show is that an additional condition on justification is that one engage in duly diligent inquiry, or something along those lines. In other words, the difference in justification in the two cases is not (directly) down to a difference in the evidence. It comes down to a difference responsible inquiry. Again, that is compatible with supervenience internalism. We've merely added a further condition on justification as opposed to established the idea that the determinants of one's justification to believe that p can "reach all the way out to the world".

It looks like the motivation for the idea that Gibbons has a responsibility to know about the note has to come from somewhere else. Gibbons agrees. What he says instead is that the motivation comes from the idea that the note is sufficiently "obvious" and therefore, whether he's aware of it or not, he should be (Gibbons 2006; 2010). According to this way of putting things, the fact that the note is on the fridge is a determinant of Gibbons' justification to believe what he's about to have for breakfast because it's so obvious. He ought to have been more attentive *because* the note is in that sense part of his evidence. This is where a further point about the motivation behind the idea that Gibbons is responsible for knowing about the note needs to be made. In particular, it's clear that the note has to be in some sense *relevant*. It's plausible that being responsible for knowing that p, entails that p is in some sense relevant to one's inquiry. For example, in the cream cheese case, the note is defeating evidence – it undermines Gibbons' justification for thinking that he will soon be eating a cream cheese omelette. We might think that this is the reason why the note is relevant, and that this in turn is at least partly why he's responsible for knowing about it. Compare: there are also a number of breadcrumbs on the countertop where he is chopping ingredients. Intuitively Gibbons does not have a responsibility to know how many breadcrumbs are on the table in this case. Even though he's in a position to know about them, they are simply irrelevant to what he thinks he's going to have for breakfast. Ultimately Gibbons himself is prepared to leave the notion of relevance at the intuitive level. He feels comfortable doing so in part because he thinks supervenience internalists in large part do so as well:

I must set aside the important question of what makes a fact, mental or otherwise, relevant to the justification of a belief. Here, I focus on the question of what kind of access we must have to the relevant facts. In this, I take it that I am following the internalist. Internalism is a supervenience thesis. It does not say that everything in the supervenience base is relevant. It says that everything relevant is in the supervenience base. If you lack some belief about how you don't feel, the introspectively accessible fact that you don't feel that way may be irrelevant to the justification of your beliefs (Gibbons 2006, p.26).

The idea, then, is that S is responsible for knowing that p, if p is sufficiently obvious – where obviousness is to be spelled out in terms of relevance and being in a position to know.

Moon (2013) objects to this way of thinking about the motivation for the claim that Gibbons is responsible for knowing about the note. Considering his objection will shed some further light on the notion of obviousness. Moon presents a case that he takes to be analogous to the cream cheese case in all respects except that we should not have the intuition that the agent is responsible for knowing about the relevant evidence. Moon asks us to consider the famous interview in which Bertrand Russell claims that, if asked by God when he got to the pearly gates why he did not believe, he would reply “there wasn’t enough evidence!” Moon adds the following alteration to the situation, in which God replies:

‘Well Bertrand, there was relevant, available evidence for my existence at various points in your life. For example, did you know that during the evening of June 16, 1956, when you were alone in your room reading your book, I had an angel silently hold up a sign with the message “God does indeed exist!” right behind your head for five minutes? If you had turned your head and checked, you would have had strong, relevant evidence for my existence; indeed, you’d have come to *know* that I exist! Hence, you should have known that I exist, and so your belief in my nonexistence was unjustified’ (Moon 2013, p.147).

The idea is that Russell is not justified in believing that God does not exist, since he should have known about the sign that was above his head for five minutes—a sign that would have provided strong relevant evidence for believing in the existence of God. The thing to ask about this case is whether it is analogous to Gibbons’ cream cheese case in the right way. I don’t think it is. In particular, the sign above Russell’s head just wasn’t sufficiently obvious to generate the intuition that he should have known about it. After all, he was busy reading a book and the sign was completely out of sight. It’s true that all he would have had to do to see the sign is turn his head upwards, but it’s not plausible that while one is reading a book, a sign that one would have no reason whatsoever to expect to materialize (i.e. brought into existence

suddenly by an *angel*) above one's head is in any sense an obvious bit of evidence for or against some standing belief (about the existence of God) that one happens to have. Meanwhile, in the cream cheese case, the note is right there in front of Gibbons on the fridge, in the place where they usually keep such notes. It's plausible that, even if Gibbons is in no sense aware of the note, the existence of the note is sufficiently obvious such that he should have known about it.

But we have to be careful here. In particular, Gibbons is walking a fine line between the first sort of motivation for the idea that he ought to have known about the note that I said was not available to him – i.e. that he is in some sense aware of it but hasn't noticed it – and the second sort of motivation—namely, the idea that he failed as a responsible inquirer. Gibbons seems to invite further scrutiny along the first sort of line when he says:

Oddly enough, not everyone is moved by my original version of the story. So I will show you how to construct a range of stories, and you can use the one that suits your intuitions. We begin with a standard-sized notepaper stuck to the door of the refrigerator. Now gradually increase the size of the note. Make the letters larger, brighter, and more colourful until you get a billboard-sized note with the letters written out in bright lights. Somewhere along the line, probably well before you get to the billboard, you will get to the point where you think that whether I noticed or not, I should have” (Gibbons 2006, p.24).

Appealing to billboard-sized notes is not doing Gibbons any favours in this context.¹¹⁷ But his point is that somewhere well before you get to the billboard, you should have the intuition that a subject can be responsible for knowing about something that they may or may not have noticed (nor are aware of in any sense). As the note gets smaller, however, pressure builds on Gibbons to acknowledge that the motivation for the idea that he was responsible for knowing about the note comes from intuitions about responsible inquiry. And this is just because the smaller the note gets, the less inclined we'll be to think that, even though he is in no sense aware of the note, he ought to be in the sense that it is right now a determinant of his justification. Perhaps Gibbons can claim that the notion of obviousness does the work he needs for walking this fine line. Recall that we spelled out obviousness in terms of being in a position to know and relevance. In a sense, perhaps we can think of the fact that notes are usually kept on the fridge as part of what makes the note *relevant* in this case (as opposed to driving the intuition that he has a responsibility to inquire about it). If he can build relevance

¹¹⁷ Again, because he can't motivate the intuition that he responsible for knowing about the note in terms of the idea that he is in some sense aware of the note.

into the notion of obviousness in this way, then perhaps Gibbons' appeal to obviousness is an appeal to something in between awareness and diligent inquiry which allows him to walk the line.

This needs further defence, and of course there are further objections to Gibbons' view we could pursue here. Moreover, we haven't even begun to explore whether Gibbons' appeal to cases in rejecting supervenience internalism is consistent with holding on to an internalist condition of accessibility on our evidence. Gibbons has a story to tell about this which relies on the notion of an "epistemic situation." In particular, he claims that we must relativise our judgements about whether an agent has privileged access to the facts about what they have justification to believe to the agent's actual epistemic situation. As I read Gibbons, he claims that two importantly different claims need to be distinguished. The first is that Gibbons (in the cream cheese case) could not make a non-culpable error about whether he has justification to believe that he is about to have a cream cheese omelette, in *any epistemic situation*. The idea is something like: we consider all the possible worlds in which Gibbons is in his kitchen and believes he is about to make a cream cheese omelette and ask whether Gibbons could make a non-culpable error about what he has justification to believe. The second claim is that he could not make a non-culpable error about what he has justification to believe *in his particular epistemic situation*. The idea is something like: holding certain epistemic features of the case fixed, we consider all the possible worlds in which Gibbons is in his kitchen and believes he is about to make a cream cheese omelette and ask whether Gibbons could make a non-culpable error about what he has justification to believe. Gibbons claims that the note should only be considered a determinant of his justification in his *particular epistemic situation*. It would not be a determinant of his justification were Gibbons to be in *any* epistemic situation (for example, if the note was in his wife's back pocket, or if Gibbons was a BIV stationed in his kitchen being stimulated to believe he was about to have a cream cheese omelette). We need to examine in a bit more detail what Gibbons means and how he might defend this claim.

