

FOUNDATIONALISM AND THE IDEA OF THE EMPIRICAL

CIARA FAIRLEY

**University College London
Submitted for the degree of PhD, May 2007**

ABSTRACT

This thesis is about foundationalism in epistemology. It distinguishes between different forms of foundationalism and defends one particular version of this doctrine. Chapter 1 gives an account of the motivations for foundationalism, including the so-called epistemic regress argument. It criticizes recent accounts of the core doctrines of foundationalism, such as those of Michael Williams and Ernest Sosa, and proposes a different account according to which foundationalism is the view that (a) some of our beliefs must be non-inferentially justified, (b) perception is a source of non-inferential justification, and (c) perception is a basic source of such justification. Chapter 2 gives an account of traditional foundationalism and tries to identify both what is right with it and what is wrong with it. It argues that the basic insight of traditional foundationalism can be detached from some of the other doctrines with which it was associated by the traditional foundationalists. That insight concerns the role of perceptual awareness or acquaintance as a regress-terminating source of epistemic justification. Chapter 3 exploits this idea in defending a more modest form of foundationalism according to which ordinary perceptual beliefs may be foundational. Chapters 3 and 4 focus on two influential arguments against the view that ordinary beliefs about the world around us can be non-inferentially justified by perception. The first argument trades on the alleged fallibility of perceptual justification, the second on its defeasibility. It is shown that neither argument poses a genuine threat to the more modest version of foundationalism that I defend. Chapter 5 compares perception with other sources of non-inferential justification such as memory and testimony. It defends the view that perception is a privileged source of non-inferential justification, even if it isn't the only source of such justification. It also contrasts foundationalism with traditional forms of externalism such as reliabilism and explains why the latter should not be counted as a form of foundationalism.

ACKNOWLEDGMENTS

Thanks to Mike Martin for getting me started on this topic and for his input in the early stages of my research. In the last two years I have been helped by Quassim Cassam whose support, encouragement, and own example made it possible for me to finish this thesis.

I would also like to thank Marcus Giaquinto and Paul Snowdon for useful discussions and feedback on earlier drafts of some of the following chapters.

I gratefully acknowledge the financial support of the Arts and Humanities Research Council and the Royal Institute of Philosophy.

This thesis is dedicated to those who have helped me most; you know who you are.

TABLE OF CONTENTS

PREFACE	5
1. WHAT IS FOUNDATIONALISM?	16
1. Introduction	16
2. The Epistemic Regress Argument	27
3. Conclusion	53
2. THE HISTORICAL FOUNDATIONALISTS	55
1. Introduction	55
2. Historical Foundationalism	57
3. Revisionary Epistemology	68
4. The Idea of the Empirical	81
5. Conclusion	87
3. THE ARGUMENT FROM FALLIBILITY	89
1. Introduction	89
2. Propositional Perception	93
3. Non-Propositional Perception	101
4. Having Reasons	116
5. The Argument from Fallibility	129
6. Conclusion	143
4. THE ARGUMENT FROM DEFEASIBILITY	149
1. Introduction	149
2. Perceptual Defeasibility	150

3.	The Argument from Defeasibility	159
4.	The World Made Manifest	169
5.	Historical Foundationalism Reconsidered	174
6.	Indefeasibility and the Mental	183
7.	Conclusion	185
5. THE PRIMACY OF PERCEPTION		188
1.	Introduction	188
2.	Concept Empiricism	194
3.	Basic Knowledge	196
4.	Generating Knowledge	202
5.	Perception and Explanation	219
6.	Conclusion	230
BIBLIOGRAPHY		233

FOUNDATIONALISM AND THE IDEA OF THE EMPIRICAL PREFACE

When I first encountered foundationalism as a topic in epistemology I had a certain understanding of what it was supposed to be. I thought that foundationalism was roughly the view that not all of our justification could be inferential and that perception is the basic source of justification that is not inferential. So foundationalism, as I understood it, was a doctrine both about the structure of human knowledge or justification, and its sources. This initial impression was strengthened when I read Contemporary Theories of Knowledge by Pollock and Cruz. They claim:

The simple motivation for foundations theories is the psychological observation that we have various ways of sensing the world and that all knowledge comes to us via those senses. The foundationalist takes this to mean that our senses provide us with what are then identified as epistemologically basic beliefs. We arrive at other beliefs by reasoning (construed broadly). Reasoning, it seems, can only justify us in holding a belief if we are already justified in holding the beliefs from which we reason, so reasoning cannot provide an ultimate source of justification. Only perception can do that. We thus acquire the picture of our beliefs forming a kind of pyramid, with the basic beliefs provided by perception forming the foundation, and all other justified beliefs being supported by reasoning that traces back ultimately to the basic beliefs (Pollock & Cruz 1999: 29).

Though Pollock and Cruz ultimately mishandle this insight they remain virtually the only commentators in this area to acknowledge the importance of perception. Their book had a lasting impact on my thinking.

I thought that if I had correctly understood foundationalism it was pretty obviously true. Although some of our beliefs are justified because we have inferred them from other things that we are justified in believing it seemed obvious to me that they cannot all be justified in this way. Some of our beliefs must be non-inferentially justified and perception is the obvious source of such justification. How then could foundationalism – which says just this – fail to be true?

It came as something of a surprise, then, when I discovered that not only is foundationalism not a position that most philosophers think is true. It is a position that most philosophers think is false. That left me puzzled: how could so many philosophers be so critical of a position that seems to get so much right? This thesis is to a large extent a direct product of that puzzlement and a more or less direct attempt to resolve it.

The first thing that I discovered when I starting reading more widely was that other people don't all understand foundationalism in the way that I did. The historical foundationalists – people like C. I. Lewis, Roderick Chisholm, and A. J. Ayer - agreed that our knowledge of the world rests upon a foundation of basic beliefs and that those beliefs are not justified in the way that the rest of our beliefs are justified. But while I took basic beliefs to include ordinary beliefs about objects and events in the world around us, the historical foundationalists took them to be beliefs about our own minds.¹ And

¹ It sounds odd to say the foundational beliefs include 'ordinary beliefs' about the world. Aren't ordinary beliefs supposed to be the beliefs that are supported by more basic beliefs? If the former are basic, what would be an example of a non-basic belief? This problem arises where we assume we can individuate the foundational beliefs in terms of their subject-matter. We will see later this is a mistake: foundational beliefs are not 'about' any particular subject-matter. What distinguishes these beliefs is the source of their justification. Here, all I mean is that the foundational beliefs might include beliefs like this: 'the squirrel is

while I thought that the basic beliefs are justified by perception the historical foundationalists claimed those beliefs are infallible and so either justify themselves or are justified by some sort of special introspective awareness. This seemed a long way from the rather commonsensical doctrine that I had always taken foundationalism to be.

In recent years foundationalism has enjoyed something of a renaissance. Unfortunately, the form of foundationalism that is popular nowadays is a long way from the position that I call foundationalism. Sometimes called ‘formal foundationalism’, the new foundationalism is a bland doctrine that normally involves no more than a commitment to the idea that epistemological properties like justification supervene on non-epistemological ones. Since this makes just about everyone a foundationalist it’s not a position that I felt very excited about defending.

I think that foundationalism is a substantive doctrine though not the very unattractive doctrine the historical foundationalists made it out to be. Foundationalism, as I understand it, has got three basic components. The first is that there must be such a thing as non-inferential justification and there must be because otherwise we face a vicious epistemic regress. This is the least contentious of what I regard as the three basic elements of foundationalism. I think that the so-called epistemic regress argument for foundationalism is a good one and I will explain why in chapter 1. Later in later chapters 3 and 4 I will counter various arguments that are supposed to show that there could not be any such thing as non-inferential justification.

on the fence’ ‘Ross is at the party’ and so on. This is something more traditional foundationalists denied.

The second component of foundationalism is that one of the sources of non-inferential justification is perception. Standard undergraduate textbooks on foundationalism tell us that the basic or foundational beliefs are either self-justified or need no justification. Such beliefs are obviously not justified by perception. Robert Audi, Keith Lehrer, and Jonathan Dancy all think of basic beliefs in this way. So what I am representing as the second essential component of foundationalism is one that very few standard discussions acknowledge.

One reason for this discrepancy is an excessive focus on Descartes and the idea that he is really the paradigm foundationalist. I think that is a mistake. If you want some paradigm foundationalists you should look to people like Locke and Hume – philosophers who on the face of it couldn't have less in common with Descartes. In fact, it's not at all obvious to me that Descartes is a foundationalist. What is important to foundationalism isn't the idea that the foundations of our knowledge are self-justifying, but that they are perceptually justified and Descartes certainly didn't think that.

The idea that perception is a source of non-inferential justification will be the focus of Chapters 2 and 3. When I say that perception is a source of non-inferential justification I take it that the beliefs perception can non-inferentially justify are, or include, beliefs about non-psychological reality, like the belief that the squirrel is on the fence. So I just mean perception in the ordinary sense. In contrast when the historical foundationalists said that 'perception' is a source of non-inferential justification they didn't mean perception in the ordinary sense. What they meant by perception is closer to what we

would call introspection and what Locke called ‘inner perception’. The beliefs that they thought perception can non-inferentially justify are beliefs about psychological reality rather than beliefs about the world around us.

These further claims are neither sensible nor essential to foundationalism. Historical foundationalism therefore represents a perversion of an otherwise sensible thought. This perversion will be the topic of chapter 2. Chapter 3 will explore a sensible version of the sensible thought. I will show how we can hang onto what is right about traditional foundationalism without committing ourselves to its less attractive features.

Many people might agree with me up to this point. They could think, yes – there is such a thing as non-inferential justification and yes - perception is a source of such justification. But there is nothing special about perception; there’s also testimony for example. Foundationalism as I understand it, denies that there is nothing special about perception. It is the view that among our sources of non-inferential justification perception is a basic source. What I mean by this is, roughly, that the other sources could not function as sources of justification at all unless perception is also a source of justification.² Although this view seems to me to be pretty obviously correct it has recently come under attack. I will be responding to some of these attacks in chapter 5.

This is the view I want to defend and these are the places I will be defending it. It should

² Notice that is weaker than the claim Pollock attributes to foundationalism in the passage quoted. He claims perception is the *only* source of justification other than reasoning. I claim merely that it is ‘a’ source of justification distinct from reasoning and, in some sense yet to be explained, a basic source. Still, it is not the only source.

be clear that the position I want to defend is ultimately very different from traditional foundationalism. Nonetheless, it bears important similarities to that view. An obvious question, therefore, is whether the standard objections to that view also apply to my view. That depends on what the standard objections are. One is that there are no self-justifying beliefs. Clearly, this is not a good objection to my view since on my view foundationalism is not committed to thinking that there are any such beliefs. The basic beliefs are justified by perception and beliefs that are justified by perception are obviously not self-justified.

A different objection is that foundationalism is committed to the ‘myth of the given’. What is that? If the ‘myth’ just involves thinking that perception is a source of non-discursive justification, then it is not a myth. It is true. But if what people who press this objection are really saying is it is only in the context of certain social practices that perception is a source of justification, then I am not denying that. I am not claiming that perception is an autonomous source of justification in *that* sense.

I think that perception can non-inferentially justify beliefs about non-psychological reality. These beliefs do not draw their justification from other justified beliefs. The justification that perception provides is belief-independent in this sense. It might depend on beliefs in some other sense.

This is a distinction we should draw even in cases in which one’s justification does derive from other beliefs. So people who press this objection are either they are denying

something that is obviously true; or they are insisting on a point that I don't need to dispute.

So that's it. In short I am really just going back to the very simple idea with which we began and which got me thinking about foundationalism in the first place. I think that foundationalism is still a live option in epistemology. One thing that can make this difficult to see is a commitment to externalism. Unlike many externalists I understand perception to involve conscious access to the world. So insisting on the centrality of perception just means insisting on the centrality of consciousness. If we only interested in reliable belief forming mechanisms there is no reason why we should take consciously so seriously.

This approach has interesting parallels with John Campbell's view of though. His view, very roughly, is that it is consciousness of the world that makes it possible for us to think about it. My view, very roughly, is that it is consciousness of the world that makes it possible for us to know about it. What could possibly be more obvious?

FOUNDATIONALISM AND THE IDEA OF THE EMPIRICAL

CHAPTER 1: WHAT IS FOUNDATIONALISM?

1. Introduction

This thesis will defend a form of foundationalism in epistemology. I think that foundationalism is an overwhelmingly natural view about the structure and sources of epistemic justification – that is, a view about what it is to have reasons for our beliefs about the world in such a way that these beliefs can constitute knowledge. Not everyone agrees. In the literature a tradition has grown up according to which foundationalism is a much less attractive doctrine than I will claim. So one thing that I am doing in this thesis is taking on a certain tradition of interpretation.

The tradition that I am opposing is long standing and still has very much the status of orthodoxy. It takes foundationalism to be a doctrine involving a commitment to certain characteristic claims. One central component of that doctrine is the idea that there are ‘epistemically basic beliefs’. By epistemically basic I mean beliefs that are not inferentially justified. Here is an example of an inferentially justified belief: I believe that England can no longer win the Ashes and my justification for that belief derives from my justification for believing they have performed poorly in the past three tests. This is an example of an inferentially justified belief. So when I talk about beliefs that are not inferentially justified - or beliefs that are ‘non-inferentially’ justified, as I will often say - I just mean beliefs that are not justified in *that* way. This is not a positive account of what

does justify these beliefs and I am not going to give you a positive account at this stage. All that is important is that such beliefs are not inferentially justified. According to the tradition that I am opposing, the foundationalist is at the very least someone who thinks that there are such beliefs.

This is one thing that any foundationalist has got to think, but it is not all she needs to think according to this tradition. Michael Williams makes this point well.

Basic beliefs are the stock in trade of epistemological foundationalists... According to foundationalism all justification starts with basic beliefs and flows “upward” from them. However, there is more to foundationalism than this. If foundationalism were no more than a structural-descriptive account of everyday knowledge it is hard to see why everyone would not be a foundationalist. Didn’t we just agree that there are lots of things we “just know”? So aren’t we all foundationalists? The answer is “No”. The theoretical commitments of traditional foundationalists are extensive (Williams 2005: 203)

So what are these further commitments? Williams goes on to mention the following four:

(1) Traditional foundationalism is *substantive*, rather than merely formal. According to substantive foundationalism, the class of basic beliefs is *theoretically tractable*. In particular, there are non-trivially specifiable kinds of beliefs, individuated by broad aspects of their content, that are fitted to play the role of terminating points for chains of justification. The distinction between basic and non-basic beliefs is thus ontological rather than merely methodological. (2)

Traditional foundationalism is *strong*. Basic beliefs, or terminating judgements are *indubitable* or (a slightly weaker notion) *incorrigible*. Basic beliefs are always basic knowledge. (3) Traditional foundationalism is *atomistic*. Basic beliefs provide absolute terminating points for justificatory chains. To do so, basic beliefs must be independently both *epistemically* and *semantically* of other justified beliefs. Since basic beliefs constitute *encapsulated* items of knowledge, there is no objection in principle to the idea of a first justified belief. (4) Traditional foundationalism is *radically internalist*. The justification-making factors for beliefs, basic and otherwise, are all open to view, and perhaps even actual objects of awareness. At the base level, when I know that P, I am always in a position to know that I know that P, and perhaps even always *do* know that I know that P (Williams 2005: 203-4).

In a similar vein, Ernest Sosa claims that:

Classical foundationalism in epistemology is the view that:

- (i) every infallible, indubitable belief is justified
- (ii) every belief deductively inferred from justified beliefs is itself justified,
and
- (iii) every belief that is justified is so in virtue of (i) or (ii) above (Sosa 2000: 14).

Williams and Sosa are not alone. The view that they describe, according to which there is a layer of epistemically basic beliefs, distinguished in terms of their content and supporting everything else that we know, is how most people in the literature understand

foundationalism. I have said that I want to defend foundationalism, but I do not want to defend the very unattractive position that Williams and Sosa describe. As others have pointed out it is doubtful that we have so many indubitable, incorrigible, beliefs or that these beliefs would provide enough of a foundation for the rest of what we know. Most ordinary beliefs about the world – like the belief that there is squirrel on the fence or the belief that Ross is on the mend - can be mistaken, they can be doubted, and they can be rationally revised. Such beliefs therefore cannot provide the foundations we are alleged to need. It is doubtful any can. Foundationalism is therefore a form of scepticism and that is not a very attractive position to want to defend.

I am not going to be defending what people like Williams and Sosa call foundationalism. My conception of foundationalism is different from and better than the standard view one finds in the literature. When I say that I want to defend foundationalism, I mean I want to defend what I call foundationalism. I am therefore proposing a distinctive account of what foundationalism really is, as well as a defence of the doctrine so defined.

So what do I call foundationalism and how does it differ from how these other authors understand that doctrine? As mentioned above, I take foundationalism to be a view about the *structure* of epistemic justification (a) and a view about its *sources* (b) & (c). With respect to (a) I take foundationalism to be a view about the structure of justification that is motivated by something traditionally called ‘the epistemic regress argument’. I will be spelling out that argument shortly. The important point for now is that foundationalism is a view about how - in very general terms - you have to conceive of epistemic justification

in order to respond satisfactorily to that argument. The foundationalist claims that in order to respond to the problem that argument raises we have to distinguish between inferentially and non-inferentially justified beliefs. As previously explained the latter beliefs are justified, but they do not draw their justification from other beliefs in the way in which my justification for believing that England cannot win the Ashes draws its justification from my belief they lost the first three tests, or the way my justification for believing that Socrates is mortal derives from my justification for believing that he is a man and that all men are mortal. The foundationalist claims that as well as inferentially justified beliefs such as these there are also non-inferentially justified beliefs. This is what Williams and Sosa mean when they talk about 'epistemically basic beliefs'. So I am agreeing with them to at least this extent. Epistemically basic beliefs are beliefs that are non-inferentially justified and according to (a) the foundationalist is someone who thinks there are such beliefs.

This is not yet the full-blown characterisation of foundationalism that one normally finds in the literature, although it may also be familiar. Nowadays some people think that (a) is all there is to foundationalism. Even those who don't, like Williams and Sosa, often give the impression that is what is really essential to foundationalism. Thus, Jonathan Dancy writes:

The claim that there are two forms of justification, inferential and non-inferential, is the core of any form of foundationalism in the theory of justification (Dancy 1985: 56).

This deflationary reading is gaining in popularity. It appeals to those who find traditional

foundationalism excessively baroque, but who can't quite bring themselves to believe that all justification is inferential.

I think the deflationary reading is too deflationary; I think that is not all there is to foundationalism. I am rejecting Dancy's deflationary reading just as firmly as I am rejecting the traditional reading.

I am rejecting the deflationary reading because foundationalism is also essentially a view about the *sources* of epistemic justification - about where justification comes from. Or so I claim. This is where (b) and (c) come in: with respect to (b) the foundationalist claims that perception or observation is a *distinctive* source of non-inferential justification; with respect to (c) she claims, further, that perception is a *basic* source of such justification - that perception is a basic way in which we come to know about the world around us.

These commitments will be spelt out further in due course, but this is the core of foundationalism as I understand it. It is a commitment to these three claims that I claim really marks foundationalism out as a philosophically interesting position and which distinguishes it from its historical rivals, rather than those traditionally focused on in the literature.³

I will return to the deflationary reading at the end of this section. How does my

³ That does not make foundationalism equivalent to 'empiricism'. The latter is a view about concepts on one important reading, whereas what I call foundationalism is a view about the structure and sources of epistemic justification. It says nothing about concepts and indeed, is perfectly compatible with the denial of empiricism in that sense.

conception of foundationalism differ from the standard view with which we began? Consider the characterisation Williams offers. On his view, foundationalism isn't just an abstract view about the overall structure of justification or the claim that there are inferentially and non-inferentially justified beliefs as (a) claims. The foundationalist is committed to highly substantive doctrines about the sorts of beliefs that can be non-inferentially justified and what it is about them that enables them to be so justified. As Williams emphasizes, foundationalism isn't merely 'formal'. In contrast, I claim that foundationalism is a generic style of response to a certain argument. As far as (a) goes, foundationalism is compatible with lots of more specific views about how our beliefs fit into the abstract structure of inferential and non-inferential justification dictated by the regress argument. It certainly doesn't require that the non-inferentially justified beliefs be indubitable, infallible, or theoretically tractable, and semantically encapsulated items of knowledge.

With respect to (b) I claim the foundationalist is someone who holds that perception is a source of non-inferential justification. The traditional definition, by contrast, makes no mention of perception at all. Far from requiring that perception be the basic source of such justification as (c) goes on to claim, the views of Williams and Sosa are compatible with thinking that it is not a distinctive source of justification at all. On their views there must be non-inferential justification but there is no requirement that it must derive from perception or that perception enjoys any other sort of epistemological privilege as I claim it does. I think that such a position wouldn't be recognizable as a form of foundationalism; it is certainly not what I understand by that name.

So there can be little doubt that the two characterisations differ in a fundamental way. If that is how they differ, what makes my view any better? Three things make my view preferable to the view one standardly finds in the literature. First, on my view foundationalism has at least some chance of being true. That seems a pretty remote prospect if foundationalism is understood in the way that Williams and Sosa understand it. I think there is no point in saddling people or positions with commitments that inevitably have the consequence that what they say is false. But even if foundationalists actually did think what Williams and Sosa suggest they thought, they *needn't* have done. I am offering them a philosophically respectable alternative. I am telling you what they should have thought, regardless of what they actually thought.

Second, as I characterise it foundationalism is an intuitively appealing position. This will become clearer in due course. The basic point though, is very simple: some of our beliefs really do seem to depend for their justification on that of other beliefs, and some of them do not, and there is nothing more ordinary or 'naïve' than the idea that perception is a basic way of acquiring knowledge of the world around us. This is basically what the foundationalist claims. My view therefore gives foundationalism roots in our commonsense thought about knowledge and justification. In contrast, Williams specifically aims to rule out this possibility. As he characterises it foundationalism is a distinctively philosophical position. His "most fundamental point" is that:

even if, at some level of abstraction, ordinary justification appears to accord with formal foundationalism this is no reason to suppose that we ought to be

foundationalists of the traditional kind...The question of whether there are basic beliefs cannot be decided by appeal to commonsense examples. This is because ‘basic belief’ is a theoretical concept, subject to stringent theoretical requirements. These do not derive straightforwardly from the desire to understand everyday justification. Rather, they are set by certain explanatory goals that are distinctively philosophical (Williams 2005: 204).

Thirdly, only my characterisation carves up the debate in a historically meaningful way. This is important since foundationalism is a label with a partly historical basis. On my view it is clear why coherentism and reliabilism, for instance, do not count as versions of foundationalism. They do not count because these positions do not hold that perception is a basic source of justification. According to the coherentist, all justification derives from coherence among one’s beliefs (BonJour 1985: esp. 87-222). According to the reliabilist, it derives from the fact that one’s beliefs are formed with a reliable process (Goldman 1986). The former is not a view on which there is such a thing as distinctively perceptual justification all; neither is a view on which perception is a basic source of justification. Such views therefore won’t count as forms of foundationalism as I understand it and that is as it should be.

It is also clear who does qualify as a foundationalist. On my view it is clear why Locke and Aristotle qualify, as well as more self-conscious foundationalists like C. I. Lewis (1946) and A. J. Ayer (1956). They qualify because they all have views on which the relevant inferential / non-inferential distinction is drawn and on which perception is a

basic source of non-inferential justification, despite the otherwise significant differences between them. These thinkers may not all understand the deliverances of the senses in the way in which we would, but all allow that perception or observation (in some sense) plays an irreducible justificatory role. This is a distinctively foundationalist idea as I conceive of foundationalism.

So my position enables us to carve up the debate in a historically sensitive way. The same is not true on the standard view. As Williams and Sosa characterise it, foundationalism is a highly specific doctrine. I will argue later that it is doubtful whether there are *any* foundationalists in Williams's sense. Even paradigm foundationalists like Russell (1912) and C. I. Lewis (1946) did not think that we could individuate the class of epistemically basic beliefs by aspects of their content. For thinkers such as these it was the epistemic source of these beliefs, rather than their subject-matter that was important. But even if these particular thinkers did somehow come out as foundationalists, it is obvious that the position Williams describes is not one that has been very widely held. It applies at best to a very small minority of thinkers. This sits oddly next to the central role that foundationalism plays in epistemological discussions and makes it hard to see why we should be interested in the position so characterised.

At this point proponents of the standard view are apt to fall back upon the deflationary characterisation touched upon earlier. For those familiar with the literature the following objection may have been brewing for some time. "It is true", they will say "that the view just characterised accurately describes what we might call 'substantive foundationalism'.

But that was always just intended as a label for a particular historical position, roughly, the view held and debated at the beginning of the 20th century by thinkers like Russell and Ayer. There is, however, a more general view that deserves the name foundationalism and which we can use when we want to understand why foundationalism enjoys the broader historical and philosophical significance that it does. This position we can call ‘formal foundationalism’”.

Formal foundationalism is closer to the commitment that I have labelled (a). As Williams characterises it, it just involves a commitment to the existence of ‘epistemically basic’ or non-inferentially justified beliefs. Sosa is more explicit: he takes formal foundationalism to embody a commitment to the supervenience of epistemic justification on non-epistemic features. He writes:

We need to distinguish, first, between two forms of foundationalism: one formal, the other substantive. A type of formal foundationalism with respect to a normative or evaluative property F is the view that the conditions (actual and possible) within which F would apply can be specified in general, perhaps recursively. Substantive foundationalism is only a particular way of doing so and coherentism is another (Sosa 2000: 14).

We needn’t worry about exactly how these formulations of formal foundationalism relate to one another since the basic strategy fails in either case. The substantive characterisation proved too narrow to be useful in delineating the essential nature of foundationalism, the more formal characterisations in contrast are far too broad. It is unclear which historical positions would fail to count as versions of foundationalism so

characterised. Coherentism certainly counts, as Sosa himself acknowledges. And while he is happy to accept that, it is a very paradoxical result. Foundationalism is a label with a partly historical basis and to apply it to positions that have so little in common with those historically called ‘foundationalist’ ultimately leads only to scepticism about the significance of the label.

I think we should avoid characterisations that have this consequence. My characterisation provides a way of understanding foundationalism on which it is both genuinely substantive and true.

2. The Epistemic Regress Argument

Now that we have some sense of what foundationalism is supposed to be we can discuss whether there are any good arguments for that view. I said that I would be taking foundationalism to be a view about justification that is motivated at least in part by the ‘epistemic regress argument’. What is that argument and how does it motivate the position I’m calling foundationalism?⁴ While there is broad agreement on the importance of this argument in motivating foundationalism there is less consensus on the form it should take.⁵ In the literature several different arguments purport to be the epistemic

⁴ Regress arguments occur in other philosophical settings. Gilbert Ryle uses one to draw a conclusion about the nature of voluntary action viz. that voluntary acts can’t be acts caused by a prior act of will if acts of will are themselves voluntary (Ryle 1949). Searle uses one in connection with intentionality (Searle 1983).

⁵ Evidence of its importance is legion. Thus, Alston claims that the main reason for being a foundationalist is “the seeming impossibility of a belief’s being mediately justified without resting ultimately on immediately justified belief” (Alston 1976: 182); Pryor calls it “the most famous argument in favour of non-inferential justification” (Pryor, 2005: 184); BonJour claims “the main reason for the impressive durability of foundationalism is not any overwhelming plausibility attaching to the main foundationalist thesis in itself, but rather the existence of one apparently decisive argument, which seems to rule out all non-sceptical alternatives to foundationalism” (BonJour 1978:1); and Bernecker and Dretske maintain “The driving force behind foundationalism has always been the threat of an infinite regress” (Bernecker and Dretske 2000: 231).

regress argument.

Sometimes people have in mind a dialectical regress in which subjects are invited to defend their beliefs in the context of an argument. This is how BonJour presents the argument in the following passage:

The most obvious, indeed perhaps the only obvious way to show that an empirical belief is adequately justified (in the epistemic sense) is by producing a justificatory argument: the belief that *p* is shown to be justified by citing some other (perhaps conjunctive) empirical belief, the belief that *Q*, and pointing out that *P* is inferable in some acceptable way from *Q*. Proposition *Q*, or the belief therein, is thus offered as a reason for accepting proposition *P*...[But] for the belief that *P* to be genuinely justified by virtue of such a justificatory argument, the belief that *Q* must itself already be justified in some fashion; merely being inferable from an unsupported guess or hunch, for example, can confer no genuine justification. Thus the putative inferential justification of one empirical belief immediately raises the further issue of whether and how the premises of this inference are justified...empirical knowledge is threatened with an infinite and apparently vicious regress of epistemic justification. Each belief is justified only if an epistemically prior belief is justified, and that epistemically prior belief is justified only if a still prior belief is justified, and so on, with the apparent result, so long as each new justification is inferential in character, that justification can never be completed indeed can never even really get started – and hence that there is no empirical justification and no empirical knowledge (BonJour 1985: 18-

19).

The dialectical regress, as I will call it, is concerned with what it takes for subjects to show that they are justified. You challenge me to defend my belief and in order to respond successfully to that challenge I must adduce some considerations in its defence thereby showing you that my belief is justified.⁶

This is not the only target of epistemic regress arguments. Other presentations of the argument focus more on what it takes for subjects to *be* justified in believing what they do. They aim to find out whether or not a subject is justified in a given belief and if so what makes it the case that she is justified.⁷ This is how DePaul presents the argument in the following passage:

According to this ancient argument, when we consider a belief that is justified and ask how it is that the belief is justified, we are typically led to another belief that supports the first. When we ask about the second belief, we may well be led to a third. The third may in turn lead to a fourth, and so on. But how long can things

⁶ This is even more explicit in Peter Klein's presentation of the argument. He asks us to imagine Fred and Doris in conversation "Fred asserts some proposition, say p. Doris says something – who knows what – that prompts Fred to believe that he had better have reasons for p in order to supply some missing credibility. So, Fred gives his reason, r1, for p. Now Doris asks why r1 is true. Fred gives another reasons, r2. This goes on for a while until Fred...arrives at what he takes to be a basic proposition, say b'. (Klein 2005: 133). Presentations of the regress do not often explicitly claim it concerns what it takes to show that a belief is justified. The argument is merely presented in such a way that is what it concerns, whether or not that is acknowledged.

⁷ The regress argument is often thought to be an ancient argument. Ancient presentations are equally ambiguous. Thus, Sextus Empiricus asks whether reasoning can ever legitimately lead to assent and writes: "the mode based upon regress ad infinitum is that whereby we assert that the thing adduced as a proof of the matter proposed needs a further proof, and this again another, and so on ad infinitum, so that the consequence is suspension [of assent], as we possess no starting point for our argument...we have the mode based upon hypothesis when the Dogmatists, being forced to recede ad infinitum take as their starting point something which they do not establish but claim to assume as granted simply and without demonstration. The mode of circular reasoning is the form used when the proof itself which ought to establish the matter of inquiry requires confirmation derived from the matter; in this case being unable to assume either in order to establish the other we suspend judgement about both" (Sextus Empiricus 1967: 166-9). In contrast, Aristotle employs a version of what I am calling the justification-making regress to show that some justification must be non-inferential (Aristotle, 1993: A3).

go for in this fashion? There would seem to be only three possibilities: the chain of beliefs either goes on forever, circles back upon itself, or stops. Finding the first two possibilities unacceptable, Foundationalists opt for the third, holding that there are some beliefs that are justified, but that are not justified by any further beliefs. The regress stops with such basic or foundational beliefs, and any other beliefs that are justified must be supported by the foundational basic beliefs (DePaul 2001: vii).

I'm going to call this the justification-making regress to distinguish it from the dialectical regress.⁸ Very often, however, one finds elements of both in the context of a single presentation. This is evident in the quotation from BonJour and it is also clear in the following passage from Dancy:

Suppose that all justification is inferential. When we justify belief A by appeal to beliefs B and C, we have not yet shown A to be justified. We have only shown that, it is justified if B and C are. Justification by inference is conditional justification only; A's justification is conditional upon the justification of B and C. But if all inferential justification is conditional in this sense, then nothing can be shown to be actually, non-conditionally justified. For each belief whose

⁸ Similarly, Susan Haack claims: "Suppose A believes that p. Is he justified in believing that p? Well, suppose he believes that p on the basis of his belief that q. Then he is not justified in believing that p unless he is justified in believing that q. Suppose he believes that q on the basis of his belief that r. Then he is not justified in believing that q, and hence not justified in believing that p, unless he is justified in believing that r. Suppose he believes that r on the basis of his belief that s. Then he is not justified in believing that r, and hence not justified in believing that q, and hence not justified in believing that p unless.... Now either (1) this series goes on without end; or (2) it ends with a belief which is not justified; or (3) it goes round in a circle; or (4) it comes to an end with a belief which is justified but not by the support of any further beliefs." Haack claims that if (1-3) is the case, then A's belief that p is not justified and goes on: "If (4), however, if the chain ends with a belief which is justified but not by the support of any further belief, A is justified in believing that p. So, since (4) is precisely what Foundationalism claims, only if Foundationalism is true is anyone ever justified in any belief. (Foundationalism is the only tolerable - non-sceptical- alternative.)" (Haack 1993: 22). For other statements of the justification-making regress see (Quinton 1973: 119) and (Pryor 2005).

justification we attempt, there will always be a further belief upon whose justification that of the first depends, and since this regress is infinite no belief will ever be more than conditionally justified...The regress argument therefore drives us to suppose that there must be some justification which is non-inferential if we are to avoid the sceptical consequence of admitting that no beliefs are ever actually justified (Dancy 1985: 55-6).

On the face of it, Dancy confuses the two different things just distinguished. He starts with a claim about what it takes to show that a belief is justified and concludes with a claim about whether or not the belief in question is justified. This is an easy mistake to make, however, since the word 'justify' is ambiguous and can be used to refer to both. Pryor calls attention to these different uses of the verb in the following passage:

On the first construal, 'justifying' a belief in P is a matter of proving or showing the belief to be just (or reasonable or credible). (Here we can include both arguments whose conclusion is P, and arguments whose conclusion is that your belief in P is epistemically appropriate, or is likely to be true.) By extension, we can also talk about things justifying beliefs; in this extended sense, a thing counts as justifying a belief if it's something you're in a position to use to prove or show your belief to be just...There's also a second way to construe the verb 'justify', which sees it as akin to the verbs 'beautify' and 'electrify'. When a combination of light and colour beautifies a room, it's not proving that the room is beautiful; rather, it's making the room beautiful. Similarly, on this understanding, justifying a belief is a matter of making a belief just or reasonable, rather than a matter of showing the belief to be just (Pryor 2005: 194).

We have seen that it is possible to initiate a regress argument using either notion. The important point is that the regresses thereby initiated will differ. One will concern whether a given belief has been shown to be justified; the other, whether a given belief is justified.⁹

Given that they differ, which do we have reason to prefer? When I talk about foundationalism as a position that is motivated in part by the regress argument what I mean is that it's a position motivated by the justification-making version of that argument. This is the more fundamental version of the argument for several reasons. It is the more prevalent, and we'll see shortly that it is by far the more plausible of the two arguments, but the most important reason is that it is only the justification-making version of the argument which threatens us with a truly unacceptable epistemological conclusion. It is only the justification-making regress which threatens to show that none of our beliefs is justified. The dialectical regress, in contrast, only promises to show that we cannot, or have not, shown that our beliefs are justified. The latter conclusion, while somewhat counter-intuitive, is not one we must do everything we can to avoid.¹⁰

⁹ In requiring that the subject respond to the question about what makes it the case she is justified it is easy to miss the fact that the questioning effectively goes higher-order: if she replies, she gives expression to a belief about what she thinks makes it the case that she is justified. To ask what justifies this belief is therefore to ask what justification she has for believing that she has a given justification. This is easy to miss, since it's easily confused with the case in which one appeals to beliefs as a *source* of justification. In the latter case, the beliefs one appeals to must be justified. To ask what justifies those beliefs is therefore not to ask a higher-order question. But that is that quite unlike the dialectical case; in that case I merely give expression to a belief in saying what makes it the case that I am justified.

¹⁰ This is especially plausible when one thinks about what it takes to show that something is the case. Alston draws attention to this in the following passage "showing by its very nature requires the exhibition of grounds. Furthermore, grounds must be different from the proposition to be shown. (This latter follows from the pragmatic aspect of the concept of showing. To show that p is to present grounds that one can justifiably accept without already accepting p. Otherwise showing would lack the point that goes towards making it what it is" (Alston 1976: 178-9).