First of all, we need to spell out in at least some detail what Gibbons means by an "epistemic situation." Here's what Gibbons says about epistemic situations:

So how do we individuate epistemic situations? Why, epistemically of course. Whatever an epistemic situation is, the following should hold true. If a has justification for believing that p in e and b does not have justification for believing that p in e', then e is not the same epistemic situation as e'. You just don't get for free the claim that if a and b are the same on the inside then a and b are in the same epistemic situation. This is to beg the question. But I do not have to beg the question against the internalist (Gibbons 2006, p.34).¹¹⁸

In other words, Gibbons says nothing positive about what an epistemic situation is. He merely says that it is *not* to be understood in terms of how things are on the inside for an agent. But I'm not sure that Gibbons needs to rely on much more than rough intuitions about what an epistemic situation is. This just isn't where the real controversial action in his view is (that's coming shortly). So I'll suggest on Gibbons' behalf that we think of an epistemic situation in the following way: an epistemic situation (with respect to proposition p) is individuated in terms of the method of belief formation one can use in forming the belief that p and the obstacles (if any) that stand in the way of successfully using it. So for example, right now my epistemic situation is such that I can come to know by visual perception that there's a computer and a phone on my desk, not to mention all manner of things in the room around me. Nothing stands in the way of me coming to know all of these things as long as I turn my head and look. In short I don't think there's any more content to the idea than the notion of what I'm in a position to know (which is of course itself a vague idea). Recall what Williamson said about being in a position to know:

To be in a position to know, it is neither necessary to know p nor sufficient to be physically or psychologically capable of knowing p. No obstacle must block one's path to knowing p. If one is in a position to know p, and one has done what one is in a position to do to decide whether p is true, then one does know p. The fact is open to view, unhidden, even if one does not yet see it" (Williamson 2000, p.95).

Gibbons claims that the only way to imagine him making a non-culpable error about whether he is justified in believing that he is about to have a cream cheese omelette is by *changing the epistemic situation*—again, for example, by changing the epistemic situation such that the note is in his wife's back pocket, or Gibbons is a BIV stationed in his kitchen being stimulated to believe he is about to have a cream cheese omelette. In other words, only in cases in which he is in a position to know about the note does the note count as a determinant of his justification. The cases in which a non-culpable error about his justification is ostensibly

¹¹⁸ What he means by this last claim is that he's offered counterexamples to supervenience internalism – for example, the cream cheese case.

possible would co-vary with cases in which he is not in a position to know about the note. In other words, holding the relevant features of the case fixed – i.e. the features such that he is in position to know about the note – he could not make a non-culpable error about the note. And so, the idea that the note (something in the external world) is a determinant of his justification is compatible in this way with the idea that he cannot make a non-culpable error about what he has justification to believe.

I'm not sure what to make of Gibbons' employment of the notion of epistemic situation. It seems perilously close to getting him the claim he wants about the compatibility of an access condition with a rejection of supervenience internalism at the cost of being trivial. The notions of being in a position to know, privileged access, and epistemic situation all seem to sit so close together that it's difficult to get a sense of a genuine explanation of the compatibility of an access condition with a rejection of supervenience internalism. Regardless, I won't pursue this critique here. My main point in appealing to Gibbons is to show that there is theoretical space worth exploring when it comes to considering whether the idea that we are always in a position to know what our evidence is implies supervenience internalism. Ultimately, although I think Gibbons' access externalism is one option available, it leaves a variety of important questions open.

5.4 Contextualism and Luminosity

There are other possibilities for making the case that privileged access to justification is compatible with supervenience externalism. In particular, it will be helpful to turn to recent work from Greco (2013). Greco (2013) connects considerations about access with a form of contextualism about justification. There's a bit more stage-setting to follow before seeing how Greco's discussion is relevant for our purposes.¹¹⁹

Greco calls his particular brand of contextualism “contextualism about foundations”.

According to Greco, a foundational proposition is one that is justified even in the absence of further support from other beliefs. The non-contextualist foundationalist claims that a certain

¹¹⁹ Greco's discussion, like Gibbons', is in terms of justification. But he also makes all of his main points in terms of propositional justification. Once again, at this stage in the thesis we are assuming that having evidence that *p* is necessary for having propositional justification that *p*. So, even if Greco does not endorse that assumption, I think I can safely interpret his discussion in this way when I want to talk about evidence. Although he doesn't frame his view this way himself, I think Greco offers a way of rejecting supervenience internalism about evidence while holding on to a condition of access to our evidence.

class of propositions enjoy the status of being foundational in virtue of properties like being indubitable, or self-evident. Meanwhile, the contextualist about foundations claims that while in any given context, certain propositions count as justified even in the absence of further support from other beliefs, there is no cross-contextual line to draw around the kinds of properties foundational propositions must have in virtue of which they enjoy the status of being foundational.¹²⁰ Rather, foundational propositions—in a given context—simply have the formal property of being *presupposed* by parties to an inquiry or conversation.¹²¹ In different contexts, different propositions are presupposed. These may be propositions about the internal world or the external world, they may be what the non-contextualist foundationalist calls self-evident or indubitable, but they certainly need not be. They can be ordinary empirical propositions such as that one has hands, for example. Greco motivates contextualism about foundationalism in virtue of its ability to provide a promising middle way through a dilemma for traditional (non-contextualist) forms of foundationalism—namely, the dilemma of making foundations too restrictive and inducing scepticism, or too broad and licensing an unsatisfactory “epistemic permissivism” (Greco 2013, p.9). But he also notes that there’s an interesting upshot when it comes to preserving internalist intuitions about access.¹²²

Here is how Greco understands access internalism:

(Access) If S has justification to believe that p, then S is in a position to know that S has justification to believe that p.

Thus, he understands access in precisely the same way that we’ve been discussing it. Indeed, he says: “(Access) just amounts to the claim that the condition of having justification to believe that P is “luminous” in the sense of luminosity for which Williamson (2000) argues that no

¹²⁰ Greco claims that the broad point he argues for is compatible with a variety of different ways of individuating contexts, including the ways commonly associated with sensitive invariantism, relativism, expressivism, and what he calls orthodox contextualism. Since his discussion proceeds on the “orthodox” framework, I’ll note what he says this view amounts to here: Sentences of the form “S’s belief that P is justified” (as well as sentences attributing foundational justification more narrowly) express different propositions when uttered in different conversational contexts. Features of conversational contexts that affect which propositions such sentences express may include the presuppositions made by the participants to a conversation, the purposes of the conversation, and the practical situation faced by the participants to the conversation. In general, a proposition P’s being presupposed by the parties in a conversational context C tends to make sentences of the form “S’s belief that P is foundational,” true, when uttered in C (Greco 2014, p.6).

¹²¹ Note, on Williams’ (1991; 2001) version of a similar view about justification, Williams also requires a reliability condition. That is, our presuppositions in any given context must in fact be reliably formed beliefs in order to count as justified.

¹²² He also applies his view to a problem about updating evidence in a Bayesian framework (Greco 2013, p.11-22).

non-trivial conditions are luminous” (Greco 2013, p.24).¹²³ Greco recognizes that the thought that we are always in a position to know what our evidence is, or whether we have justification to believe that *p*, can easily seem to lead to a problematic position according to which our evidence or justification must supervene on facts that are reflectively or introspectively accessible to one. He puts the point nicely:

A natural, almost inescapable thought, is the following. (Access) could only be true if there were some distinctive realm of propositions that constituted the supervenience base for facts about what we have justification to believe, such that concerning claims in this distinctive realm, one is always in a position to know the truth (Greco 2013, p.25).

The natural thought is that justification, or what counts as our evidence, supervenes on the reflectively accessible, or introspectible—in short the thought is that it supervenes on the internal. But Greco is sympathetic to an access condition on justification.¹²⁴ He thinks he can show that an access condition does not lead to supervenience internalism. The basic idea is that, once we endorse contextualism about justification, the claim that being justified in believing that *p* is a luminous condition is a much less controversial idea.