Things would be different if we had some reason to link these two notions of justification and tie the conditions under which a belief can be justified with the conditions under which one has shown that it is justified. On certain views of justification one cannot be justified unless one can show that one is justified. Dancy's mistake may not then be a mistake but a substantive claim about the nature of justification. This is a view about justification with eloquent exponents. Thus McDowell describes 'the time honoured connection in our discourse between reasons the subject has for believing as she does and reasons she can give for thinking that way' and criticises writers like Peacocke for having to 'sever' that connection (McDowell 1994: 162-166).

There is clearly something to the picture of justification that McDowell recommends here. It is certainly true that when we have reasons for our beliefs, we can very often give them. It is also true that we normally expect other people to be able to give us their reasons (the very young think that we should be able to do that indefinitely) and we frequently take the fact that others can't give us their reasons as a sign that they do not really have any. That is partly why, as Austin pointed out, the question of why the subject believes as she does can be asked not just out of "respectful curiosity", but pointedly; her inability to answer can only reveal that she ought not to have been so bold (Austin 1979: 78). These are no doubt some of the reasons the dialectical regress can strike us as plausible. When one has reasons, the normal expectation is that one will be able to give them.

However, this is just an expectation and it is defeasible. At least, that is how I will be

talking about justification. When I talk about justification I am going to allow that a subject can have reasons or justification for believing what she does without necessarily being able to show that she has them, much as a person can be honest or funny without necessarily being able to defend the claim that she is honest or funny when under attack. This seems the more natural usage and the more plausible. At the most basic level talk of justification is a way of appraising someone who is doing well in her thoughts; there is no reason to think that necessarily brings with it the ability to show that you are justified.¹¹ The latter requires that you recognise that you are justified and have the ability to articulate and perhaps even defend the grounds upon which that is so. This looks like a more sophisticated cognitive achievement. Even where we are able to defend the claim that we are justified it is still important to distinguish between what it is that shows that we are justified and what it is which makes it the case that we are justified. Not everything that plays one role may be capable of playing the other. In order to show that you are justified you have to adduce claims to that effect, which claims are expressive of your beliefs. Showing that you are justified may therefore always involve appeal to beliefs. We will see shortly that it is crucial to foundationalism that being justified does not.

None of this to deny that there is something gripping about, what we might call, the ‘internalism’ McDowell here expresses - the idea that when one is justified one’s justification ought to be somehow ‘available’ to one. It is very plausible to think that for something to be *your* reason as opposed to just a reason, or for it to be what makes *you* justified in believing, it has to be accessible to you: it has to be a basis upon which you can

¹¹ Alston claims this as an “elementary point” (Alston 1976: 178).

justifiably form a belief, not just a basis upon which a belief could be so formed. But since it is not in virtue of your ability to articulate your reasons as your reasons that they count as yours we can acknowledge what is right about this line of thought without going as far as McDowell. They can be your reasons or what makes you justified in believing what you do even if you cannot state that fact when pressed, much less defend it under questioning.¹² It is also plausible to think that reasons must possess a certain sort of ‘evidence’ or perspicuity. They must be the sorts of things of which you are somehow aware: if they do not consciously reveal the world as being a certain way, what leads you to believe it to be one way rather than another? Again though that doesn’t warrant going as far as McDowell. What reasons make manifest, first and foremost, is the layout of the world, not their own status as reasons. It is the latter, however, that one would require for showing that one is justified.

As I understand it, then, justification is the sort of thing you can have without necessarily being able to show that you have it however often the two may in fact accompany one another. And when I talk about foundationalism as a position motivated in part by the epistemic regress argument I mean it’s a position motivated by the justification-making version of that argument.

So how does that argument motivate what I am calling foundationalism? Recall that the argument starts from reflection on cases of inferential justification – that is, cases in

¹² If one models accessibility as, in effect, belief – so that for a given fact to be accessible is for you to believe that fact obtains –there would be a much tighter connection between what is accessible and what is capable of being articulated (assuming beliefs are capable of being articulated). But we have no reason to model accessibility in that way.

which one's justification derives from other things that one has justification for believing. This is how many of our beliefs do seem to be justified. For instance, I may believe that Tony Blair will not win another election because I believe his policy in Iraq has been so unpopular with the electorate. That belief may be what makes it the case that I am justified in my belief about his electoral prospects. Or I may believe that Socrates is mortal because I believe that Socrates is a man and that all men are mortal. Some people call that 'mediate justification' since other beliefs, in this case my beliefs about foreign policy and mortality, mediate my justification. I'm going to stick with inferential though, since it is an important part of the reason why those further beliefs justify me that there is an acceptable inference between them and the belief they are claimed to justify. The fact that Blair's policies have been unpopular with the electorate makes it likely that he will not win: my beliefs thus stand in a relation of probabilification. In other cases the relation will be one of implication: that Socrates is a man and all men are mortal implies that Socrates is mortal.¹³

According to the epistemic regress argument where we have a belief that is justified we can ask what makes it the case that belief is justified: we can ask why the subject believes as she does or what justifies her in that belief. So in answer to the question what justifies my belief about Blair's electoral prospects we can appeal to my belief about his foreign policy. But what justifies this belief? If I don't have any reason to believe his policy has been unpopular then I won't have any reason to believe that he will not win. So what

¹³ It is important to what I am calling inferential justification that the inference be available between one's beliefs. Some people think that all justificatory relations obtain in virtue of inferential relations between a subject's attitudes, whether or not the attitudes involved are beliefs. This will not make all justification inferential in the sense in which I am interested. I'll return to this issue in chapter 3.

justifies me in the belief about his foreign policy? Here again the answer may involve appeal to beliefs and what justifies these beliefs may be yet further beliefs still. But how far can things carry on in this fashion?

On the face of it there seem to be only four ways in which the justificatory regress can pan out:

1. The regress ends with a belief that is not justified. While it is not justified it is still able to justify other beliefs.
2. The regress goes on forever: the belief that p is justified by the belief that q, and the belief that q is justified by the belief that r, and the belief that r is justified by....and so on, ad infinitum.
3. The regress circles back upon itself: the belief that p is justified by the belief that q, and the belief that q is justified by the belief that r, and the belief that r is justified by the belief that p.
4. The regress ends with a belief that is non-inferentially justified: while it is justified, it does not draw its justification from other justified beliefs.

If this is the choice with which the regress presents us, what would be a good response? According to the position that I am calling foundationalism it is only if the fourth option is correct, and some of our beliefs are non-inferentially justified, that any of our beliefs are justified at all. This is how the regress argument motivates foundationalism. According to this argument foundationalism is the only alternative to the view that none of our beliefs is justified. This ought to make it irresistible to all but the most sceptically minded.¹⁴

¹⁴ It is also possible to frame the argument in terms of knowledge. Just as we ask why one believes p, so too

Now we know what the regress argument is supposed to be, we can ask whether it is any good. The argument is an argument by elimination so the case in favour of foundationalism is only as good as the case against scepticism and the other options (1-3). I am not going to defend the rejection of scepticism, but I will now defend the rejection of the three alternatives. While they all enjoy support in some quarters, a strong case can be made against each of them. The case against them is strong insofar as it relies upon assumptions about justification that it is overwhelmingly plausible to make. It may be possible to give them up, but why do so unless we really have to?

The first option claims that the regress ends with an unjustified belief. While it is unjustified it is still able to justify other beliefs. How can a belief that is not itself justified, justify other beliefs? Many writers take that to be obviously impossible. Thus, Susan Haack simply states without further ado:

If A believes that p on the basis of his belief that q, then he is not justified in believing that p unless he is justified in believing that q (Haack 1993: 22).

Haack is not alone. The idea that beliefs must be justified in order to justify other beliefs is intuitive and it is grounded in a picture of the way in which beliefs confer justification that makes a lot of broader sense.

Take any central case in which one belief justifies another belief and it seems to do so in virtue of inferential relations between the propositions believed. The reason why my

we also ask how one knows p. It is easy to use the latter to initiate a regress: much of what we know we know because we have inferred those things from other things that we know, but could all our knowledge be like this? Not according to the foundationalist.

belief about Blair's unpopular foreign policy justifies my belief that he will not win another election is that it stands in a relation of probabilification to the latter. Similarly, the reason why my beliefs that Socrates is a man and all men are mortal justifies me in believing that Socrates is mortal, is that the former imply the latter. Inferential relations are therefore an important part of the story as far as beliefs go. But it is obviously not enough for one belief to justify another that it merely stand in inferential relations like these to it. That would make it far too easy to be justified since every belief stands in an infinite number of such relations to all manner of other beliefs (including itself).

Of course, such relations might explain why I believe certain things, given what else I believe. They might in that sense make it rationally intelligible that I believe as I do. But they do not, by themselves, give me any justification to believe those things, since they do not, by themselves, give me any reason to suppose that things actually are as my beliefs represent them as being. This is what Laurence Bonjour is getting at in the following passage. He writes:

For the belief that *p* to be genuinely justified by virtue of such a justificatory argument, the belief that *q* must itself already be justified in some fashion; merely being inferable from an unsupported guess or hunch, for example, can confer no genuine justification (Bonjour 1985: 18).

Bonjour's thought is that merely appealing to inferential relations will not do. This seems right; we want justification to be a guide to how things actually are in the world. We want it to have a connection with truth and mere inferential relations do not secure that.

So we need a further constraint and this is precisely the role played by the requiring the beliefs involved be justified. By specifying that the ‘inputs’ to this potential inference be justified we plug the intuitive justificatory gap that merely believing something leaves open.¹⁵ This gives us a picture of the way in which beliefs confer justification according to which justificatory status is *inherited*. Beliefs justify other beliefs to which they are suitably related and they do so by passing on their own justification. This is a compelling picture and it explains why an unjustified belief cannot make another belief justified. A belief cannot pass on justification it does not itself possess, just as I cannot inherit your car if you do not yourself possess one.¹⁶ If so then the first option according to which the regress ends with an unjustified belief is a non-starter.

Up until now I have talked rather loosely about the way in which beliefs confer justification. Actually the claim being made here is a claim about inference and the conditions under which the fact one’s beliefs stand in inferential relations to one another is capable of conferring justification. I am claiming that inference, in this sense, is a *conditional* vehicle of justification; it only confers justification where the input beliefs are already justified. Dancy makes a similar point in the following passage:

Inference is basically a matter of moving from premises to conclusion along an acceptable path. If the premises are unjustified there will be no justification for the conclusion - at least not by this inference (Dancy 1985: 55).

¹⁵ In other cases perhaps the sort of ‘attitude’ involved might plug that gap. But this is not plausible in the case of belief: the mere fact that one believes something, together with the fact that what one believes implies or probabilifies something else, isn’t enough to confer justification on the latter. Here, we must specify that the attitude (*viz.* belief) have a certain additional property, namely that of being justified.

¹⁶ In a similar vein, Jose Zalabardo claims “When a proposition *p* obtains warrant inferentially, it inherits it from other propositions to which it is suitably related. And *p* cannot inherit from other propositions warrant that the latter don’t possess” (Zalabardo, unpublished).

This makes inference an essentially *dependent* form of epistemic justification: its existence and functioning as a source of justification depends upon the existence and functioning of some other source of justification.¹⁷

This is why we need to distinguish between the conditions under which justification is inferential, and the conditions under which it derives from beliefs. Beliefs may be capable of furnishing us with a different sort of justification. In that case it may not be true that they must be justified in order to confer justification.¹⁸ However, this justification will not be inferential justification. This response will then no longer be a version of the first option; it will be a version of the fourth. It is hard to imagine what that role might be and that makes it tempting to frame the conclusion that I have drawn about inference as a conclusion about beliefs more generally and the conditions under which they are capable of furnishing us with justification. Strictly speaking, though, I have only argued for the claim about beliefs insofar as they confer justification by standing in inferential relations to other beliefs. I claim that in those circumstances, beliefs can only make other beliefs justified where they are themselves justified.

This is a very plausible idea and it rules out the first option according to which the regress ends with an unjustified belief. Notice though that is not part of a general claim

¹⁷ See also (Ginet 2005: 148-9).

¹⁸ Suppose I have the unjustified belief that Fino is matured in contact with air. Can't that belief still make me justified in believing that I have at least one belief? If so it is not true that only justified beliefs can make other beliefs justified, since my belief about Fino is not justified. Perhaps this makes sense; either way it would not be a case of inferential justification since there are no appropriate inferential relations between these two beliefs. There are appropriate inferential relations between the belief that I believe that Fino is matured in contact with air and the belief that I have at least one belief. But it is not obvious the former belief is unjustified; it is my views about Fino that are unjustified, not my views about what I believe about Fino.

about justifiers; it is a specific claim about inferentially justified beliefs. It says they can only confer justification where they are themselves justified. For all that's been said there may be other things which can make a subject justified – other things which can confer justification upon a belief – about which it doesn't even make sense to wonder whether or not they are themselves justified. As we will that possibility is central to foundationalism.

What about the second option? Is there anything wrong with supposing that the regress goes on ad infinitum? Since we do not actually have an infinite number of beliefs supporting each and every one of our beliefs, this response cannot be the only non-sceptical way of terminating the regress.¹⁹ Assuming it is unacceptable to claim that none of our beliefs is justified, the possibility of an infinite regress cannot show that foundationalism doesn't accurately describe the actual justificatory structure of our beliefs. This may be all the foundationalist needs, but I think we are entitled to a stronger conclusion in any case since 'infinetism' does not succeed in articulating a justificatory structure our beliefs could possibly enjoy (whether or not they actually enjoy it). Or so I will now argue.

The reason is very simple and takes us back to the point made at the end of the previous discussion in connection with the proposal that the regress ends with an unjustified belief. We are now considering a response according to which the regress goes on ad infinitum with each belief inferentially justified by some further belief. However. I have already

¹⁹ Some self-styled 'infinetists' are not committed to thinking we must actually have an infinite number of beliefs. Peter Klein thinks our beliefs merely become more justified, the greater the number of beliefs we have in support of them. This is not a version of 'infinetism' as I understand that position; it is a version of coherentism (as Klein himself acknowledges) (Klein 2005).

argued that inference is an essentially dependent source of justification: its existence and functioning as a source of justification depends upon the existence and functioning of another source of justification. Given that is so, it is not possible that inference could be the only source of justification. It is not intelligible that it could be the only way in which justification is conferred upon a subject's beliefs as the present response envisages since it is not an autonomous source of justification in that sense. If that is right, then the second option according to which the regress goes on ad infinitum is no good either.²⁰

What about the third option according to which the regress circles back upon itself? According to this option what justifies my belief that it rained last night might be my belief that the grass is wet and what justifies my belief that the grass is wet might be my belief that it rained last night. Does this make sense? It looks like a case in which a belief effectively justifies itself and that seems to defeat the whole point of requiring that beliefs be justified in the first place.

The obvious inadequacy of this response brings out even more clearly the underlying inadequacy of inference as a source of justification and helps explain why we should think that it is a fundamentally dependent form of epistemic justification. If inference were not a dependent source of justification as I have claimed it is unclear why circularity of this sort wouldn't be acceptable. After all, one of the beliefs that every belief stands in

²⁰ Here is another way to see that: arguments are only as good as their starting points. This is what I mean when I say that inference is a conditional vehicle of justification - whether it succeeds in conferring justification depends on whether the starting points are any good. An infinite regress is compatible with the starting points being all good or being all bad. The fact that, for every belief there is some further belief that would support it doesn't suffice to determine whether they do support it. All it rules out is the possibility that if any of the beliefs in the series is justified, then they aren't all justified given that each has successors that would justify it. But if inferential relations don't suffice to determine whether any is good though then something else must be necessary.

inferential relations to is itself; the belief that p stands in a relation of implication to the belief that p. So if inferential relations between one's beliefs were sufficient for justification (if inference could in that sense be one's sole source of justification) then one could be justified in believing anything whatsoever provided only that one does believe it.

This is totally unacceptable: beliefs are not justified simply in virtue of being held and they do not in that way justify themselves.²¹ Indeed, this misses the whole point of requiring that beliefs be justified in the first place. The original idea was that there should be something which functions as a reason why the belief is likely to be true and makes it something you ought to believe, something over and above the mere fact you do believe it. Thinking of inference as a dependent source of justification as I have done enables us to explain why circularity of this sort is unacceptable. Given that is so the third option according to which the regress goes round in such circles is also no good.

Having just said that I want to acknowledge that this is not the only way that people have understood the suggestion that the regress 'circles' back upon itself. According to a position known as coherentism the objection just raised goes wrong in assuming a 'linear' conception of conception. According to the coherentist's holistic alternative we are not to think of justification being passed from one belief to the next, eventually landing up back with the belief with which the series begun. Rather, justification is a property of an individual's entire set of beliefs. Specifically, it is that property the set enjoys when it's

²¹ Thus Quinton writes: 'For a belief to be justified it is not enough for it to be accepted, let alone merely entertained: there must also be good reason for accepting it' (Quinton 1973: 119).

individual members 'cohere' with one another and it accrues to each individual belief in virtue of its membership of such a set of beliefs.

This is a very implausible account of what justifies our beliefs, but I do not want to take issue with it here. The important point is that this strategy effectively appeals to beliefs that are *non*-inferentially justified in terminating the regress. In fact, it says that all of our beliefs are non-inferentially justified, since it says that they are all justified by the fact they belong to a coherent set of beliefs and that is not where justification comes from in standard cases of inferential justification.

What I mean by this is that the relations the coherentist appeals to are not inferential relations of the usual sort - they are not like the relation that Socrates is a man and all men are mortal stand in to the belief that Socrates is mortal - they are far more extensive and encompass all of an individual's beliefs. There is also a difference in the sort of access the subject has to those relations and the role they play in getting her to be justified. In the Socrates case, there's a fairly robust sense in which I believe that Socrates is mortal *because* I believe that he is a man and that all men are mortal; the latter beliefs really are operative in getting me to believe as I do, and if I had different beliefs on that front, I'd adjust my beliefs about his mortality accordingly. This is quite unlike the holistic case, where it is doubtful whether I am aware of the relevant facts about all my beliefs and their inter-relations (or could even easily become aware of them), and doubtful these facts are in any way operative in getting me to believe as I do. This may be the basis on which certain coherentist epistemologists form their beliefs, but it is clearly

not the basis upon which most of us do so.²²

Finally, though, even the coherentist doesn't just appeal to facts about the inferential relations between our beliefs, however extensive we take the set to be. She also has a story to tell about why the fact that one's beliefs cohere makes them likely to be true, and hence why it should be a justification that they provide. For instance, Davidson appeals to the fact that beliefs are by nature veridical (Davidson 2000). This ought to remove any remaining temptation to call this a case of inferential justification, in the ordinary sense. We certainly do not ordinarily appeal to facts about the nature of beliefs in making sense of inferential justification. Coherentism is therefore a version of the fourth strategy which I have associated with foundationalism, rather than the third.

I am not claiming that coherentism is a particularly plausible version of the fourth strategy or that there is nothing awkward about so characterising it. I am just claiming that overall it is better seen as a version of the fourth strategy which appeals to a non-inferential source of justification, than to one which maintains that all justification is inferential in the relevant sense. In the end, though, it may be more accurate to say the coherentist simply rejects the framework of inferential and non-inferential justification within which the regress is set. After all, that position is one according to which all beliefs have the *same* source of justification; they are all justified by the fact they are

²² Intuitively to be justified in believing *p* is to believe *p* on the basis of the facts which give you justification to believe *p*. This is hard for the coherentist to make sense of, since it is very hard to see how a belief's inter-relations to all other beliefs could be the 'basis' on which you adopt it. The only obvious way of making sense of this is to suppose you believe *p* on the basis of a coherentist meta-argument claiming that the belief that *p* coheres with the rest of one's beliefs, and that beliefs which cohere are likely to be true. This makes sense of how such facts could intelligibly be the basis upon which one believes but it is clearly very implausible as a description of the basis upon which most people form their beliefs. Coherentism therefore leaves most people's beliefs unjustified and that isn't much of a recommendation.

members of a coherent set. Ironically, it may turn out that the best way of bringing out the difference between coherentism and more traditional versions of the fourth response is not by stressing the idea that some justification is non-inferential (since that is something the coherentist thinks is true of all beliefs) but by stressing the idea that some justification is actually inferential (some of it really does derive from other things that we have justification for believing). Intuitively, not all beliefs have the same source of justification and this is something the coherentist denies.

So I have claimed that inferential justification is a fundamentally dependent form of epistemic justification and I have rejected options (1-3) on that ground. However, I have also suggested that subtle variations of those options ought to count as versions of the fourth strategy. It is now time to consider this option – the option according to which the regress ends with a belief that is non-inferentially justified. Would this be a good response to the regress?

On the face of it, yes it would. Unlike the other three options there is nothing immediately implausible about the suggestion that the regress ends with a belief that is non-inferentially justified.²³ Some people deny that a belief can be justified by anything other than its inferential relations to other justified beliefs. But there is nothing intuitive about this view. It rests entirely upon philosophical arguments that we will later find wanting. Given that options (1-3) are no good, and that there is nothing intuitively problematic with the fourth option, the regress argument does look like a good argument

²³ Of course, it is much harder to see how one would stop the dialectical regress in this way. How can one show that a belief is justified other than by producing a justificatory argument in its favour? See also (Pryor 2005: 193-4).

in favour of thinking that there must be such a thing as non-inferential justification.

But is it a good argument in favour of foundationalism? It is certainly not a good argument in favour of foundationalism in the *traditional* sense. Nothing in this argument supports the demand for a layer of non-inferentially justified or epistemically basic beliefs distinguished in terms of their content or enjoying the sorts of strong epistemological privileges that Williams and Sosa describe. All this argument supports is the claim that there must be some beliefs that are non-inferentially justified; there must be some beliefs that do not draw their justification from their inferential relations to other justified beliefs. It doesn't tell us anything about the sorts of beliefs that can be non-inferentially justified or what it is about them that enables them to play that role. Any substantive claims on that score are totally unmotivated. The regress argument is therefore not a good argument in favour of foundationalism as traditionally conceived.

Is it a good argument in favour of foundationalism in my sense? I have claimed that foundationalism is a view about the structure of justification (a) and a view about its sources (b) & (c). With respect to (a) I have claimed that foundationalism is a view about the structure of justification that is motivated by the epistemic regress argument. The foundationalist claims that we must acknowledge the existence of non-inferentially justified beliefs in order to respond satisfactorily to that argument. This is precisely what we have just seen the regress does establish, given that the other three alternatives are no good and that scepticism is false. So the regress argument is a good argument in favour of (a). That is hardly surprising though since (a) simply commits the foundationalist to

whatever it is one needs to solve that problem. Once one sees foundationalism as a response to the regress argument it is easy to see that my response is a better response to that argument than the traditional view, since my view is actually motivated by that argument. The traditional view, by contrast, commits itself to all sorts of things that just aren't relevant to solving that problem.

Of course someone might now say: why call that view 'foundationalism'? In a way this is a good question. As we have just seen the regress argument only gets you as far as thinking there must be some non-inferential justification. This cannot be all there is to foundationalism; it is far too permissive. We have seen that slightly modified versions of options (1-3) all appeal to non-inferential sources of justification on one reading, yet it would be wrong to think of those positions as forms of foundationalism in any serious sense. This is why (a) is only one component of the position I call foundationalism. Foundationalism isn't just a view about the overall structure of justification. It is also a view about the sources of justification. This is the point of (b) and (c). A foundationalist also thinks that perception is a distinctive and privileged source of non-inferential justification.

So is the regress argument a good argument in favour of these claims? No. The regress argument doesn't say anything about where justification actually comes from; it just tells us where it *doesn't* come from. So the regress argument is therefore not a good argument in favour of foundationalism in my sense since it doesn't on its own get you all three components of that position. It is a good argument for (a) but not for (b) or (c).

Of course, once we get as far as acknowledging the existence of non-inferential justification an overwhelmingly natural question presents itself: where does such justification come from? There is then a very natural progression from this view to full-blown foundationalism in my sense, since the natural answer to this question is to advert to the senses. This is an independently plausible claim about where justification actually comes from and the beauty of this response is that it looks like it provides just the kind of non-inferential justification that we need – the sort of justification, that is, which doesn't land us back with a form of the regress. Unlike inference, perception isn't an essentially dependent form of epistemic justification. It doesn't merely spread around justification that is already there or require antecedently justified beliefs as inputs. It can be what, in the first instance, gives us justification to believe; it is a source of new justification.

But this response, however natural, is not mandated by the regress; it is not something that argument establishes. This is not to deny that the story one tells about how or why the senses are a privileged source of justification might not be tacitly informed by the picture of justification underlying the regress, or that rivals like coherentism might not run counter the spirit of that argument. I have already pointed out that coherentism sits oddly next to the idea that there are both inferential and non-inferential sources of justification, and it repudiates the linear conception of justification underlying that argument by denying that we can meaningfully ask after the justification of an individual belief without settling the status of all the rest. We will see in later chapters how foundationalism is more in keeping with the spirit of the regress. My point is merely that

the foundationalist's claims about the primacy of perception and about where non-inferential justification comes from are not an inevitable consequence of the regress argument.

This is a strength rather than a weakness in the foundationalist's argument since it depends upon considerations that even those who claim not to be moved by the regress or its linear conception of justification ought to take seriously. Views like coherentism, which fail to do so, and which deny that there is any such thing as distinctively perceptual justification are therefore doubly wide of the mark. The fact that they aren't ruled out by the regress is not enough to save them given that this is not the only thing to be said in favour of foundationalism. Foundationalism is also motivated by the desire to accommodate the obvious fact that perception plays a justificatory role and that is something the coherentist denies. This is a reason not to take that position seriously quite aside from concerns about the regress.

3. Conclusion

In conclusion, I have argued that the epistemic regress argument doesn't establish as much as some have thought since it doesn't establish foundationalism. It is a good argument in favour of one component of that view, but not the other. This is not a problem for my view, however, since I am not taking foundationalism to be motivated solely by the regress argument. The regress argument motivates the foundationalist's claims about the overall structure of justification. But foundationalism is also a view about its sources. Specifically, it is the view that perception or observation is the basic

source of justification. This is not motivated by the regress argument and nor is it meant to be. It is an independently plausible claim about where justification actually comes from. The next chapter will explore two very different ways in which a foundationalist can try and hold onto it.

FOUNDATIONALISM AND THE IDEA OF THE EMPIRICAL

CHAPTER 2: THE HISTORICAL FOUNDATIONALISTS

1. Introduction

In chapter 1 I offered a certain characterisation of foundationalism. Foundationalism is a view about the structure of epistemic justification and a view about its sources. Specifically, it's the view that (a) some of our beliefs must be non-inferentially justified, and that (b) perception is not only a source of such justification, but (c) a basic source of such justification. If you look at what actual historical foundationalists have said though, their position seems to be quite different from what I'm calling 'foundationalism'. They don't appear to be committed to (b), let alone (c). And they are committed to, or at least endorse, a series of further claims that I don't talk about at all. Historically, foundationalists had distinctive views about the content or nature of the basic or non-inferentially justified beliefs, the source of our knowledge of them, and the epistemic credentials of such beliefs. I am silent about these further claims. I do not commit the foundationalist to them; and, in fact, my view is that these further claims are almost certainly false. So there's a discrepancy; while I commit the foundationalist to (a), (b), and (c), foundationalists historically endorsed a very different set of claims.

Someone might therefore say: 'you can call your position 'foundationalism' if you like, but that is not a label that makes any historical sense. Your position is just too different from those historically called foundationalist so to insist on calling it that is just

anachronistic'. My reply to this objection is that we have to distinguish between what the historical foundationalists thought and what they should have thought – that is, between those claims they actually endorsed and those they were committed to endorsing. What I'll be arguing in this chapter is that historical foundationalists needn't have committed themselves to these further claims and, what's more, they shouldn't have done. They needn't have done because you don't need to think these further things in order to deal with the epistemic regress argument. And they shouldn't have done for lots of familiar reasons; most obviously, you end up with an implausible and unworkable account of what actually justifies most of our beliefs.

Insofar as the historical foundationalists really did commit themselves on this score, their position is no good. This is the moral of the first half of this chapter. But this chapter is not all negative. Foundationalism isn't just important assuming we ignore everything that foundationalists ever actually said. Some of the considerations motivating these thinkers were understandable and a lot of what they thought was almost right. Or so I will argue. In particular, I will argue that although traditional foundationalists might not have explicitly asserted (b) or (c) they did, in effect, think something like that. They thought that observation or perception (albeit of a funny sort) was a source of justification fundamentally distinct from reasoning or inference. There is something importantly right about that. Unfortunately in their case other commitments got in the way and prevented them from seeing that straight. Under the weight of these additional commitments that idea got perverted and ultimately transposed into something importantly different. Nonetheless, the original insight – the real driving force – is one that we should hold onto

and in chapter 3 I'll show how we can do so without being foundationalists in the traditional sense.

What I really want to do in this chapter, then, is to save historical foundationalism from itself. I want to show why all the things which traditionally got these thinkers into trouble are things they didn't need to think. And I also want to show why my own position, which doesn't endorse these further claims, still deserves the label 'foundationalism'. It's not just that the label fits if you leave aside everything the foundationalists historically thought. It fits in a historically meaningful sense since the two positions are ultimately driven by the same basic intuition. In both cases that intuition is a good one.

2. Historical Foundationalism

So who are these 'historical foundationalists' and what did they actually think? While this label is perhaps most famously associated with early twentieth century thinkers such as Russell, Lewis, Chisholm, and Ayer, it can also be applied to early figures like Locke. They all thought that some of our beliefs must be non-inferentially justified and they all had a distinctive view about what could be justified in this way and why. It's those views which distinguish historical foundationalism from the version of foundationalism that I'll eventually be arguing we should adopt. Three claims, in particular, stand out. They are all part of the positive account of the non-inferentially justified or basic beliefs. I am going to call them the Subject-Matter, the Source, and the Status, Proposals respectively.

First, historical foundationalists had a distinctive view about the nature or content of the

epistemically basic beliefs. They thought that only beliefs about our own psychological states could be non-inferentially justified. In particular, they thought that ordinary beliefs about non-psychological reality - like the belief that there is a squirrel on the fence - could not be non-inferentially justified. For them the foundational knowledge was limited to knowledge of our own minds. Any belief about how things are in the world external to one's own mind depends for its justification on inferences one can make from beliefs that are about one's own mind. I'm going to call this the Subject-Matter Proposal since it's a claim about the sorts of things that we can be non-inferentially justified in having beliefs about.

Second, historical foundationalists had a distinctive account of how we are in a position to form justified beliefs about our own minds. To claim that such beliefs are epistemically basic or non-inferentially justified is merely to say what doesn't justify them; it says that they do not get their justification from their inferential relations to other beliefs. But we're also owed a positive account of what does justify these beliefs. The traditional answer was 'observation'. You are in a position to justifiably make judgement about your own psychological states because you 'observe' or are 'acquainted' with or are otherwise somehow aware of those states, and form your beliefs about them on that basis. This is why such beliefs are not inferentially justified since observation is not a form of inference. It offers a fundamentally different model of how we can be in a position to make judgements. I am going to call this the Source Proposal since it is a claim about where the justification for basic beliefs comes from, given that it doesn't come from other justified beliefs.

The third defining feature of historical foundationalism is an account of the epistemic status of beliefs about our own psychological states. They had impressive epistemological credentials according to the historical foundationalists and were commonly held to be both infallible (that is, incapable of being mistaken) and indefeasible (incapable of being rationally revised). This is what I'm calling the Epistemic Status Proposal.

These three proposals help to define historical foundationalism. What is the relationship between them? In the secondary literature it is common to stress the third proposal concerning the epistemic status of beliefs about our own psychological states (Audi, 1998, Bonjour 1985, Dancy 1985, Lehrer 2000, Pollock & Cruz 1999. Traditional foundationalists it is claimed had a particular view about non-inferential justification according to which only beliefs that were incapable of being mistaken or rationally revised could be justified in this way. After all, if a belief could be mistaken it would be as much in need of epistemic support as any other belief and so could hardly provide the secure foundations upon which the others rest. Beliefs about non-psychological reality, like the belief that there is a squirrel on the fence, can be mistaken. But beliefs about psychological reality cannot and it is because these beliefs can't be mistaken that they can provide the secure foundations that we need. Thus Dancy writes:

How is it that beliefs about our present sensory states need no support from others, while all other beliefs require such support? The answer comes from the third element of classical foundationalism: this is that our beliefs about our present sensory states are infallible. It is because of this that they can play the role

ascribed to them in this form of empiricism; beliefs about our present sensory states can be our basis – can stand on their own two feet and support the rest – because they are infallible (Dancy 1985: 53-4).

On this reading the overall account is driven by a certain view about non-inferential justification and what it would take for any belief to be justified in this way. This view explains why beliefs about our own psychological states are epistemically basic, as the Subject Matter proposal claims, since it's only those beliefs which can plausibly be thought to enjoy such an exalted status. And it also explains why such beliefs have the epistemic source that they do, since observation is the basis upon which these beliefs are held on the traditional account. I am therefore going to call this the epistemological route to historical foundationalism.

This explanation is probably the right one in some cases.²⁴ It seems to be the picture that Russell has – he has a certain epistemological view about what it would take for any belief to be non-inferentially justified and that drives him to look for things which meet these requirements. The things he comes up with are so-called sense-data and ourselves. We can be mistaken in our beliefs about how things are in the world around us, but it is less obvious that we can be mistaken in our belief about the things that Russell picks on. The idea that we have privileged access to the self is widespread and not without appeal, and it was thought equally hard to make sense of error in the case of sense-data since the latter are mind-dependent objects that have exactly the properties they appear to have. Such objects were said to be 'self-intimating' to the subject upon whom their existence

²⁴ The precise nature of the epistemological view is open to question. In chapter 4 I will argue that it is indefeasibility rather than infallibility, which really drives this account. This is not the standard view one now finds in the literature.

depends. According to Russell these are the things with which we are ‘acquainted’. Moreover, it is our acquaintance with them that then explains how we are in a position to know about them. We are in a position to know or justifiably form beliefs about them insofar as we are aware of them in this special way and believe what we do on that basis. So acquaintance or awareness is the source of our justification, where that is essentially observation of a special sort of object (Russell 1912).

In other cases, however, this is not the right order of explanation. In the case of someone like Locke it’s not so much that he has a special view about non-inferential justification, which then leads him to look for beliefs which could be justified in this way. Rather, he takes it for granted that observation is a source of justification and just happens to have a particular metaphysical view about what it is that we actually observe. This view is independently motivated. It claims that we do not perceive objects in the external world, at least not directly. We only really perceive our own ideas. For someone like Locke this is why it’s only beliefs about the latter that can be non-inferentially justified; we do not observe the former and so observation trivially isn’t available to justify our beliefs about them. The epistemic status of basic beliefs then falls out of this, since the things which Locke thinks we actually perceive happen to be such that we cannot be mistaken or rationally revise our beliefs about them when we form those beliefs on the basis of observation.

I’m going to call this the metaphysical route to historical foundationalism since it starts from a particular metaphysical picture of what it is that we observe. In both cases beliefs

about mind-independent objects depend on inferences we can make from beliefs about our own psychological states. But in Locke's case there is no prior commitment to the idea that if we did observe such objects, observation still couldn't non-inferentially ground our beliefs about them. He just has other reasons for thinking we don't observe mind-independent objects. The overall account is therefore driven by a metaphysical view about what it is that we actually observe, rather than an epistemological view about the conditions under which observation, per se, can be a source of non-inferential justification.

These are two different routes to the position that I'm calling historical foundationalism. Clearly in the case of many historical figures one can find elements of both. I am not claiming that these two motives operate entirely independently of one another, merely that in certain cases one is more pronounced than the other.

At first glance historical foundationalism looks more like the position that Williams and Sosa describe and which I rejected in the first chapter than what I call foundationalism. In fact, however, not even historical foundationalists thought that the class of epistemically basic beliefs are 'theoretically tractable' as Williams requires. In particular, they did not think that there were 'non-trivially specifically kinds of beliefs individuated by broad aspects of their content that are fitted to play the role of terminating points for chains of justification' (Williams 2005: 203). The emphasis on content is absolutely crucial to the orthodox view one finds in the literature, but it is completely foreign to historical foundationalism. On the historical account what is important about beliefs about one's

own psychological states isn't some aspect of their content; it is not that there's a certain subject matter 'the self and its states' about which one is guaranteed not to be mistaken. These beliefs are special because - and only because - we form them on a certain basis, namely on the basis of acquaintance with (or observation of) those very states.

According to the historical foundationalists, acquaintance or observation consciously presents us with the facts that make what we believe true: it makes those facts 'manifest' to us by consciously revealing them to us and that is why we are in an epistemically favourable position to make judgements about them. For the historical foundationalists, then, it's the epistemic source of belief about our own psychological states rather than their content that is important in explaining why such beliefs can be justified in a way that doesn't derive from the justification of other beliefs.²⁵ The standard view one finds in the literature leaves this out altogether.

The easiest way to see why that is so important is to think about cases in which we hold beliefs about our own psychological states on non-observational grounds. In such cases we won't get the same explanation of why the beliefs in question are justified if I am right. If it is in virtue of the *circumstances* in which one comes to entertain the relevant belief that one is in a position justifiably to judge, then we shouldn't expect beliefs formed in different circumstances to be justified in the same way, despite the fact that they concern the same topic. In fact, this is exactly what we find.

²⁵ I am not claiming these two things are totally distinct. Clearly it's because such beliefs are 'about' what they are about (viz. our own minds) that they can be justified in the way that they are justified, that is, by observing our own mind. If they were about something else altogether, observation could not justify them. Still, it is the latter which explains why they are justified, not some aspect of their content.

Consider the following case from Pollock and Cruz:

Consider shadows on snow. Because shadows on white surfaces are normally grey, most people think that shadows on snow are grey. But a discovery made fairly early by every landscape painter is that they are actually blue. A person having the general belief that shadows on snow are grey may, when queried about how a particular snow-shadow looks to him, reply that it looks grey, without paying any serious attention to his percept. His belief about how it looks is based upon his general belief rather than inspection of his percept, and is accordingly wrong. This shows that the belief is not incorrigible... Suppose further you're your inductive evidence is faulty and you are unjustified in believing that shadows on white surfaces are grey. Then you are unjustified in believing that the snow shadow looks grey (Pollock & Cruz 1999: 58-60).²⁶

Cases like are often thought to be a problem for the foundationalist; and they would be if the orthodox view were right and the foundationalist really was trying to delineate basic beliefs in terms of their content as Williams and others assume.

I'm going to call that the Propositional Assumption since it claims that the foundationalist aims to articulate a set of propositions that we are non-inferentially justified in believing.

This assumption goes unchallenged in the literature, but it is not an assumption to which

²⁶ Here is a different example (also based on an case of Pollock's): suppose I have an alarm clock and I notice that every time the alarm goes off a light red light flashes in the lower left hand corner of my visual field and it appears red to me. A big bee now suddenly comes on the scene and hovers perilously close to my nose. Not surprisingly I forget all about the alarm clock and focus on the bee in the middle of my visual field. But I then hear the alarm go off and that gives me inductive grounds for believing that 'it appears to me as if there is a red light flashing' even if, having been so preoccupied with the bee, I fail to notice that is how things appear to me.

the historical foundationalists were committed, as my characterisation makes clear.²⁷ They don't think that any old beliefs about one's own psychological states must be non-inferentially justified. It is only beliefs that one holds on the basis of observing those states, and not all beliefs that are 'about' one's own psychological states need to be held on that basis.²⁸

Pollock and Cruz consider that response, but duly reject it. They write:

There is a response to this counter-example which has considerable intuitive pull, at least initially. This is to agree that not all beliefs about how things appear to us are prima facie justified, but those based upon being appeared to in that way are. Taken literally, this makes no sense. Prima facie justification is a logical property of propositions. A proposition cannot have such a property at one time and fail to have it at another (Pollock and Cruz 1999: 60-1).

In fact this does make sense - provided that one is not committed to the Propositional Assumption. If I am right the historical foundationalists were not committed to this assumption; they think exactly what Pollock and Cruz think they cannot think and this is possible precisely because they do not aim to articulate a set of propositions we are non-inferentially justified in believing in the way that he assumes.²⁹

²⁷ Bill Brewer appears to be committed to this assumption, though not in the name of foundationalism. Brewer claims that in entertaining the content of a perceptual state one thereby necessarily recognises its truth, and hence has reason to endorse that content. This is how he thinks that our perceptual experiences provide reasons for our judgements about the world. (Brewer 1999: esp. 204-6).

²⁸ How does this relate to what I earlier called the Subject-Matter Proposal? The latter lists the sorts of things we can be non-inferentially in having beliefs about. According to the traditional foundationalists that meant our own minds. This is not what the Propositional Assumption claims; it claims we can identify a set of propositions we *are* non-inferentially justified in believing. But whether we are inferentially or non-inferentially justified in believing some proposition depends on more than just what basis propositions like that *can* be held; it depends on what basis the proposition in question are held and that requires looking at more than the content of the beliefs involved.

²⁹ Unfortunately they draw the wrong conclusions about foundationalism from their own examples. This is because he falsely assumes that foundationalism is a 'doxastic' theory. In the passage just quoted he goes

This is easy to overlook since we don't normally hold beliefs about our own minds on other grounds. According to the historical foundationalists observation of one's own mind is the canonical ground for beliefs about it. Not surprisingly therefore observation turns out to be the ground upon which these beliefs are in fact always held since there is no reason not to hold those beliefs on observational grounds, if such grounds are genuinely available. Still, it is not impossible to hold such beliefs on other grounds, even on the traditional account, and that tell us something important about historical foundationalism.

This is a much more attractive account of the source of our justification for beliefs about the mind, since there is no aspect of their content to which we can plausibly appeal in any case. In the literature it is often objected that beliefs about one's own psychological states aren't 'self-evident'. Unlike the thought that whoever is tall is tall, say, the mere entertaining of propositions about one's own psychological states does not put one in a position to know whether or not they are true.³⁰ This is very plausible. However, it is only an objection to the foundationalist if we assume that they must appeal to a subject's grasp of the content of basic beliefs to explain why such beliefs are justified in a way that

on: "the claim actually being made here is presumably a different one, viz., that when we are appeared to in a certain way, that in and of itself can make us at least defeasibly justified in believing that we are appeared to in that way...later in this book we will endorse a theory providing such a foundation for epistemic justification, but notice that such a theory is no longer a doxastic theory. The justifiedness of beliefs is no longer determined exclusively by *what* we believe. What percepts we have is also relevant. Thus this is not a way of saving doxastic foundationalism' (Pollock 1999: 61). I don't know how they first came to the conclusion that foundationalism is a doxastic theory.

³⁰ One could try and argue that the relevant content is only available in the presence of its presented subject-matter, so that it is impossible even to frame thoughts about one's own psychological states in other circumstances. But that is implausible and it has problematic consequences in other areas; if the foundations are really that limited, it's even less plausible to suppose they suffice to support everything else that we know. Moreover, it's unclear to what extent this account succeeds in preserving a pure subject-matter account. What is special about the particular circumstances in which one comes to entertain these thoughts, if not the fact that one is consciously presented with the items about which one judges? In that case though do we really have anything more than a mere re-labelling of a source based account?

doesn't derive from the justification of other beliefs.³¹ This is just a mistake. Self-evidence isn't the only alternative to inferential justification. Observation is another and that I suggest is precisely what the historical foundationalists did appeal to.

So I have now explained what the historical foundationalists actually thought and why what they thought is different from what they are usually portrayed as having thought. Broadly speaking, they thought that basic beliefs are beliefs about one's own psychological states; what justifies such beliefs is the fact that one observes or is acquainted with one's own psychological states; and that beliefs held on such a basis are epistemically privileged. This makes historical foundationalism a lot more plausible and interesting than it is often thought to be. Observation does seem to be a fundamental way in which we get in a position to know things. So if it is true that we do observe our own mental states this will constitute an appealing account of what justifies our beliefs about them.

But 'more plausible' does not mean 'plausible' and in other respects historical foundationalism remains distinctly odd. Even if we do observe our own mental states they are presumably not the only things that we observe or, indeed, the most obvious. Intuitively, we also observe objects and events in the world external to our minds, like squirrels sitting on fences. Historical foundationalists deny that this is so and that looks a long way from the truth. The next section will look at just far away they got.

³¹ Some historical figures we often think of as foundationalists may have thought that; perhaps Descartes did. I don't think that these thinkers – that is, those who give no weight to observation – deserve to be seen as foundationalists. If that means that Descartes does not count as a foundationalist, so much the better. I don't see that there is anything problematic about that.

3. Revisionary Epistemology

The previous section looked at what the historical foundationalists actually thought. In any event it is clear this is not what they should have thought. This is so for a number of familiar reasons. The most obvious and the most powerful concerns what we are to say about the justification of beliefs that aren't about our own psychological states. Even if traditional foundationalism does provide a satisfactory account of what justifies beliefs that are about our own minds, we have as yet no account of what justifies beliefs about the world external to our minds. Most of our beliefs fall into the latter category, so the foundationalist had better have a good answer to this question. The answer they traditionally gave was inference: beliefs about the world draw their justification from inferences we can make from beliefs about our own minds. This is hardly surprising. Foundationalists are committed to the view that all of our beliefs are either inferentially or non-inferentially justified. Beliefs about the world don't qualify as non-inferentially justified on the traditional view, so they must be inferentially justified (if they are justified at all), and it is hard to see what else to appeal to as premises other than beliefs about one's own mind.

This view is subject to fatal objections. It is implausible as a description of the psychological process by which we form beliefs about the external world. Such beliefs don't appear to be formed on the basis of any kind of reasoning or conscious inference; they normally seem to be formed directly on the basis of perception or other's say-so. Further, it is implausible as an account of the justificatory status of such beliefs. The

precise nature of the problem here will depend on the details of the particular account given. On one view, our own mental states (including our ideas and ‘sense-data’) are held to be distinct from mind-independent objects in the physical world, and to stand in for or represent them. On another view, they partly constitute them.

The latter view, known as ‘phenomenalism’, seems to call into question the existence of genuinely mind-independent objects; while the former, commonly called ‘indirect realism’, will have to rely on general bridging principles linking our ideas or sense-data with the objects in the world that they are held to represent. In the second case, the problem concerns the justification of the bridging principles. One possibility is that they are general causal principles stating that ideas or sense-data of such-and-such a sort are reliably correlated with (or caused by) physical objects of such-and-such a sort. Given such general principles, one’s own psychological states can serve to indicate how things are in the world around one; they can be a sign that they exist for one who knows the relevant principles. However, it is a familiar point that our justification for these principles appears not to be independent of our justification for beliefs about the world. Michael Martin makes this point in the following passage:

If we inquire into our reasons for believing that certain kinds of object are normally responsible for certain kinds of experience then we cannot avoid appealing to past perceptual beliefs concerning our encounters with particular objects of that kind as part of our justification for these beliefs. If this is so, it suggests that perceptual beliefs about particular objects must ground our general beliefs about the causal connections between types of experience and the types of

things in the world, which cause them rather than vice-versa (Martin 1995: 42). This is a problem assuming that the relevant bridging principles figure as premises in an inference to some conclusion about the world. If Martin is right, our justification for these premises is not independent of the conclusion it is meant to establish. This is a possibility we earlier rejected in connection with the epistemic regress.

Another possibility, though, is that the relevant inference is abductive: the existence of such and such state of affairs in the world being held to be the best explanation of our enjoying the mental goings-on that we do. Here, the problem is to say why that hypothesis - that things in the external world are thus and so - should really be the 'best' explanation of the course of one's mental life. Ruling out other explanations, like the hypothesis of a Berkeleyan God keeping the perceptible world in existence from one moment to the next, is a familiar problem which it unclear how the traditional foundationalist can hope to solve.

These questions lack satisfactory answers. The literature documenting why that is so is vast and it is not one to which I intend to add. My claim is that even if the traditional foundationalists did have a satisfactory account of what justifies beliefs about our own psychological states, they lack a satisfactory account of what justifies beliefs that aren't about our own psychological states. This makes traditional foundationalism a fairly radical form of scepticism since most of our beliefs fall into the latter category and that is a position we have reason to reject.

However, historical foundationalism doesn't just fail on its own terms. Even if that position could successfully reconstruct our justification for beliefs about the mind-independent world using only the resources it permits itself viz. premises about the mind-dependent world, the very idea of such a reconstruction is independently objectionable. It is objectionable since it is not as if we do not ordinarily have views about what actually justifies us in such beliefs. These views are wildly at odds with what the historical foundationalist says. So that account is an essentially revisionary account. This is another reason to reject it unless we have some special reason to suppose our ordinary views are mistaken on this score.

To see this, take a case in which I stop you on the street to ask the whereabouts of a shop that I am interested in finding. Suppose, being a local, you know and give expression to that knowledge, and that I simply take your for word for it and believe the shop to be where you have told me. In this case, and others like it, it is very natural to think that what justifies me is simply the fact that you have told me where the shop is.³² We could try and reconstruct my justification in inferential terms, perhaps using premises about what I take it you have said, and the likelihood that you are not lying and so on. But in many cases it is doubtful that we could do so satisfactorily. It is not often that I will have much justification for the relevant premises and if I am not justified in believing the premises of my argument then I will not be justified in any conclusion that I infer on their basis.³³ But even if we could do this it still wouldn't change the fact that the natural view

³² Or, if one prefers, my having learnt from you where it is (McDowell 1998b).

³³ In a similar vein McDowell writes "if we make the ancillary premises seem strong enough to do the trick, it merely becomes dubious that the tourist has them at his disposal; whereas if we weaken the premises, the doubt attaches to their capacity to transmit, across the argument, the right sort of rational acceptability for believing its conclusion to amount to knowledge...Let it be the most favourable case we can imagine. Let the hearer have all kinds of positive evidence that the speaker is speaking his mind: a steady honest-looking

is that my justification is not inferential. The natural view, and the view we all give in these cases, is that my justification derives from your having told me. This is to appeal to a non-inferential source of justification.

The point I am making here is really very simple. We do have intuitions about how our beliefs are justified and it is normally pretty easy to get people to acknowledge when their justification is tacitly inferential – even in cases in which it mightn't initially seem to be so. There is, however, no reason to suppose that this must be so in the present case; very often the intuition that my justification derives from having been told and not from the tacit availability of an inference is stubborn.

This intuition is even more robust in other cases. Consider a case in which I look out of the window and see that there's a squirrel sitting on the back fence. This is the sort of thing that most of us think we are sometimes in the best possible position to judge – the light is good, there is nothing blocking my view, and so on. Again it might be possible to reconstruct my justification for believing that the squirrel is sitting there using only premises about how things are for me, psychologically speaking, and some general principles linking things being that way for me psychologically and the likelihood things

gaze, a firm dry handclasp, perhaps years of mutual reliance. Surely it is always possible for a human being to act capriciously, out of character? And even if the speaker is speaking his mind, how firm a hold can the hearer possibly have on the premises, needed on this view, that the speaker is not somehow misinformed about the subject matter of the conversation? However favourable the case, can the hearer really be said to know that his informant can be relied on now, in such a way that his verdict can be used in a non-question-begging certification that what he has acquired is an epistemically satisfactory standing? The supposition that the informant is, perhaps uncharacteristically, misleading the hearer or, perhaps surprisingly, misinformed about the topic is not like the typical suppositions of general sceptical arguments (e.g. 'Maybe you are a brain in a vat'), where, it is at least arguable that no real possibility is expressed. In Simon Blackburn's phrase, mistakes and deceptions by putative informants are 'kinds of things that happen'. It is not clear that the approach I am considering can make out the title to count as knowledge of any beliefs acquired from someone else's say-so. And too much overturning of intuitions must surely make it questionable whether the general account of knowledge is a good one' (McDowell 1998b: 418-20).

really are that way in the world external to my mind. This is what the historical foundationalist thinks we must do. But even if we could, it still wouldn't change the fact that this is a very strange and baroque account of what justifies that belief. The natural view (and the one that we all actually give when asked) appeals to what we perceive. Intuitively, it's the fact that I can *see* that the squirrel is sitting on the fence that gives me reason for believing as I do, not beliefs which, in all likelihood I don't actually have about the way things look and how that makes it likely things really are so. The latter is totally unmotivated. The obvious and simple view of these matters is just to stick with what we actually all say when asked justificatory questions and what we say is things like 'because I can see that squirrel is sitting on the fence'.³⁴

So, even if the historical foundationalists could successfully reconstruct our justification for beliefs about the mind-independent world in inferential terms this still would still not be a reason to think that what they say is actually true; that it accurately describes what, in actual fact, justifies those beliefs. Their account is revisionary of our ordinary views about what grounds or justifies those beliefs and that is a reason to reject it all else being equal.

I have now given you some reasons why historical foundationalists shouldn't have

³⁴ When pressed we often retreat to claims about looks. Faced with the question 'But how do you know that it's really a squirrel you see?' we will often say something more guarded, like 'well it at least looks like a squirrel'. But the fact that we retreat to claims like this when under attack does not show they are the grounds upon which we initially judged. Thus, Williamson writes "it is a fallacy to assume that retreats in the face of doubt always reveal a pre-existing structure of justification. Someone may be simultaneously disposed to retreat to premises about appearances if put under pressure by idealists about the external world and to retreat to premises about brain scans if put under pressure by eliminativists about the mind. In responding to a doubt, we look for ground that it does not undermine, but where that ground is depends on the doubt. That we can be made to retreat to a place does not show that it is where we started from" (Williamson, forthcoming).

thought what they actually thought. An obvious question is: did they need to think these things? The epistemic regress argument certainly doesn't commit them to such a restrictive view. The regress argument demands that some of our beliefs be non-inferentially justified, but it doesn't say anything about which beliefs can be justified in that way or what it takes for a source to be capable of furnishing us with such justification. The problem with inference is that it is an essentially dependent source of epistemic justification; in order for it to furnish us with justification it requires antecedently justified beliefs as inputs. So the same had presumably better not be true of a non-inferential source; it must furnish us with justification that does not derive from the justification we have to believe other things. But observation or 'acquaintance' certainly looks like it can play that role. The sort of justification it provides us with does not derive from other justified beliefs; it offers a different model of what it is to be in a position to know about the world.

So if the historical foundationalist are right to think that we do observe our own mental states, they'll be right to conclude that our beliefs about them are non-inferentially justified. Many philosophers think there is no reason to suppose we do observe our own mental states. But there is certainly no reason to think that we only observe our own mental states; or that observation is necessarily the only source of non-inferential justification. If not, there is no reason to believe the historical foundationalist when they say that only beliefs about our own psychological states can be non-inferentially justified.

Things might be different given a different understanding of the regress argument.

Sometimes people think the problem that argument raises is *temporal*. If for every belief that is justified there must be some further belief that one is be justified in believing before one can be justified in the first, how could justification ever get started.³⁵ They conclude that there must be some beliefs that can be justified prior to any other beliefs being justified; there must be some beliefs that we can in that way start from. The problem with beliefs about objects in the world external to our minds is that it doesn't look like a subject can have those beliefs (let alone be justified in them) unless she has lots of other beliefs. I can't believe that the squirrel is on the fence unless I have the concept squirrel and that plausibly requires me to have certain sorts of beliefs about squirrels. I may have to believe that squirrels are animals. This requires I have the concept animal and that plausibly requires me to have yet further beliefs. Beliefs about objects in the world external to our minds therefore do not look like they can stop the regress, if the regress is understood in temporal terms. They aren't beliefs that we can in that way start from.

Traditional foundationalists tended to think that the same was not true of beliefs about our own psychological states (Lewis 1946). They had a certain view of how our psychological concepts get their meaning according to which one could possess them without having to have lots of other beliefs. That meant they could claim that a subject didn't even need to *have* other justified beliefs in order to be justified in believing things on the basis of acquaintance or observation. One could be justified in believing of the red-y brown squirrel-ish shaped sense-data that one currently perceives that that looks brown to me irrespective of what other beliefs one happens to have. This belief would be

³⁵ For a presentation of the argument along temporal lines see (Moore 2002: 122-3).

‘semantically free-standing’ as well as epistemically freestanding. Not only would it not draw its justification from other beliefs, it wouldn’t even require their existence.

This is what Michael Williams is getting at when he says that according to the traditional foundationalists, basic beliefs represent ‘semantically encapsulated items of knowledge’ (Williams 2005: 204). I agree with Williams that this is what lots of traditional foundationalists thought. I also think it would be a reason to follow them in privileging beliefs about the mind if we had any reason to accept this view of concepts or this view of the regress. But we have no reason to accept this view of concepts and no reason to take the temporal regress seriously. Maybe you can’t be justified in believing anything about the world around you until you believe lots of things about it – maybe justification, to that extent, emerges en masse. So what? In order for there to be any justified beliefs there don’t need to be any beliefs that are justified before all the rest. This is just a separate issue from the issue of whether or not all your justification could derive from the justification you have to believe other things. Your justification can be non-inferential, even if it’s not possible to be justified in that way without also being justified in believing lots of other things.³⁶ In that case your justification will not derive from other beliefs, though it will in a sense depend on them (since it will require their existence). This is not a problem though, since it is only the former that we have reason to worry about. It is the

³⁶ The following example from James Pryor nicely illustrates this point: ‘Consider: in order to have the concept of a unicorn I may need to believe (i) that unicorns have hooves, and (ii) that unicorns have horns. Now suppose I acquire evidence that a virus has killed all hoofed creatures. Since I believe unicorns to be hoofed creatures, I form the belief (iii) that no unicorns currently exist. It is clear that (ii) plays no role in justifying this belief. This shows that there can be propositions that you need to believe in order to have certain concepts (you need to believe (i) in order to have the concept of a unicorn) without those propositions mediating your justification for every belief involving the concept. Now (iii) is not an immediately justified belief. But it serves to make my point. We can see the same phenomenon with beliefs that are good candidates to be immediately justified like (iv) If any unicorn exists, it is identical with itself. (ii) plays no more role in justifying that belief than it plays in justifying (iii)’ (Pryor 2005: 198).

derivation of justification that the epistemic regress worries about and we avoid that by requiring that your justification not derive from other beliefs (whether or not it depends upon them).³⁷

So the further commitments of historical foundationalism are not needed to deal with the epistemic regress problem. What, then, could have lead these thinkers away from the obvious view of these matters just sketched and towards their own peculiar alternative? In the previous section I mentioned two possible sources of motivation for that view – one metaphysical, the other epistemological. The metaphysical route to historical foundationalism claims that we don't actually perceive objects in the world around us. Contrary to what we all ordinarily think, we only really perceive things which stand in for or represent those objects - mind-dependent objects, commonly called 'ideas' or 'sense-data'. This would explain why perception cannot non-inferentially justify our beliefs about non-psychological reality. If perception does not make us aware of objects in the world external to our minds, why should it give us any reason to believe that they are one way rather than another? What explanation could we possibly give? An intuitive justificatory gap remains, which gap can only be plugged by appeal to general bridging principles of the sort to which proponents of this view really did appeal.

Despite what proponents of this view thought, however, there is no good reason to think that perception doesn't put us in touch with the mind-independent objects in the world that it seems to, or therefore that we are only ever aware of mind-dependent replicas.³⁸

³⁷ Cf. (Burge 2003: 503-48).

³⁸ I am not denying that they had arguments for this view. I am denying that they had *good* arguments for that view. Most of the considerations adduced in support appeal to facts about illusions and hallucinations

Moreover, even if perception did work in the way they thought, that still leaves open the possibility that there are sources of non-inferential justification other than perception – sources like testimony. As yet, nothing has been said to block that move.

The metaphysical route to historical foundationalism is therefore a non-starter. If we really do perceive objects in the world around us, why isn't perception available to non-inferentially justify our beliefs about them, just as it justifies our beliefs about mind-dependent objects on the traditional account? For those of us willing to accept that we do perceive objects in the world external to our own minds, only the epistemological route to historical foundationalism remains.

The epistemological route appeals to a certain view about non-inferential justification and what it would take for any belief to be justified in that way. As it is standardly reconstructed in the literature the reasoning goes like this: (1) if a belief can be mistaken, it can't be non-inferentially justified, (2) beliefs about the world can be mistaken. So (3) beliefs about the world can't be non-inferentially justified. According to what I'll call the Simple Reading this is precisely the sort of reasoning that lead the historical foundationalists away from the intuitively appealing idea that perception is available to non-inferentially justify our beliefs about the external world. This establishes the negative part of the traditional thesis: beliefs about the world can't be basic. The positive claim

and try to generalise something from that case to the normal case; the time-lag argument and the argument from illusion are both examples traditionally appealed to in this connection (Ayer 1956 is a prime example). For a detailed discussion of the argument from illusion in all its forms see Michael Martin's forthcoming book 'Uncovering Appearances' (forthcoming, OUP). In contrast, Locke thought that introspection revealed that the immediate objects of perception were our own ideas. He writes: '*What Perception is, ever one will know better by reflecting on what he does himself, when he sees, hears, feels, etc. or thinks, than by any discourse of mine. Whoever reflects on what passes in his own Mind, cannot miss it.*' (Locke 1975: 143).

that beliefs about one's own psychological states can be basic then goes through provided only that one assumes that the same is not true of the latter and this is precisely what thinkers like Russell and Ayer thought.

The Simple Reading dominates the literature and the view it gives expression has the status of orthodoxy. I think that it is mistaken as a reading of traditional foundationalism. Thinkers like Ayer were ultimately more concerned with the conceivability of mistakes, than their possibility. Nonetheless, there is clearly some point to the idea that what motivates denying that beliefs about non-psychological reality can be non-inferentially justified is the thought that non-inferentially justified beliefs would have to be peculiarly epistemically privileged. The Simple Reading is just the most popular way of spelling out what these privileges would amount to.

In chapter 4 I will argue that the relevant privilege is indefeasibility: it's the fact that beliefs about our own psychological states are indefeasible rather than the fact that they are infallible which accounts for their special historical status. Properly diffusing the epistemological motivation for historical foundationalism will thus have to wait till later. At first glance, however, this line of thought does little to motivate the move away from the obvious sounding thought that perception is available to non-inferentially justify beliefs about the world. Why should we accept that the mere fact a belief is capable of being mistaken means that it must be inferentially justified? Even if we were to assume that what justifies a belief must rule out the possibility the belief is mistaken, and so conclude that (insofar as perception fails to do that) perception fails to justify our beliefs

all by itself, it still wouldn't follow that perceptual justification is tacitly inferential. To say that a belief is inferentially justified is to make a claim about the nature of the positive support it enjoys and where that support comes from; it says that it comes from its inferential relations to other justified beliefs. But even if we'd conceded that perception cannot justify our beliefs about the world all by itself, so that something else must be necessary, it wouldn't follow that what more is necessary is other justified beliefs. Much less that perception itself plays no role at all, and is entirely supplanted by beliefs about what one seems to perceive and inferences one can make from these to beliefs about the world. This is what would need to be true for one's justification to be inferential. To be inferential one's justification must derive exclusively from the inferential relations one's belief bears to other justified beliefs. This is where the traditionalist foundationalist assumes that one's justification must come from in this case, but so far no good argument for that view is in sight.

We will return to the epistemological roots of historical foundationalism in chapters 3 and 4. At this stage I merely hope to have put the traditionalist on the defensive and shown that theirs is a position we have good reason to reject all else being equal. Given that this is so, someone might well ask: why bother taking it seriously to begin with? This is a good question. The short answer is that traditional foundationalism gets something fundamentally right. Spelling this out is the aim of the final section of this chapter.

4. The Idea of the Empirical

What does historical foundationalism get right? What it gets right, I suggest, is the idea

that perception or observation is a source of justification that is genuinely distinct from inference. That is to say, that one is in a position to justifiably make judgments about things in virtue of being aware of them, and that we cannot explain why that is so via any sort of analogy with the inferential case.

But is it true that perception is a source of justification that is distinct from inference? It might be objected that this couldn't possibly be true because perception itself always involves inference and that it follows from this that perceptual justification can't be distinct from inferential justification. In what sense, however, is inference always a component of perception? One possibility is that perception involves conscious inference from beliefs, say beliefs about how things appear to one. This really would be a threat to the claim that there is a sharp distinction between perceptual and inferential justification but there is no reason to think that perception always, or even commonly, involves inference in this sense.

A different possibility is that perception always involves some form of unconscious processing or 'binding'. This is not something that either the historical foundationalists or I need to deny. Perception might be 'inferential' in *that* sense, but this has little bearing on the intuitive distinction between inferential and non-inferential justification. A genuinely inferential justification is one that proceeds from premises that the subject doing the inferring believes. There is no suggestion that sub-personal binding or information processing proceeds in this way. To describe perceptual justification as inferential solely on the basis of the involvement of sub-personal processing in perception

is to deprive the notion of inference of its usual connotations. If one is prepared to do that then one can of course insist that there is no real difference between perceptual and inferential justification. The problem, however, is that there is a perfectly intuitive distinction between these two forms of justification and that is a good reason not to think of what our perceptual systems do as ‘inferring’. When I see that there is a squirrel on the fence and I don’t in any interesting sense infer that there is a squirrel on the fence inferring is the sort of thing that one does when one is presented with signs of the presence of a squirrel not when one is presented with the squirrel itself.³⁹

The view that perception provides a fundamentally different model of what it is to justifiably make judgements in this sense is very intuitive. Some prominent theories deny that is so, but they are very hard to believe. Coherentists think that our only real model of epistemic justification is provided by inference, broadly conceived (Davidson 2000: 154-163). They think our grasp of what it is to be justified is provided by the thought that there exist logical relations within our system of beliefs. Such theories seem incredible precisely because there is nothing intuitive about restricting justification in this way or ignoring the obvious fact that we do have other models of justification one of which is provided by perception. Any case in favour of a more restrictive view must therefore rest upon philosophical argument. As we will see in the next chapter the arguments just aren’t that good and that just leaves the presumption in favour of perception standing.

This presumption is one to which the historical foundationalist ultimately fails to do justice. Still, it remains one to which they try to do justice. As I have characterised

³⁹ For a different view see (Harman 1973: esp. Ch. 11).

foundationalism, it is one to which any foundationalist must try to do justice since that is just what (b) commits them to. As I characterise it, foundationalism isn't just a view about the structure of epistemic justification as (a) claims, it is also a view about its sources. Specifically, it is the view that *perception* is a source of non-inferential justification. (a) merely draws a distinction in abstract space – it claims there is a distinction between inferential and non-inferential sources of justification. (b) says the second class is not empty; it says that perception is a non-inferential source of justification.

I have claimed that this is an overwhelmingly natural view that is as important in motivating foundationalism as the regress argument. Indeed in one sense it is simply the flip side of that coin, since it is an account of the source of the justification that the regress says must exist. Without such an account foundationalism looks half-baked or incomplete. That is why in chapter 1 I said that foundationalism is an intuitive view. As I characterise it, the desire to acknowledge the seemingly obvious fact that perception is a source of justification distinct from reasoning is essential to foundationalism. I am going to call this idea the Idea of the Empirical.

This idea has shaped epistemological reflection since Aristotle, and it is crucial to motivating foundationalism. Yet it is almost universally ignored in the literature. Pollock and Cruz are one of just a handful of commentators to recognise the role it plays in motivating foundationalism. They write:

The simple motivation for foundations theories is the psychological observation

that we have various ways of sensing the world, and that all knowledge comes to us via those senses. The foundationalist takes this to mean that our senses provide us with what are then identified as epistemologically basic beliefs. We arrive at other beliefs by reasoning (construed broadly). Reasoning, it seems, can only justify us in holding a belief if we are already justified in holding the beliefs from which we reason, so reasoning cannot provide an ultimate source of justification. Only perception can do that. We thus acquire the picture of our beliefs forming a kind of pyramid, with the basic beliefs provided by perception forming the foundation, and all other justified beliefs being supported by reasoning that traces back ultimately to basic beliefs (Pollock and Cruz 1999: 29).⁴⁰

The most straightforward way to accommodate this insight is to think that perception is a source of non-inferential justification and that what it non-inferentially justifies us in believing depends on what we actually perceive. Hence, given that we do perceive mind-independent objects in the world around us, perception is available to non-inferentially justify our beliefs about them.

I have claimed that the historical foundationalists sort of saw that but in their case that insight was twisted. Under the weight of their extraneous metaphysical and epistemological views the original idea was transformed into something importantly different and much less plausible. In their case it is not perception as we ordinarily think of it that plays the ultimate grounding role; it is a special sort of perception ('acquaintance'), or the perception of a special sort of object (our own 'ideas').

⁴⁰ Notice that in this passage they are also making a stronger point: viz. that perception is the 'only' source that is privileged in this way. This is not a view to which I commit the foundationalist. (c) merely claims that perception is 'a' basic source of non-inferential justification.

Still, even here the centrality and importance of ordinary perception is not completely lost sight of. This is straightforwardly so on the metaphysical view, since that is just meant to be an account of ordinary perception, but one can see it even on the more epistemologically motivated versions of the view.⁴¹ Russell is perhaps the thinker whose view is maximally unfavourable to the one that I describe. He certainly eschews talk of ‘perception’ in favour of talk about ‘acquaintance’. But while this ought to be a label for a special epistemological relation that one can stand in to a state of affairs, it is clear that Russell effectively models acquaintance upon ordinary perception. The latter provides his only really worked out picture of what acquaintance might be.⁴² Moreover, features of ordinary perception play an indispensable role in explaining why it is that acquaintance should justify a subject in believing anything whatsoever. It is because acquaintance affords us conscious awareness of those states of affairs it purports to justify us in judging about in the way that we naively suppose perception does; and because we form our beliefs on that basis, that we are justified in judging as we do. Ordinary perception therefore functions as the paradigm upon which these other notions are modelled, even for someone like Russell.

This is importantly different from the way in perception may be a paradigm for a reliabilist or on other non-foundational theories. It may well be true that among the reliable belief forming mechanisms that we actually have perception is the one with

⁴¹ Clearly, they also took themselves to be offering an account of how one is in a position to know about what one believes, say.

⁴² Russell also thought that we were acquainted with universals and my analysis is harder to apply to that cases since Russell certainly did not think that we perceived universals (Russell 1912: esp. 28). But universals are normally thought to be a problem for his view. The fact they sit rather awkwardly may well be a reflection of the fact that Russell is effectively working with a perceptual model.

which we are most familiar, or best equipped to understand, or know the most about. It may be for that reason that it functions as a ‘paradigm’ in reliabilist theories of knowledge; it is certainly the example most of them tend to use. However, the reliabilist wants to abstract away from features of perception other than its reliability. Perception is only a paradigm because of what it illustrates about reliability – other features of it are irrelevant to saying why it should be a source of justification. This is possible because perception does not provide us with our basic picture of what justification is like for a reliabilist. Our basic picture of justification is furnished by the idea the idea of a reliable belief forming mechanism and that is an entirely general claim (Goldman 1986). It is one that we could grasp independently of thinking about perception or any other specific source, even if (in the genesis of understanding), it is one that we only actually come to see is true by reflecting upon particular cases.

This is fundamentally different from the role that I suggest perception plays for the foundationalist. For the foundationalist, perception is not just illustrating a general moral. Its status as a source of justification is *sui generis*. That is part of what I mean when I say that perception is not just *a* source of non-inferential justification but a *basic* source of non-inferential justification. This is what (c) claims and that is the third component of the position I call foundationalism. Further discussion of (c) will have to wait until chapter 5. The important point at this stage is that we can make sense of the idea that perception is a source of non-inferential justification (and a potentially basic source at that) even as far as the historical foundationalists are concerned.

5. Conclusion

The next chapter will offer a positive account of how perception can be a source of non-inferential justification for beliefs about the world external to our minds. You don't have to be a reductionist to think that a foundationalist must say something further about why perception should be a source of justification for beliefs about what we perceive. Here, the historical foundationalist may be thought to be at a distinct advantage since on their view it is impossible for you to be mistaken in the beliefs that you form on the basis of 'perception'. The same presumably isn't true if we take perception to involve a relation to mind-independent objects. Beliefs about non-psychological reality can be mistaken. According to the Simple Reading that is precisely what launches the retreat inwards towards beliefs about our own psychological states with all its attendant difficulties. This is the orthodox view one finds in the literature. The next chapter will look at what a more modest foundationalist can say in response. As we will see, we needn't be quite as modest as the orthodox would have us believe.

FOUNDATIONALISM AND THE IDEA OF THE EMPIRICAL

CHAPTER 3: THE ARGUMENT FROM FALLIBILITY

1. Introduction

The take home message of the previous chapter was that there's a good idea underlying traditional foundationalism, but that traditional foundationalism goes about developing this idea in the wrong way. The good idea is that perception is a distinctive source of non-inferential justification. Where traditional foundationalism goes wrong is in its account of the sorts of beliefs that perception can justify. Traditional foundationalism has a narrow conception of the scope of perceptual justification; it claims that perception can only justify beliefs about psychological reality – that is, beliefs about our own ideas or sense-data - and this is what ultimately gets that position into all the trouble previously discussed.

If it's a good idea to think that perception is a source justification but a bad idea to adopt a narrow conception of the scope of perceptual justification, there's an obvious way of holding onto the good idea without holding onto the bad idea. The obvious way is to adopt a broad conception of the scope of perceptual justification. The Broad View, as I am going to call it, says simply that among the beliefs that perception can non-inferentially justify are beliefs about the world around us. When I talk about beliefs about the world around us I mean beliefs about non-psychological reality - for example, the belief that the squirrel is on the fence, or that the toast is burning.⁴³ Although this seems

⁴³ Berkeley also thought that perception was a source of justification for beliefs 'about the world around

obviously true, it is something lots of philosophers have denied.