Structurally, Greco’s thought is that whenever *S* has justification to believe that *p*, in a given context, *S* will also be in a position to know that *S* has justification to believe that *p*. The only way to make it the case that *S* is not in a position to know what *S* has justification to believe, is by changing the context, and thereby changing the facts that *S*’s justification supervenes on. An example will be helpful to spell this out. Imagine that Jim is hosting a party and he either tells Susan that the party starts at 8pm or that it starts at 9pm. Now imagine some people are discussing Susan’s epistemic credentials and someone raises the question of what justification Susan has to believe about when the party starts (in this context it’s a shared presupposition that Susan heard what Jim said). The natural answer, Greco suggests, is that she has justification to believe whatever Jim told her. In the given context, the facts about what Susan has justification to believe supervene on the facts about what Jim said (as Greco puts it, they

¹²³ We can translate this into a claim about what accessibility of evidence would amount to.

¹²⁴ He spells out the potential motivation in terms of considerations about guidance (which I examined last chapter), and also in terms of an access condition’s ability to explain the felt badness of “Moorean conjunctions” about justification (Greco 2013, p.24). This is an interesting area for further exploration of the motivation for an access condition which I would have liked to examine last chapter, but it would have taken us too far afield (see also Gibbons 2013; Smithies 2012).

determine the “local supervenience base”). Now, since we’re assuming that Susan heard what Jim said, it’s also true that Susan is in a position to know what she has justification to believe.

What’s important to note is that facts about what Jim said do not make a plausible supervenience base for facts about what an agent has justification to believe across *all* contexts (as Greco puts it, they do not determine a good “global” supervenience base for justification). Imagine, for example, we come to believe that Jim had mumbled or that Susan is hard of hearing (in this context, it’s no longer presupposed that Susan heard what Jim said). If we continue to think that the facts about what Jim said determine what she has justification to believe, then we’ll be claiming that she is not in a position to know what she has justification to believe (i.e. that Jim said the party starts either at 8pm or 9pm). Greco suggests, however, that in this context, when we ask about what justification Susan has to believe about when the party starts, a more natural answer is not going to be facts about whatever Jim said—rather, it may be facts about when Jim’s parties have started in the past (if we presuppose that Susan remembers this), or what Jim’s email said (again, presupposing that she remembers what it said). The point is that the facts we now take her justification to supervene on are once again facts that she’s in a position to know. So, when the context changes (such that what we are presupposing in assessing what Susan has justification to believe changes), the facts that what Susan has justification to believe supervenes on are always facts that she’s in a position to know. But this is not to say that there is some realm of facts which could form a supervenience base for justification such that in *any* context Susan would be in a position to know what she has justification to believe. As Greco puts it:

...[B]ecause we do not believe in cognitive homes of the traditional sort, we'll admit that whatever B is [i.e. the supervenience base for S's justification], we will be able to raise reasonable doubts about S's ability to know which are the truths in B. To the extent that we take such doubts seriously, the context will shift, and in our new context C', B will not be the local supervenience base for facts about what S has justification to believe. Rather, there will be some new supervenience base B', such that S does count as in a position to know which are the truths in B' (Greco 2013, p.28).

There are obviously a lot of details to be filled in here, and I haven’t even begun to provide any sort of support for Greco’s contextualism about justification. I think an interesting question is whether we would *need* to endorse Greco’s kind of contextualism in order to defend the idea that privileged access to justification is compatible with supervenience externalism (for

example, it's not clear to me exactly how to understand whether Gibbons' view about epistemic situations could be understood as a non-contextualist way, or whether it is just some other kind of contextualism about justification). In any case, I think Greco provides another potential way of showing that the idea that we are always in a position to know what our evidence is does not simply entail accessibilist supervenience internalism. There are options on the table for resisting that move. I take this to suggest that it would be hasty to conclude that accounting for natural reaction in the first way (that is, by rejecting evidential externalism and allowing that we are always in a position to know what our evidence is) makes it impossible in-principle to reject the sameness of evidence thesis.¹²⁵

5.5 Sameness of Evidence and the Closure Problem

All of this said, I haven't come up with a way of showing how this could be done. Indeed, even if there is no in-principle reason to think it cannot be done, I find it difficult to see how a strategy that holds on to the idea that one is always in a position to know what one's evidence is will be able to meet the explanatory challenge. In particular, if we acknowledge that the good and bad cases are indistinguishable, it's simply very difficult to see how one could claim that one's evidence in the good and bad cases is different, while holding on to the accessibility condition. The best way to reject the sameness of evidence thesis in light of the explanatory challenge seems to be some form of evidential externalism. However, of course, I've claimed that doing so comes with a significant burden, namely accounting for the natural reaction to evidential externalism in the second way – that is, showing why (A) or the accessibility condition on attributability (or both) is false, and explaining away the intuitive plausibility of these theses.

I haven't tried to do this on behalf of the evidential externalist. That strikes me as a book-length project in its own right. While I cannot engage further in that direction here, I think it's important to take a step back and note that the dialectic I've ended up with raises a question: perhaps rejecting the sameness of evidence thesis is not a promising way to respond to scepticism about evidential justification after all. Despite the motivation I've put forward for

¹²⁵ It's important to point out that the contextualist views about justification or evidence that Greco's view is inspired by have their own particular ways of responding to the sceptic (Annis 1978; Neta 2003; Williams 1991; 2001). I haven't discussed these responses in this thesis.

pursuing this line, we can see now that it faces considerable challenges – at least as troubling as the challenges faced by all of the other anti-sceptical strategies I’ve considered in this thesis. This invites a reassessment of the approach. In light of this, I want to close this chapter by examining a further reason we might think that rejecting the sameness of evidence thesis cannot provide an adequate response to scepticism about evidential justification.¹²⁶ I want to examine whether the approach has adequate scope for dealing with scepticism about evidential justification generally – in other words, whether there are any obviously important forms of scepticism remaining that rejecting the sameness of evidence thesis would leave untouched.

We started the discussion of scepticism with the closure problem for knowledge. This quickly led to the claim that the really interesting forms of scepticism challenge our evidential justification to believe ordinary empirical propositions. But I noted briefly that the closure argument can be reformulated in terms of evidential justification. Pritchard (*forthcoming*) shows how this can be done, although he does it in terms of “rationally supported knowledge”. For our purposes, we can translate “rational support” to “evidential justification.” The argument goes as follows:¹²⁷

- 1**) One cannot have evidential justification to believe that one is not a BIV.
- 2**) If one cannot have evidential justification to believe that one is not a BIV, then one cannot have evidential justification to believe that one is seated at a desk.
- 3**) Therefore, you cannot have evidential justification to believe that you are seated at a desk.

While the anti-sceptical strategies we’ve considered that proceed by denying the sameness of evidence thesis —if correct—clearly have a way of dealing with the underdetermination argument (and I hinted at the possibility that they might have a way of undermining the plausibility of Wright’s and Pryor’s sceptic’s appeal to (SPJ)), an important desideratum for

¹²⁶ The considerations to follow are of course similar in spirit to the issue we saw Owens and Schiffer raise in Chapter 3 for Williamson’s anti-sceptical strategy. That is, they claimed that even if we can reject the sameness of evidence thesis, this isn’t enough to defeat the sceptic. I noted that Williamson can respond at least to Owens’ objection by simply allowing that our justification would be different in good and bad cases.

¹²⁷ Pritchard formulates the argument as a paradox instead of an argument. That is, instead of the conclusion (3**), he simply points out that it’s of course highly intuitive that, imagining you are in fact seated at a desk, you do have rationally supported knowledge that you are seated at a desk. This claim stands in conflict with (1**) and (2**). Hence, according to Pritchard, we have a paradox.

these strategies is that they are able to deal with this sort of argument as well.¹²⁸ After all, this is just another way of formulating scepticism about evidential justification.

The most obvious way for a disjunctivist to challenge the reformulated closure argument (if they do not want to deny a closure principle for rationally supported knowledge) is to challenge premise (1**).¹²⁹ On the face of it, it seems that they are well positioned to do so. That is, they can claim that the only motivation for the claim that we do not have evidential justification to believe we are not BIVs is a sceptical appeal to the favouring claim that's at work in the underdetermination argument. That is, the sceptic proceeds by claiming that one's evidence does not favour that one is not a BIV over that one is a BIV, and thus, via (UP),¹³⁰ that one does not have justification to believe that one is not a BIV. If this is right, then disjunctivist strategies can (in principle) resist this claim in the same way that they resist the sceptical appeal to underdetermination of evidence concerning one's ordinary empirical beliefs. They will have a story to tell about how we can be justified in believing that we are not BIVs and reject premise (1**) of the reformulated closure argument on that basis (of course, they will also need a diagnostic story to tell about why this seems so implausible).¹³¹ The trouble is, this sort of response is conditional on its being true that the only motivation for premise (1**) is the sceptic's appeal to the favouring claim and (UP). If the sceptic has some *other* motivation, which strategies that proceed by challenging the sameness of evidence thesis don't speak to, then this approach to premise (1**) will at best challenge only one way of motivating that premise, and thus leave an important argument for evidential scepticism open for

¹²⁸Concerning Pryor's (SPJ), the way I said there is potential for undermining the plausibility of the sceptic's appeal to (SPJ) was as follows. Pryor's explanation of the plausibility of (SPJ) was that radical sceptical scenarios are "bad" in the sense that even if they obtained we'd have the same grounds to believe ordinary empirical propositions – this was the explanation for why it's plausible we should have antecedent justification to believe they do not obtain if we are to have justification to believe ordinary empirical propositions. I said that we might try to undermine the plausibility of the sceptic's appeal to (SPJ) by claiming that it's *false* that one's evidence (or "grounds") would be the same if one were a BIV. In effect, we'd be denying that our grounds E always "allow" the obtaining of a radical sceptical scenario in Pryor's sense, and thus that a condition on the hypothesis as counting as "bad" is not met.

¹²⁹ This closure principle for rationally supported knowledge goes as follows: "If *S* has rationally grounded knowledge that *p*, and *S* competently deduces from *p* that *q*, thereby forming a belief that *q* on this basis while retaining her rationally grounded knowledge that *p*, then *S* has rationally grounded knowledge that *q*" (Pritchard *forthcoming*, Ch.1, p.13). It's not clear why we would find this principle any less plausible than the closure principle for knowledge (CK3) that I formulated in Chapter 1. That is, it seems to satisfy the same intuition that one's knowledge is preserved through competent deductions. Why would evidentially grounded knowledge be any different? But considering this issue further here would take us too far afield.

¹³⁰Recall that the underdetermination principle (UP) says: For all *S*, *p*, *q*, if *S*'s evidence for believing *p* does not favour *p* over some hypothesis *q* which *S* knows to be incompatible with *p*, then *S*'s evidence does not justify *S* in believing *p*.

¹³¹ Pritchard, at least, is deeply engaged in this project (2012; *forthcoming*).

consideration. Again, we might take this as further reason – in addition to the problems I’ve discussed so far – to abandon the project of rejecting the sameness of evidence thesis.

There are of course alternative possible explanations of what motivates (1**).¹³² For example, Cohen (1998) argues that the claim that we don’t have evidential justification to believe that we’re not BIVs is motivated by a principle that says:

(Z) For all S, p, if the truth of p would explain S’s evidence, then S’s evidence does not justify not- p

(Z), like (UP), is meant to tell us why, according to the sceptic, some body of evidence E does not justify S’s belief that some radical sceptical hypothesis is false. Roughly, it is because that evidence would be *explained* by the radical sceptical hypothesis. Cohen’s point is that, if the truth of the BIV hypothesis can explain E, then E cannot justify the claim that I am not a BIV. According to Cohen, premise (1**) depends on the thought that every bit of evidence one might appeal to to justify ~BIV would be explained by the BIV hypothesis. Cohen claims that this principle does not appeal to claims about underdetermination of evidence and that it underwrites a very natural way of understanding the plausibility of premise (1**). “My claim is not that this reasoning will move everyone to skeptical doubt. But I do think it represents a very natural motivation for premise [(1**)]” (Cohen 1998, p.147).¹³³

We might think that if Cohen is correct, and we can motivate (1**) with (Z), then it’s not clear that we can challenge (1**) via a rejection of the sameness of evidence thesis. That is, we might wonder whether, in seeking to establish that our evidence is not the same in the good and bad cases, our anti-sceptical strategy will stand up to the idea that our evidence would be *explained* by the BIV hypothesis in the good case. The obvious point to make here is that Cohen’s principle seems to work as motivation for (1**) only insofar as we assume that our evidence consists in non-factive mental states. After all, the only way it would be possible for a BIV hypothesis to explain the evidence one has for believing, say, that they have two hands, is if that evidence is non-factive. As a BIV, one does not have two hands, so a BIV hypothesis cannot explain one’s having evidence that entails that one has two hands. In other words, it

¹³² Some of the material in this section draws heavily on my Boulton (2013).

¹³³ Cohen appeals to (Z) in the context of challenging Brueckner’s (1994) claim that we *must* appeal to (UP) in motivating (1**). Cohen acknowledges (in a footnote) that the “order of justification” between (Z) and (1**) is unclear to him.

seems the principle is only plausible insofar as we *presuppose* a highest common factor conception of evidence. If so, then the disjunctivist is just as well-positioned to deal with this motivation for (1**) as she is for dealing with motivation that directly appeals to (UP) and the favouring claim.

That said, it is interesting to consider whether Cohen might take the intuitive plausibility of the idea that a BIV hypothesis could explain one's evidence as reason to resist the idea that our evidence could possibly consist in factive mental states. In other words, we might take his principle as *motivation* for a highest common factor conception of evidence. The thought would be that if it seems highly intuitively plausible that a BIV hypothesis would explain all of one's present evidence for believing ordinary empirical propositions, and if the truth of a BIV hypothesis would entail that one's evidence must be non-factive, then one's evidence must (always) consist in non-factive mental states.

I'm not sure how a disjunctivist would want to deal with this claim. For one thing, it seems open to the disjunctivist to insist again that the only reason it seems intuitively plausible that a BIV hypothesis would explain one's evidence is if we are already taken in by the highest common factor conception of evidence. Indeed, they'll note that what's intuitive here is simply the idea that all of one's present *phenomenal experiences* would be explained by a BIV hypothesis, not one's *evidence*. But again, this may be a sort of stand-off that is similar to the stand-off that Pritchard seemed to be in when it came to dealing with the initial highest common factor argument I presented in Chapter 3. In particular, that stand-off had to do with whether someone who makes the move from

(P2) The 'good' and 'bad' cases are phenomenologically indistinguishable

to the conclusion that

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

owes the disjunctivist a justification for why this is legitimate move, or whether the disjunctivist owes the proponent of such a move a justification for why it is not. And again, I noted that that the proponent of the move from (P2) to (C1) *does* have a kind of justification for why it is a legitimate move to make in the form of an explanatory challenge. The idea was

that the best explanation of the indistinguishability of one's evidence in the good and bad cases is that one's evidence *is* the same. In light of that challenge, I claimed that the way for the disjunctivist to approach this stand-off—if it is one—is by endorsing some form of evidential externalism. It seems likely that something similar might apply when it comes to a disjunctivist response to any claim that Cohen might make about the intuitive plausibility of his principle (Z). That said, I think it's worth digressing for a moment and considering whether Cohen's principle is really all that plausible as motivation for (1**), independently of the issue of whether our evidence can consist in factive mental states.