There are two important components of the Broad View. One is its conception of the *scope* of perceptual justification, and the other is its insistence that the kind of perceptual justification it is talking about is *non-inferential*. What I want to do in this chapter is to defend both elements of the Broad View. I want to do that because I think that the Broad View is correct. But it's also worth pointing out some of the other advantages of that view.

One is that it is at least arguably the ordinary or naïve view of perceptual justification. To the extent that we ordinarily have views about such matters it would seem that we have no difficulty with the idea that among the beliefs that perception can non-inferentially justify are ordinary beliefs about the world around us. So we could also call the Broad View the 'naïve view'. By that I mean it's the view that seems most natural and obvious to us prior to philosophical reflection. Obviously the fact that it is the naïve view is not a knock down argument in its favour. However, it does mean that the Broad View is much less revisionary than the traditional foundationalist conception of these things, and that's a good thing.

A further advantage of the Broad View is that it doesn't end up positing excessively narrow foundations – that is, foundational beliefs that are so restrictive in their scope as to be incapable of supporting the rest of what we know or justifiably believe. It thereby

us', he just thought that the world around us was mental or psychological (Berkeley 1964). I here mean to rule out that possibility.

promises to form the core of a more modest form of foundationalism – one that doesn't have the sceptical consequences of traditional foundationalism. Again that is an advantage not to be sniffed at.

Having said all of this the most important question about the Broad View is whether it is actually true, not whether it's what we ordinarily think anyway. I think that the Broad View is correct and I'm going to explain why I think that in this chapter. But many philosophers – including those who are not foundationalists at all – have thought that the Broad View just can't be right and that there are decisive philosophical reasons for revising what we ordinarily think.

What are these allegedly decisive objections to the Broad View? One is that perception only fallibly justifies beliefs about non-psychological reality and that perceptual justification therefore can't be non-inferential. I'm going to call that the Argument from Fallibility. Another, is that perception only defeasibly justifies beliefs about non-psychological reality, and that perceptual justification therefore cannot be non-inferential. I will call this the Argument from Defeasibility. According to these arguments if one wants perceptual justification to be both infallible and indefeasible then it had better pertain to a very special subject matter. This is just what the traditional foundationalists thought, the special subject matter being psychological.

Given these objections to the Broad View anyone who wants to defend that view is going to have to do several different things. The first thing which needs to be done, and which

I'll be doing in the next section, is to explain in what sense perception can non-inferentially justify beliefs about the world around us. The second thing that one would need to do in order to defend the Broad View is to rebut the philosophical arguments against it. In the final part of this chapter I will tackle the Argument from Fallibility against the Broad View. One way of tackling this argument is to deny its premise: that is, to deny that perceptual justification is fallible, even if its subject matter is non-psychological. A more familiar strategy would be to accept its premise but to dispute that the conclusion follows; in other words, to argue that perceptual justification can be both fallible and non-inferential. I have some sympathy for both these responses to the fallibility objection and will explain why later on in this chapter.

The Argument from Defeasibility against the Broad View will be the topic of the next chapter. Again the options are to deny that perceptual justification is defeasible, or to deny that it's being defeasible entails that it's inferential. In this case, only the second option looks plausible. So to recap: the object of the exercise in this chapter is to spell out and defend a broad conception of the scope of perceptual justification that is different both from traditional foundationalism's narrow conception of perceptual justification, and from non-foundationalist conceptions of perceptual justification, such as those of Donald Davidson and Laurence Bonjour.

The plan for this chapter is as follows: in the next section I'll give a more detailed account of the Broad View and explain the sense in which it delivers non-inferential perceptual justification. In the following part I will tackle Davidson and Bonjour, and

show that their conception of perceptual justification, to the extent that they have one, is inferior to mine. Then, in the final part, I will deal with the Argument from Fallibility.

2. Propositional Perception

In the last chapter I claimed that a good, everyday answer to the question of why some subject believes as she does, or what justification she has for that belief, will often appeal to the fact she can see that things are so. Suppose Ann is doing the washing up and glances up to look out of the back window. When she looks out she sees that there is a squirrel sitting on the fence. If Ann believes that there's a squirrel on the fence, on the basis of what she can see, a perfectly acceptable answer to our justificatory question will cite the fact that Ann can *see that* there is a squirrel on the fence: that is what makes her justified in believing that there is a squirrel on the fence. In all likelihood that is the answer that Ann herself would give us if we asked her.⁴⁴

Seeing that the squirrel is on the fence is a case of what Fred Dretske called 'epistemic seeing' (Dretske 1969: Ch.3). Epistemic seeing, in turn, is a case of what we might call 'epistemic' or 'propositional' perception.⁴⁵ One can see that the squirrel is on the fence, but one can also *hear that* Ross is at the party - as when one overhears him talking to the host - and similarly for the other three sense modalities: one can *feel that* the dog is wet, *smell that* she is wearing perfume, and *taste that* the water is salty. Intuitively, just as we can explain what justifies Ann in believing that the squirrel is on the fence by citing the fact she can see that the squirrel is on the fence, so too can we explain my justification for

⁴⁴ The reply, 'Because, I can see that there's a squirrel' is often given in response to what many sees as a more demanding epistemological question, viz. 'How do you know that there's squirrel on the fence?'

⁴⁵ This label is Cassam's, not Dretske's (Cassam 2007: esp. 27-70).

believing that Ross is at the party by appealing to the fact I can hear that he is at the party. These sorts of explanations are completely commonplace. So the naïve or ordinary view at least, seems to be that epistemic perception in this broad sense, and not just epistemic seeing, is a way in which we can have justification for our beliefs about the world.

How does that bear on the question of whether or not perception can be a source of non-inferential justification for beliefs about non-psychological reality as the Broad View maintains? Well, pretty directly - since epistemic perception is a form of perception (even if it is not the only form), and the sort of justification that it provides is non-inferential. What makes it non-inferential is that it appeals to the fact that the subject is *in* a certain perceptual state – a state in which she perceives that something is the case. It does not appeal to her beliefs about what sort of state she is in, or inferences she can make from those beliefs to a belief about the world. Ann is justified because she can see that the squirrel is on the fence, not because she believes she can see that the squirrel is on the fence. Indeed, in order to see that the squirrel is on the fence, or, more generally perceive that *p*, it is neither necessary nor sufficient that one believe that one perceives that *p*. So the sort of justification that seeing that *p* - and epistemic perception more generally – provide is not inferential; it does not derive from the justification the subject has to believe other things.

This is often overlooked in the literature. Two things make that easy to do. One is that talk of “the fact that” the subject sees that *p* sometimes misleads people – as if that fact were somehow different from, or over and above, the subject’s simply seeing that *p*. It is

not. The second and more compelling explanation for the oversight appeals to the dialectical context in which questions about justification normally get raised. Thus, suppose I ask Ann what justification she has for believing that the squirrel is on the fence, and she tells me, saying 'I can see that the squirrel is on the fence'. In making that assertion Ann thereby gives expression to one of her beliefs - namely, a belief about what she thinks makes it the case that she is justified. This is an inevitable consequence of making sincere assertions, but it leads some people to think that, if Ann speaks truly, what makes her justified isn't really her seeing that the squirrel is on the fence but rather her belief that she can see that the squirrel is on the fence.

Again this is a mistake: if Ann speaks truly, what justifies her in believing that the squirrel is on the fence is the fact she can see that the squirrel is on the fence, not her belief that she can see that is so. This is what her explanation actually cites and any impression to the contrary is just an artificial product of the dialectical context. Indeed, if this sort of example were sufficient to show that Ann's justification derives from her beliefs, then all justification would trivially so derive since in responding to justificatory challenges and making claims about what justifies one, one trivially gives expression to one's beliefs about what justifies one. If this is the reason for thinking that all justification is 'inferential' it is not a very interesting one.

So we shouldn't be misled by these facts into supposing that the sort of justification that epistemic perception provides is tacitly inferential. Are there any other grounds for doubt? Epistemic perception is conceptual: in order to see that the squirrel is on the fence

one needs the concepts that figure in the that-clause. One cannot see that the squirrel is on the fence if one lacks the concept squirrel. I am not claiming that all perception is conceptual in this way, or even that all perception insofar as it is a source of justification must be conceptual. I am merely claiming that epistemic perception is conceptual and that ought to be relatively uncontroversial.⁴⁶

Does this raise a problem for the idea that epistemic perception can be a source of non-inferential justification? Many have thought that it does. After all, having concepts often involves having beliefs. Plausibly, I do not count as possessing the concept squirrel unless I have certain kinds of beliefs about squirrels; I may have to believe that squirrels have tails. So if Ann cannot see that the squirrel is on the fence without having the concept squirrel, then she cannot see that the squirrel is on the fence without believing that squirrels have tails. Any justification that Ann gets for the belief that the squirrel is on the fence by seeing that the squirrel is on the fence therefore depends upon the fact that Ann also believes that squirrels have tails. Ann's justification therefore isn't independent of her beliefs, but depends upon them, at least to this extent.

Lots of people think this a problem, since they assume it means Ann's justification must be inferential. This is a mistake, but it is an easy one to make given how some people

⁴⁶ Epistemic perception is not necessarily conceptual on all readings of 'conceptual'. I am claiming that it is conceptual because you need concepts in order to be in a state in which you perceive that p; I think that is pretty uncontroversial. A different reading of 'conceptual' has it that for a state to be conceptual it must be 'composed' of concepts. I am silent about whether or not epistemic perception is conceptual in that sense. On the face of it, it is hard to see how it could be. One's mental state itself is presumably not composed of concepts. A different suggestion is that the mental state that is 'perceiving that p' is an attitude to a proposition, and that propositions are composed of concepts. In that case what you perceive is composed of concepts, even if your mental state itself is not. Whether epistemic perception is conceptual in any of these further senses depends on whether it really is an attitude to a proposition and if so, whether propositions are composed of concepts in the relevant sense. These are issues I cannot hope to resolve in this thesis. For more on this debate, see (Peacocke 1992).

define ‘inferential’ and ‘non-inferential’. However, in cases of inferential justification one’s justification doesn’t just depend upon one’s beliefs, it derives from them. When I believe that England can no longer win the Ashes because I believe they have already lost 3 of the 5 tests, my justification for believing they can’t win doesn’t just depend on the fact I believe they’ve lost 3 of the 5 tests; it comes from that belief. The latter belief is the source of my justification – it is where I inherit, or get, my justification from - and if it is not justified then my belief that they can no longer win will not be justified either.

This is clearly not what is going on in the case in which Ann is justified in believing that the squirrel is on the fence because she can see that the squirrel is on the fence. While it has to be true that she believes that squirrels have tails this belief plays no role in justifying her belief that the squirrel is on the fence; it is not part of what confers justification on that belief. Intuitively, it is no more part of the source of Ann’s justification than any of the very many other things that also have to be true for Ann to see that the squirrel is on the fence. Ann must exist, she must have properly functioning eyes, and a fence must have been erected in the garden at some stage. These are also all things which have to be true for Ann to see that the squirrel is on the fence but nobody would think that they play any role in justifying her in believing that there is a squirrel on the fence. Exactly the same is true of her belief that squirrels have tails.

More generally, we can distinguish between sources of justification (those things which, intuitively speaking, confer justification upon a subject), and those things which merely have to be true for the subject to be justified, or background conditions as I will call

them. The latter enable the subject to take advantage of the justification on offer to her, without themselves being what is justifying her. In Ann's case the source of her justification is her seeing that the squirrel is on the fence. In the cricket case, the source of my justification is my belief about England's poor performance in the first three tests, and the fact it supports a further belief. That is where the subject's justification comes from in these cases; they are its sources. On the other hand, there are all those things which merely have to be true for those subjects to be justified; the squirrel must have somehow found its way onto the fence and cricket must be a game that England can just about play. These aren't things which confer justification on their beliefs; they aren't sources of justification, they are mere enabling conditions.⁴⁷

What cases of inferential justification make clear is that beliefs can play the first role. Beliefs can be one's source of justification and one can inherit one's justification from them. In contrast, what the present example of Ann makes clear is that beliefs can *also* play the second role: they can be mere background conditions. In that case, there is no reason to think the justification involved is tacitly inferential.

There is nothing tricky about this distinction. I am not trying to pull the wool over anyone's eyes. It is a perfectly obvious and intuitive distinction to draw and one that we draw even in cases of inferential justification. To be justified in believing that England have lost 3 of the 5 tests, I need the concept England. That may require I have certain sorts of beliefs about England; I may have to believe that England is a nation. Still, this belief does not justify me in believing that England cannot win the Ashes, it's my belief

⁴⁷ For further discussion of this distinction see (Burge 1993) and (Cassam 2007: esp. 1-50).

they've lost 3 of the 5 tests that does that.

So epistemic perception may depend upon beliefs (since it is conceptual and having concepts plausibly requires beliefs); but that doesn't mean that it derives from beliefs, or, therefore, that it is inferential. Moreover, there is nothing objectionable about dependence per se. As we saw in the first chapter it is the derivation of justification from one belief to the next with which the epistemic regress is concerned. This is what the regress argument focuses on since it is only in these cases that we are threatened with a vicious regress of justification. If one's justification for believing *p* derives from the belief that *q*, then not only must *q* be justified, but any justification one has for believing *q* must be antecedent to one's justification for believing *p*. If it is not, then one's justification will be vitiatingly circular and this is precisely what makes the regress vicious.

The same is not true in cases of mere dependence. Suppose I can't be justified believing that I am in pain unless I also believe that someone is in pain (perhaps, as some have claimed, having the latter belief, or being disposed to infer it, is a condition on possession of the concepts requisite for believing that I am in pain.) In that case, being justified in believing that I am in pain depends upon my believing that someone is in pain. It must be true that I have the latter belief. But it is not plausible to require both that this belief be justified and that any justification I have for it be antecedent to my justification for believing that I am in pain. There is absolutely nothing objectionably circular about getting justification for believing that someone is in pain by inferring that is so from one's own case. In that case, unlike in the genuinely inferential case, one's justification

will not be vitiatingly circular and the regress will not extend viciously backwards.

So not only is the sort of justification that epistemic perception provides not inferential. There is no other reason for a foundationalist, or anyone else interested in stopping the epistemic regress, to find it objectionable either.⁴⁸

Epistemic perception is therefore one source of non-inferential justification for beliefs about the world around us, and that is enough for the purposes of defending the Broad View with which this chapter began. It is enough, since epistemic perception can be a source of justification for beliefs even where their subject matter is non-psychological. As we have seen it can be a source of justification for beliefs about squirrels, and people at parties, and other objects and events in the world around us.

This is one way in which to defend a broad view of the scope of perceptual justification. This view, in turn, promises to constitute the core of a more modest form of foundationalism. It is not obvious though that this is the only way in which to defend the Broad View. In the next section I'm going to consider an alternative strategy by considering an objection to the present line of thought. The objection, which derives from an argument of Davidson's, is that the present line of thought is still too concessive to views on which all justification is inferential.

⁴⁸ I am not suggesting that this is an obvious mistake or one that no one has ever made. On the contrary, this is undoubtedly one of the reasons traditional foundationalists wanted to privilege beliefs about our own psychological states. They thought they had to appeal to semantically encapsulated items of knowledge since any sort of dependence upon beliefs would be problematic. I am merely claiming that it is a mistake. It is certainly not essential to foundationalism.

3. Non-Propositional Perception

Davidson famously once claimed that only a belief can justify another belief. Writing in defence of that theory he claims:

What distinguishes a coherence theory is simply the claim that nothing can count as a reason for holding a belief except another belief. Its partisan rejects as unintelligible the request for a ground or source of justification of another ilk (Davidson 2000:156)

What Davidson is saying here is that all justification is doxastic – it all derives from beliefs. On a strict reading of what he is saying, I have already refuted his view. Seeing that the squirrel is on the fence is not a belief, but it can still justify Ann in believing that the squirrel is on the fence. So what Davidson says is false: not all justification does derive from beliefs.

There is, however, a more charitable reading of what Davidson is saying. On this more charitable reading what he is saying is that justification is a relation that can only obtain between mental states with propositional content. Propositions can stand in what he calls ‘logical relations’ to one another – that is, inferential relations like entailment and probabilification. On this more charitable reading of Davidson’s view, justification is all still inferential, since it all derives from the inferential relations between propositions. It is just that it needn’t all be doxastic since it needn’t necessarily all derive from relations between propositions that are believed.

This is what Davidson is saying on the more charitable reading of his position. In fact

other remarks make it clear this really was his view, despite the provocative and oft-quoted slogan with which we began. An obvious worry about my position therefore is that although what I say is inconsistent with Davidson's view on a strict reading, it is not inconsistent with what he says on the more charitable reading. So Davidson could just accept everything that I have so far said about epistemic perception and its being a source of justification, but still insist that that's consistent with thinking that all justification is inferential, contra the Broad View.

I think that this is not a genuine worry about my position. There are at least three reasons why it is not. First, when I say that perceptual justification is a source of non-inferential justification I do not mean that it is not a relation between mental states with propositional content. What I mean when I say that perceptual justification is non-inferential is that it is not a relation between propositions believed and that it doesn't have the form of a move from premises to conclusion. This can be true even if epistemic perception is a mental state with propositional content.

On my view it is not the proposition that Ann sees that the squirrel is on the fence which makes her justified in believing that the squirrel is on the fence. Further, seeing that the squirrel is on the fence is not a premise in an argument giving Ann justification. It is the experience itself which makes her justified (whether or not this experience is a mental state with propositional content). Of course, someone else reporting Ann's epistemic position would have to use a proposition to specify Ann's epistemic predicament and what it is that she thus sees. They could point out that what Ann's sees in this sense

entails what she believes since what she sees is what she believes viz. that the squirrel is on the fence. But when Ann sees that the squirrel is on the fence she does not make a transition of that sort, from a premise about what she sees to a conclusion about the world. It is the perception itself – her seeing that the squirrel is on the fence - and not some description of it, which justifies her in believing that the squirrel is on the fence.⁴⁹

In the sense in which Davidson can agree with me, then, that is not a threat to the claim that perception is a source of non-inferential justification - as I understand it. Moreover, it is only as I would have us understand that claim that we have any reason to think not all justification can be inferential and that some of it must be non-inferential. This is what the regress argument teaches; there is no parallel reason to think that all justification cannot be inferential in Davidson's sense, so this is not a possibility we need to rule out.

Second, there is in any case a more obvious thing to say about all this. I don't think it is obviously correct to think that epistemic perception is propositional. Still, if it is propositional the sense in which that's so is very different from the sense in which *belief* is propositional. The real point of Davidson's picture though, is to assimilate them; to represent perceiving that p as relevantly similar to believing that p from an epistemological point of view. On my view they are really very different, so it is a mistake to think that our two pictures are fundamentally alike.

⁴⁹ Any philosophical explanation of why seeing should be a source of justification will presumably appeal to facts about seeing, and these can be expressed using propositions and arranged in the form of a philosophical argument like the one I am now giving. This is true of anything that might be said to justify a subject (whether or not it is itself propositional) provided only that some minimal philosophical explanation is possible. This is not a reason for thinking that all justification is propositional in any interesting sense.

What I mean is this: beliefs represent the world as being a certain way and they are capable of conferring justification insofar as we have some reason to believe that the world is as our beliefs represent as being. The world might not be the way our beliefs represent it after all and that is precisely why reasons are required. Nonetheless, representing the world in one way will entail or make it probable the world is a whole host of other ways. If I believe that Ross will be out of surgery before 1pm, what I believe entails that Ross will be out of surgery before 2pm. So if I have reason to believe the world is as I initially represent it to be, I will also have reason to believe it is those further ways it can truly be described as being. This is how beliefs extend the reach of what we are justified in believing.

But epistemic perception is not a source of justification because it represents the world as being a certain way, or represents it in anything like the way that beliefs do. Intuitively, perception is a source of justification because it actually puts us in touch with the objects about which we judge – it presents them to consciousness in a way that mere thought or imagination fail to do. So it is no good trying to understand why perception is a source of justification along the lines that we have for beliefs, even if they are both mental states with propositional content.

In fact, when you start thinking about the differences between, say, seeing that the squirrel is on the fence and believing that the squirrel is on the fence, even the idea that the former is a mental state with propositional content starts to come under pressure. Seeing that *p* is often described as a so-called ‘propositional attitude’; where other

examples of such attitudes are: hoping that p, fearing that p, and of course, believing that p. Presumably, the point of describing seeing in these terms is to capture the idea that what you see can (in some sense) be the same as what you believe. This is certainly true in the following sense: we can often describe what is seen using a proposition that could equally well express what is believed, hoped, or feared. You see that the squirrel is on the fence and that can also be something that you believe, hope, or fear. In another sense, though, what you see is not at all the same as what you believe. What you see is the squirrel on the fence – that is, some concrete 3-D scene in the physical world around you - and that is not in any obvious or natural sense ‘what’ you believe.

I am not denying that epistemic perception is propositional in any sense; I am just claiming the sense in which that is so is very different from the sense in which the more canonical attitudes like belief are propositional. Given these differences it is hardly surprising that there is an epistemological difference between them. This is enough to mark a significant difference between my picture and the Davidsonian picture, since his picture is essentially an assimilationist one: he wants to assimilate perception and belief. On my view, that is a mistake. They are really very different - so much so in fact, that even lumping them together under the single heading ‘propositional attitudes’ is somewhat suspect.

This is the second reason not to worry about whether or not what Davidson and I say is consistent. Lastly, and most obviously, when I say that perception is a source of justification I don’t just mean that epistemic perception is a source of justification. On my

view there is nothing wrong with thinking that *non-epistemic* perception is a source of justification. This kind of perception is not propositional, whatever we say about the epistemic case.

What I mean is this: Ann can be justified in believing that the squirrel is on the fence because she sees that the squirrel is on the fence, but she can also be justified in believing that the squirrel is on the fence because she *sees* the squirrel on the fence. Seeing the squirrel on the fence does not merely cause her to believe the squirrel is on the fence – it gives her justification for believing that the squirrel is on the fence. But seeing the squirrel on the fence is not a mental state with propositional content; it is not an attitude canonically ascribed by means of a that-clause. I ascribe it just by saying ‘and Ann sees/saw/is about to see, the squirrel on the fence’. So Davidson is wrong to think that justification is always a relation between mental states with propositional content. It is not and seeing is a counterexample.

Of course, seeing will not always justify a subject in her beliefs about what she sees. If Ann sees a squirrel on the fence this will only justify Ann in believing that there is a squirrel on the fence if certain other things are true. Ann must be able to recognise squirrels when she sees them and she must be such that she could not easily have gone wrong in believing what she sees to be a squirrel. So seeing will only justify a subject in her beliefs in the right circumstances. This does not mean that the obtaining of these circumstances is what is really doing the justifying. Those further things are certainly enabling conditions, but what this means is that they help spell out the conditions under

which *seeing* will justify Ann. If that is true, then it remains true that what is doing the justifying is her seeing and this - to repeat - is not a mental state with propositional content.

Thus, it is not just that I reject the assumption that only beliefs can justify other beliefs as the strict reading of Davidson claims. I am also rejecting the idea that justification is a relation that can only obtain between mental states with propositional content. This is what the charitable reading says Davidson says. So even if you read Davidson charitably my position is still inconsistent with his.

Someone wanting to defend Davidson has basically got three options available to them. The first option is to deny that non-epistemic perception can be a source of justification. The prospects for this response look pretty bleak since the idea that it is, is one with which we are ordinarily perfectly happy. It is as natural to appeal to the fact that Ann sees the squirrel on the fence, as it is to appeal to the fact she sees that the squirrel is on the fence. Anyone wanting to reject the first of these options had therefore better have some pretty good arguments.

Are there any such arguments? Here is what Davidson has to say in defence:

The relation between a sensation and a belief cannot be logical, since sensations are not beliefs or other propositional attitudes. What then is the relation? The answer is, I think, obvious: the relation is causal. Sensations cause some beliefs and in this sense are the basis or ground of those beliefs. But a causal explanation

of a belief does not show how or why the belief is justified (Davidson 2000:157)

This is not an argument. Davidson just denies that non-propositionally structured items like sensations, can stand in justificatory relations to beliefs; he doesn't refute that view. Moreover the embedded suggestion - that we cannot really make sense of the idea such items might justify beliefs (also implicit in the earlier quote where he claims that idea is "unintelligible") is wholly unconvincing. The idea that Ann is justified in believing the squirrel is on the fence because she sees it on the fence is not one that most people struggle to understand, and there is no evidence we secretly translate that explanation into one that appeals to something propositional in form.

Of course when philosophers try and explain to Ann why she is justified – when I give her chapter 3 of this thesis, say, and she reads all about non-epistemic perception and how it gives us reasons for our beliefs about non-psychological reality because it puts us in touch with the objects about which we judge and so on - my explanation is propositional in form; it has the form of premises and conclusions (I hope). All explanations have that form. So if that is all Davidson is saying then what he is saying is unobjectionable. But this is not a good reason for thinking that what justifies Ann is *itself* something propositional - any more than the fact that Gordon Ramsay can explain to me what makes something taste delicious by saying something propositional in form means that what makes it taste delicious it itself something propositional in form.

So Davidson is either saying something true but completely unremarkable or he is saying something philosophical interesting and substantive but false.

Williamson offers a different defence of the same basic idea. He claims that some of the central functions of what he calls ‘evidence’ can only be sub-served by things which are propositional in form. He writes:

Only propositions which we grasp serve the central evidential functions of inference to the best explanation, probabilistic confirmation, and the ruling out of hypotheses (Williamson 2000: 196-7).

Suppose this is true and we assume that what Williamson means by ‘evidence’ lines up with what I am calling ‘justification’. Still, it is unclear why everything that can be evidence must be able to play all the roles that evidence is capable of playing – even all its central roles. A central function of games is to promote team spirit and a sense of collective endeavour. This is a function that cricket sub-serves well, though it is not one that chess, let alone solitaire, do much to promote. This does not mean that chess and solitaire are not games or that we have any great difficulty in understanding how they can be games. On the face of it exactly the same is true in the case of evidence or justification: something can be among the central functions of evidence without its being the case that actually fulfilling that function is essential to something’s being evidence in the first place.⁵⁰

I am not claiming that nothing could lead us rationally to revise the view that non-propositional perception can be a source of justification, merely that we had better have good reasons for doing so, given the presumption in its favour. That restriction is, after

⁵⁰ Matters would be different if ‘central function’ meant ‘essential function’. In that case all the philosophical work remains to be done in showing that the selected functions really are essential. That now looks as hard to defend as the original claim that simple seeing cannot be a source of justification.

all, not one that we ordinarily feel compelled to make. My point is just that so far the reasons offered aren't that good.

The second option for someone wanting to defend Davidson is to accept that non-epistemic perception is a source of justification, but to insist that is only because it is epistemic perception in disguise. Of course, even if that is true it is not a threat to my position for the reasons I went into earlier in the discussion of epistemic perception. But it is not true in any case; there is no good reason to think non-epistemic perception is only a source of justification because it is un-obviously epistemic. For all that has been said so far, the circumstances in which the two confer perceptual justification need not even coincide. However, even if they did, the conclusion still would not follow. On the contrary, we might as well just argue the converse: why isn't all talk of 'perceiving that something is the case' merely elliptical for talk of perceiving in certain circumstances? If so then perceptual justification is *never* a relation between mental states with propositional content since perceiving is not such a state.

If anything, that is the more plausible line to take. As we saw earlier there is something independently odd about thinking that even epistemic perception is propositional. Unlike the more canonical propositional attitudes, seeing seems much too passive for talk of its being an 'attitude' to seem at all natural. Intuitively, seeing that the squirrel is on the fence is not a stance that I actively adopt towards the squirrel's being on the fence in the way in which believing or even hoping that the squirrel is on the fence are. It is simply something that I take in. Moreover while epistemic seeing is relational, what it relates us

to is intuitively not something abstract like a proposition. When one sees that the squirrel is on the fence what one is related to is a concrete state of affairs in the physical world that is literally made up of the squirrel, and the fence, and their relation. As we saw earlier, this is not in any obvious sense ‘what’ one believes.

Finally, epistemic constructions describe a specific way in which the world is perceived: one sees that the squirrel is on the fence. Yet it is natural to think that one and the same episode of seeing can justify an indefinite range of beliefs. Normally when I see the squirrel on the fence I take in more than just the squirrel’s being on the fence and what I see may accordingly justify me in a range of further beliefs – that is, beliefs other than that the squirrel is on the fence. When I see that the squirrel on the fence that very episode may also justify me in believing that Mr Squirrel Nutkins is on the fence, assuming I am sufficiently familiar with Mr Nutkins, the family pet. If so, it seems to be the non-epistemic construction - talk of seeing and the circumstances in which we see – which ground attributions of the epistemic, and not vice versa. Again this makes it natural to think of the former locution as the more basic of the two.

So the second strategy, which says that non-epistemic perception is only a source of justification because it is un-obviously epistemic, is also no good. Only the third option remains for someone wanting to defend Davidson. This option agrees that simple perception isn’t epistemic perception in disguise and so agrees that non-epistemic perception isn’t a source of propositional justification. Nevertheless, it insists that non-epistemic perception cannot justify our beliefs about the world all by itself. According to

this response, non-epistemic perception merely functions as a sign which, in the presence of other justified beliefs, the subject can use to infer a belief about the world, in much the same way that if I am at sea and see a lighthouse flashing, I can infer that the surrounding water is shallow - provided I know that flashing lights indicate shallow water.

In the lighthouse case, what I see does not justify my belief about the depth of the water all by itself. It merely functions as a sign which, in the context of my knowledge or justified belief that flashing lights mean shallow water, I can use to establish that the water is shallow. According to the final response this is precisely how non-epistemic perception works and it is therefore not the autonomous source of justification that I claim.

What I mean is that the sort of justification it provides isn't genuinely independent of beliefs. This is easy to see in the lighthouse case. In this case, it is true that my justification does not derive exclusively from what I believe (the fact I see the lighthouse flashing is also meant to be relevant, and this is not a belief); but it is not as if my belief that flashing lights means shallow water is a *mere* enabling condition or that it plays no role in justifying me in believing that the water is shallow. On the contrary, this belief is playing a justificatory role. My justification seems to partly derive from it, even if it doesn't play that role by figuring as a premise in an argument from which the conclusion follows.⁵¹

⁵¹ This shows that there is a gap between whether or not a justification is 'inferential' in the traditional sense (i.e. exclusively belief-based) and whether or not it partly derives from beliefs. In the lighthouse case my justification isn't inferential in the traditional sense: I do not infer that the water is shallow from my belief that I see that the lights flashing and that flashing lights mean shallow water; part of my justification comes from the fact that I actually see the lights flashing, and seeing the lights flashing is not it a belief. Still, my justification is also partly belief-based. This seems to be more widespread than is normally

While this sort of dependence is not inferential in the traditional sense, it seems to be objectionable in much the same way that regular inferential justification is objectionable.⁵² So it would be worrying if non-epistemic perception worked like that. This is precisely what the third response alleges.

Non-epistemic perception does not work like that though and the lighthouse analogy is a bad one. It's a bad analogy since there is nothing in the case in which I believe that there is a squirrel sitting on the fence when I see the squirrel sitting on it that plays a role analogous to the role that seeing the lighthouse plays. Seeing the lighthouse flashing is a 'sign' that the water is shallow and if I know that it is a sign I can be justified in believing that the water is shallow when I see the lighthouse flashing. This is something I can conclude given what I see. But seeing the squirrel on the fence is not a *sign* that there is a squirrel on the fence; it is not something I am permitted to conclude on the basis of what I see. This is just a misuse of what we mean by something's being a sign.

It would be appropriate to talk of signs if I saw a pile of nut husks, or some bristly red hairs caught in the slats of the fence. In that case I might be justified in concluding that a squirrel had been sitting on the fence - they would be 'signs' that a squirrel had been

recognised. Take a standard case of enumerative induction: having observed the sun rising 'n' number of times I conclude it will rise tomorrow. As it is normally reconstructed, I reason from *beliefs* about what I have observed to a conclusion about the world. But why must we reconstruct my reasoning in this way? The natural view, surely, is that, in this case, as in the lighthouse case, part of my justification comes from my observations themselves, not beliefs about what I have observed. This suggests we may need to modify the notion of 'inference' and what it is from which we can 'infer'.

⁵² As we saw previously, if my justification partly derives from the belief that flashing lights means shallow water, then that belief had better be justified and justified in a way that is independent of the belief it is meant to justify. Otherwise my justification will be vitiatingly circular in the very same way that it would be in more traditional cases of inferential justification.

about for one who knows a bit about squirrels. But it would be bizarre to say the same in the case in which I literally *see* the squirrel sitting there. Similarly, I might look at the toaster and see smoke pouring out. That might be a sign that the toast is burning. But if I look at a piece of toast in flames, carbonizing before my very eyes, I do not see something that is a sign that the toast is burning: I literally see the burning toast.⁵³ This is quite unlike the lighthouse case: in this case I needn't even be looking at the water, and even if I do see the water, I needn't see its shallowness. This may just not be something I can see from where I stand. This is why it is plausible to think that the fact that the lighthouse is flashing does function as a genuine sign and that my beliefs about what flashing means do play a justificatory role. This is how I have access to the fact that the water is shallow. In the squirrel case though, it is not; I have access to the fact that the squirrel is on the fence because I actually see it sitting there.

I have claimed that in the case in which I see the squirrel sitting on the fence it also has to be true that I have certain recognitional capacities and beliefs. If I didn't have them, the fact I see the squirrel on the fence would not make me justified in believing the squirrel is on the fence. But it doesn't follow from this that it's my possession of those capacities that is doing the justifying or that those beliefs and capacities are playing anything more than a mere enabling role as I have claimed.

I think that this is the most intuitive thing to say in these cases and if that is right there is no reason for a foundationalist, or anyone else interested in stopping the epistemic regress, to find non-epistemic perception any more problematic in principle than its

⁵³ For a line of thought which I think is similar in some respects to this one, see (Travis 2004).

epistemic analogue.

To sum up: in this section I have argued that my picture of perceptual justification is fundamentally different from Davidson's both because my conception of how epistemic perception can be a source of justification is fundamentally different from his and because I allow that non-epistemic perception, that's to say non-propositional perception, can be a source of justification. This is not something Davidson could say even on a charitable reading of his position. In the next section I am going to consider another influential challenge to my view. This time the challenge comes from Laurence Bonjour.

4. Having Reasons

I have claimed that perceiving is a way of acquiring justification for our beliefs about the world that is fundamentally distinct from inference or reasoning. Crudely, Bonjour's view is that there is no such thing as distinctively perceptual justification in this sense. Like Davidson he thinks that our only real model of what it is for a subject to have justification for one of her beliefs – in his terms, for her to be 'in cognitive possession of a reason' (Bonjour 1985: 31) – is for her to believe the premises of an argument from which it follows that what she believes is (or is likely) to be true. This is very clear in his central anti-foundationalist argument.

Here is how Bonjour formulates that argument (and here I quote):

- (1) Suppose that there are basic empirical beliefs, that is empirical beliefs (a) which are epistemically justified, and (b) whose justification does not

depend on that of any further empirical beliefs

- (2) For a belief to be epistemically justified requires that there be a reason why it is likely to be true.
- (3) For a belief to be epistemically justified for a particular person requires that this person be himself in cognitive possession of such a reason.
- (4) The only way to be in cognitive possession of such a reason is to believe with justification the premises from which it follows that this belief is likely to be true.
- (5) The premises of such a justifying argument for an empirical belief cannot be entirely apriori; at least one such premise must be empirical.

Therefore, the justification for a supposed basic empirical belief must depend on the justification of at least one other empirical belief, contradicting (1); it follows that there can be no basic empirical beliefs (BonJour 1985: 32).

Here, BonJour just assumes that the only way in which a subject can have a reason for what she believes is by possessing ‘a justificatory argument’ in its favour (Premise 4). If she lacks such an argument, or fails to believe its premises, the belief won’t be justified for her. That is why BonJour thinks that no belief could be ‘epistemically basic’ since to be epistemically basic a belief would have to be both justified and such that the subject’s justification for it does not consist in her possession of a justificatory argument.

This is just to deny what I have so far claimed. I have claimed that perception provides a fundamentally different model of what it is for a subject to have reason or justification for her beliefs. Ann has reason for believing that there is a squirrel on the fence because she

sees or sees that there is a squirrel sitting on the fence. This is not a matter of Ann's believing the premises of some argument from which it follows that there is a squirrel sitting on the fence. It cannot be reduced to that or explained in those terms. If that is right then it is false that we can only have reason or justification for our beliefs by possessing a justificatory argument of the sort Bonjour describes. Perceiving something (in the context of certain abilities and environmental facts) or perceiving that something is the case is another way in which we can be justified.