For one thing, it is not immediately obvious that sceptical hypotheses are equally good explanations of E as other anti-sceptical hypotheses, even when E is construed non-factively. For example, the BIV scenario does not immediately come across as a satisfactory competitor with what we can call the “real world hypothesis”. Of course, it is difficult to say why. But one possible way of spelling this out is due to Vogel (1990). Vogel claims that on the real world hypothesis the spatial properties of objects already do real explanatory work (Vogel 1990, p.664). “The fact that something is spherical explains why it behaves like a sphere (in its interactions with us and with other things). If something that is *not* a sphere behaves like one, this will call for a more extended explanation” (Vogel 1990, p.664). Vogel's point is that spatial behaviour comes “for free”, so to speak, with real spheres but it must be supplied by additional BIV programming when the ‘spheres’ in question have no spatial properties at all.¹³⁴

To be sure, insofar as it has been offered as an anti-sceptical strategy, Vogel's argument has attracted criticism (Fumerton 2005). But regardless of whether or not this line of thought is successful as an anti-sceptical strategy, it does articulate the intuition that the BIV hypothesis has shortcomings as an explanation of our evidence (even construed non-factively). That is, something along the lines articulated by Vogel gives expression to an inchoate, intuitive thought about the explanatory shortcomings of the BIV hypothesis. In this respect, even if Vogel's argument can ultimately be challenged as an anti-sceptical strategy, it provides a useful

¹³⁴ Here is a very brief illustration of the point. It is a necessary truth about physical objects that two objects cannot occupy the same location at the same time. “We do not need any empirical law or regularity to explain this” (Vogel 1990, p.664). Because the BIV scenario posits objects with mere pseudo-locations, some explicit principle will be required to explain why no two objects can have the same pseudo-location. In other words, an additional regularity is required to explain why objects that are not genuinely spatial behave as though they are. According to Vogel (given certain assumptions about goodness of explanation—presumably these are assumptions of a “unificationist” variety (Kitcher (1981)) these are grounds for taking the real world hypothesis to be the better explanation.

way of bringing out the claim that (Z) is at best a sort of motivation for (1**) that puts the sceptic on less than plausible ground at the outset. If one appeals to the claim that the BIV hypothesis would *explain* one's evidence in order to motivate the claim that one does not have evidential justification to believe that the BIV hypothesis doesn't obtain, it seems both natural and epistemically relevant to wonder whether the BIV scenario is a very *good* explanation of one's experiences. The explanatory goodness of the BIV hypothesis may or may not be a decisive concern—or even, in the final analysis, anti-sceptically relevant—but it is an issue, nevertheless, that makes for implausible ground upon which to *motivate* or in other words explain the intuitive plausibility of (1**).

I've been focusing on Cohen because I think he presents a particularly elegant and interesting alternative to the idea that (UP) and the favouring claim motivate (1**). Of course there are other arguments to consider concerning what motivates (1**). Other discussions on this issue include Wright (2004), Greco (2007), Weatherson (2007), and Briesen (2010). Briesen and Weatherson both think that an argument which Weatherson calls the “exhaustive argument” best supports (1**). The argument essentially proceeds as follows: S is not justified in believing that they are not a BIV by way of empirical evidence; nor is S justified in believing that they are not a BIV by way of non-empirical evidence; since all evidence is either empirical or non-empirical, S is not justified in believing that they are not a BIV, and thus does not know that they are not a BIV. I don't want to place too much weight on disputing Briesen's claim that this argument “does not refer to the underdetermination principle” (Briesen 2010, p.9). I will note, however, that a natural question to ask Briesen (and Weatherson) is: what motivates the claim that S is not justified by way of empirical evidence in believing that they are not a BIV? Is this to be understood as primitively plausible? Or is there some further explanation for why it is that S isn't empirically justified in her beliefs about not being a BIV. A good candidate is of course the plausibility of claims about underdetermination of evidence.

In a recent discussion, Dodd argues that “the closure argument is better than the underdetermination argument” (Dodd 2012). Dodd has a particular version of the closure argument in mind which appeals precisely to (1**). What he means by the foregoing claim is that his closure argument is harder to respond to. Dodd (following Cohen) distinguishes between the “refutability” and “cogency” of an argument.

That being said, there are different ways to evaluate the relative strengths of different arguments for a single conclusion. Cohen (1998, p. 156) distinguishes between *the refutability* of an argument and *the cogency* of an argument. One argument may be less refutable than another argument because it relies on weaker premises. But that doesn't mean it's more cogent—i.e., that it makes a better case for the conclusion (Dodd 2012, p. 343).

He notes that he hasn't said anything about the relative cogency of the closure and underdetermination arguments:

I've argued that the Closure Argument relies on weaker premises. In Cohen's terminology, it's better in the sense of being *less refutable*. But it doesn't follow that the Closure Argument is also *more cogent* than the Underdetermination Argument. Why think there's anything to [(1**)]—the claim *that S's evidence doesn't provide S with propositional justification for believing not-sk* [the proposition that they are a BIV]? Maybe we find this premise plausible only insofar as we're thinking that S's evidence doesn't discriminate between *h* [the proposition that they have hands] and *sk*. S's evidence supports *h* no better than it supports *sk*. If it's thoughts like these that underlie our finding [(1**)] plausible, then there's another sense in which the Closure Argument is no better than the Underdetermination Argument. That would be because we find the key premise of the Closure Argument compelling only to the extent to which we find the key premise of the Underdetermination Argument compelling (Dodd 2012, p.343).

The basic idea is that, for all Dodd is concerned, the *reason why* we find premise (1**) of the closure argument plausible is *because* we find the claims about underdetermination in the underdetermination argument plausible. So Dodd's discussion does not present any reason to think that a disjunctivist strategy that proceeds by challenging the claims about underdetermination in the underdetermination argument does not have scope for dealing with the reformulated closure argument. Thus, we can see that a variety of recent discussions of the motivation for (1**) place it within the scope of a rejection of the sameness of evidence thesis. As such, this is good reason to think that a disjunctivist strategy has at least as much scope as any other strategy I've considered in this thesis for dealing with scepticism about evidential justification generally (at least when it comes to evidential justification concerning our ordinary beliefs about the external world). If the argument in Chapters 1 and 2 is on the right track, the disjunctivist plausibly has *more* scope for dealing with scepticism about evidential justification than any of the other strategies I've considered in this thesis. The question, of course, is whether such a strategy is ultimately plausible. But the foregoing is motivation for taking the issue I've sharpened when it comes to rejecting the sameness of evidence thesis seriously, as opposed to abandoning the project of rejecting the sameness of evidence thesis.

5.6 Conclusion

I started this chapter by looking at the options when it comes to accounting for the natural reaction and rejecting the sameness of evidence thesis. Williamson's diagnosis of the natural reaction suggests that the only way of accounting for it is by rejecting the claims that I've said it amounts to—namely the (A) principle and the condition of accessibility on attributability. His diagnosis came down to the idea that an internalist condition of accessibility on our evidence amounts to endorsing accessibilist supervenience internalism. I claimed that this is tantamount to the idea that our evidence is the same in the good and bad cases. But I appealed to the work of Gibbons and Greco to show that there are theoretical resources available for arguing that an internalist condition of accessibility does not entail supervenience internalism. As such, I claimed that—at least insofar as these considerations are concerned—it would be hasty to conclude that we only have one option when it comes to accounting for the natural reaction and rejecting the sameness of evidence thesis. That said, I then noted that we are no further along in understanding how such a view would go, or what the best way to reject the sameness of evidence thesis is. In light of the challenges this strategy faces, I concluded the chapter by reconsidering the potential scope of disjunctivist strategies in responding to scepticism about evidential justification generally. I concluded that the prospects on this score are good.

Conclusion

The central claims of this thesis can be put as follows.

- 1) The best (or at least a particularly promising) way of responding to the most interesting (or at least a particularly interesting) argument for external world scepticism involves rejecting the sameness of evidence thesis.
- 2) There are two broad ways of rejecting the sameness of evidence thesis. One way holds on to a condition of accessibility on our evidence; the other way rejects that condition.
- 3) Holding on to the accessibility condition makes it very difficult to meet the explanatory challenge (although we can endorse a weak form of the accessibility condition – but I argued that this seems to undermine the motivation for holding on to an accessibility condition).
- 4) Rejecting the accessibility condition generates the natural reaction. That is, granting that there is a normative requirement such that agents ought to proportion their beliefs to the evidence (which, again, is something I've assumed in this thesis – along with the main parties to the debate I'm engaging in), rejecting the accessibility condition implies that agents can be normatively required to \emptyset when they are not in a position to know that they are normatively required to \emptyset . The natural reaction expresses the idea that this is impossible. I filled out the natural reaction in terms of the (A) principle:

(A): S ought to \emptyset only if, if S doesn't \emptyset , then S is attributable for not \emptyset -ing

and an accessibility condition on attributability, which said:

Necessarily, an agent is attributable for \emptyset -ing only if they are in a position to know that they are \emptyset -ing (and they have not culpably failed to put themselves in a position to know that they are \emptyset -ing).

These two claims together (along with the truth of the evidential norm) imply that a view about evidence which allows that we are not always in a position to know when the triggering conditions in our normative requirements obtain must be mistaken.

- 5) Any anti-sceptical strategy that rejects the sameness of evidence thesis must account for the natural reaction. We can do so in two broad ways: a) by holding on to the accessibility condition; or b) by showing why either (A) or the accessibility condition on attributability or (or both) is mistaken and explain away their intuitive plausibility.
- 6) Both a) and b) are live options, but I've left it an open question what the best way to go is.

There is still of course much to be discussed about the plausibility of the central ideas I've sketched out above. Given the seeming impossibility of meeting the explanatory challenge while endorsing an accessibility condition on our evidence, I am inclined to think that rejecting (A) or the accessibility condition on attributability (or both) is where the most promising action is. Indeed, my aim in this thesis has not been to show that there is no way out of this issue. It has been to sharpen the issue and explore some related issues in the process. The fact that I've been focusing so heavily on the natural reaction should not be taken to imply that I have no sympathy or hope for evidential externalism. Rather, my discussion of the challenges for rejecting the sameness of evidence thesis has been guided by the idea that it's important to understand what we are doing when we approach a rejection of the sameness of evidence thesis in one way or another. I think doing this via evidential externalism comes with underappreciated difficulties, and I've spelled out what I take these to be.

Of course, there are many other potential ways of responding to scepticism besides rejecting the sameness of evidence thesis. I've explored a good number of these in this thesis. I found that close examination of the issues that these strategies give rise to suggests that rejecting the sameness of evidence thesis is a particularly promising way of responding to external world scepticism. But it is worth reviewing the results of that discussion.

I started off with strategies for dealing with the closure-based argument—rejecting premise (1), rejecting premise (2), and semantic contextualism. The overarching dialectic that emerged when examining those strategies was the following. First, when considering rejecting premise (1), there are of course well-known difficulties with rejecting the closure principle. But more interestingly when this strategy is carefully examined it seems to suggest that a more natural approach to closure-based scepticism is to challenge the second premise (this idea emerged when modifying Nozick’s modal conditions on knowledge to basis-relative formulations). But rejecting premise (2) requires explaining how we can know the falsity of sceptical hypotheses. An important early example of such an explanation is Sosa’s safety-based approach. The denial of (2) strikes many as simply too implausible, or at the very least raises important questions about externalist responses to scepticism. DeRose takes the issues that arise when denying either premise of the closure-based argument to motivate contextualism, which seeks neither to deny the first nor second premise outright. But I argued that going contextualist (at least in the paradigmatic way represented by DeRose) seems to require presupposing that the sceptical problem turns on a high-standards conception of knowledge. Of course, DeRose also has resources in the form of externalist conditions on knowledge at this point. But, again, we saw that there are important issues that arise with externalist responses to external world scepticism. I argued that these issues turn on considerations about what an adequate response to external world scepticism should hope to achieve. I argued that one condition of adequacy in responding to external world scepticism is being able to account for radical forms of scepticism which challenge not only that our beliefs enjoy the epistemic status of knowledge (however that status is conceived) but also that our ordinary empirical beliefs are justified, or that we are reasonable in holding them.

I introduced the underdetermination problem as an example of a sceptical argument that challenges the idea that we are justified or reasonable in holding ordinary empirical beliefs. But in Chapter Two I examined two anti-sceptical strategies that deal with a different form of scepticism about evidential justification. In particular, these strategies engage with a sceptical argument that makes central use of a claim that we need antecedent justification to believe that we’re not BIVs in order to have justification to believe ordinary empirical propositions. I argued that Pryor’s and Wright’s views are caught in a dialectical circle. I took the difficulties faced by Pryor and Wright, along with further considerations about the relationship between the sceptical argument they engage with and the underdetermination argument, to motivate

engaging directly with the underdetermination argument, and in particular the sameness of evidence thesis.

All of that said, the issue I've tried to sharpen in this thesis need not be understood as an issue that must be resolved insofar as we want to resolve the problem of the external world. Rather, it may simply be an issue that we need to resolve if we want to approach that problem (understood as an underdetermination problem) successfully via a rejection of the sameness of evidence thesis. Still, again, the first two chapters of this thesis contain a detailed discussion which suggests that approaching external world scepticism in this way is a particularly interesting and promising approach. Moreover, at the end of the last chapter I provided a further defence of the claim that rejecting the sameness of evidence thesis has adequate scope when it comes to responding to scepticism about evidential justification generally.

The problem of our knowledge of the external world is a particularly resilient problem. Part of this may be due to the fact that there's arguably not really *one* problem of our knowledge of the external world. As Williamson puts it, "There is no silver bullet against scepticism. It is robust, in part because sceptical arguments form a complex terrain across which sceptics show great facility in subtly shifting their ground as they come under pressure at point after point" (Williamson 2009, p.357). An advantage of bringing the underdetermination argument to centre stage is that the plausibility of its central commitment—the sameness of evidence thesis—is particularly susceptible to scrutiny. But challenging the sameness of evidence thesis requires showing that the view about evidence we end up with does not itself lead to unacceptable consequences. A promising way of challenging the sameness of evidence thesis—evidential externalism—does lead to counterintuitive consequences when it comes to the issue about access and normative requirements. But I said in the introduction that the question of what sort of access we have to our normative requirements is the focus of an increasingly sophisticated discussion in contemporary epistemology. One advantage of the foregoing discussion is that it brings the problem of external world scepticism directly within the scope of that debate. This is a promising result for future prospects of responding to external world scepticism.

Bibliography

- Alston, W. 1985. Concepts of epistemic justification. *The Monist*, 68:57-89.
- _____ 1989. *Epistemic Justification*. Cornell University Press, Ithaca.
- Annis, D. A contextualist theory of epistemic justification. *American Philosophical Quarterly*, 15:213-219.
- Becker, K. 2012. Methods and how to individuate them. In Becker, K. and Black, T. eds., *The Sensitivity Principle in Epistemology*, Cambridge University Press, Cambridge.
- Bergmann, M. 2004. Epistemic circularity: malignant and benign, *Philosophy and Phenomenological Research*, 69:709–727.
- _____ 2006. *Justification without Awareness*. Oxford University Press, New York.
- _____ 2008. Externalist responses to skepticism. In Greco, J. ed., *The Oxford Handbook on Skepticism*, Oxford University Press, Oxford.
- Berker, S. 2013. The rejection of epistemic consequentialism. *Philosophical Issues*, 23:363-87.
- Black, T. 2002. A Moorean response to brain-in-a-vat skepticism. *Australasian Journal of Philosophy*, 80: 148-163.
- Boult, C. 2013. Epistemic principles and sceptical arguments: closure and underdetermination. *Philosophia*, 41:1125-1133.
- Briesen, J. 2010. Reconsidering closure, underdetermination, and infallibilism. *Grazer Philosophische Studien*, 80: 221-234.
- Brogaard, B. 2010. Disjunctivism. *Oxford Annotated Bibliographies Online*.
- Brown, J. 2013. Infallibilism, evidence and pragmatics. *Analysis* 73:626-635.
- Brueckner, A. 1994. The structure of the skeptical argument. *Philosophy and Phenomenological Research*, 54:827-835.
- _____ 2010. $\sim K \sim SK$. In *Essays on Skepticism*, Oxford University Press, Oxford.
- Byrne, A. and Logue, H. Introduction. In Byrne, A. and Logue, H. eds., *Disjunctivism: Contemporary Readings*. MIT Press.
- Chrisman, M. 2007. From epistemic contextualism to epistemic expressivism. *Philosophical Studies*, 135:225-254.
- _____ 2012. 'Ought' and control. *Australasian Journal of Philosophy*, 90:433-451.
- _____ 2012b. The normative evaluation of belief and the aspectual classification of belief and knowledge attributions. *Journal of Philosophy*, 109:588–612.