In other words I am claiming that we should be permissive when it comes to justification. I am not just claiming that we have two different models of the way in which perception can be a source of justification. I am claiming that we have two different models of what it is to be justified at all – that is, two different models of what, at the most basic level, it is to have reasons for one's beliefs. This is to deny what Bonjour assumes. On this view inference does not furnish us with our only understanding of justification; perception provides another model. Moreover, just as inferential justification does not demand supplementation via perception in order to be intelligible as a source of justification on Bonjour's view, the same is true of perception on mine. Perception does not demand the support of what Locke called 'concurrent reasons' though we frequently have such reasons (Ayers 1991: 166-72); and it is not secretly constituted by such reasons in the way Bonjour assumes.

This is what I mean when I say that perception offers a fundamentally different model of what it is to have reasons for one's beliefs. An obvious question therefore is this: why

should these both be models of what it is to have reasons or justification? If inference and perception are as different as I have claimed, how can we make sense of them both being models of the same underlying thing? This is an important question and it is one that I will return to in chapter 5.⁵⁴ My point at this stage is that we do succeed in making sense of them as such, whatever the difficulties of doing so may be. So it must be possible to do so, whatever people like Bonjour say. In actual fact, of course, we have no more trouble with the idea that there is more than one way in which to acquire justification than we do with the idea that there is more than one way to play a game. It might be difficult to say exactly why different games – games perhaps as diverse as cricket and chess – both count as games, but it is not as if there is any doubt about whether or not they do. I think exactly the same is true in the case of justification. In both cases, it is the philosophical project of describing our practice that is difficult, not the practice itself.

I have claimed that this insight – that is to say, the idea that perception does furnish us with a distinctive model of justification – is central to foundationalism. This is partly why it is so odd to think of Descartes as a foundationalist, since he did not think of perception in this way. For Descartes, it is just as true as for Bonjour, that ‘the senses’ per se lack independent authority. They are useful signs for one who has concurrent reasons, but they have no intrinsic epistemic authority (Descartes 1996: Sixth Meditation). This is fundamentally opposed to what the foundationalist thinks on my account of foundationalism.⁵⁵ Of course, Descartes may count as a foundationalist because of the

⁵⁴ It might be thought that a Davidsonian - who countenances only epistemic perception - has less trouble on this score since he can appeal to the fact that all justification is ‘inferential’ (in some suitably broad sense) in accounting for the unity of the concept. It would be a mistake to think this makes things any easier. It would merely shift the problem to inferential justification: what makes that a concept with any significant internal unity, given than variety of cases that it covers?

⁵⁵ It might be argued that Locke ought not to count as a foundationalist on that basis either since there is

structural similarities his view bears to more traditional forms of foundationalism, or because of his quasi-perceptual view of the source of apriori justification; so perhaps he is a foundationalist of sorts. But he certainly shouldn't be seen as a paradigm foundationalist. This is not the reductio of my characterisation that some may suspect. Descartes is hard to classify as a historical thinker in more respects than one and we should no more balk at refusing to call him a foundationalist, than we should at refusing to think of him as a so-called Cartesian about the mind.⁵⁶

I have claimed that a foundationalist is someone who thinks that perception is a genuine source of justification and that is what Bonjour denies. He claims we cannot really make sense of that idea and that is very hard to believe. Again I am not claiming nothing could lead us to rationally revise that view, merely that we had better have good reasons for doing so given the presumption in its favour. This is what Bonjour has so far failed to provide and that just leaves him denying the seemingly obvious.

In fact, Bonjour does have one argument for thinking we can't make sense of a non-inferential model of justification. This is what I will call the argument from clairvoyance.

Bonjour asks us to imagine Norman, who:

some sense in which the senses also function as mere signs for claims about the world on his view too. But that had more to do with Locke's metaphysical picture of the objects of perception than any epistemological qualms about perception per se. Unlike Descartes, Locke thought that the direct objects of perception were our own ideas and he certainly didn't think that the senses were mere signs for beliefs about them. Moreover, it's not obvious Locke did think the senses functioned as 'signs' for beliefs about the world. He certainly did not on the whole think we needed concurrent reasons for believing they were a reliable or accurate guide to reality. On the contrary, he thought the status of so-called 'sensitive knowledge' was sui generis. This may not be sustainable, given Locke's overall metaphysical picture. But it is a sign of the pressure that he evidently felt, and that Descartes did not, to do justice to the idea that the senses are a source of knowledge or justification for our beliefs about the world, unaided by reason. This view simply struggled to receive its full and proper expression under the weight of his other metaphysical commitments.

⁵⁶ Cassam argues for the latter view in (Cassam, forthcoming).

...under certain conditions which usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power under circumstances in which it is completely reliable (BonJour 1985: 41)

The argument from clairvoyance claims that in this case Norman does not have any reasons for the beliefs he forms on the basis of his clairvoyant powers. If Norman had such reasons they would be genuinely non-inferential, but he does not. So clairvoyance isn't a source of non-inferential justification.

This is meant to be an argument against the possibility of non-inferential justification more generally, on the assumption that any picture of what the source of such justification would be, or what such reasons would be like, would have to be relevantly like clairvoyance, and hence (given our assumptions) not a source of justification at all. Although BonJour does not make that assumption explicit it is clearly his view.

One familiar response to this argument is to question the opening assumption – the assumption that Norman does not have any reason or justification for the beliefs he forms on the basis of his clairvoyant powers. So called 'externalists' about justification often complain that in denying Norman has justification BonJour merely parades his own

opposing intuitions and in doing so simply begs the question against their externalist alternative. A more concessive externalist response is to accept that Norman does lack reasons or justification for his clairvoyant beliefs, but to deny that 'reasons' or 'justification' are required for knowledge. According to this line of thought, 'reasons' in the sense in which Bonjour is interested in them, are of little epistemological interest to the rest of us.

Clearly, the second response is only more effective as a strategy if our intuitions that Norman knows are any less congenial to Bonjour than our intuitions that Norman has reasons or justification. I doubt whether that is so. Either way, neither of these two externalist options is an entirely cost-free strategy and both leave us having to deny something seemingly intuitive.

A much more obvious response is to agree with Bonjour that Norman lacks reasons or justification for the beliefs he forms with his clairvoyant powers, but to deny that that shows anything about the possibility of non-inferential justification in general. It doesn't show any such thing because clairvoyance is a just bad picture of what the alternative must look like. Bonjour's whole analogy therefore fails to get off the ground and the spreading step fails.

That is surely especially plausible in the present context, where we are talking about perception and whether *it* can furnish us with an alternative model of justification. Perception is nothing like clairvoyance. When one perceives an object or state of affairs

one is related to it in a very special way – a way that mere thought or imagination do not provide, let alone clairvoyance. In perception, one is made consciously aware of the objects and events in the world around one that one perceives and their properties: it is that very squirrel with which one is presented and of which one is aware. This simply has no analogue in the case of clairvoyance; our best gloss on what clairvoyance is like is that it involves a strong hunch.

This is why it is so hard to believe that clairvoyance is a way in which to acquire justification for our beliefs.⁵⁷ But we can agree with Bonjour that clairvoyance isn't a source of justification and agree that seems genuinely intuitive, contra the externalist. Yet still that that shows anything about whether there could be other sources of non-inferential justification, in particular, whether perception could be one of them.

I think that this is a much more plausible line to take. How plausible it is in the final analysis will no doubt depend on how we think about perception. On certain ways of thinking about perception, it is much less readily intelligible why it should be a source of justification or provide a truly distinctive or basic model of what it is to have reasons.

Against those views, Bonjour's complaints have more force.

Consider what James Pryor has to say about the epistemological role of perception:

My view is that our perceptual experiences have the epistemic powers the

⁵⁷ I mean a source of *non*-inferential justification. Clearly, we could have excellent reason for thinking that clairvoyance is reliable, in which case 'it' may very well give us reasons for the beliefs about the world that we form on its basis. In that case, however, our justification will not be non-inferential; it will derive at least in part from the belief that clairvoyance is reliable.

dogmatist says they have because of what the phenomenology of perception is like. I think there's a distinctive phenomenology: the feeling of seeing to ascertain that a 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. We find this phenomenology in perception...When you have a perceptual experience of your hands, that experience makes it feel as though you can just see that hands are present. It feels as those hands are being shown or revealed to you. This phenomenology may be present in other mental episodes too...My view is that our perceptual justification comes from that phenomenology. Having the phenomenology of seeming to ascertain P is what makes us have prima facie justification to believe P (Pryor 2004: 356-7).

I think that this view is incredible. Here, Bonjour's complaints really do seem to have force. Why should states in which you merely 'seem to ascertain' that a proposition is true give you any reason to believe that things actually are so? That seems a perfectly legitimate question to my ear.

I am not denying that we couldn't have a model like Pryor's, or that we couldn't try and explain why perception is a source of justification in these terms. Faced with Bonjour, Pryor should say precisely the sort of thing that I say; namely, there is nothing more to be said about why states in which you 'seem to ascertain that a proposition is true' are reason-giving. He should just insist that those sorts of states provide one of our most fundamental grips on what it is to have reason. FULL STOP. But it is striking just how

implausible this sounds. Intuitively, Bonjour's question really does get a grip; it's just not at all obvious or natural to say that these sorts of states just are cases in which you have reason. End of story. On the contrary, that fact cries out for further explanation.

The view that I defend does better in that respect. My view appeals to the fact that the subject perceives or perceives that things are a certain way in explaining what justifies her beliefs. These states consciously relate the subject to the objects and events in the world that her beliefs concern and make her consciously aware of their properties; it is the very squirrel that Ann believes is on the fence which she actually sees sitting there. And it is surely no great mystery how *that* can put Ann in a position to know something about the squirrel. Intuitively, that just is what it is to have a reason for believing something about the world.

This view is so natural that even Pryor eventually falls back on it. This is effectively what he appeals to when it comes to explaining why the phenomenology of perception should be reason giving: it is because "that experience makes it feel as though you can just see that hands are present. It feels as those hands are being shown or revealed to you as he claims" (Pryor, *ibid.*). On my view those hands really are being shown or revealed to you and that is precisely *why* you have reason for judging. So Pryor's explanation is ultimately parasitic on mine.⁵⁸

⁵⁸ I think that is significant given how intuitive Pryor-style views are often thought to be. This assumption places a crucial role in framing the contemporary debate. It is used to put pressure on views like mine and suggest that somehow the burden of proof in this area really lies with me – my view is the counter-intuitive view and so I am the one that need to provide all the reasons. I think that gets things exactly the wrong way round. What I hope to have brought out above is how very counter-intuitive the Pryor style view is at the most fundamental level. That should help us redress the dialectical balance.

Pryor is a so-called ‘common factor’ theorist about perception. So does this make me a ‘disjunctivist’ about perception? Some will suspect that it does and some will want to reject my view on those grounds. Even those who don’t object to disjunctivism, per se, may find it hard to believe that our common sense views about the epistemological role of perception could commit us to such substantive claims about its underlying metaphysical nature.

I agree. But I am not committed to disjunctivism. I am claiming that perceiving or perceiving that something is the case is a source of justification for our beliefs about the things we perceive and that perception does that by consciously relating us to those very things; when one sees the squirrel on the fence, it is that very squirrel of which one is aware. It should be relatively uncontroversial that that very state (your seeing) couldn’t exist unless the squirrel existed and you were related to it in that very way.

This is not what disjunctivists claim; even people as opposed to disjunctivism as Searle or Davidson or Pryor could agree with that. What they dispute is whether a mental state of the same fundamental kind could occur in the absence of the squirrel. They think it could and indeed would occur, even if you were hallucinating. That is what the disjunctivist denies. He denies that there is some ‘narrow’ psychological state in common between cases of illusion and hallucination and cases of veridical perception in virtue of which these three states count as being the fundamental kind of psychological states that they are (Snowdon 1981, 2005; Martin 2002).

My view is simply silent about this. It might be true, it might not, and what I say is compatible with either since even if there is a common factor of the sort just described it does not follow that whether or not one is justified is a function solely of the presence or absence of that factor. Indeed, even if a subject's justification is just a function of what mental or psychological states she is in, it may be function of more than what fundamental kind of psychological state she is in. This is especially plausible when we consider the grounds upon which philosophers of perception tend to individuate the mental or carve out its fundamental kinds. They think that what 'fundamental' kind of psychological state a subject is in is fixed by how things seem to the subject of that state: by that state's 'phenomenal character' (Martin 2002; Soteriou 2005). There is no obvious reason to think that what is epistemically important – that is, what confers justification – is just a factor of how things seem to the subject in this sense.

Having said that all that, it may turn out that disjunctivism is ultimately the best way of holding onto the view I defend. Perhaps once we concede the existence of a common factor between the case in which one perceives and merely seems to perceive, it would be impossible to believe that factor wouldn't also fix the epistemological facts or leave any genuine explanatory work for perceiving to do; or perhaps it's implausible to suppose that what is epistemically significant isn't determined by how things seem to the subject in the sense in which philosophers of perception are interested.⁵⁹ If so, we will have to adopt a disjunctive view of the phenomenal character of perceiving and seeming to perceive in order to hold onto the view that perceiving the world is, as I maintain, a genuinely

⁵⁹ Williamson argues against a parallel claim in the case of knowledge and belief (Williamson 2000: esp. 2.4).

distinctive source of justification for beliefs about it.

However, establishing any of these claims would require lots more work and all go far beyond what I am claiming. I am merely claiming that perceiving is a source of justification for beliefs about the world and that is so because it consciously relates us to the objects and events in the world about which we judge. On the face of it, that is compatible with lots of different views about what it is to be related to an object in this way and whether that state has anything significant in common with a state in which one merely seems to be related to the world.

So if you don't like the idea that perceiving is a source of justification in this sense, you should go ahead and reject my view. But don't reject it because you think it commits you to disjunctivism.

To sum up: in this chapter I have claimed both that perception is a source of justification, genuinely distinct from inference, contra Bonjour; and that the most natural and straightforward explanation of this fact appeals to the relational nature of perceiving, contra Pryor - the fact that perception, unlike clairvoyance, say, consciously relates us to objects and events in the world around us about which we form beliefs. I have claimed that states in which we perceive and perceive that things are the case are states in which we are actually in touch (sometimes literally) with the things in the world that our beliefs concern; we are made aware of those very objects and their properties. Given that this is so, it is no mystery why perception is a source of justification for our beliefs about them.

That is what I have claimed. As we have seen, not everyone agrees. I have already discussed Bonjour's view. In the final section I will look at what people like Pryor have to say. Unlike Bonjour, Pryor's view is not completely lacking in all motivation. Pryor and others like him want to do justice to that idea that a subject unknowingly suffering an illusion or hallucination has the same justification as one who is perceiving the world around her. On the face of it, this is an idea to which my view will struggle to do justice. If one is justified because one actually perceives that things are a certain way, why should one be justified when one is not, and merely seems to be related to the objects and properties in the world about which one endeavours to judge? Many philosophers claim to find that counter-intuitive and opt for a Pryor-style view on these grounds.⁶⁰ This further challenge to my view is one I will address in the next and final part of this chapter.

5. The Argument from Fallibility

I have explained what the Broad View of perceptual justification is and why that is a pretty plausible, commonsensical thing to think. It solves the epistemic regress problem and it doesn't have the restrictive consequences of traditional foundationalism. So why is it that philosophers haven't on the whole seen this? Historically, two main arguments have prevented them taking that view sufficiently seriously. The Argument from Fallibility claims that perception only fallibly justifies beliefs about non-psychological

⁶⁰ These are not the only options in this area. Tyler Burge has a view that is intermediate between Pryor's and mine in many respects. He wants to allow that subjects can be justified when they suffer certain sorts of perceptual illusions provided that perception is in general reliably veridical. This might be thought less implausible: maybe it is more readily intelligible that a source of representations, which is reliably veridical, should be a source of reasons for beliefs about what is represented. For more on Burge's view, see (Burge 2003).

reality and that it therefore cannot non-inferentially justify them, contra the Broad View. The Argument from Defeasibility claims that perception only defeasibly justifies beliefs about non-psychological reality and that it therefore cannot non-inferentially justify them, contra the Broad view. Unlike the views discussed so far these are actual arguments; they don't just deny that perception could non-inferentially justify beliefs about the world, they give concrete philosophical reasons for thinking that this can't be so. Moreover these are reasons that many philosophers have been moved by. According to what I earlier called the 'Simple Reading', it is precisely because beliefs about non-psychological reality can be mistaken that traditional foundationalists were lead to privilege beliefs about our own minds. So these are the arguments we really need to engage with. In the rest of this chapter I'll tackle the Argument from Fallibility, defeasibility will then be the topic of the next.

The Argument from Fallibility claims that it follows from the fact that perception only fallibly justifies beliefs about non-psychological reality that it must inferentially justify them. Faced with that argument there are two basic moves available to those of us wanting to hold onto the Broad View: we can either deny that perception does only fallibly justify beliefs about non-psychological reality; or we can deny that the conclusion follows, that is, deny that it follows from the fact that perceptual justification is fallible that it must be inferential. I'm going to start off by looking at the second response, that is, questioning the argument's validity. This will strike many as the more obvious of the two responses and it leads very naturally into questions the first response addresses. As we are about to see in thinking about why validity might be a problem it will emerge that the

argument's premise is much less straightforward than most people assume.

Why should anyone think that it follows from the fact that perceptual justification is fallible that it is inferential?⁶¹ Let's agree straight off that there are fallibly justified beliefs and that some of those beliefs are also inferentially justified – inductively justified beliefs seem to be a case in point. If my justification for believing that Captain Molski will win derives from the fact I believe that she is the fastest dog on the track and that this makes it probable she will win, then my justification for believing she will win is inferential. It is also fallible; intuitively, I may be justified in believing that Captain Molski will win on those grounds and yet, for all that, she does not.

However, it is not enough for a defender of the Argument from Fallibility merely to point out that there are fallibly justified beliefs that are also inferentially justified. She needs to show us that there is some connection between these two things – that it somehow follows from the fact that a belief is only fallibly justified that it is inferentially justified, or that it is inferentially justified in virtue of being fallible.

So the question remains: what has fallibility got to do with inference in this sense? To see the intuitive connection here, think again about the inductive case. In this case what it means to say that my belief that Captain Molski will win is only 'fallibly' justified is that

⁶¹ There is something slightly odd about concluding that 'perceptual' justification is inferential. In that case, what makes that justification *perceptual*? To call a justification perceptual is to make a claim about the epistemic source of that justification; not just its causal origin or the causal origin of the beliefs it is claimed to justify. Thus, in the event this argument succeeds in showing that perceptual justification cannot be non-inferential. I think it would be more natural to conclude that there is no such thing as properly perceptual justification. On this view, the epistemic role that I claim perception plays is instead played by beliefs about what we perceive and how that makes it likely the world is arranged in such and such ways.

my justification for that belief leaves open the possibility that I am mistaken. I can be justified in believing that she will win and it still be false that she does win. If that is possible though, am I really justified in believing that she will win? I might still be justified in believing that she will probably win in those circumstances, but I surely can't be justified in believing that she will actually win if my grounds leave open the possibility that she will not.⁶² Intuitively, to be justified in taking this further step and believing that she really will win I need to exclude the possibility that she will not and I can only do that by inferring that is so from other justified beliefs.⁶³

I think that is the intuitive thought connecting inference and fallibility. In the case of induction, however, I cannot rule out the possibility of mistake – that is precisely makes my inductive reasons 'inductive' rather than 'deductive'. What someone who wants to press this worry really ends up calling into question is how there can be such a thing as inductively justified beliefs – that is, how inductive reasons can be genuine *reasons* given that they leave open the possibility of mistake.⁶⁴ If so, this argument isn't really an argument for thinking that fallible justification must be inferential. Really, it is an argument for thinking that there is no such thing as fallible justification, since the only

⁶² This is an intuition McDowell claims to be moved by (McDowell 1998a, 1998b). It is certainly not without appeal though it is hard to see how to reconcile it with the equally intuitive idea that induction is a genuine source of knowledge.

⁶³ Presumably even that is no good if those beliefs in turn leave open the possibility of mistake. Why should it be acceptable to allow that one's grounds for ruling out the possibility of mistake, do not themselves rule out the possibility of mistake, when it is not acceptable to claim that is so in the original case? Thus this line is immediately regressive. It could only be stopped by appealing to infallibly justified beliefs.

⁶⁴ This line of thought is familiar from Hume (Hume 1975: esp. sec. IV). I am assuming that induction needn't be grounded in metaphysical necessities and hence that there is genuinely room for the possibility of mistake on a intuitive understanding of what one's grounds are in the inductive case. Perhaps some will deny that. They need to give us a different understanding of the distinction between 'inductive' and 'deductive' reasons, and a different account of what exactly one's grounds are in the inductive case. None of this is to deny that good inductive arguments rule out the possibility I could easily have been wrong. That is a different matter to the issue of whether or not they logically exclude the possibility of error.

grounds for thinking that fallible justification must be inferential turn out to be grounds for thinking that there is something funny about fallible justification. What started out as an argument for a pretty commonsensical thought has therefore ended up as an argument for ‘infallibilism’.

In that case, maybe a better question is whether the premise of the Argument from Fallibility is true, not whether the argument is valid.⁶⁵ So far I’ve assumed that perception does only fallibly justify beliefs about non-psychological reality. This is the standard view in the literature, but is it true? Not on the face of it. What justifies Ann in believing that there is a squirrel on the fence is the fact Ann sees or sees that there is a squirrel on the fence. Those grounds do not leave open the possibility that Ann is mistaken: Ann cannot see a squirrel on the fence unless it is true that there is a squirrel on the fence and she cannot see that there is a squirrel on the fence if she could be mistaken in believing that there is a squirrel on the fence. ‘See’ is a ‘success’ verb: one cannot see the squirrel on the fence unless the squirrel on the fence exists. Similarly, ‘sees that’ is ‘factive’: one cannot see that the squirrel is on the fence unless the squirrel is on the fence. So if that is what justifies Ann in believing that there is a squirrel on the fence, and I claim it is, then Ann’s justification is not fallible.

Exactly the same is true of perception more generally – I cannot perceive ‘a’ unless ‘a’

⁶⁵ Notice, I am not endorsing this line of thought or suggesting that the argument *is* valid. One could as easily claim that this line of thought has therefore failed to establish *any* link between inference and fallibility since the only grounds for thinking that fallible justification *must* be inferential turn out to be grounds for thinking there is no such thing as fallible justification. And there is such a thing as fallible justification. So, the argument is no good. In that case we can allow perceptual justification is both fallible and non-inferential. This is the response Pollock and Cruz favour (Pollock and Cruz 1999: 43-44). On this model, perceptual reasons need not be inductive reasons. They simply figure alongside such reasons, as equally basic constituents of our ratiocentive framework.

exists and I cannot perceive that 'a is F' unless a is F. So the justification that perception more generally provides is not such as to leave open the possibility of mistake. If this is right then the Argument from Fallibility is unsound and we can reject it on those grounds, quite aside from any worries we might have about its validity.

This response claims that perception is a source of infallible justification, and that the beliefs that it infallibly justifies are or include beliefs about non-psychological reality, like the belief that the squirrel is on the fence. On this view – which I have also claimed is the naïve view – ordinary perceptual beliefs about the world around us count as infallible, and not just beliefs about our own psychological states as the traditional foundationalists maintained. This will strike many philosophers as so obviously incredible that I had better say a bit more.

Perhaps the first thing I should say is this: I am not claiming that what you believe when you believe something about the world around you on the basis of perception – that is, the proposition believed – is incapable of being false or mistaken. What Ann believes is that the squirrel is on the fence and clearly it could be false that the squirrel is on the fence. It is a contingent truth, after all, that the squirrel is anywhere near the fence. So ordinary perceptual beliefs are not infallible or incapable or being mistaken in that sense. This is how some philosophers characterise the notion of infallibility. Bernard Williams is a case in point. He claims that the belief that p is 'in corrigible' (in his terms) iff S believes that p, entails p (Williams 1978: 306). As we have just seen, that is a test that ordinary perceptual beliefs do not pass.

However, that is not the relevant notion of infallibility. What I mean by this is that people who traditionally appealed to the notion of infallibility wanted it to be *epistemically* relevant. But the fact that a belief is infallible was meant to figure in the explanation of why that belief was justified. This is certainly what the traditional foundationalists were after. The mere fact that a belief is infallible in Williams's sense tells us nothing about how or whether the subject is justified in that belief. If I believe any necessary truth my belief will be infallible in his sense, but I can certainly be unjustified in believing a necessary truth.

Williams's characterisation therefore doesn't serve the epistemological function that infallibility was historically designed to serve. This has led most contemporary philosophers to jettison the notion of infallibility altogether, but this is a mistake. The problem is that infallibility is a modal notion: to say one believes *p* infallibly is to say one cannot be mistaken in believing *p*. The problem with Williams's formulation is that the modality attaches to the wrong thing; on his formulation it attaches to the proposition believed. Intuitively, though, we want the modality to attach to the subject. We want *her* to be such that whatever she believes, what she believes cannot be mistaken. Unlike Williams's notion this is not a test that all necessary truths pass, and it does seem to be genuinely epistemic relevant. So it is a mistake to dismiss the whole notion of infallibility – we just need to understand it in the right way.

This is how the notion of infallibility was traditionally understood. Moreover, this is the

sense in which I am claiming that ordinary perceptual beliefs are infallible: Ann's belief that the squirrel is on the fence is infallible because Ann is such that whatever she believes on that score what she believes cannot be mistaken. That is what I am claiming at least, but I can see that some will require more convincing.

Perhaps it will help to start by saying something about how any subject could be infallible in this sense. One possibility is that the subject's belief concerns a special subject matter: perhaps, as with Descartes' cogito, what the subject believes is such that merely entertaining the thought it expresses is sufficient for the subject to determine whether or not what she believes is true. This is what many philosophers have been tempted about so-called self-evident truths, like whoever is tall is tall. I am going to call that the Subject-Matter Reading.⁶⁶

Another possibility is that the subject's belief has a certain source and that that explains why she cannot go wrong. If the Pope believes that p because God has told him that p, then the Pope is equally such that what he believes cannot be mistaken. That is not due to some special feature of the content of his beliefs: the Pope might know that the end is nigh because God has told him, but it is not a self-evident truth that the end is nigh. If he knows that the end is nigh, it is because (and only because) his belief has a special source. I am therefore going to call that the Source Reading.

⁶⁶ This is somewhat closer to the Williams characterisation, since it at least appeals to an aspect of the proposition believed. Even so, it is not identical with that characterisation. Necessary truths all pass Williams's test but most of them are not self-evident in this sense (mathematics would be a lot easier if they were). Conversely, many things which are self-evident are not necessarily true. Unless you are Timothy Williamson you probably will not think that 'I exist' is a necessary truth, but many philosophers think it is a self-evident truth.

We now have two different explanations of how a subject's beliefs could be infallible in the sense that I have claimed is relevant to our discussion. Both have played a role in the history of philosophy: Descartes favoured the first and traditional foundationalists the second. It is perhaps not surprising, then, that it is in the second sense that I think that ordinary, perceptually justified beliefs are infallible. Ann's belief that there is a squirrel on the fence is infallible not because it is self-evident that there is a squirrel on the fence, but because her belief has a certain source, namely her seeing or seeing that there is a squirrel on the fence. Those grounds exclude the possibility that what Ann believes is mistaken and they are the grounds upon which she judges.

That is what I am claiming. I hope it now seems less obviously false than before, but it is still not trivial. I have not made it true at the expense of its philosophical interest. Not everyone will agree that our perceptual beliefs are infallible in even this sense, since not everyone will agree that these are the grounds upon which we are justified in holding those beliefs. We saw earlier that people like James Pryor think that what justifies Ann in believing that the squirrel is on the fence is the fact she is in a state in which 'she seems to ascertain' that it is true that there is a squirrel on the fence. Those grounds do not rule out the possibility of mistake; Ann can seem to ascertain that the squirrel is on the fence and yet fail to actually ascertain that the squirrel is on the fence since the squirrel is elsewhere. In that case Ann will be justified in believing that the squirrel is on the fence even though it is false that the squirrel is on the fence. So Ann's belief is at best fallibly justified.

I have not rejected that view because I think that the Argument from Fallibility is valid and that this view would make perceptual justification inferential, given that all fallible justification must be inferential if that Argument is valid. I have rejected it because I think it is a very strange and baroque account of what actually justifies our perceptual beliefs. I think there is no good philosophical reason for abandoning the naïve view that what justifies our perceptual beliefs is the fact that we are in states that consciously relate us to the things about which we judge. Those states are flatly incompatible with things not being as we judge.

So I am not arguing for a generalised ‘infallibilism’. I am not saying that any justification worth its salt must exclude the possibility of mistake. There is certainly something intuitive about that, but it is hard to reconcile with the equally intuitive idea that induction can be a genuine source of knowledge or justification despite the fact it does not thus exclude the possibility of mistake.⁶⁷ You do not have to be an infallibilist in any more general sense to accept my story. I am merely claiming that it is an interesting fact about our perceptually justified beliefs that what actually justifies those beliefs does, in fact, rule out the possibility of mistake - quite aside from whether or not if it didn’t that

⁶⁷ McDowell claims otherwise. He writes: “induction can have a confusing effect here: it can seem to be a counter-example to the principle. But demanding that an argument be conclusive is not the same as demanding that it be deductive.” (McDowell 1998b: 421). That is certainly true on some readings of ‘conclusive’. (It is true on Dretske’s theory of conclusive reasons. Indeed, this is the whole point of his theory, see (Dretske 2000b). But it doesn’t look like a live option on McDowell’s own reading. The principle he appeals to says, (and here I quote) “the argument would need to be conclusive. If you know something, you cannot be wrong about it” (ibid.). Only deductive arguments ensure that you cannot be mistaken; inductive arguments do not. Whether or not I will be mistaken is a different matter; good inductive arguments may certainly rule that out but even the highest common factor theorist whom McDowell claims to be opposing can require that our epistemic standings rule out the possibility we *will* be mistaken or that we *could* easily be mistaken. Of course, McDowell might just be making the point that if you *know*, you cannot be mistaken (whether or not your grounds make it the case that you couldn’t be mistaken). This is certainly true; but trivially so. Anyone can agree with that. Indeed, anyone must agree with that since knowledge is ‘factive’: there is a deductive argument from S knows that p, to p. So at best what McDowell says here is highly misleading.

justification would still be worth the name.

Of course even on my view having such grounds doesn't rule out every way in which the subject could go wrong. Suppose Ann sees the squirrel on the fence but believes irrationally that there is no squirrel on the fence or that squirrels do not exist. Or suppose she is careless and forms the belief that there is a cat on the fence, rather than the squirrel that she quite plainly sees. In those cases, what Ann believes may well be mistaken. Clearly these are logical possibilities; seeing the squirrel on the fence does not exclude the possibility that Ann's belief could be mistaken for these reasons. But this sort of error is not peculiar to perceptual beliefs. If Ann is negligent in forming her beliefs, doesn't properly attend to what she perceives, or if her beliefs are not appropriately related to what it is that she perceives, then Ann's beliefs may very well be mistaken. But all beliefs are subject to mistakes of these kinds and no notion of infallibility could reasonably be expected to exclude them. Even so-called self-evident truths are not infallible in this sense and nor are beliefs about one's own psychological states. One can believe irrationally that it does not follow from the fact that if A is taller than B, and B taller than C, that A is taller than C; and if I do not properly attend to my sensations I may mistake a tickle for an itch. These are not reasons for thinking that beliefs about our own psychological states are not infallible or that belief in self-evident truths is not infallible – at least not when one properly attends, is not being irrational, and one's beliefs are appropriately related to one's grounds. Why can't exactly the same caveat be allowed to apply in the perceptual case?

Perceptual beliefs are vulnerable to a certain sort of mistake to which the others are often thought to be immune. Ann might believe that the squirrel is on the fence because she seems to see a squirrel on the fence but where, because she is hallucinating, there is really no squirrel there to be seen. In that case, if Ann believes that there is a squirrel on the fence her belief will be mistaken. However, this is not a case in which Ann is justified in believing that the squirrel is on the fence because she sees that there is a squirrel on the fence and in which her belief is mistaken.⁶⁸ In this case Ann does not see the squirrel on the fence. So this fact has no power to show that Ann's belief that there is a squirrel on the fence is not infallible in the sense in which I claim it is. That belief is infallible because of its source and in this case the source is different.⁶⁹

Exactly the same is true of any belief that is infallible because of its source. If I believe that I am depressed because of what my analyst tells me, or that I am in pain because that is what the doctor has told me I will feel when the needle goes in – then my belief may also be mistaken.⁷⁰ Beliefs about one's own psychological states are therefore no more immune from error in this sense than more ordinary beliefs about non-psychological reality. If they seem different it is only because we typically form beliefs about psychological reality on a certain basis, namely on the basis of being presented with the

⁶⁸ I am not even committed to thinking this is a case in which Ann is *justified* in believing that the squirrel is on the fence. After all she doesn't see the squirrel, she merely seems to see one. For all that has been said so far Ann therefore has no reason at all to believe that there is a squirrel on the fence. Of course, she might have inductive grounds for believing that things which look like squirrels normally are squirrels, but that is a different matter altogether.

⁶⁹ One might try a similar move in the case of induction, though it would be a lot harder to pull off. In that case it's much harder to come up with a plausible alternative for one's grounds, in such a way that they do not leave open the possibility of mistake.

⁷⁰ Some philosophers think it is sufficient for one to be in pain that one sincerely believe that one is in pain. I think this is a mistake. It may be that we never believe that we are in pain unless we are in pain. That doesn't show that the belief that one's in pain is sufficient for one to be in the pain. It may just be that we never form the belief that we are in pain other than because we feel our own pain.

facts these beliefs concern. What normally justifies me in believing that I am in pain is the fact that I feel my own pain and believe I am in pain on that basis.⁷¹ In that case it remains true that I cannot be mistaken. But there is no more guarantee in this case than in the perceptual case that if I believe that I am in pain on some other basis my belief will be similarly immune from mistake. Beliefs about non-psychological reality therefore have at least as good a claim to be infallible as beliefs about psychological reality, contra the traditional foundationalist.⁷²

This is what I am claiming, at least. Having said that I don't want to go to the stake for the view that beliefs about psychological reality really *are* infallible. Perhaps if you really press matters we cannot isolate a subset of mistakes as those that involve purely 'verbal slips' or that result from irrationality, insufficient attention, or that are not properly grounded in what it is that one perceives. Maybe these mistakes are all substantive, as Austin liked to claim (Austin 1962: 112-3).⁷³ What I do want to go to the stake for is the

⁷¹ At least, that is the sort of thing that people in the literature say. As a statistical claim it is surely false. Indeed, it gets things precisely the wrong way round. I think it is very common to form beliefs about one's own psychological states on others grounds – I normally form them on the basis of others say-so. In contrast, for those of us who didn't ground up in the 1960's, it is very rare that we ever form beliefs about the world on the basis of hallucinations.

⁷² Of course one can be ignorant of the layout of non-psychological reality in a way in which many think one cannot be with respect to psychological reality. But ignorance is not a form of error. If Ann refuses to believe anything when presented with the squirrel on the fence she is not in error; she merely fails to know something she might otherwise have known. That is not a reason to think that ordinary perceptual beliefs aren't incapable of being mistaken.

⁷³ Austin poured scorn on the idea that such mistakes involve merely 'verbal slips' He writes: "Ayer tries, as it were to laugh this off as a quite trivial qualification; he evidently thinks that he is conceding here only the possibility of slips of the tongue, purely 'verbal' slips (or of course of lying). But this is not so. There are more ways than these of bringing out the wrong word. I may say 'Magenta' wrongly either by a mere slip, having meant to say 'Vermilion'; or because I don't quite know what 'magenta' means, what shade of colour is called magenta; or again, because I was unable to, or perhaps just didn't really notice or attend to or properly size up the colour before me. Thus, there is always the possibility, not only that I may be brought to admit that 'magenta' wasn't the right word to pick on for the colour before me, but also that I may be brought to see, or perhaps remember, that the colour before me wasn't magenta. And this hold for the case in which I say, "It seems, to me personally, here and now, as if I were seeing something magenta", just as much as for the case in which I say, 'That is magenta.' The first formula may be more cautious, but it isn't incorrigible." (Austin 1962: 112-113)

claim that beliefs about non-psychological reality have *as good* a claim to be infallible as beliefs about psychological reality. That is not something most philosophers think is true, let alone trivially so.⁷⁴

I think this view does most justice to our ordinary view of perceptual justification. I also think it does something at the philosophical level to explain why it is that it should be a ‘justification’ that perception affords one. What I mean is that this view makes the connection to truth – the connection that any justification has to secure – perfectly straightforward. Perception consciously presents us with the very things about which we judge and makes their nature manifest. In believing what we do on those grounds, we cannot go wrong. That is why it is no mystery that it is the source of justification we so readily take it to be. The same cannot be said for the fallibilist view of Pryor, considered earlier. Not only is it odd to appeal to states in which we merely seem to perceive that things are a certain way, it also leaves a big gap at the level of explanation – that of saying why seeming to perceive is a source of reasons. This is another good reason for preferring my view.

So I have now defended what I will call an Infallibilist view of perceptual justification. This represents a further respect in which my position is continuous with more traditional forms of foundationalism. But my view enables us to hold onto what is gripping about

⁷⁴ Traditional foundationalists might claim it is not the mere availability of other methods that makes for difficulties. The problem only arises where we cannot keep track of which method we are using. In the case of psychological reality that is meant to be unproblematic: it is easy to determine when we believe something on the basis of ‘inner perception’ as opposed to others say-so. The same is not true in the perceptual case: we are not always in a position to know when we are perceiving as opposed to merely seeming to perceive and can we rationally be employing the latter ‘method’ without noticing. I will be exploring these issues in the next chapter.

traditional foundationalism without committing us to its less plausible features. If I am right the beliefs that perception infallibly justifies include beliefs about non-psychological reality. I think this is the view which strikes us as most plausible prior to philosophical reflection, but is not one that many philosophers have been drawn to in the last 100 years. So why is that? Why has ‘fallibilism’ suddenly come to seem inevitable, when previously infallibilism seem so natural?

6. Conclusion

I want to conclude this chapter with a speculation about why that is so, though I hope what I say will not be of mere sociological interest. The basic reason I suggest is this: philosophers are very impressed by the idea that subjects unknowingly suffering an illusion or hallucination have the same justification as those who are veridically perceiving the world around them. This intuition now has its very own argument; commonly referred to as ‘the new evil-demon problem’ it is basically just an updated version of Descartes’ malicious demon argument (Descartes 1996: First Meditation). It claims that my ‘twin’ and me have exactly the same justification for our beliefs about the perceivable world, despite the fact that my twin is the victim of an elaborate deception and merely seems to perceive the world around him (Sosa 1991: 281). If this is right, it’s hard to see how perceptual justification can be a function of something that is only present in the case of veridical perception as I claim viz. contact with the world. Rather, it must derive from something that cases of veridical perception have in common with the others – perhaps, as Pryor suggests, from the phenomenology of perception and the fact that in having a perceptual experience one seems to ascertain that a proposition is true.

I have already said that I think that view is perverse. It is also unmotivated, since there is nothing so very intuitive or commonsensical to recommend the idea that subjects in such circumstances really *do* have the same justification as those actually perceiving the world around them. So while I started off by saying that philosophers are ‘very impressed’ by the idea that such subjects have the same justification as their more fortunate counterparts, what I really mean is that they are *overly* impressed by that idea; I think they have simply forgotten what a strange view it really is.

Of course, there is something to the idea that subjects who form beliefs about the world when unknowingly suffering from hallucinations cannot be ‘blamed’ for the beliefs that they form on that basis. As Williamson points out they have “a cast iron excuse” for having formed those beliefs (Williamson, forthcoming). That is how things seemed to them and that makes it perfectly intelligible why they believe what they do. Moreover, they needn’t be irrational in these beliefs. They needn’t have any special reason to suspect that they are hallucinating. The important question is why we should we think it follows from this that they have any reason for the beliefs they form on this basis? If I jump out of a window while under the misapprehension that the building is on fire it is perfectly explicable why I act as I do and I certainly needn’t be irrational in jumping. It may even be a perfectly ‘reasonable’ thing for me to do in some sense of ‘reasonable’. Still, from a commonsense point of view it remains true that I don’t actually have any reason for acting as I do. I certainly *seem* to have a reason and that explains why I act as I do. But I don’t *actually* have any reason. That is precisely why we say to people who

jump out of windows when there is no fire: oh dear, you really shouldn't have jumped...

I think that this is the most intuitive thing to say in this case, as it is in the epistemic case.⁷⁵ This is a way of making the point that we are naturally realists about reasons – at least until philosophy gets hold of us. We think that in the realm of reason there is a distinction to be drawn between how things seem and how they really are – between what one seems to have justification for believing and what one really has justification for believing. What I am claiming is that there are no good philosophical reasons for revising this view. In the case of perceptual justification, we should remain the realists about reasons that we naturally all are.⁷⁶

⁷⁵ Indeed it's arguably even more plausible in the epistemic case. In the cases sceptics focus on the beliefs we form result from *deception* (I am the victim of an evil demon...etc etc). This needn't be so in the practical sphere.

⁷⁶ None of this is to say that how things seem to us perceptually (e.g. how they look) is irrelevant. It's by looking the way they do that objects consciously reveal themselves to us. Moreover, when pressed we often retreat to claims about how things look. My point is just that this doesn't show that 'looks' are all that is important, or that they are the grounds upon which we must have been judging all along.

FOUNDATIONALISM AND THE IDEA OF THE EMPIRICAL

CHAPTER 4: THE ARGUMENT FROM DEFEASIBILITY

1. Introduction

At the end of the last chapter I argued that perceptual justification is in fact infallible. My question in this chapter is whether it is also indefeasible. I'm going to argue that it is not. So I will be arguing that the infallibility of perceptual justification does not entail its indefeasibility as many have thought. What, then, is the significance of the fact that, as I am going to argue, perceptual justification is defeasible? The significance is that the defeasibility of perceptual justification might be thought to create problems for my view on at least two different fronts. On the one hand it might be thought to imply that perceptual justification is inferential. This is what I have been calling the Argument from Defeasibility against the view that perceptual justification is non-inferential. On the other hand the defeasibility of perceptual justification might be thought to call into question the idea that perception really puts us in touch with non-psychological reality in such a way as to make its layout manifest to us.

I'm going to argue that both these claims about the alleged consequences of perceptual defeasibility are unfounded. I think it does not follow from the fact that perceptual justification is defeasible that it is inferential or that it fails to put us in touch with reality in the way in which I have been maintaining in this thesis. So perceptual defeasibility therefore poses no threat at all to my view.

The plan in this chapter is this: in the next section I'm going to defend the claim that perceptual justification is defeasible, despite being infallible. I'm going to distinguish between several different senses of 'defeat' and argue that perceptual justification is defeasible in each of these senses. Then, I'm going to address the Argument from Defeasibility. I will show that this argument fails since, although it is sound, it is invalid; perceptual justification can be defeasible without being inferential. In the following section I will tackle the other big worry about accepting perceptual defeasibility – that is, the worry about manifestation – and explain why that worry is also unfounded. In the final section of the chapter I will draw attention to an important structural difference between my position and traditional foundationalism with respect to defeasibility. I will argue that once we see how perceptual justification can be defeasible, we will also see that traditional foundationalists were wrong to think that we have indefeasible access to psychological reality.

2. Perceptual Defeasibility

I have said that I will be arguing that perceptual justification is defeasible. To say that a given justification is 'defeasible' is to say that it is capable of being defeated. But what does that mean and why should we think that it is true of perceptual justification? In the literature there is more than one way of understanding the notion of defeat and it's not always obvious what they all have in common. As a rough, first stab we might say that in cases of defeat the subject is deprived of knowledge she might otherwise have had. Perceptual justification is therefore defeasible insofar as the knowledge it grounds is

capable of being defeated in this way. This is not completely useless, but it's not that helpful either. If I shoot you in the head I deprive you of knowledge you might otherwise have had, but it would be wrong to call my shooting you a 'defeater' in the sense in which we are interested.

We can get a clearer fix on what is at issue by looking at some concrete examples of perceptual defeat. Goldman's famous barn example is often thought to be a classic case in which a perceptual justification is defeated (Goldman 1992: 86). Seeing a barn in good light and from a reasonable distance is normally a basis on which I can come to know that what I see is a barn. But suppose, as in Goldman's example, that the barn I see in good light and from a reasonable distance is the only real barn in an environment full of fake barn facades - indistinguishable in good light and from a reasonable distance from the real barn that I actually see. In that case, I do not acquire knowledge that there is a barn in front of me despite the fact that I would know that it's a barn if I were seeing it in more favourable circumstances.

So the justification that seeing a barn (in good light and from a reasonable distance) provides for believing that something is a barn is capable of being defeated; while it will sometimes enable one to acquire knowledge that what one sees is a barn, it will not always do so. The presence of the fakes can prevent one acquiring knowledge that one might otherwise have had and it's in that sense in which their presence is said to 'defeat' one's justification.⁷⁷ Let's call this case BARN.

⁷⁷ At this stage we can leave it open whether or you one would still be justified in believing that there is a barn in front of you. That depends on whether you still sees that there is a barn n front of you. This is an issue to which I return below.

BARN is an example in which I never actually have the knowledge I might otherwise have had. But it is easy to imagine cases in which I am deprived of perceptual knowledge I intuitively did once have. Suppose I am looking at the real barn at t_1 and all is normal; intuitively I know that there is a barn in front of me. But what if, while I am looking, a series of fakes is constructed around me. At t_2 , when the construction is finished I do not know that there is a barn in front of me since at t_2 I am in precisely the same situation that I was in in BARN and we have already conceded that I do not know the truth of the proposition in that case. So while I knew that I was facing a barn when I first started looking, that is something I no longer know. So I am deprived of perceptual knowledge I did once have.

These are both examples in which I am deprived of perceptual knowledge by facts about the situation of which I am unaware; I don't know anything about the fakes, but they still prevent me acquiring knowledge. There are also cases of perceptual defeat in which I am deprived of knowledge by evidence that I do possess. Consider the following example from Michael Martin:

Suppose you know that I have a system capable of causing perfect hallucinations of oranges in subjects, and that I regularly run tests where I alternate the actual viewing of an orange with a perfect hallucination of one. You subject yourself to my machine. Unknown to you the machine has developed a serious fault and is incapable of causing hallucinations: if it looks to you as if there is an orange there, then that could only have been because you are seeing one. Nonetheless, you have

information which seems sufficient to make rational a doubt on your part as to whether there really is an orange before you when it looks to you as if that is what is there...you have reason sufficient to undermine the warrant that experience provides for judgement (Martin 2001: 444-5).

Let's call this case ORANGE. ORANGE fits in better with some of the characterisations of defeat that one now finds in the literature. Thus Williamson claims that:

Define a way of having warrant to assert p to be defeasible just in case one can have warrant to assert p in that way and then cease to have warrant to assert p in that p merely in virtue of gaining new evidence (Williamson 2000: 265).

This fits in better with ORANGE, since in ORNAGE I presumably would have known that there is an orange before me had I looked before having heard anything about the hallucination machine. But when I am told, I then cease to know in virtue of the evidence that I gain.⁷⁸ As we saw earlier, though, not all cases of defeat work in this way. We can also be deprived of perceptual knowledge by evidence that we do not possess.

Despite this difference, these two cases do have something important in common. In both cases I am deprived of perceptual knowledge by certain epistemological facts about the situation that obtain independently of what I happen to believe and that operate irrespective of these beliefs. I am therefore going to call these epistemological defeaters since they represent genuine epistemological obstacles to perceptual knowledge.

Notice, the fact that perceptual justification is capable of being defeated in this way is

⁷⁸ This is intuitively quite unlike BARN since in that case the presence of the fakes deprives me of knowledge whether or not evidence they exist is ever in my possession. For more on defeat via evidence one does not possess, see (Harman 1973: Ch. 9).

perfectly compatible with my claim that perceptual justification is infallible. In these cases it is still true that I cannot be mistaken in believing there is a barn or an orange when I judge that there is a barn or an orange because I see that there is a barn or an orange before me. It is just that in these cases I could easily have formed that belief on a different basis (e.g. by looking at a fake barn façade) and in which case I would have been mistaken. However, it's unclear that one really sees that there is a barn or an orange before one in these cases and if I do not see that those things are so, I will not be justified in believing that they are so. We saw earlier that to have perceptual justification, I have to be in the right circumstances and these facts (that is, the presence of the fakes and the hallucination machine) may make it the case I am not in such circumstances. If that is right then the presence of the fakes and the hallucination machine don't just deprive me of perceptual knowledge I might otherwise have had. They deprive me of perceptual justification I might otherwise have had; in turn, this may be what explains why I lack the knowledge I might otherwise have had.

So in cases of epistemological defeat I am deprived of perceptual knowledge by the existence of some genuine epistemological obstacle. There are also cases of defeat in which I am deprived of knowledge by the presence of an obstacle that is purely psychological. In these cases I do not know simply because I do not believe, not because there is any genuine epistemological obstacle to believing. In these cases I deprive myself of knowledge and I do that by refusing to believe what I do in fact have the best possible grounds for believing. I am therefore going to call these cases of psychological defeat since they function by bring about a certain psychological result, viz. lack of belief.

Cases of psychological defeat are also clearly possible where perceptual knowledge is concerned. Suppose you ask me why I believe that I have hands and I answer by citing the fact that I can see that I have hands; this is what makes it the case I am justified in believing I have hands. Now imagine you have been busy reading Vogel's paper 'Are there any Counterexamples to the Closure Principle?' (Vogel 1990) over the weekend and you point out to me that knowing that I have hands entails I also know I am not a handless brain in a vat being artificially feed all my experiences, including the experience I now have of seeming to see that I have hands. So, you pointedly ask: am I *really* claiming to know that I have hands? All of a sudden I feel terribly flustered – somehow you always manage to get the better of me in these sorts of arguments. So I conclude that I don't really know that I have hands, since I don't really know that I am not a handless brain in a vat and cease to believe that I have hands on that basis; agnosticism, I think, is the safest policy for me.

Let's call this case HANDS. Clearly we could elaborate HANDS in such a way that it is just an example of epistemological defeat. Conceivably, you might give me excellent reasons for supposing that I am a handless brain in a vat. We might live in a world very unlike the actual world in which 1/5 people are in fact secretly envatted and you might point out how relatively high the probability is that I am too and how, if I were, things would seem just the same to me as they now seem. Or, you might coolly remark that you only asked as the experiment is now drawing to a close and all vats are to be unplugged, before manipulating the course of my experience in such a way as to make it extremely

plausible that you are the mad scientist who has cruelly envatted me. Clearly, these are both ways in which we could have developed HANDS and in both of these cases it would have been plausible that I do not know that I have hands. However, these are not cases of psychological defeat. In these cases I don't just fail to know because I fail to believe. In these cases I fail to know because there is a genuine epistemological obstacle to my knowledge.

I have deliberately not elaborated HANDS in this way. In HANDS as I describe it, you merely raise the possibility that I could be a brain in a vat and I am so moved by this speculation – so overly moved – that I conclude I do not know that I have hands, and so refrain from believing that I have any. I am assuming, in other words, that not any old consideration in favour of *p* counts as a genuine reason to believe *p*. In this case I am also deprived of perceptual knowledge I might otherwise have had: my justification is defeated. But that is because, and only because, I no longer believe that I have hands, not because there is any genuine reason for me to believe that.⁷⁹

Cases of psychological defeat are even more obviously compatible with the fact that, as I have claimed, perceptual justification is infallible. In these cases I don't believe anything

⁷⁹ Martin's case is therefore not a case of psychological defeat as I am employing that label. In the passage from which the quote is taken, he does stress the fact that the subject refrains from believing that there is an orange before her and it is clear he thinks that is relevant to whether or not she has knowledge. He writes (in the section I omitted above) "Given that doubt, you do not endorse appearances, and despite the fact that it looks to you as if that (the thing before you) is a certain way (the way an orange can look), you refrain from making any judgement about the matter. So, the experience you have is independent of your beliefs – it can look to you as if something is that particular way without you so believing it to be. And your failure to believe is a reflection of the defeasibility of perceptual justification – you have reason sufficient to undermine the warrant that experience provides for judgement." (Martin, *ibid.*) As the last line makes clear, you have *reason* sufficient to undermine the warrant that experience provides for judgement. If so, there is a genuine epistemological obstacle to your knowledge; the obstacle being that you lack perceptual justification. Whether this is a plausible description of the case he gives is a separate question.

so there is no question of my belief being mistaken. I merely refrain from believing something I might otherwise have believed, but ignorance is not a form of error. Moreover, what I might otherwise have believed is that I have hands and that is not a belief that could have been mistaken in any case, given that I really do see that I have hands. Your speculation does not interfere with that. Unlike epistemological defeat psychological defeat needn't necessarily deprive me of perceptual justification; I still see that I have hands, and so still have justification for believing that I have hands. It is just that I fail to believe that I have hands and so trivially lack knowledge that I have hands.

So we now have two sorts of perceptual defeat – epistemological and psychological. Some people will probably want to question whether the examples that I have given are all cases of actual defeat. Whether that's so, depends on whether the subjects actually lack perceptual knowledge in these cases, and here opinions may differ. One might think that I do know that there is a barn in front of me in BARN. Maybe I am lucky, but knowledge does not exclude every kind of luck imaginable.⁸⁰ I think that is not completely implausible.⁸¹ Conversely, hard-core 'externalists' will probably maintain that the mere fact that it would be unreasonable for me to believe there is an orange before me, once you've told me about your hallucination machine, is not enough to deprive me of knowledge if I persist in believing that there is an orange before me despite you. This looks less plausible to me, though still clearly possible. Still others will reject the way that I have described HANDS. They will insist the mere fact that I no longer believe that

⁸⁰ For some nice examples see (Sainsbury 1997).

⁸¹ Clearly, intuitions are heavily effected by the vagaries of description. For instance, the fact that barns are actually pretty big objects seems to make a difference. Lots of people's intuition that you know the coin you have just picked out of your pocket is a real 50p, despite the overwhelming preponderance of fakes, is much less robust.

I have hands does not prevent me knowing that I have hands, since belief is not a condition for knowledge.⁸² How plausible that is will depend on how firmly entrenched is the link between knowledge and belief. Philosophers have tended to assume it is pretty deep, though the old adage ‘I knew it! I just didn’t believe my eyes’ suggests it may be less so.⁸³ Finally, even those who don’t deny that these particular cases are all cases of actual perceptual defeat may disagree about the sort of defeat they represent and whether it is properly regarded as epistemological or psychological in nature.

These are all things someone could say by way of response. I do not intend to take issue with such an opponent here. My point isn’t that these particular cases must be cases of perceptual defeat and my aim is not to give you a long list of considerations that really do defeat a given perceptual justification.⁸⁴ My point is just that there *are* cases are which perceptual justification is defeated, whether or not these cases are among them. There really are cases in which we are deprived of perceptual knowledge that we intuitively might otherwise have had. This weaker claim is much harder to reject and it is enough to show that perceptual justification is capable of being defeated.⁸⁵

⁸² There are also cases where I persist in believing *p*, despite believing I do not know *p*. For instance, I might claim not to *know* that my car is parked where I left it since I don’t know that it hasn’t just been stolen, yet still *believe* it is parked where I left it. In that case I won’t know that I know since I don’t believe that I know, but I will still know where my car is parked.

⁸³ There are also the famous Radford examples (Radford 1966). For more recent discussion, see Williamson (Williamson 2000: esp. Ch.1).

⁸⁴ Giving such a list may not even be possible – at least not in any kind of general way. Perhaps we can only sort cases by reference to our intuition to count them as cases of knowledge. For helpful comparison, see (Austin 1962: esp. Lecture X).

⁸⁵ There are further distinctions one can draw here. One that is relevant to our purposes is between ‘rebutting defeat’ and ‘undercutting defeat’ – that is, between cases in which defeat gives you reason to believe *p* is false and cases where it gives you reason to suppose that the belief that *p* is inadequately grounded. I will focus largely on the latter in what follows. For helpful further discussion, see (Pollock and Cruz 1999: 36-8)

3. The Argument from Defeasibility

What unwelcome consequences can we draw from the fact that, as I have just argued, perceptual justification is defeasible? According to the Argument from Defeasibility it follows from the fact that perception only defeasibly justifies beliefs about non-psychological reality that it can only inferentially justify them. This is a direct threat to the Broad View of perceptual justification that I have defended. According to the Broad View, the beliefs that perception non-inferentially justifies include beliefs about non-psychological reality. So what should we make of this challenge?

Unlike the Argument from Fallibility, we cannot claim the Argument from Defeasibility is unsound; as we have just seen, perceptual justification is defeasible. But is the argument valid? That is, does it really follow from the fact that perception only defeasibly justifies beliefs about non-psychological reality that perceptual justification must be inferential? Notice that, as in the fallibility case, it is not enough for a defender of this argument merely to point out that there are justifications that are both defeasible and inferential. Once again, induction seems to be a case in point. I can be justified in believing that Captain Molski will win because I am justified in believing that she is the fastest dog on the track. Here my justification is inferential, but it is also capable of being defeated: if I acquire good reason to believe that Captain Molski has been doped I will no longer be justified in believing that she will win, even if (as it turns out) she does. So there certainly are examples of justifications which are both defeasible and inferential. That doesn't show that there is any connection between these two things though or that such justifications are inferential in virtue of being defeasible.

This is what the Argument from Defeasibility needs to establish. How might this be done? I can think of only two strategies. To see the first, take a case in which perceptual justification isn't defeated. Suppose I'm looking at the barn in good light and from a reasonable distance and that there are no fakes in the vicinity. In these circumstances I will ordinarily come to know that there is a barn in front of me. But I will not know this (even in those circumstances) if I happen to believe that I am in fake barn country surrounded by facades that I cannot distinguish from the real thing. If I believe *that*, then it would surely be irrational for me to persist in believing that I am confronted by a barn. Similarly, if I believe that there are excellent reasons for thinking I am currently hallucinating, since you've just told me all about your marvellous machine, it would normally be irrational for me to persist in believing that there is an orange before me when that is how things look to me. And I will not acquire knowledge if it would be irrational for me to persist in that belief.⁸⁶

We can put the point here very simply: we can be deprived of knowledge because we believe that we are in circumstances that really would deprive us of knowledge. In other words the *belief* that we are in such circumstances (or what I will henceforth call 'defeating circumstances') is itself a defeater. This belief can also prevent us acquiring perceptual knowledge we might otherwise have had.⁸⁷ Given that that's so, someone might argue that defeasible justification therefore cannot be non-inferential since it

⁸⁶ Not everyone accepts even that much is true. For a powerful defence of such a view, see (Ayers 1991: 170-1)

⁸⁷ This is not a psychological defeater, it is an epistemological defeater since it represents an epistemological obstacle to knowledge viz. irrationality. In other respects though the case is more similar to cases of psychological defeat. The obstacle is consequent upon something psychological and to remove it one just needs to refrain from believing

depends upon the fact the subject believes that she is not in defeating circumstances.

This is the first strategy for someone wanting to show that defeasible justification must be inferential. It looks distinctly unpromising. At best it threatens to show that my knowledge that I'm confronted by a barn depends upon the fact that I also believe that I am not surrounded by indiscriminable barn facades.⁸⁸ It doesn't show that my justification for that belief derives from the belief that I am not surrounded by fakes or therefore that my justification is inferential. No doubt it would strike us as odd if most subjects who believe that there is a barn didn't also believe that they aren't in fake barn country. This doesn't show their justification for the former derives from the latter any more than the fact it would be odd for me to believe that I am in pain without also believing that someone is in pain shows that part of my justification for believing that I am in pain derives from the belief that someone is in pain. So the most this line of thought promises to show that it is a normal concomitant of a subject's knowing such things as that there is a barn in front of them, that they also believe they are not in fake barn country.⁸⁹ This has no power to show their justification is tacitly inferential.

However, it is doubtful this line of thought even succeeds in showing that much. The starting point, recall, was the observation that it would be irrational for a subject both to believe that they are confronted by a barn and that they are in defeating circumstances.

⁸⁸ Clearly, it only shows that in cases where it would be irrational for me to persist in the belief that I am confronted by a barn. That needn't always be so; it won't be in a case in which I know I am looking at the only real barn and all the others are fakes. In that case I can rationally believe both that I am confronted by a barn and that I'm surrounded by visually indistinguishable barn facades.

⁸⁹ That may be less plausible on a more demanding conception of what belief involves. I am trying to be maximally favourable to my opponent at this point. If belief just involves the disposition to sincerely assent when prompted (not having been put off by being asked so obvious a question etc.) it would be pretty unusual to find someone who has one of these beliefs without the other.

Why should that show the subject must believe that she is not in defeating circumstances? If the obstacle to her knowing is just the belief that she is in such circumstances, we remove that obstacle by removing that belief. In other words, it must be the case that the subject does *not* believe that she is in such circumstances. That does not entail she must *believe* she is not in such circumstances. Beliefs admit of both ‘internal’ and ‘external’ negation (Evans 1982: 226 n. 36). Only the latter, that is, the absence of the belief that one is in defeating circumstances, appears necessary.

Once again it would no doubt strike us as strange if conceptually competent subjects didn’t believe that they were not in fake barn country whenever they believe themselves to be confronted by a barn. But it cannot be a requirement for knowledge that they have that belief. We are not irrational in failing to grasp every entailment of what we believe or draw all the conclusions we are committed to, even comparatively obvious ones.⁹⁰ It is even harder to see why that requirement should hold in the case of less conceptually sophisticated subjects – that is, subjects who lack the concepts necessary to frame thoughts about fake barns or grasp their rational bearing on one’s ability to spot a barn when one sees one. We surely do not want to prevent these subjects having ordinary perceptual knowledge just on the grounds that it would be irrational for them to lack a belief they aren’t even capable of framing. Further, the mere fact they lack those concepts is not an objection in its own right. To claim otherwise would be incredible, especially when we consider the number and variety of different circumstances that would defeat

⁹⁰ Sometimes pointing out an entailment to a subject will lead the subject to suspend her original belief. I might suspend my belief that my car is parked where I left it, when you point out this entails it hasn’t been stolen in the last 5 minutes. Clearly, there is nothing irrational in combining that with the belief that one is in defeating circumstances.

any given perceptual justification. We presumably don't want to require that we all believe with respect to each and every one of those circumstances that it does not obtain; but it seems equally implausible to attribute to us all the more general belief that defeating circumstances do not obtain since that belief is composed of concepts that even philosophers struggle to articulate.⁹¹

So we have as yet no reason to think that perceptual knowledge depends on anything more than the fact that the subject lacks the belief that she is in defeating circumstances obtain.⁹² Strictly speaking, of course, this is not enough; lacking that belief also has to be epistemically appropriate for the subject. We cannot acquire perceptual knowledge just by stubbornly refusing to believe that things are amiss when we have excellent reason for thinking otherwise. In the ordinary case, though, what makes one's lack of belief reasonable is just one's lack of grounds. I normally have no reason to believe that I am hallucinating and that is what makes it acceptable for me to lack the belief that I am. I do not need to have acquired any special reasons for thinking that I am not, or anything therefore which suggests that perceptual justification isn't a perfectly good stopping point in the regress of justification.

⁹¹ Harman opts for the latter strategy: "it is very likely that there is an infinite number of different ways a particular inference might be undermined by misleading evidence one does not possess. If there must be a separate essential conclusion ruling out each of these ways, inference would have to be infinitely inclusive—and that is implausible. Therefore it would seem that the relevant inferences must rule out undermining evidence one does not possess by means of a single conclusion, essential to the inference, that characterises all such evidence. It is not at all clear what distinguishes evidence that does undermine from evidence that does not...since I am unable to formulate criteria that would distinguish among these cases, I will simply label cases of the first kind "undermining evidence one does not possess" (Harman 1973: 150). The objection raised here obviously has less force on less demanding conceptions of what is required for one to count as believing things like defeating circumstances do not obtain. My main point still stands though; even if you do need to have this belief, it is not part of the source of the justification for one's ordinary perceptual beliefs. It therefore has no tendency to show one's justification is tacitly inferential, contra the Argument from Defeasibility.

⁹² To lack the belief that p one doesn't need the concepts that figure in the belief that p. So there is no parallel worry about hyper-intellectualisation on this account.

This highlights an important asymmetry between beliefs and their absence. It would not be epistemically appropriate for me to *believe* that defeating circumstances do not obtain just because I lack reasons for thinking that they do. Beliefs are not justified by ‘default’ or until and unless reasons transpire to the contrary. They require positive support, which the mere fact a belief is true does not provide; it is something I must actively go out and acquire. The same is not true of the absence of belief. It can be epistemically appropriate for me to lack the belief that defeating circumstances obtain just in virtue of the fact I lack reasons for believing that they do; I needn’t have any special reason for believing that they do not.

In this respect, it is the belief that not-p that is the true contrary of the belief that p. This is hardly surprising since only the former is an attitude. The absence of belief is not a stand one actively adopts on the world; it is just the absence of one. My point is that while we are required to have reasons for the stands that we do take, whether pro or anti, we are not required to take a stand on every issue. Sometimes agnosticism is acceptable and unlike positive stances, mere agnosticism – that is, the lack of belief either way – does not call for reasons in the very same way.

This undermines the first strategy for showing that defeasible justification must be inferential. The second strategy appeals to the sort of explanation we can give of why subjects lack knowledge in cases of defeat. The thought here is very simple. Suppose we did think that in cases of defeasible justification part of the subject’s justification came

from the belief that defeating circumstances do not obtain. This would make the subject's justification inferential and it would also offer a neat explanation of why the subject doesn't acquire knowledge in cases of defeat. In these cases the belief that such circumstances do not obtain is false and it is a widely accepted principle about knowledge that subjects can't acquire knowledge where what they believe rests essentially upon a false belief. The 'No False Lemma's Requirement' as Harman calls it, is a principle about knowledge that we have independent reason to accept. So this explanation is parsimonious and that is an explanatory virtue.⁹³

This is an idea that many philosophers have been moved by. The assumption that defeat is to be explained in these terms informs all the early work in this area; this is why Lehrer and Paxson begin their seminal paper by assuming that defeat can only operate in cases of 'non-basic' knowledge: that is, cases where:

a man knows that a statement is true because there is some other statement that justifies his belief (Lehrer & Paxson 2000: 31).

Later on, recognising that defeat also applied in the perceptual sphere, Harman claimed:

I shall argue that we cannot easily account for perceptual Gettier examples unless we assume that even simple perceptual knowledge is based on inference. In that case, the perceiver can be assumed to infer that the explanation of there seeming to be a candle ahead is that he is seeing a candle there. He comes to know that there is a candle there only if he is right about why it seems to him that there is a candle there (Harman 1973: 23).

⁹³ This is obviously a more general version of the strategy pursued in relation to Gettier's original examples, which claimed they involve inferences resting essentially upon false beliefs. For a comprehensive survey of response to the Gettier examples, see (Shope 1983).

So the idea that defeat is to be explained in terms involving beliefs is one with an illustrious history.

Nevertheless, it is not the only explanation of why subjects lack knowledge in cases of defeat, or indeed the best explanation. As we are about to see there are other explanations that do not involve appealing to beliefs at all. Consider BARN: why should seeing a barn in fake barn country not be a way of acquiring knowledge that there is barn before you? The obvious explanation, surely, is that you could easily have gone wrong in believing there is a barn in front of you.⁹⁴ This explanation does not appeal to beliefs, so it will not make the subject's justification inferential. Yet it appeals to a principle about knowledge that we have quite as much independent reason to accept as the No False Lemmas requirement viz. the principle that subjects cannot acquire knowledge where what they believe could easily have been false (Sainsbury 1997; Sosa 1999; Williamson 2000).

The 'Safety Requirement' as it is called can also explain why some (if not all) of the cases in which evidence we do possess deprives us of knowledge. Consider ORANGE: if the hallucination machine very rarely fails to function properly, I could easily have gone wrong in believing that there is an orange before me and ignoring the evidence to the contrary.

However, perhaps not all of the relevant cases can be handled in this way. As Martin describes ORANGE it is meant to be a case in which I couldn't easily have gone wrong

⁹⁴ When we change the example so that the subject could not easily have gone wrong (e.g. suppose she has a guide who know the area and that he wouldn't have brought her to see a fake barn) the intuition that she no longer knows is correspondingly less robust.

in believing that there is an orange before me, but in which I still lack knowledge. As he sets things up:

...the machine has developed a serious fault and is incapable of causing hallucinations: if it looks to you as if there is an orange there, then that could only have been because you are seeing one (Martin, *ibid.*).

If cases like this are genuinely possible – and I here leave it open whether or not they are - then we need another explanation of why they aren't cases of knowledge.⁹⁵ But even here alternatives aren't impossible to find. So at worst we'll be left without a unitary account of why all the different cases count as cases of defeat count as such. This is not an objection.

A different explanation might appeal to the fact that there are certain canons of reason or rationality, and that knowledge is not compatible with the violation of these canons. If you tell me you have a machine capable of causing perfect hallucinations of oranges and we have a long and trusting relationship on the basis of which I know you are very likely to be well-informed about the topic at hand, it would not normally be reasonable for me to simply disregard what you have said and persist in believing what I would otherwise have believed. On the face of it what other people tell us, particularly those we trust, can be a genuine reason for us to believe what they say, despite the fact that in ignoring them we might not easily have gone wrong. This may be enough to explain why we lack

⁹⁵ What I mean is this: something might not qualify as a 'genuine' reason to not to believe-p, unless it is true that in believing p you could easily go wrong. Whether that's so depends on how we finesse the idea that one could easily have gone wrong. We could just fix it in such a way that it automatically covers all these cases and so insist that if we lack knowledge it must be because we could easily have gone wrong. But what would that prove? In any event, Martin's case might still represent a case of purely psychological defeat.

knowledge.⁹⁶

This explanation isn't the same as one which appeals to the Safety Requirement, but it doesn't essentially appeal to beliefs either. What makes it the case that I lack knowledge in these cases is the fact it would be unreasonable for me to believe the proposition in question and there need be no further explanation of why that is so in terms of some other requirement. The idea that certain things run counter to reason in this way is just as basic a part of our thought about knowledge as any. There need be no force to the demand we explain that it in other terms, let alone, terms that involve beliefs.

This explains epistemological defeat. Only psychological defeat now remains. But psychological defeat is easy to explain since in cases like HANDS the reason one lacks knowledge is simply that one lacks belief and belief is I am assuming an essential component of knowledge. Clearly, this explanation does appeal to a fact about what the subject believes (or fails to believe): to know that I have hands I must believe I have hands. So if I do not believe that I have hands, I will not know that I have hands either. However, this can hardly be thought to render my justification inferential. The belief that I have hands is not part of what gives me justification for believing that I have hands on anyone's view.

So I have now explained why in all three cases of perceptual defeat the subject lacks

⁹⁶ What those canons dictate will presumably be different in the case of different subjects. Information it is unreasonable for me to ignore might not be unreasonable for you to ignore. For instance, even if I cannot rationally ignore the possibility I am hallucinating in ORANGE, it is hard to believe the same is true of subjects who do not grasp the rational bearing of hallucination-machine's on one's ability to know an orange when one sees them. I here leave undone the difficult job of spelling out these conditions in any kind of general way (assuming such specification is even possible).

knowledge she might otherwise have had. In none of those cases does the explanation appeal to the fact the subject's justification derives ineliminably from what she believes. So in none of these cases does it follow from the fact that perceptual justification is defeasible, that it is inferential. This disposes of the Argument from Defeasibility. Unlike the Argument from Fallibility, the Argument from Defeasibility has true premises; but like the Argument from Fallibility, it is invalid.

4. The World Made Manifest

I have now shown that it does not follow from the fact that perceptual justification is defeasible that it is inferential. Perhaps, however, perceptual defeasibility poses a different sort of threat to my view. I have claimed that perception is a source of justification because it puts us in touch with the objects and events in the world about which we judge. On my view, perception brings the world to consciousness and in doing so makes the layout of reality manifest. How can that be if, as I have argued, perceptual justification is defeasible?

The second challenge that I'll be looking at under the heading of defeasibility claims that these two features of perceptual justification are incompatible: if a subject can rationally refrain from judging that she has hands even when she sees that she has hands, then what she sees cannot really make it manifest to her that she has hands. So my view of perceptual justification which attempts to combine these two features is incoherent.