- _____. *forthcoming. The Meaning of 'Ought': Beyond Expressivism and Descriptivism in Metaethics.* Oxford University Press, Oxford.
- Chuard, P. and Southwood, N. 2009. Epistemic norms without voluntary control. *Nous*, 43: 599-632.
- Cohen, S. 1998. Two kinds of skeptical argument. *Philosophy and Phenomenological Research*, 58: 143-159.
- _____. 2002. Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research*, 65:309-329.
- _____. 2010. Luminosity, reliability, and the sorites. *Philosophy and Phenomenological Research*, 81:718-730.
- Coliva, A. 2010. *Moore and Wittgenstein: Scepticism, Certainty, and Common Sense.* Palgrave Macmillan.
- Conee, E. and Feldman, R. 2004. *Evidentialism.* Oxford University Press, Oxford.
- Comesaña, J. 2013. On an argument against Immediate Justification. In Sosa, E., Steup, M. and Turri, J. eds., *Contemporary Debates in Epistemology, 2nd Edition,* Wiley-Blackwell.
- Davies, M. 2004. Epistemic entitlement, warrant transmission and easy knowledge. *Aristotelian Society Supplementary Volume*, 78:167-245.
- DeRose, K. 1995. Solving the skeptical problem. *Philosophical Review*, 104: 1-52.
- _____. 2006, 'Bamboozled by our own words': semantic blindness and some arguments against contextualism. *Philosophy and Phenomenological Research*, 73:316-338.
- Dodd, D. 2012. Evidentialism and skeptical arguments. *Synthese*, 189:337-352.
- Dorsch, F. 2008. The diversity of disjunctivism. In Haddock, A. and Macpherson, F. eds., *Disjunctivism: Perception, Action, Knowledge.* Oxford University Press, Oxford.
- Dretske, F. 1970. Epistemic operators. *The Journal of Philosophy*, 67: 367-397.
- Feldman, R. 1999. Contextualism and skepticism. *Philosophical Perspectives*, 13:91-114.
- Fish, W. 2009. Disjunctivism. *Internet Encyclopedia of Philosophy.*
- Foley, R. 1987. *The Theory of Epistemic Rationality.* Harvard University Press, Cambridge.
- Foot, P. 1972. Morality as a system of hypothetical imperatives. *Philosophical Review*, 81:305-316.
- Fumerton, R. 1995. *Metaepistemology and Skepticism.* Rowman and Littlefield.
- _____. 2000. Williamson on skepticism and evidence. *Philosophy and Phenomenological Research*, 60:629-635.

- _____ 2005. The challenge of refuting skepticism. In Steup, M. and Sosa, E. eds., *Contemporary Debates in Epistemology*, Wiley-Blackwell.
- Gibbons, J. 2006. Access externalism. *Mind*, 115: 19-39.
- _____ 2010. Reply to Moon. *Mind*, 119:153-158.
- _____ 2013. *The Norm of Belief*. Oxford University Press, Oxford.
- Ginet, C. 1975. *Knowledge, Perception, and Memory*. Springer.
- Goldman, A. 1979. What is Justified Belief? In Pappas G. ed., *Justification and Knowledge*, D. Reidel, Dordrecht.
- _____ 1986. *Epistemology and Cognition*. Harvard University Press, Cambridge, MA.
- _____ 2009. Williamson on knowledge and evidence. In Greenough, P., Pritchard, D. and Williamson, T. eds., Oxford University Press, Oxford.
- Greco, D. 2013. Cognitive mobile homes. *unpublished*.
- _____ *forthcoming*. Could KK be OK? *Journal of Philosophy*.
- Greco, J. 1999. Agent reliabilism. *Philosophical Perspectives*, 13: 273-96.
- _____ 2000. *Putting Skeptics In Their Place*. Cambridge University Press, New York.
- _____ 2004. Externalism and skepticism. In Schantz, R. ed., *The Externalist Challenge*, De Gruyter.
- _____ 2005. Justification is not internal. In Steup, M. and Sosa, E. eds., *Contemporary Debates in Epistemology*, Wiley-Blackwell.
- _____ 2007. External world skepticism. *Philosophy Compass*, 2:625–649.
- _____ 2010. *Achieving Knowledge*. Cambridge University Press, Cambridge.
- Haddock, A. and Macpherson, F. 2008. Introduction: varieties of disjunctivism. In Haddock, A. and Macpherson, F. eds., *Disjunctivism: Perception, Action, Knowledge*. Oxford University Press, Oxford.
- Hawthorne, J. 2005. The case for closure. The challenge of refuting skepticism. In Steup, M. and Sosa, E. eds., *Contemporary Debates in Epistemology*, Wiley-Blackwell.
- Hofweber, T. 1999. Contextualism and the meaning-intention problem. In Korta, K., Sosa, E. and Arrazola, X. eds., *Cognition, Agency and Rationality*, Kluwer, London.
- Joyce, J. 2004. Williamson on evidence and knowledge. *Philosophical Books*, 45:296-305.
- Jenkins, C.S. 2007. Entitlement and rationality. *Synthese*, 157:25-45.
- Kallestrup, J. 2012. Bootstrap and rollback: generalizing epistemic circularity. *Synthese*, 189:395-413.

- Kelly, E. *Forthcoming*. What is an Excuse? In Tognazzini, N., and Coates, J. eds. *Blame: New Essays*, Oxford University Press, Oxford.
- Kelly, T. 2003. Epistemic rationality as instrumental rationality: a critique. *Philosophy and Phenomenological Research*, 66:612–640.
- Kitcher, P. 1981. Explanatory Unification. *Philosophy of Science*, 48:507-531.
- Kornblith, H. 2000. The Contextualist evasion of epistemology. *Noûs*, 34:24-32.
- _____ 2001. Epistemic obligation and the possibility of internalism. In Fairweather, A. and Zagzebski, L. eds., *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, Oxford University Press, Oxford.
- Kratzer, A. 1977. What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, 1:337-55.
- _____ 1981. The notional category of modality. In Eikmeyer, H., and Reiser, H. eds., *Words, Worlds, and Contexts*, de Gruyter, Berlin.
- _____ 1991. Conditionals. In Stechow, A., and Wunderlich, D. eds., *Semantik: ein internationales Handbuch der zeitgenössischen Forschung*, De Gruyter, Berlin and New York.
- Leite, A. 2006. Epistemic instrumentalism and reasons for belief: A reply to Tom Kelly's "epistemic rationality as instrumental rationality: A critique". *Philosophy and Phenomenological Research*, 75:456–464.
- Lewis, D. 1996. Elusive knowledge. *Australasian Journal of Philosophy*, 74:549–567.
- Littlejohn, C. (2011). Evidence and knowledge. *Erkenntnis*, 74:241-262.
- _____ (2011b). Evidence and armchair access. *Synthese*, 179:479-500.
- _____ *unpublished*. Infallibilism, evidence, and infelicity.
- Owens, D. 2000. *Reason Without Freedom*. Routledge, New York.
- _____ 2004. Williamson on scepticism and rationality. *Philosophical Books*, 45:306-312.
- Macarthur, D. 2003. McDowell, scepticism, and the 'veil of perception'. *Australasian Journal of Philosophy*, 81:175-190.
- McDowell, J. 1995. Knowledge and the internal. *Philosophy and Phenomenological Research*, 55: 877-93.
- McGinn, M. 1989. *Sense and Certainty*, Basil Blackwell, Oxford.
- McHugh, C. 2012. Epistemic deontology and voluntariness. *Erkenntnis*, 77:65–94.
- McLaughlin, B.P. 1996. Lewis on what distinguishes perception from hallucination. In Akins, K. ed., *Perception*, Oxford University Press, New York.
- Moon, A. 2010. Gibbons on epistemic internalism. *Mind*, 119:143-151.