In response to this challenge I am going to argue for a version of what I will call

‘compatibilism’. I will argue there is nothing incoherent in thinking both that perception does make the layout of reality manifest, and that the justification it provides is capable of being defeated. The reason they are compatible is very simple. What perception makes manifest is the world, but it is not always manifest to us whether or not we are perceiving. There are states subjectively indiscriminable from genuine perceptions (cases of hallucination, say) in which the world is not made manifest though it seems to be. If one were in such a state, one would have no reason at all for judging. On my view, unlike some of the views discussed previously, to have reason to believe that you have hands you must see that you have hands; it must be those very hands, right there in front of you, of which you are aware and merely seeming to see that you have hands does not give you that.⁹⁷ Thus, to think that you are in such a state (that is, a state in which you merely seem to see that you have hands) is to think you are in a state in which you would have no reason at all for judging that you have hands. Clearly, a rational subject who thinks that will refrain from judging that she has hands. Rational subjects, after all, will not judge where they take those judgements to be unfounded.⁹⁸

In a case like HANDS the subject is wrong to think that she does not see that she has hands. Still, this needn’t make her irrational. It is just a fact about the nature of rationality, as opposed to justification, that we can be moved by doubts that we have no

⁹⁷ I’m obviously not claiming that this is the only way a subject can have reason to believe she has hands. Blind people have also reasons for thinking that they have hands but their reasons do not come from visual perception. My point is that you do not have any reason to believe you have hands merely in virtue of seeming to see you have hands.

⁹⁸ The same is not so obviously true if we substitute ‘could’ for ‘is’: I may have reasons for thinking that I could be hallucinating (I might think that is a logical possibility, however remote), but that doesn’t automatically mean that I have any reason to think that I am hallucinating or therefore that I cannot rationally judge that I have hands. We can rationally judge where we think our judgements *could* be unfounded even if (in those circumstances) we can also rationally refrain.

genuine reason to be moved by and do so without being irrational. This, in turn, can make it rational for us to believe things for which we have no genuine justification.⁹⁹

This explains how it is possible for a subject rationally to refrain from judging that she has hands even though she plainly sees that she has hands. She can rationally refrain since she can see that she has hands without knowing that she sees that she has hands. If she does not know that she sees that she has hands, any doubt about whether or not she sees may rationally lead her not to believe that she has hands, despite the fact that *is* something perception puts her in a position to know.

This is all perfectly compatible with my view. It is no part of my view that perception makes it manifest to us whether or not we are perceiving; it is the world, not our epistemic access to it, which perception makes manifest.

This does not mean that we must *know* that we are perceiving in order to acquire perceptual knowledge. Clearly, this would be a threat to my view that perception genuinely makes knowledge of the world available. But it's simply not true. Even in HANDS, the subject need not abandon the unfounded suspicion she is hallucinating and believe that she has hands in order to know that she has hands. Nothing more is required and that is precisely why the knowledge is genuinely available to her.¹⁰⁰

⁹⁹ This is not a problem vis-à-vis the regress since that argument is concerned with the conditions under which one's beliefs are justified. The regress assumes that justification is a positive epistemic status. In contrast, a belief might be 'rational' (or not irrational) just as long as it doesn't conflict with other things you believe.

¹⁰⁰ Strictly speaking, I suppose perception might be accused of failing to make the world 'manifest' on these grounds: one has got to lack the belief that one isn't perceiving. If that is right, I can't think of any good reason for thinking perception must make the world manifest in this demanding sense.

Of course, people who suspect they are hallucinating often require more convincing. This does not mean that they are right to demand more, or that the rest of us who do not, and who persist in our ordinary perceptual beliefs are somehow in the wrong or being wantonly irresponsible. The ‘high’ standards of the epistemically cautious do not make them any more principled, than do the ‘high’ standards of the person who refuses ever to cross the road. In neither case are we rationally obliged to be so cautious; so in neither case can we be faulted in failing to be.¹⁰¹

This is not to deny that there is any sense in which we are ‘lucky’. We are lucky we are not so risk averse, but there is nothing objectionable about that. Being lucky enough not to be moved by certain unfounded doubts and not to suspend one’s ordinary perceptual beliefs in their face is not an epistemic vice. It is an epistemic virtue since the beliefs one otherwise abandons are ones for which one has the best possible justification. This is not like running out into the middle of the road, eyes-closed; it is like crossing having looked left, right, and left again. Since this is all we are required to do, we cannot be blamed for having done more.

At the same time, though, as the possibility of psychological defeat makes clear we are not rationally obliged to be bold and to persist in our beliefs in the face of doubt. It is an

¹⁰¹ It is just a mistake to think having extra assurance always makes one’s original grounds better. If I know I locked the front door it doesn’t matter how many times I go back and check. Double-checking may put my mind at rest but it doesn’t make me ‘know’ any better. Of course, I may be in a better position to defend my belief that it is locked when you ask me. But that is just reflection of the fact one is better placed in arguing with an opponent the broader the range of considerations one can adduce in one’s defence. That way one is more likely to find common ground. That does show anything interesting about what it is to have justification; it just tells us something about what makes for success in arguments.

epistemic virtue, no doubt, to persist in the beliefs that one has justification for believing since that way lies knowledge. At the same time, however, to fail to be bold – that is, to fail to believe all that one has justification for believing – does not necessarily make one irrational.¹⁰² Extreme caution is an epistemic vice in some sense, but it is not a vice of reason. It is more like being stuck on the wrong side of the street when the shop one wants to get to is on the other side; we do not have to cross, it's just a bit annoying.

I have claimed this is what makes room for the possibility of subjects – rational subjects – whose extreme caution prevents them knowing everything they might otherwise have known. The world is as manifest to them as it is to us. They simply choose not to take advantage of it. This is not something perception can force them to do though. So our theory cannot be blamed for failing to do so. Given that is so, perceptual defeasibility poses no threat at all to the view that perception genuinely makes the world manifest to us.

I have now explained the sense in which perceptual justification is defeasible and explained why that is not a threat to my view. It does not make perceptual justification inferential, contra the Argument from Defeasibility. And it does not mean that perception isn't a way in which the world is made manifest. So we can accept that perceptual justification is defeasible. This is something traditional foundationalists were much more reluctant to accept. In the rest of this chapter I want to try and explain why that is so: how is it that something which I claim is perfectly acceptable could once have seemed so very unacceptable?

¹⁰² If one persisted indefinitely of course, eyebrow might start to be raised.

5. Historical Foundationalism Reconsidered

In the previous section I claimed that the world could be manifest to us in perception without our access to the world being similarly manifest, and that we could in consequence fail to know things about non-psychological reality that perception put us in a position to know. Traditionally, the same was thought not to be true of psychological reality. Historical foundationalists, in particular, thought that the layout of psychological reality could not rationally remain a mystery to us did we but properly attend to it. This is what Hume was getting at when he writes:

For since all actions and sensations of the mind are known to us by consciousness, they must necessarily appear in every particular what they are, and be what they appear. Every thing that enters the mind, being in reality a perception, 'tis impossible any thing shou'd to feeling appear different. This were to suppose, that even where we are most intimately conscious we might be mistaken. (Hume 1978: 190)

There is certainly something philosophically gripping about the idea that psychological reality is self-intimating in this way. Unlike the case of non-psychological reality where we can appeal to the possibility of hallucination, say, there is no obvious explanation of how the mental can ultimately remain hidden from view.

Suppose this is right. If so, there is a difference between beliefs about psychological reality and beliefs about non-psychological reality after all, and a difference between traditional foundationalism and the more modest form of foundationalism that I have

defended. But the difference is not in whether or not the foundational beliefs are infallible as the Simple Reading would have us think. I have also claimed that the foundational beliefs are infallible, since beliefs about non-psychological reality are infallible where those beliefs are justified by perception. The difference is that traditional foundationalists thought that the foundational beliefs are *indefeasible*. They thought that these beliefs are incapable of being rationally revised and unlike beliefs about non-psychological reality, only beliefs about psychological reality can plausibly be though enjoy that privilege.

I have already shown that traditional foundationalists were wrong to think that the foundational beliefs have to be indefeasible because I have already shown that beliefs about non-psychological reality can be both defeasible and non-inferential: such beliefs can provide acceptable stopping points in the regress of justification, despite the fact that they are capable of being defeated. The interesting question is why these thinkers were so keen on indefeasibility in the first place. There are lots of bad reasons for thinking that it matters – the Argument from Defeasibility being one. But are there any good reasons for thinking that indefeasibility matters?

One reason is that we desire knowledge - that is, the freedom from ignorance and not just the freedom from error. The latter is easy: believe nothing and you certainly won't be mistaken. But you won't know anything either and it is not an appealing epistemic policy for precisely that reason. As epistemic agents, we want our ignorance removed. It is a consequence of my view, however, that we can fail to know things that perception genuinely puts us in a position to know. We can rationally remain ignorant of certain facts

about non-psychological reality.

This is not a live option with indefeasibly justified beliefs. If the historical foundationalists are right there are no 'ringers' for inner perception to which we can plausibly appeal in making sense of the possibility we might fail to know all that we are in a position to know. In the case of our own minds mistakes aren't just impossible, they are inconceivable and any failure to believe must be a failure of reason.

So it is clear why indefeasibility would be nice: it buys us a special sort of security that I have claimed perception genuinely does not secure. Still, this is not a good reason for thinking that foundational beliefs must be indefeasible or for privileging beliefs about our own minds at the expense of all others as the traditional foundationalists were lead to do. The price of our absolute knowledge of psychological reality turns out to be ignorance of non-psychological reality. This ought not to constitute a good bargain by anyone's lights.

Moreover, while it might be desirable to know what we are in a position to know it is not an obligation: nice does not mean necessary, and sometimes all is not equal. I might be in a position to know that the pain in my leg is more like an itch than it is like a tickle. But there are more important things than that to worry about in epistemic life. In failing to know that, and focusing my attention on more important matters, I don't do anything wrong from an epistemic point of view; I do something right.

A second reason for cleaving to a picture of indefeasible foundations springs from

‘internalism’ of a certain familiar sort. As we have seen it is a consequence of my view that I can see that I have hands without thereby knowing that I see that I have hands. So I can be in a position to know that I have hands, without knowing that I am in a position to know that I have hands. The world can be evident to me without my epistemic access to the world also being evident.

This strikes many philosophers as uncomfortably close to denying the so-called KK Principle that defines a familiar sort of internalism. The KK Principle says that if one knows that p, then one knows that one knows that p. But this principle is false: I can know that there are no typographical errors in my thesis because, having checked each page individually, I know of each that it contains no errors and so infer that the thesis as a whole contains no errors. Valid deductive reasoning from known premises is normally a way in which to extend our knowledge and there is no reason why this case should be an exception. So I can know that my thesis contains no errors. Still, I do not know that I know that it contains no errors. The fact that most documents that long do contain typographical errors may make it the case I do not believe that I know that it contains no errors; if you were to ask me whether or not that is something I know I would in all likelihood deny that is so. Nonetheless, I might still believe it contains no errors and so might still know that it does. Much of our knowledge has this status – this is how we establish that there are exceptions to rules we once thought were universal and it is hard to see how science could get by if that were not so.

Although the KK Principle is false it has residual philosophical appeal. Part of that appeal

derives from our tendency to over-assimilate the case of beliefs and procedures or methods for arriving at beliefs. As we have seen beliefs require reasons. This is just the moral of the regress argument: beliefs do not justify themselves. It is tempting to think the same must be true of procedures or methods; to think that unless we have some positive assurance that we are correctly employing a given procedure we cannot acquire justification by relying on it, just as in the parallel case of beliefs, they do not assume any positive epistemic standing merely in virtue of being held or employed.

Some philosophers think that this is obviously true: Crispin Wright is a case in point. He writes of an analogous principle, which he calls the ‘Proper Execution Principle’ that “it is apt to impress as barely more than a platitude” (Wright 1991: 99). This principle claims that:

If the acquisition of warrant to believe a proposition depends on the proper execution of some procedure, then executing the procedure cannot give you any stronger a warrant to believe the proposition in question than you have independently for believing that you have executed the procedure properly (Wright 1991: 99).

From this perspective defeasibly justified beliefs pose a problem. In the case of beliefs justified by perception we have seen that we do not automatically have any assurance that we are perceiving the world – that we really do see that we have hands and do not merely seem to see them. How, then, can it be that one acquires justification for believing that one has hands? This can look like the merest good fortune – a reckless, irresponsible gamble of a sort we normally think incompatible with knowledge.

This is no doubt where the appeal of indefeasibility lies. If a belief is indefeasibly justified we cannot rationally doubt whether or not it is well founded. In such case it isn't just the world which is evident to us, so too is our epistemic access to the world. From the 'internalistic' perspective now being considered this alone can seem to provide the sort of assurance that a properly responsible epistemic subject would demand.

But while this is where the appeal of indefeasibility lies, it is also clear where the blame should lie. Though tempting, internalism in this sense is hopeless and the analogy from which it springs a bad one. It is a bad analogy since beliefs aren't in the final analysis anything like methods for forming them, any more than they are like their own absence. This is so because of all the obvious differences between them: what is puzzling are not the differences, but how anyone could think otherwise.¹⁰³ Beliefs are unusual in more respects than one and there is simply no good reason to think they are central in the way this line of thought suggests or that they provide a useful paradigm upon which other epistemological categories can all be modelled.

One difference concerns the need for prior vindication. In the case of a 'procedure' like perception we do not require independent reasons to believe that we are perceiving before we can acquire the knowledge that perception vouchsafes for us. This fact is just as basic a fact about justification as they get. The same need not apply equally or without qualification to all methods. Perhaps clairvoyance or wishful thinking (could we but

¹⁰³ Most obviously, only belief is an attitude. I am certainly not claiming one can be justified in believing one can rely on a method in the absence of reasons. I am merely claiming that one can rely on it in the absence of such reasons.

regard them as procedures for finding out about the world) would require such certification in order for the beliefs they deliver to be ones for which we have good reasons. However, this has less to do with ‘procedures’ in general and more to do with the specific nature of clairvoyance or wishful thinking and the fact that it is not really intelligible that the justification they deliver is anything *other* than inferential. Clairvoyance only gives us reasons insofar as we have some independent assurance that our clairvoyant powers are reliable or are properly functioning – that is how we make sense of it as a way of finding out about the world at all.

This is fundamentally different from what I have claimed is true of perception. If I am right, perception offers us a fundamentally different model of what it is - at the most basic level - to have reasons for one’s beliefs. And if we have reasons for our perceptual beliefs what more could we possibly require?

This is not to deny that we ever want more, or even that more is impossible to get. We can and sometimes do have reasons for believing that we are not being deceived in certain ways. My point is merely that such assurance is not an indispensable part of what it is to have perceptual justification in the first place: it is additional assurance. Moreover, where we do know that we are not being deceived, that knowledge is not basic knowledge: if we know that we are not hallucinating, it is not because it is evident to us that we are not hallucinating in the way in which my hands are evident to me when I see them. If we know that we’re not hallucinating or being deceived it is on the basis of other things that we know.

A foundationalist can hardly object to that. The starting point for that position was the idea that there are two fundamentally different sources of knowledge or justification: there is what we know because of what we are presented with via sources such as perception; and there is what we know via inferences from that knowledge - things which, while not themselves evident, are knowable on the basis of what is evident. My view places the knowledge that we are perceiving – that is, knowledge about our epistemic access to the world – at the second level. It does not place it out of reach altogether.

This makes room for the possibility of a certain sort of self-vindication as far as perception goes. If you do not need independent reasons to believe that you are not hallucinating in order to acquire the perceptual knowledge that there is a hand before you then you can establish that you are not hallucinating by inferring that you are not hallucinating from what you know on the basis of perception, since you can reason that ‘if *that* is a hand, then I cannot be hallucinating, and that is a hand; so, I cannot be hallucinating’. This gives us a different way of holding onto the KK principle latterly rejected since this argument (or one just like it) is, in principle, available to anyone in possession of ordinary perceptual knowledge. Thus, even if we do not always have independent assurance that we are not hallucinating (as proponents of this principle presumably always wanted) assurance is in a sense always available.

Many philosophers reject the idea that any source of knowledge can vindicate itself in

this way. It is certainly not something we leave room for in the case of beliefs and there is, admittedly, something artificial about such arguments. It misrepresents the situation we are in to suppose the first time that any of us comes to know such things as that we are perceiving (or that ‘the external world exists’) is when, as philosophy undergraduates, we run through ‘Moorean’ arguments such as these. But if what I have said is right, the possibility of self-vindication is one that we must learn to live – at least in the case of perception. Indeed, if what I have said is right, it is what makes room for the possibility of any sort of vindication at all. Those philosophers who reject it and cling to KK are therefore hankering after something they simply cannot have.¹⁰⁴

6. Infeasibility and the Mental

Up until now I have pretended to go along with the traditional foundationalist and have assumed that there is a significant asymmetry between beliefs about psychological reality and beliefs about non-psychological reality with respect to their defeasibility. I have conceded that infeasibility would be nice if we could get it, but claimed there is no reason to think the fact that we can’t has the unfortunate epistemic consequences the traditional foundationalists suspected. Perceptual justification is defeasible, but it is still non-inferential and perception is still a way in which the layout of reality is made manifest. This should be enough. Indeed it is enough for our purposes. But there will always be those greedy souls who think we should be able to have our cake and eat it – every last crumb. So I want to end by encouraging a little optimism among the

¹⁰⁴ I am not denying that my opponent has arguments in favour of the alternative view; I am denying that he has good arguments. For Wright’s own response to the problem he raises, see (Wright: 1991; and previously 1985). For a more recent response, in some ways similar to Wright’s, see (Zalabardo, forthcoming).

avaricious.

To see why optimism might be warranted on this front recall that the basis of my claim that perceptual justification is defeasible was the idea that we cannot always tell whether or not we are perceiving. I can see that I have hands and yet rationally refrain from believing that I have hands, since it is not evident to me that I see that I have hands. But whether or not I see that I have hands is arguably part of psychological reality. Seeing that I have hands is a psychological state I may or may not be in. So if I cannot, by introspection alone, determine whether or not I am in it, I cannot determine by introspection alone all there is to know about the layout of psychological reality. I can rationally refrain from believing something that is true about psychological reality, and hence that reality cannot force itself upon the rational subject in the way that Hume and others thought. In principle we can remain as ignorant of certain aspects of our own minds as we can of the world outside our minds.¹⁰⁵

This suggests a more radical conclusion since in HANDS I am not merely ignorant of a fact about psychological reality that I am genuinely in a position to know via introspection. I don't just refrain from judging something I have introspective reasons for judging. In this case I *cannot* know by introspection alone whether I genuinely do see that p, or just seem to see that p. So if that is a psychological fact about me, then introspection alone is not in a position to tell me all the psychological facts, even in the most

¹⁰⁵ Of course there might be other reasons to doubt that beliefs about psychological reality are indefeasible. It is hard to put limits on which beliefs philosophical reflection can lead us rationally to revise: perhaps a clever article in *Mind* persuades me that the mental is not luminous, and I revise some of my psychological self-ascriptions on that grounds.

favourable circumstances – let alone, as traditional foundationalists supposed, all the facts tout court.¹⁰⁶

Ironically, then, the mental may turn out to be much more hidden from view than the rest of reality. And for any of us who have ever tried ‘introspecting’ that, I suggest, ought not to come as much of a surprise.¹⁰⁷

7. Conclusion

According to the literature, traditional foundationalists were driven by the desire to find foundations for our knowledge that are infallible. I have found such foundations but not in the place the traditional foundationalists locates them. Unlike the existing literature though, I have argued that traditional foundationalists were ultimately motivated by a different ideal – an ideal that perception genuinely does not secure. This is the ideal of total transparency – that is, transparency both in the world and in our epistemic access to it – or what, following Williamson we might call the ideal of luminosity (Williamson 2000: esp. Ch. 4). This idea has deep roots in the history of philosophy and it is an idea with continuing appeal. Given what I have been arguing in this chapter though, it is not an ideal that we can or should seek out at the expense of all others. If what I have been saying is right, our access to the world is often much more straightforward and unproblematic than our access to our access to the world. If that were not so,

¹⁰⁶ Clearly, a traditional foundationalist may question whether that is a psychological fact about me. That introduces difficult issues about the scope of the mental which I cannot here hope to resolve.

¹⁰⁷ This is compatible with the idea there are some facts about the mind which we cannot doubt; pain is often brought forward in this regard. In contrast, all of our beliefs about non-psychological beliefs may be thought to be revisable in principle. My claim is merely that some claims about our own minds are also rationally revisable. So even if there is an asymmetry, it is not the one the traditional foundationalists insisted upon.

epistemology would be as easy as deciding whether or not we have hands.

FOUNDATIONALISM AND THE IDEA OF THE EMPIRICAL

CHAPTER 5: THE PRIMACY OF PERCEPTION

1. Introduction

In this thesis I have claimed that a foundationalist is someone who thinks that (a) some of our beliefs must be non-inferentially justified; (b) perception is a distinctive source of non-inferential justification; and (c) perception is a basic source of such justification. In chapter 1 I defended (a) and in chapters 2, 3, and 4, I defended (b). What about (c)? What does it even mean to say that perception is a ‘basic’ source of epistemic justification and why should anyone think that this is true?

In this chapter I am going to give you an account of the sense in which it is true. I will argue that perception is basic in exactly the sense in which a foundationalist must think that it’s basic. In other words, I’m going to argue that the things you have to think to be a foundationalist are the things that are actually true.

A different question is why (c) matters. I think it matters for two reasons. The first is historical. I have claimed that foundationalism is a label with a partly historical basis. Doing justice to that position means trying to do justice to what particular thinkers actually thought, and it is undeniable that foundationalists have always thought that perception is a special source of justification with a privileged role in yielding us knowledge of the world.

(c) also matters because it is part of what make foundationalism a philosophically interesting position and distinguishes it from its rivals. Clearly (a) and (b) also do important work on this front; a coherentist can perhaps agree that some of our beliefs are non-inferentially justified as (a) claims, but coherentists do not think that perception is a source of non-inferential justification. They think that all justification derives from the fact one's beliefs belong to a coherent set of beliefs. This is not a view on which there is such a thing as distinctively perceptual justification as (b) claims. On my account coherentism therefore does not qualify as a version of foundationalism and that is at it should be.

While (a) and (b) rule out some philosophical alternatives to foundationalism, they do not exclude them all. Consider reliabilism: someone like Goldman could certainly accept both (a) and (b). But do we really want to say that reliabilism is a form of foundationalism? No doubt this is less obviously objectionable than suggesting that Davidson is a foundationalist, but it is still somewhat unsatisfying. Reliabilism fundamentally has little in common with those positions historically called foundationalist and it is hard to believe that we're not losing sight of something important by obscuring these differences. This is also part of the point of (c).

Ideally, then, (c) will enable us to hold onto what is *distinctive* about foundationalism in contrast to positions like reliabilism and preserve what is *true* about it in contrast to its rivals. Showing how we can achieve both of these aims is the purpose of this chapter.

So far I've claimed that (c) perception is a basic source of justification. What I mean by this is that the knowledge it grounds is special in relation to other kinds of knowledge. For this reason I'll sometimes refer to (c) as the view that perception is a 'basic source of knowledge' or an 'epistemically basic' source, but nothing important turns upon these differences. However, there are different ways of understanding the idea that perception is 'basic' in any of these senses and important things do turn on these differences.

Some people take that claim in an incredibly strong way. They think that what it means to say that perception is 'basic' is that perception is privileged in relation to all other sources. To this end, perception is held to be: (i) the only source of concepts; (ii) the only source of non-inferential justification; or (iii) the only generative epistemic source.

The first thing to notice about each of these claims is that they are not very plausible. There are non-perceptual sources of concepts, there are non-perceptual sources of non-inferential justification, and there are even non-perceptual epistemic sources that are generative; sources like 'reason' and introspection come to mind in this connection. The second thing to notice about these uniqueness claims is that they are unnecessarily strong. For a start, they are not claims to which the traditional foundationalists were committed. Further, they are not claims to which I am committed. I have only claimed that perception is a basic source, and this does not commit me to the view that it is 'the' basic source, or privileged in relation to all others. This is not to deny that the stronger claim has in each case had exponents, but in this chapter I will argue that these claims are not defensible.

Perception isn't basic in any of the senses described by (i), (ii), and (iii).

Although this is not a problem for the foundationalist it does raise the following important question: in what sense *is* perception basic? In this chapter I will suggest that the fundamental contrast is not between perception and all other epistemic sources, but between perception on the one hand, and memory and testimony on the other hand. When philosophers talk about where empirical knowledge of the world comes from these are the three sources they typically identify. So what I am saying, and what I think foundationalists are saying when they say that perception is 'basic', is that these three sources are not all on a par and that perception is privileged in relation to the other two. These are the sources I will be focusing on in this chapter, and when I talk about perception as a basic source what I mean, unless otherwise stated, is that it is basic in relation to memory and testimony.

Even if we confine our attention to these three sources it is still not plausible that perception is the only source of concepts, or the only source of non-inferential justification. I will argue that memory and testimony are not secondary in these respects. But what is plausible is that perception is the only one of these three sources that is truly generative. This is the reading of (c) that I am going to defend. I think that's a very intuitive view. It is one that seemed obvious to traditional foundationalists. However, it has seemed far from obvious to some recent commentators. In a series of influential articles Jennifer Lackey has argued that memory and testimony are also generative. One thing that I will be doing in this chapter is showing why Lackey is mistaken.

Showing that perception is special relative to memory and testimony is a way of showing that perception is a basic epistemic source and that is all I've taken foundationalists to be saying. It might be tempting to go on to claim that perception is the only generative epistemic source and that it is therefore not just a basic source, but the basic generative source. I don't want to go that far. Like the traditional foundationalists I don't want to rule out the possibility that reason and introspection are also generative sources. Nor do I want to say that reason and introspection are just forms of perception as some empiricists have claimed.

This is why my claim is only that perception is 'a' basic source. It's worth noticing, however, that while reason and introspection might be generative sources, the knowledge that they generate is in the one case 'a priori' knowledge, and in the other case 'self-knowledge'. So if you were interested in arguing for a uniqueness thesis you could maintain that perception is the only generative source of non-inferential empirical knowledge of non-psychological reality. I am perfectly happy with this claim, though it is a bit of a mouthful. My point is merely that we shouldn't abbreviate it by saying that perception is the only generative epistemic source, period. This would be to ignore a priori knowledge and self-knowledge and so is either false or, at best, misleading.

Since my official formulation of (c) is that perception is a basic source of non-inferential knowledge or justification, a real threat to my position would be one that identifies another source (i.e. a source other than perception, memory or testimony) that is also a

generative source of non-inferential empirical knowledge of non-psychological reality. In this thesis I have taken it that there are no such sources. This is not an unreasonable assumption. Barry Stroud makes a similar point in the following passage.

What happens in the case of the external world is that we want to understand how any propositions about an independent world are known to be true by anyone. But we must explain that knowledge in the light of other facts about human beings which we feel we cannot deny: in particular, that human beings get their knowledge of the world somehow from sense perception – ‘either from the senses or thorough the senses’, as Descartes put it. No divine messages from on high, no extra-sensory access to things around us, are to be assumed to be at work. So far, that is simply a very general ‘anthropological’ fact about the human condition. The question is how knowledge of the world is possible in the light of that fact (Stroud 2000:128-9).

Notice, like Stroud, I am not claiming that there couldn’t be other sources relevantly like perception in their ability to generate knowledge. When presented with putative examples like extra sensory perception, clairvoyance, or telepathy, the foundationalist has to think as hard as anyone else about what to say about them. But foundationalism also has descriptive aims; it aims to describe the structure and sources of human knowledge as we know it. If that is right then it is simply irrelevant adverting to sources that do not exist.

This all concerns the reading of (c) which I want to defend, but before we get that far we need to know why the other readings are no good.

2. Concept Empiricism

The first thesis that I am going to discuss claims that perception is privileged because of the special role that it plays in furnishing us with concepts. What sort of role is that? One reading has it that perception is special because it is a source of concepts. This is undoubtedly true, and without concepts we couldn't so much as frame the thoughts necessary for knowing. So this thesis bears an intelligible relation to the idea that perception is a basic source of knowledge. Yet it's hard to believe this is what the foundationalist has in mind in claiming that perception is basic. This reading leaves it open that there are lots of other sources of concepts and that sits oddly next to the traditional idea that perception enjoys a special privilege in relation to the rest. Moreover, it's hard to find people who disagree with the very weak claim that perception is a source of concepts; so it is hard to see it as a distinctively foundationalist commitment.

A stronger reading would claim that perception is the only source of concepts. Locke and the 18th century British Empiricists were famous for thinking that 'experience' was the source of all concepts, and they were certainly foundationalists.¹⁰⁸ While this reading has more bite, it suffers from the defect of being false. It is not true that for a subject to have acquired the concept of an F, she must have previously perceived F's. She may have acquired that concept from someone else and that does not in general require her to have perceived instances of the concept.¹⁰⁹ This is how a great many of our concepts are in fact

¹⁰⁸ Strictly speaking, that is not equivalent to my claim that 'perception' is the source of all concepts, since by 'experience' Locke meant to include both sensation and reflection. Nonetheless, reflection was held to be a form of 'inner perception' - in all respects just like 'external sense' but directed inwards at the mind's own ideas and operations (Locke 1975: Bk. 2 Ch. 1).

¹⁰⁹ I am assuming that similar objections could be raised against proposals that do not require the subject to have previously perceived an instance of the relevant concept, but under which she has still acquired the concept by means of perception. That may not hold true if one is sufficiently liberal about what it is to have acquired a concept 'by means of' perception. But then the proposal ceases to be interesting since it is one

acquired: I acquired the concept electron in science lessons, not by perceiving electrons.

We might try claiming that, even so, the person from whom I acquired the concept must have perceived instances of the concept, so that perception explains how the concept came into being, even if not the history of any given individual's acquisition of it. Or we might try claiming that the concept electron is a 'complex' concept composed of simpler concepts which are ultimately derived from perception, as the empiricists themselves claimed.¹¹⁰ But even if we revise the original thesis in this way and thereby render it more plausible (and that seems doubtful in the electron case) it would not in the end help save it as a gloss on what is distinctive about foundationalism.

The reason is that the thesis that perception, broadly conceived, is the only source of concepts is the defining thesis of concept empiricism. Yet concept empiricism and foundationalism are not the same thing; foundationalism is a thesis about knowledge or justification, not about the origin of concepts. Concept empiricism may or may not be true, but one can be a concept empiricist without being a foundationalist, and one can be a foundationalist without being a concept empiricist, regardless of how many foundationalists may, in fact, have been concept empiricists.¹¹¹

that most people will endorse.

¹¹⁰ Empiricists like Locke and Hume thought there could be complex ideas which were not directly derived from perception. Nonetheless, these ideas had to be composed of simple ideas which were derived from perception: in Hume's terms, simple ideas are 'copies of impressions' (Hume 1978: Bk. 1 Part 1 sec. 1). It's not obvious that electron is a complex concept in this sense. And simple ideas needn't be copies of impressions in any case: that is the point of Hume 'missing shade of blue' (Hume 1975: 16).

¹¹¹ 'Empiricism' is often used as a label for both doctrines so the confusion is perhaps not surprising. The important point is that foundationalism is distinct from empiricism - conceived of as a doctrine about concepts. In fact, it is perfectly compatible with so-called 'nativism' about concepts.

Moreover, concept empiricism only promises to show that perception is an enabling condition for knowledge. If perception is the source of all concepts, and there can be no knowledge without concepts, then there can be no knowledge without perception. This does not show that perception is a basic source of knowledge or justification. Indeed, it does not show that perception is the source of any of our knowledge. This is surely not what foundationalists are trying to capture when they claim that perception is basic.

I will come back to the distinction between sources and enabling conditions later; we will see that it is significant.¹¹² The important point to take from the discussion at this stage is that concept empiricism simply casts too wide a net to be useful in delineating the essential nature of foundationalism.

3. Basic Knowledge

The second thesis that I will be discussing claims that perception is basic because only perception is capable of giving us non-inferential justification. It is basic because it alone gives us so-called ‘basic knowledge’. Is that true and do you need to think it’s true to be a foundationalist? I will argue that it is false and that perception is not the only source of non-inferential justification. It’s a good job therefore that you don’t need to think this in order to be a foundationalist. You would think it was necessary given a commitment to certain prior assumptions, but those are assumptions that we have no good reason to accept. Or so I will now argue.

I just said it is false that perception is the only source of non-inferential justification. Why

¹¹² For a good discussion of this distinction see (Cassam 2007: 16-22)

is that? The obvious answer is that there are counter-examples to that claim. Consider remembering that you had toast for breakfast this morning. This can justify you in believing that you had toast for breakfast this morning. Yet remembering that you had toast for breakfast this morning is not a belief and nor is it something which needs to be believed to do its justifying work. So remembering that you had toast for breakfast this morning can be a source of non-inferential justification; it can justify you in believing that you had toast for breakfast, without that belief's necessarily drawing its justification from other justified beliefs. This is a counter-example to the claim that perception is the only source of such justification, since remembering that you had toast for breakfast this morning is not the same as perceiving that you had toast for breakfast this morning.

Here is another counter-example: suppose I believe that it is 3pm because you have told me that it is 3pm. Your having told me that it is 3pm can be what justifies me in believing that it's 3pm, despite the fact that your having told me that it is 3pm is not a belief. So testimony is another source of non-inferential justification. Again, we see that it is false that perception is the only source of non-inferential justification.

These remarks are not intended to foreclose all debate. Some will insist that memory and testimony are not really capable of giving us non-inferential grounds for believing things – not if we look closer.¹¹³ That may be so. My point is merely that it's prima facie plausible to think they do give us such grounds and there is no reason to saddle the foundationalist with prima facie implausible commitments, all else being equal. Things

¹¹³ Even some people, who agree with me about perceptual justification, think that is so: James van Cleve is a good example of someone who thinks the case of testimony is relevantly different (van Cleve 2006).

would be different if we were already convinced that this was the only way to spell out the claim that perception is basic so that a foundationalist really must deny that memory and testimony cannot also be sources of basic knowledge. But there is no reason to think this is so; we're about to see that there are at least two further ways in which to understand the claim that perception is basic.

I just said there is no reason for a foundationalist to reject the claim now being made about memory and testimony. In fact, this is a claim they have good reasons to accept. We saw earlier that it's reflection upon ordinary, commonsense examples which first motivates the claim that perception is a source of non-inferential justification. So if reflection also suggests that memory and testimony may furnish us with such grounds, why not once again take reflection at face value?¹¹⁴

Moreover, a foundationalist has much to gain by thinking that memory and testimony can be sources of non-inferential justification. The foundationalist wants to account for the justification of all our beliefs. If memory and testimony-based beliefs cannot be non-inferentially justified then they cannot figure amongst the foundations of our knowledge. Is it really plausible to suppose we can reconstruct all the rest of our knowledge if the foundations are as limited as this response suggests? Many have thought not. In contrast, if such beliefs can themselves figure amongst the foundations then the chances of affecting a successful reconstruction look correspondingly better.¹¹⁵ This advantage is not

¹¹⁴ The epistemic regress argument merely tells us that some of our beliefs must be non-inferentially justified, so there must be sources of such justification. That doesn't tell us anything about which beliefs are non-inferentially justified or what the sources of such justification are. For an answer to that question we must look to examples.

¹¹⁵ Even traditional foundationalists saw this (Lewis 1946: 336).

to be sniffed at. To fail adequately to explain what justifies all of our beliefs is to land up being overly sceptical about the extent of our knowledge, and such scepticism removes any reason we have to look favourably upon foundationalism in the first place.

I just said there is no reason for a foundationalist to deny that memory and testimony give us non-inferential justification for believing. This is not quite right; there is at least one reason, but it depends upon an assumption that we looked at in chapter 2 and previously found wanting. The Propositional Assumption claims the foundationalist must divide up beliefs into the inferentially and non-inferentially justified just on the basis of their propositional content or subject-matter so that if belief in a given proposition ever counts as non-inferentially justified, it always counts as non-inferentially justified. On this picture, it is propositions which are or aren't 'foundational'.

This is specifically problematic when memory and testimony are taken into account, since virtually any proposition can be held on their basis. So if propositions held on the basis of memory and testimony count as being non-inferentially justified, there will be no limit to the number of propositions which count as being justified in this way. This would be to abandon the foundationalist's idea that what is non-inferentially justified constitutes a privileged subset – the 'foundations' upon which everything else rests. On this picture there is simply too little left for the foundations to support. Pollock and Cruz raise precisely that worry in connection with memory. They write:

The only way the foundationalist can allow that the process of remembering can confer justification on a belief is by supposing that memory provides us with

epistemologically basic beliefs. It is important to realize that what is remembered can be a proposition of any sort at all. Sometimes there is a temptation to suppose that we can only remember facts about the past, but memory is just the process of retrieving stored information, and that information can be of any sort. For example, I can remember that $4+7=11$. This is a timeless truth. I can remember general truths e.g., that birds fly. And I can even remember facts about the future, such as that there will not be another solar eclipse in North America until 2032. By definition, epistemologically basic beliefs comprise a privileged subclass of the set of all possible beliefs, so it cannot be true that the proposition remembered is always epistemologically basic. (Pollock and Cruz 1999: 47).¹¹⁶

This would be a principled reason for a foundationalist to deny that memory and testimony can be sources of non-inferential justification, assuming they had reason to accept the Propositional Assumption. This is clearly what Pollock and Cruz are assuming. He writes:

Justification is a logical property of propositions. A proposition cannot have such a property at one time and fail to have it at another (Pollock and Cruz 1999: 61).