- Moore, G.E. 1939. Proof of an external world. *Proceedings of the British Academy* 25.
- _____ 1925. A defence of common sense. In Muirhead, J. H. ed., *Contemporary British Philosophy* (2nd series), Allen and Unwin, London.
- Nagel, T. 1979. *Mortal Questions*. Cambridge University Press, New York.
- Neta, R. 2003. Contextualism and the problem of the external world. *Philosophy and Phenomenological Research*, 661: 1-31.
- Neta, R. and Pritchard, D. 2007. McDowell and the new evil genius. *Philosophy and Phenomenological Research*, 74:381–396.
- Nozick, R. 1981. *Philosophical Explanations*. Harvard University Press.
- Pedersen, N. J. 2009. Entitlement, value and rationality. *Synthese*, 171: 443-57.
- Plantinga, A. 1993. *Warrant: The Current Debate*. Oxford University Press, Oxford.
- Pollock, J. 1986. *Contemporary Theories of Knowledge*. Rowman and Littlefield, Totowa, N.J.
- Pritchard, D. 2005. The structure of sceptical arguments. *Philosophical Quarterly*, 55:37-52.
- _____ 2005b. Wittgenstein's *On Certainty* and contemporary anti-scepticism. In Moyal-Sharrock, D., and Brenner, W. H. eds., *Readings of Wittgenstein's On Certainty*. Palgrave Macmillan, London.
- _____ 2005c. *Epistemic Luck*. Clarendon Press.
- _____ 2006. McDowellian neo-Mooreanism. In Macpherson, F., and Haddock, A. eds., *Disjunctivism: Perception, Action, Knowledge*, Oxford University Press, Oxford.
- _____ 2007. How to be a neo-Moorean. In Goldberg, S. ed., *Internalism and Externalism in Semantics and Epistemology*, Oxford University Press, Oxford.
- _____ 2008. Sensitivity, safety, and anti-luck epistemology. In Greco, J. ed., *The Oxford Handbook of Skepticism*, Oxford University Press, Oxford.
- _____ 2012. *Epistemological Disjunctivism*. Oxford University Press, Oxford.
- _____ 2012b. Anti-luck virtue epistemology. *Journal of Philosophy*, 109:247-79.
- _____ forthcoming. *Epistemic Angst*, Oxford University Press, Oxford.
- Pryor, J. 2013. There is immediate justification. In Sosa, E., Steup, M. and Turri, J. eds., *Contemporary Debates in Epistemology, 2nd Edition*, Wiley-Blackwell.
- _____ 2004. What's wrong with Moore's argument? *Philosophical Issues*, 14:349–378.
- _____ 2001. Highlights of recent epistemology. *British Journal for the Philosophy of Science*, 52:95-124.
- _____ 2000. The skeptic and the dogmatist. *Nous*, 34: 517–549.

- Reichenbach, H. 1938. *Experience and Prediction*. University of Chicago Press, Chicago.
- Rysiew, P. 2001. The context sensitivity of knowledge attributions. *Noûs*, 35:477–514.
- Scanlon, T.M. 1998. *What We Owe to Each Other*. Harvard University Press, Cambridge.
- Schiffer, S. 1996. Contextualist solutions to scepticism. *Proceedings of the Aristotelian Society*, 96:317-333.
- _____ 2009. Evidence=knowledge: Williamson's solution to skepticism. In Greenough, P., Pritchard, D. and Williamson, T. eds., Oxford University Press, Oxford.
- Shoemaker, D. 2011. Attributability, answerability, and accountability: toward a wider theory of moral responsibility. *Ethics*, 121:602-632.
- Schroeder, M. 2011. Oughts, agents, and actions. *Philosophical Review*, 120:1-41.
- Silins, N. 2005. Deception and evidence. *Philosophical Perspectives*, 19:375–404.
- Smith, A. 2008. Control, responsibility, and moral assessment. *Philosophical Studies*, 138:367-392.
- _____ 2012. Attributability, answerability, and accountability: in defense of a unified account. *Ethics*, 122:575-589.
- Smithies, D. 2012. Mentalism and epistemic transparency. *Australasian Journal of Philosophy*, 90:723-741.
- _____ 2013. Review of *Epistemological Disjunctivism*. *Notre Dame Philosophical Reviews*.
- Srinivasan, A. unpublished. What's in a norm?
- Sosa, E. 1991. *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge University Press, New York.
- _____ 1999. How to defeat opposition to Moore. *Philosophical Perspectives*, 13:137-49.
- _____ 2007. *A virtue epistemology: apt belief and reflective knowledge, volume 1*. Oxford University Press, Oxford.
- Steglich-Petersen, A. 2010. Luck as an epistemic notion. *Synthese*, 176:361-377.
- _____ 2011. How to be a teleologist about epistemic reasons. In Steglich-Petersen, A. and Reisner, A. eds., *Reasons for Belief*. Cambridge University Press.
- Stine, G.C. 1976. Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, 29:249-261.
- Strawson, P. F., 1962. Freedom and resentment. *Proceedings of the British Academy*, 48:1–25.
- Stroud, B. 1984. *The Significance of Philosophical Scepticism*. Clarendon Press, Oxford.
- _____ 1989. Understanding human knowledge in general. In Clay, M. & Lehrer, K. eds., *Knowledge and Skepticism*, Westview Press.

- Talbert, M. 2008. Blame and responsiveness to moral reasons: are psychopaths blameworthy? *Pacific Philosophical Quarterly*, 89:516–535.
- Vogel, J. 1990. Cartesian skepticism and inference to the best explanation. *Journal of Philosophy*, 87:658-666.
- _____ 2008. Epistemic bootstrapping. *Journal of Philosophy*, 105:518-539.
- Weatherson, B. 2007. The Bayesian and the dogmatist. *Proceedings of the Aristotelian Society*, 107:169-185.
- Watson, G. 2004. Two faces of responsibility. In *Agency and Answerability: Selected Essays*, Oxford University Press, Oxford.
- _____ 2011. The trouble with psychopaths. In Wallace, R.J., Kumar, R. and Freeman, S. eds. *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon*, Oxford University Press, Oxford.
- Williams, B. 1976. Moral luck. *Proceedings of the Aristotelian Society*, 50: 115-135.
- _____ 1981. Philosophy. In Finley, M.I. ed., *The Legacy of Greece: A New Appraisal*. Clarendon Press, Oxford.
- Williams, M. 1991. *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*, Princeton University Press, Princeton.
- _____ 2001. *Problems of Knowledge: A Critical Introduction to Epistemology*. Oxford University Press, Oxford.
- _____ 2008. Responsibility and reliability. *Philosophical Papers*, 37:1-26.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford University Press, Oxford.
- _____ 2000. Skepticism and evidence. *Philosophy and Phenomenological Research*, 60:613-628.
- _____ 2009. Replies to critics. *Williamson on Knowledge*. In Greenough, P., Pritchard, D. and Williamson, T. eds., Oxford University Press, Oxford.
- Wittgenstein, L. 1969. *On Certainty*. Anscombe, G. E. M. and von Wright, G.H. eds. Paul, D. and Anscombe, G. E. M. tr., Wiley-Blackwell, Oxford.
- Wright, C. 2000. Cogency and question-begging: Some reflections on McKinsey's paradox and Putnam's proof *Philosophical Issues*, 10:140-63.
- _____ 2002. (Anti-)sceptics simple and subtle: G. E. Moore and John McDowell. *Philosophy and Phenomenological Research*, 65:330-348.
- _____ 2004. Warrant for nothing (and foundations for free)? *Aristotelian Society Supplementary Volume*, 78:167–212.

Zagzebski, L. 1999. What is knowledge? In Greco, J. & Sosa, E. eds., *The Blackwell Guide to Epistemology*, Wiley-Blackwell, Oxford.