This is an assumption we have no reason to accept. Indeed, we saw earlier that even the most traditional forms of foundationalism, which hold that only beliefs about one's own psychological states are non-inferentially justified, do not endorse the Propositional Assumption. What is important in determining whether or not a belief is inferentially or

¹¹⁶ In a similar vein they later write: "There would be no epistemologically basic beliefs if this principle were true. The result would be a coherence theory rather than a foundations theory, because an essential claim of a foundations theory is that the epistemologically basic beliefs form a privileged subset of beliefs on the basis of which other beliefs are justified" (Pollock and Cruz 1999: 60).

non-inferentially justified are the grounds upon which that belief is held. There is no reason to think we can, or must, determine what grounds a belief is actually held upon, just by looking at its propositional content.

The point here is really very simple: one and the same proposition can be believed on a variety of different grounds. Testimony is just a particularly vivid illustration of this point, but even beliefs about one's own psychological states can in principle be held on more than one basis. Given that that's so we cannot hope to settle the question of whether or not a belief is non-inferentially justified just on the basis of the propositions involved.¹¹⁷ We cannot even look to canonical grounds, or grounds upon which such propositions tend to be held. We need to look to actual grounds.¹¹⁸ This is to abandon the Propositional Assumption and with it any remaining reason to think that memory and testimony-based beliefs are problematic in principle for the foundationalist.¹¹⁹

4. Generating Knowledge

The third reading that I am going to look at claims that perception is basic because it is the only 'generative' source of knowledge. A generative source of knowledge is one that increases our overall stock of knowledge: it increases the number of propositions which are actually known and not merely the number of individuals who know them, or the

¹¹⁷ Perhaps things would be different if such beliefs were justified in virtue of their content, but as we saw earlier they are not.

¹¹⁸ Traditional foundationalists may have thought we never *do* hold belief about our own minds on grounds other than observation: why would we, if such grounds are always available? This doesn't change the fact that we can hold them on other grounds or, therefore, that it isn't their actual grounds (and not their content) which is important in explaining why they are justified.

¹¹⁹ Of course, a foundationalist may have specific epistemological views about memory and testimony according to which they turn out not to be non-inferential sources of justification. Still, this is not a principled reason why a foundationalist, per se, ought to think that such sources do not yield such justification.

number of different ways in which those propositions are known. The traditional view about memory and testimony is that they are not generative sources of knowledge. Memory merely ‘preserves’ knowledge from one time to another, while testimony ‘transmits’ knowledge from one subject to another. Crucially, neither is thought to be capable of generating knowledge in the first place; both rely upon some other source’s previously having done so. Let’s call this the traditional view of memory and testimony.

Robert Audi describes this view as follows:

Just as we cannot know that p from memory unless we have come to know it in another way, say through perception, we cannot know that p on the basis of testimony unless the attester...has come to know it (at least in part) in another way...Memory and testimony...are not generative with respect to knowledge: characteristically, the former is preservative, the latter transmissive (Audi 1997: 410).¹²⁰

The same is not true of perception. In order for Ann to know that the squirrel is on the fence by means of perception, it is not necessary that Ann (or anyone else) already knows that the squirrel is on the fence; nor, *a fortiori*, is it necessary that Ann (or anyone else) knows that the squirrel is on the fence via a source other than perception. This is a way of making the obvious point that perception is a way in which things first come to be known; it is capable of generating new knowledge in a way in which memory and testimony are commonly held not to be.

¹²⁰ Indeed, Audi claims the fact they don’t generate knowledge is “the most important thing that memory and testimony have in common” (Audi 2006:44). Similarly, Michael Dummett writes “Memory is not a source, still less a ground, of knowledge: it is the maintenance of knowledge formerly acquired by whatever means” Dummett (1994: 262).

The third reading claims that perception is unique in this respect. It claims that perception is the only one of these three sources that is a truly generative source of knowledge and that that is why it is basic. This is the reading of (c) that I am going to defend. To traditional foundationalists it seemed obvious that perception was privileged in this respect and it is certainly a very natural view. However, it is one that has recently come under attack. In a series of influential articles Jennifer Lackey has argued that it is false that perception is the only generative source of knowledge. Moreover it is false not just because, as we conceded in the introduction to this chapter, there are *non-empirical* sources of knowledge that are also generative. Rather, it is false because the traditional view about memory and testimony is false. According to Lackey, they are also generative sources of knowledge; they can also create new knowledge, even if they frequently do not.¹²¹ Defending (c) requires showing this is false and that is the purpose of the rest of this section.

I'm going to start with the case of testimony since it is the more straightforward of the two cases. Here is Lackey's first example:

Suppose that a Catholic elementary school requires that all teachers include sections on evolutionary theory in their science classes and that the teachers conceal their own personal beliefs regarding this subject matter. Mrs Smith, a teacher at the school in question, goes to the library, researches the literature from

¹²¹ In fact, Lackey claims more. She thinks the examples can also be framed in terms of 'justified' or 'rational' belief' and hence that they promise to show that memory and testimony are not generative epistemic sources, more broadly speaking. This is also something that was also traditionally denied (Plantinga 1993: 61; Owens 2000: 156).

reliable sources, and on this basis develops a set of reliable lecture notes from which she will teach the material to her students. Despite this, however, Mrs Smith is herself a devout creationist and hence does not believe that evolutionary theory is true, but she nonetheless follows the requirement to teach the theory to her students. Now assuming that evolutionary theory is true, in this case it seems reasonable to assume that Mrs Smith's students can come to have knowledge via her testimony, despite the fact that she fails condition (ii) [the belief condition on knowledge] and hence does not have the knowledge in question herself. That is, it seems that she can give to her students what she does not herself have. For in spite of Mrs Smith's failure to believe and therewith to know the propositions she is reporting to her students about evolution, she is a reliable testifier for this information, and on the basis of her testimony it seems that the students in question can come to have knowledge of evolutionary theory (Lackey 1999:477).

Suppose we agree that the students in this example acquire knowledge they didn't previously have and that they acquire it from someone who doesn't themselves possess that knowledge viz. Mrs Smith.¹²² This is an attack on what I earlier called 'the traditional view of testimony' – the view that testimony is a transmission mechanism, and that a subject cannot transmit knowledge she does not possess. This is certainly how Audi

¹²² Even this is less than obvious. We can acquire knowledge by listening to what other people say ("we can come to know via their testimony" as Lackey likes to put it) without their testimony being the source of our knowledge. Other people can be mere 'mouthpieces': their words can give expression to the thoughts of another and can make the other person's knowledge available to us. Why think anything more is going on in Lackey's example? In cases like this it is irrelevant whether or not the speaker knows. Indeed it's not obvious the speaker must even understand what she says. Suppose Mrs Smith is ill and Elodie the French teacher steps in. Elodie doesn't understand English, though she does a passable impression when she reads out material to the class from the books Mrs Smith left behind on her desk. In this case, the students also acquire knowledge and they acquire that knowledge by listening to Elodie. Such cases aren't counter-examples to the traditional view of testimony: what testimony transmits is knowledge, but not every case in which we acquire knowledge by listening to x is a case in which x's testimony is itself the source of our knowledge.

presents the traditional view in the passage quoted above. He claims:

...we cannot know that p on the basis of testimony unless the attester...has come to know t (at least in part) in another way (Audi 1997: 410)

If Lackey's example succeeds, this is false.

Still, Lackey's example does not show that testimony can *generate* knowledge in the sense in which we're interested. It doesn't show that testimony can increase the total number of propositions that are known, as opposed to the number of individuals who know them. Let 'T' represent all the propositions about evolution, knowledge of which the students acquire. Lackey has certainly not described a case in which it was not known that T was true, it then comes to be known that T is true, and it comes to be known as a result of testimony. That has no plausibility in this case; presumably nobody wants to say that T comes to be known for the *first* time as a result of testimony. Darwin discovered that T was true by doing some empirical science and his results have subsequently been passed on in a chain with the students at one end and him at the other. He certainly knew that T was true and his grounds were not testimonial.

So at most this example promises to show that we can acquire knowledge from other people even if they do not themselves possess that knowledge, just as we can acquire shares in a company by relying on stockbrokers even if the brokers do not themselves own those shares. The stockbroker acts an intermediary and the same is possible in the epistemic case. This doesn't show that testimony can generate knowledge, any more than transfers on the stock market generate additional stock.

A more plausible example of a case in which testimony generates knowledge is this one. Consider Norman, BonJour's completely reliable clairvoyant: suppose Norman believes that Captain Molski will win the 3.30 as a result of his clairvoyant powers. Norman does not know that Captain Molski will win since, as we saw earlier, clairvoyance is not a source of knowledge for the clairvoyant. But if Norman tells me that Captain Molski will win and I happen to know he formed this belief with his fully reliable clairvoyant powers then I will acquire knowledge that Captain Molski will win. Here, a proposition that genuinely was not known comes to be known and it comes to be known as a result of testimony; that is, via Norman's having told me. So in this case testimony really is functioning as a generative source of knowledge.

This case is more plausible than the last, though it is still not terribly convincing. What is unconvincing in this case is the idea that Norman's testimony is really the source of my knowledge. In this case my belief that Captain Molski will win does not constitute knowledge just because Norman has told me that this is what is going to happen. It constitutes knowledge because (and only because) I know various facts about the provenance of Norman's belief. Given that knowledge, I can use what Norman says to work out how things are in the world - in much the same way that we use instruments like thermometers. But in that case my knowledge is inferential, not testimonial. My knowledge that Captain Molski will win derives from my knowledge that Norman is a clairvoyant and owes its epistemic credentials to that knowledge, just as a sailor's knowledge that the water is shallow when he sees the lighthouse flashing owes its

epistemic status to his knowledge that flashing lights mean shallow water.¹²³

So we are still no closer to finding a case in which testimony itself generates knowledge.

Lackey's entire case therefore rests upon examples like this last one:

Jane is currently in the grips of sceptical worries that are so strong that she can scarcely be said to know anything at all... That is, her belief that she could now be the victim of an evil demon is strong enough to defeat the justification she has for many of her ordinary beliefs and moreover, it is currently an undefeated defeater. Jim, a passer-by, approaches her, asks her where the café is, and she reports that it is around the corner, but does not report her sceptical worries to Jim. Now Jim has never considered any sceptical possibilities at all, and hence he does not have any doxastic defeaters for his ordinary beliefs. Furthermore, he does have positive reasons for accepting Jane's report, e.g., he has perceived a general conformity between facts and the reports of many speakers in these types of contexts and, he had inductively inferred that speakers are generally reliable when they are giving directions, and Jane does not indicate any behaviour which indicates a lack of sincerity or competence with respect to her report. So Jim forms the true belief that there is a café round the corner on the basis of Jane's testimony. Given that Jane has an undefeated defeater, which Jim does not have, he has knowledge, which she lacks. Yet at the say time it seems possible for Jim to come to know that the café is around the corner via Jane's testimony even though her sceptical

¹²³ Perhaps this is a case of testimonial knowledge in some suitably extended sense of 'testimony'. My knowledge does partly depend on Norman's having told me, just as the sailor's knowledge partly depends on what he sees. Nonetheless such cases do not represent the normal case and we should be sceptical about thinking they show anything substantive about the normal case. They certainly do not show that testimony can ordinarily generate knowledge all by itself, where such background beliefs are absent.

doubts currently undermine her knowing this...and thus it seems possible for a hearer to acquire knowledge on the basis of a speaker's testimony even when the speaker does not personally have the knowledge in question (Lackey 1999:484).

The first thing to notice about this case is that it is like previous case in the following respect: Jim *has* justified beliefs on the basis of which he can infer that what the speaker says is likely to be true, whether or not he does infer that. If that is why Jim knows, then this is just another version of the previous case and the same remarks apply to it. Jim may know, but his knowledge doesn't really derive from testimony; it is inferential knowledge. Call this scenario 1.

However, we needn't assume that Jim does rely upon those background beliefs. In this respect this case differs from the previous case: I only know that Captain Molski will win because I know that Norman is a reliable clairvoyant. Jim on the other hand may just believe the shop is around the corner because Jane has told him and her testimony may accordingly be the source of his knowledge. Call this scenario 2.

Lackey does not clearly distinguish between scenarios 1 and 2. Only scenario 2 is relevant though, since only scenario 2 promises to show that *testimony* is capable of generating new knowledge. So this is the important case to consider. Unfortunately, the case is completely implausible so construed. It is implausible either because (I) Jane does know and therefore passes on her own knowledge or because (II) Jane does not know, but does not pass on any knowledge either. In neither case does Lackey have what she needs. (I) can be true either because belief is not a necessary condition for knowledge, thereby

rendering Jane's lack of belief irrelevant; or because Jane does believe the shop is round the corner, whatever she says to the contrary.¹²⁴ Both strike me as considerably more plausible than Lackey's description. However, if you are really convinced that Jane doesn't know that still leaves option (II): Jane does not know, but does not pass on knowledge either. This is also more intuitive than Lackey's description. After all, if Jane doesn't know where the shop is, why does she tell Jim it is round the corner? She is telling him something that, by her own lights, isn't true. That sounds suspiciously like deception. Why, then, suppose that Jim acquires knowledge by listening to her?

Thus, Jane either does possess the knowledge she passes on; or she does not possess that knowledge, but does not pass it on either. In neither event do we have a case in which testimony itself creates new knowledge, or therefore a counter-example to the claim that perception is the only generative source.

To summarise: I have argued that when you press hard on the point that it is the generation of knowledge (and not whether or not a speaker must have knowledge in order to pass it on) which is relevant, the alleged counterexamples are either irrelevant or unconvincing. Hence, testimony is not a generative source of knowledge and so is not a counter-example to the claim that perception is the only such source.

That still leaves the case of memory. Is memory, at least, capable of generating new

¹²⁴ One might think that the fact she tells Jim the shop is round the corner without intending to deceive him is evidence that she thinks precisely that. Moreover, it seems as if she can believe that the shop is round the corner without necessarily believing things like the external world exists. In that sense her 'scepticism' may be compatible with ordinary knowledge.

knowledge? This case is more tricky. Let's start by considering some examples from Lackey:

While an undergraduate in college, Nora was a very careful and epistemically reliable recipient of testimony, with one notable exception: she was overly susceptible to peer pressure from two of her friends who belonged to a religious cult. After repeatedly hear them rant and rave about the corrupt minds of non-believers she eventually became convinced that the testimony of atheists is nearly completely unreliable. During this time, Nora had several conversations with Calvin, a fellow student in one of her classes who, as a matter of fact, was an extremely reliable source of information and whom she had every reason to believe was both competent and sincere with respect to his reports. Yet Nora also knew that Calvin was an atheist, and so she believed him to be a highly unreliable epistemic source. One day after class, they were discussing World War II and Calvin told Nora, much to her surprise that Hitler was raised a Christian. Being momentarily caught off guard, Nora found herself believing this proposition on the basis of Calvin's otherwise epistemically flawless testimony. Now several years after graduating from college, Nora is no longer in touch with her friends who were members of the religious cult and she has ceased believing that the majority of the testimony offered by atheists is highly unreliable - such a belief has simply faded from her memory. At the same time, however, she still believes on the basis of memory dating back solely to Calvin's testimony that Hitler was raised a Christian (Lackey 2005: 644-5).

What does this case show? Lackey takes it to show that the traditional view of memory is

false. It is false because this is a case in which a subject, Nora, first comes to know a proposition on the basis of memory: the proposition that Hitler was raised a Christian. Nora didn't know that Hitler was raised a Christian when she first acquired that belief due to the presence of relevant counter-beliefs (viz. her beliefs about the unreliability of atheists) and she cannot know that Hitler was raised a Christian as long as she has those beliefs. But since she no longer believes that atheists are unreliable, that belief no longer prevents her from having knowledge. So she now knows that Hitler was raised a Christian. Moreover, since she only continues to believe this on the basis of a memory dating back to Calvin's testimony, memory must be the source of her knowledge. So Nora now knows something on the basis of memory that she did not know when she first acquired her belief. This is incompatible with the traditional view that memory is never capable of generating new knowledge.

Here is Lackey's second example:

Two days ago, Arthur was visiting his Aunt Lola and, while they were eating lunch she mentioned to him, without disclosing the source of her information that the mayor of their city had been caught accepting bribes in exchange for political favours. Arthur unhesitatingly formed the corresponding belief. At the same time, however, there was a vast conspiracy on the part of the mayor's allies to protect his political reputation, and so they exploited their high-powered connections in the media to cover up this indiscretion. To this end they convinced all of the major newspapers and television stations to report that the mayor's political opponents had orchestrated a plan to win the upcoming election by falsely presenting him as

having been the recipient of bribes. However, because both Arthur and Aunt Lola rarely pay attention to the news, they were entirely unaware of all of the stories surrounding the mayor. Thus unbeknownst to both Arthur and Aunt Lola, every major newspaper and television network was reporting that the mayor had not accepted bribes and was instead the victim of a devious scheme at the very time that Arthur was forming the belief that the mayor had been the recipient of bribes on the basis of Aunt Lola's testimony. Now, as it turns out the mayor had in fact accepted bribes in exchange for political favours, all of the reports to the contrary were false, Aunt Lola was not only a highly reliable source of information in general, but had also heard this news directly from the mayor's epistemically reliable secretary, and Arthur's true belief about the mayor was reliably formed. Since then and, once again, unbeknownst to Arthur and Aunt Lola, the scheme to cover up the mayor's indiscretion has been exposed, and all of the major newspapers and television stations are now reporting that the mayor did accept political bribes. At the present time, then, there are no longer any vast amounts of available evidence indicating that the mayor had been framed, Throughout all of this, Arthur has remained blissfully ignorant of all of the relevant reports, and he currently continues to believe that the mayor was the recipient of bribes solely on the basis of remembering Aunt Lola's original testimony (Lackey 2005: 640).

Again, Lackey takes this example to show that the traditional view of memory is false. It is false because when Arthur first acquired his belief about the mayor he failed to acquire knowledge; Arthur did not initially know that the mayor had accepted bribes because of the presence of available counter-evidence, specifically, the newspaper and television

reports to the contrary. Since that counter-evidence is no longer available, it no longer prevents Arthur from knowing that the mayor accepted bribes. So Arthur now knows this, despite the fact that he did not know this when he originally acquired a belief to that effect. Moreover, he now knows this on the basis of memory, since he continues to believe the mayor accepted bribes solely on the basis of a memory dating back to his Aunt's testimony. If this is right then the traditional view, that memory is never a generative source of knowledge, must be wrong.

What should we make of these examples? I think they are pretty unconvincing. What is unconvincing, specifically, is the claim that memory is the source of the subject's knowledge in these cases.¹²⁵ Imagine asking how Nora knows that Hitler was raised a Christian. The obvious answer, surely, is: 'because Calvin told her'. This, rather than memory, is the source of her knowledge. Of course, it is true that Nora did not know that Hitler was raised a Christian when Calvin first told her, but this does not mean his testimony cannot be the source of her later knowledge.

Exactly the same is true in Lackey's second example. How does Arthur know that the mayor accepted bribes? Answer: because his Aunt told him. The fact he did not know this when he first acquired that information is irrelevant since it does not prevent that information being what explains how he later knows.

¹²⁵ This is not the only ground for doubt. It's not obvious that Nora does come to know that Hitler was raised a Christian; or that Arthur didn't know the mayor had accepted bribes all along. And if the subjects did previously know (or do not now know) then these cases won't be counter-examples to the traditional view of memory. I suspect we could develop Lackey's stories either way.

What I am claiming, in effect, is that the following makes perfectly good sense: a subject S now knows that p because in the past she was informed that p, even though (at that time) she did not know that p. Clearly Lackey thinks this does not make any sense; that is why she thinks memory must be the source of the subject's knowledge. Later on, writing in her own defence, she claims:

A belief that was not known (or, in Case 2, justified or rational) when originally acquired became known (and, Case 2, justified/rational) at a later time without input or assistance from any other epistemic source besides memory. Thus, even though memory did not generate the belief in question, it generated the epistemic status of the belief in question. And this is sufficient not only to falsify the [Preservation View of Memory], but also to conclude that memory has the capacity to function as a generative epistemic source (Lackey 2005: 649).

This is just a mistake. If I am right, testimony can be the source of the subject's knowledge in these cases. So there is no pressure to look for an alternative source.

Certainly, this sounds better to my ear than claiming that memory is the source of Nora and Arthur's knowledge. In these cases memory contributes absolutely nothing in its own right. It merely preserves a belief long enough that other factors cease to be relevant. This is surely not enough to warrant claiming that it 'generates' the knowledge in question. The natural thing to say is that the source of the subject's knowledge in these cases is the source of her earlier information and in these cases that source is testimony not memory.¹²⁶

¹²⁶ If a subject can know that p because she was in the past informed that p (despite not then knowing that p), then her being informed that p does not entail that she knows that p. This doesn't mean the fact she was informed can't ever explain how S knows that p, or therefore that it cannot ever be the source of her

Lackey's own examples thus fail to show that memory is a generative source of knowledge since in the cases that she describes *it* is not the source of the subject's knowledge at all. But I think there are better examples. Suppose I get in late from work and absent-mindedly put my keys down on the sideboard before going up to bed. The next morning, wanting to leave the house, I start hunting around for my keys. Where are they I wonder? I try to think back to where I might have put them - mentally retracing my steps in memory - and then I remember: I put them on the sideboard.

In this case it certainly sounds less odd to say that memory is the source of my knowledge. I come to know where my keys are *by remembering* where I left them. Such cases are not uncommon or wholly contrived. Thus, suppose there has been another murder in Peckham and the Scene of Crime Officer wants to know what I remember about that night: do I recall noticing anything odd or out of the ordinary? I certainly wasn't aware of anything suspicious at the time, but I know it's important so I keep at it - carefully going over the scene again in my mind. And yes, now I think back: I do remember having seen someone a bit suspicious. There was a man parked on the corner in a vintage blue Mercedes and I am forced to admit, that on reflection, that certainly is somewhat unusual in Peckham.

knowledge. Throwing a brick at a glass window does not entail it will smash; that doesn't mean it never explains why it does. So if that is Lackey's worry then it is unfounded. Naturally, when S has been informed that p and still fails to know that p, we expect there to be some explanation of why that is so. But there is an explanation of why the subjects don't acquire knowledge in Lackey's examples: Lola believes that atheists are unreliable, and Arthur is in the midst of a political conspiracy. These factors prevent Arthur and Lola acquiring the knowledge they might otherwise have had and this is precisely why, when those obstacles later disappear, they do acquire the knowledge that previously was closed to them. In both cases, their beliefs must be preserved in memory, but that does not make memory the source of their knowledge.

I think cases like this pose the biggest threat to the view I want to defend – that is, the view that memory is never a generative source of knowledge. In these cases there is at least some temptation to think memory is contributing something in its own right. These cases are not like Lackey’s examples where memory merely preserves a belief for long enough that other factors cease to be relevant; in these cases attending to the scene in memory seems to play a positive role in making knowledge possible. Accordingly, there is a recognizable temptation to say memory is the source of one’s knowledge.

Still, these cases are not decisive. It is often just as natural to say that I have simply *forgotten* where I left my keys, or forgotten whether I noticed anyone suspicious in Peckham that night. These are things I did previously know for however brief a time, but which (being unimportant or distracted) I soon forgot. So I didn’t first acquire the knowledge in question via memory; I acquired it via perception. Memory is merely reminding me of facts I have forgotten; it is not generating new knowledge.

This may even be the more natural thing to say on reflection; but perhaps we have mixed intuitions. What is clear is that I must have *registered* the scene I later recall and that this registration is not itself a product of memory. More often than not it is the product of perception. In the crime scene example, I remember having *seen* someone in a vintage blue Mercedes. Similarly, I presumably felt myself putting the keys down at the very instant I placed them on the sideboard, whether or not that remained at the forefront of my mind for long.¹²⁷ This makes it natural to suppose that the facts I later recall are at

¹²⁷ We presumably do not want to deny subjects can ever properly be said to perceive that which they fail to notice or actively attend to. This would rule out too much of what we ordinarily count as perception. Moreover there is normally an explanation of why subjects fail to perceive things which are in the vicinity

least ones that I could have known at the time, even if they are not ones I actually knew due to factors like lack of attention. In that case, it is more natural to say the source of my knowledge is perception, rather than memory.

So while like the keys case do pose some sort of problem for my view, they are clearly not decisive. It is simply too unclear what is really going on in them for that to be plausible. The most that is plausible is that memory can sometimes make a contribution towards the generation of knowledge in a way in which it is hard to believe testimony does. Even if this is right though (and I am not unambiguously endorsing that claim¹²⁸), memory is still not generating knowledge in its own right in the way that perception can. Intuitively, perception can be both the source of your information and the source of your knowledge and the same is not true of memory. Insofar as it has the capacity to generate knowledge at all, that capacity is essentially dependent upon information being registered via a source other than itself.

One way to put this would be to say that perception, unlike memory, is a non-dependent generative source. Memory and testimony are different from perception because their capacity to yield knowledge depends essentially upon the fact there exist other ways of

e.g. they wasn't looking in the right direction. This needn't be so in the cases I have described. Michael Martin argues the fact that a scene is later available to the subject in memory is precisely evidence the subject did once perceive it. On his view, memory involves a re-presentation of some past perception (Martin 1992). But even if this sort of registration does not always amount to conscious perception as Martin claims, it is not exactly unconscious either.

¹²⁸ As just seen it's often just as natural to say the source of the subject's knowledge is the source of her earlier information. Moreover, it is not clear in what sense 'attention' should be put down to memory. All I am claiming is that this is the most that can be said for the thesis that memory is a generative source. Part of the problem here relates to unclarity in the notion of a 'source' and what it is for something to be the source of your knowledge.

acquiring information about the world.¹²⁹

In conclusion: none of the potential counter-examples to (c) is decisive. Most are clearly no good and the one example that looks threatening, namely the keys examples, is not clearly problematic. This is not a good enough ground on which to reject the view that perception is the only generative source. So I have now defended the claim that perception is privileged in relation to memory and testimony in at least this respect; it is the only non-dependent generative source of empirical knowledge of non-psychological reality. This is a way of defending the claim that perception is a basic source of justification just as (c) claims. Spelling out the precise sense in which that is so turned out to be harder than one might have expected, but it would be even harder to believe there is no such sense. The idea that perception is privileged is an overwhelmingly natural one and it is one that we can and should hold onto.

5. Perception and Explanation

I have now argued for (c). An obvious next question is this: why is (c) true? Is there any positive account we can give of why perception is such a basic source of knowledge? Indeed, is there any positive account we can give of why perception is a source of knowledge at all?

The reliabilist has an answer to this question. The answer goes like this: knowledge is true belief produced by a reliable process and perception is a source of knowledge insofar

¹²⁹ Perception may often rely on information from other sources, and if it didn't we would certainly know as lot less than we do via perception. But it doesn't look like it absolutely must depend upon them in the same way in which memory and testimony do.

as it satisfies this condition. According to the reliabilist we can explain why perception is a source of knowledge in *more basic* terms. What makes this an explanation in more basic terms is that the notion of ‘a true belief produced by a reliable process’ is one upon which we have an independent grip. Reliabilism claims to have a generic fix on knowledge, a general framework into which it can fit perception and thereby see it as an instance of some broader phenomenon depending on its fit or lack of fit with the relevant criteria. That is the whole point of reliabilism; it aims to give us *general* criteria for knowing.¹³⁰ Whether perception lives up to those criteria is then an empirical question.

Reliabilists are not the only people who promise to explain perception’s status as a source of knowledge in more basic terms. Any theory that offers a reductive analysis of knowledge – that is, an analysis in terms of non-circular necessary and sufficient conditions - promises to do just that. The theory specifies general conditions for knowledge; whether perception lives up to those conditions is then an empirical matter. Thus a defender of the ‘tripartite’ analysis of knowledge will claim that knowledge is justified true belief, just as the reliabilist insists that knowledge is true belief produced by a reliable process. Both will claim that perception is a source of knowledge only insofar as it meets these conditions.

I think that this idea is foreign to foundationalism. For foundationalists there is an important sense in which we *cannot* explain why perception is a source of knowledge in more basic terms. When it comes to knowing, perceiving is as basic as it gets.

¹³⁰ ‘Reliabilists’ who concede that we cannot grasp the notion of reliability independently of knowledge are not ‘reliabilists’ in my sense. On my account, reliabilism is an essentially reductive project: if not, I have no problem with it.

I am going to call this the idea that perception is explanatorily basic. This offers a very different gloss on the idea that perception is a ‘basic’ source of knowledge. Unlike (c), the claim that perception is explanatorily basic does not attempt to privilege perception in relation to other sources of knowledge, like memory or testimony. It claims that perception enjoys a privilege in its own right - its status as a source of knowledge cannot be further explained – and this can be true whether the same is true of all sources, or of none.

I think that this claim is both true and defensible. The implication is that reliabilism and other forms of reductionism therefore fail to explain perception’s status as a source of knowledge. They fail either because (i) they fail to explain how it yields knowledge, or because (ii) they fail to explain how it yields knowledge in more basic terms. Consider again what a reliabilist has to say: she claims that knowledge is true belief produced by a reliable process. Yet not all such beliefs are instances of knowledge: Norman’s clairvoyant beliefs do not amount to knowledge even if those beliefs are true and reliably formed. It might be a necessary condition for a belief to count as knowledge that it be reliably formed, but it is not a sufficient condition. So we cannot explain why perception is a source of knowledge just by saying that it satisfies this condition. This is the first horn of our dilemma.

So the reliabilist must supplement her original analysis or modify the way we understand it. However, when she tries and do that – that is, tweak the notion of reliability so that it

really does cover all and only those cases that we genuinely want to count as cases of knowledge - it becomes doubtful that the explanation given genuinely is one in more basic terms. In this case the relevant notion of reliability is no longer one we can plausibly be said to grasp independently of thinking about what it is meant to explain.¹³¹ This is the second horn of our dilemma.

I think exactly parallel problems afflict all existing attempts to analyse the concept of knowledge. I think this concept can't be analysed in more basic terms. This is not a claim I can hope to defend in the last ten pages of this thesis, though it has been defended elsewhere.¹³² It is, however, part of what I am rejecting when I claim that perception is an 'explanatorily basic' source of knowledge. I am claiming that we cannot understand perceptual knowledge as an instance of some independently intelligible notion of knowledge; it is a type of knowledge that cannot be explained in more basic terms.

I don't want to defend this claim because I think that it is *essential* to foundationalism. As a matter of historical fact I think most foundationalists probably *did* think that perception was explanatorily basic in this sense. I also think that this is part of what, at the deepest level, distinguishes foundationalism from other important theories on the market. Reliabilists certainly do not think that perception is explanatorily basic in this sense. Nonetheless, it would be hard to defend the claim that this idea is really essential to foundationalism in the way in which I have claimed that (a), (b), and (c) are essential.

¹³¹ This is often known as the 'generality problem'. I am assuming that reliabilists have failed to come up with a convincing response to this problem. Clearly that's something some of them may deny.

¹³² See (Williamson 2000: esp. Ch. 1); and for problems, (Cassam, forthcoming).

It would be hard for two reasons. The first is that it is an issue that very few of them ever explicitly addressed. People like Locke certainly never *said* that the concept of knowledge could not be analysed in this sense. I suspect most of them simply never thought about it. They certainly didn't think this was what they ought to be doing. The idea that the concept of knowledge can usefully be analysed or that this is the proper task of epistemology is a modern preoccupation. More embarrassingly, it's also an idea that some canonical foundationalists in the modern period apparently did go in for. Ayer claimed that knowing was 'having the right to be sure' (Ayer 2000: 22-44) and Chisholm tried to define knowledge in terms of the directly evident (Chisholm 1966). I don't have the space to explain why these explanations fail or why they are not genuinely reductive. So I will not try and defend the claim that all foundationalists must – deep down – all really think perception is explanatorily basic.

Second, many will find such a defence unwelcome. Most people do not have the intuition that foundationalism is fundamentally opposed to positions such as reliabilism in the first place and these people will find any commitment on this score unduly restrictive. They will claim it prevents us seeing the things these positions have in common and that is more important than the respects in which they differ.

So I am not claiming that the idea that perception is explanatorily basic is one that all foundationalists either have accepted or must accept. I have deliberately refrained from describing it as an essential commitment of foundationalism alongside the other three for precisely the reasons just given. Whether any version of 'foundationalism' worth the

name must hold that perception is explanatorily basic seems to me to be a much less interesting question than the issue of whether or not perception is explanatorily basic. This is the issue I want to focus on. I think perception is explanatorily basic and this is what I want to persuade you of in the last part of this chapter.

So far my claim has been purely negative: I have claimed that we cannot explain perception's capacity to yield knowledge in more basic terms. The project of explaining perception status in more basic terms requires a reductive analysis of knowledge itself, and that is what I have claimed is not possible. This raises an obvious question: if we do not explain why perception is a source of knowledge in more basic terms, with what right do we treat it as a source of knowledge at all? What positive explanation can we give if not the one latterly rejected?

I think foundationalists had a distinctive answer to that question. Consider the following remark from Paul Snowdon:

Another alternative, though, is that the link between perception and knowledge explains the content of the concept of knowledge. Thus the idea might be that our fundamental understanding of what knowledge is as what is yielded by perception in certain circumstances (Snowdon 1998: 301).¹³³

Suppose Snowdon is right and knowledge is what is yielded by perception. In that case it would be obvious why we could give no further explanation of perception's status as a source of knowledge. On this account it is not that we first have an independent grip on knowledge and then get to the idea that perception is a source of it. Rather, we start off

¹³³ Notice, Snowdon does not himself endorse this suggestion.

with the idea that perception is a source of knowledge. Knowledge is then just what perceiving gets you. So there can be no explanation in more basic terms of why perception is a source of knowledge; on the contrary knowing just is the kind of thing that perceiving can give us.

What, if anything, is wrong with that suggestion? The most obvious problem concerns seemingly non-perceptual sources of knowledge. It would be very implausible to have to claim that all knowledge is perceptual knowledge or that perception is the only source of knowledge. Memory and testimony are not ways of perceiving the world yet they are sources of knowledge. In addition, mathematical knowledge surely cannot be understood as that which is yielded by perception. Nor, on the face of it, can the knowledge I have that I intend to go to Lebanon this summer. So it looks like it is false that perception is the only source of knowledge.

One response to this objection would be to simply bite the bullet and insist that these are all forms of perceptual knowledge. Some people have argued that memory is a form of perception, namely, perception of the past. Traditional foundationalists certainly had no problem with the idea that self-knowledge was a form of perceptual knowledge. And the history of debate about understanding-based knowledge or the so-called *a priori* bears ample witness to the power of the perceptual model and the grip it has on our thinking about knowledge. One 'sees' by the natural light of reason, and even for Descartes it clear and distinct 'perception' that enables one to verify that God exists and is no deceiver. Indeed, even those hostile to a priori knowledge apparently succumb to the power of

perception - the very idea that a priori knowledge is 'mysterious' seems to imply little more than a prejudice in favour of perception.

So this response has certainly had famous defenders. However, it is implausible in its conception of what perception can explain and unnecessarily hard-line. Snowdon's original claim after all was merely that perception furnishes us our 'most fundamental grasp' on what it is to know. This doesn't entail that perception furnishes our only grasp upon knowledge.

At this point, we can decide how liberal we want to be. We might claim that there are at least two different sources of knowledge – perception and reason - both of which non-derivatively put subjects in a position to know about the world and both of whose capacity to do so cannot be explained in more basic terms. This is what many of the traditional foundationalists thought. On this reading, perception still isn't 'the' basic source of knowledge if that implies that it is the only source with these privileges, but it is one of just two.

However, we needn't be as restrictive as even this line of thought suggests. We cannot be if, unlike the traditional foundationalists, we are not convinced that self-knowledge is a form of perceptual knowledge or that memory and testimony can be explained in wholly perceptual terms.

What is there to prevent us acknowledging that there are lots of different ways of

acquiring knowledge? One problem is this: why do we count them all as sources of knowledge, if not because we grasp what they have in common? Without some common factor it can seem that knowledge itself would lack unity. This problem is based on a false assumption however since we *do* grasp what they all have in common: they are all sources of knowledge. That is the feature they all share and that can be the ground upon which we count them as such, just as we count cricket and chess and solitaire together on the basis that they are all games. All I am denying is that we classify them all as cases of knowledge by grasping something else they all have in common. To assume otherwise would be question-begging.

The simple fact is that different sources of knowledge - perception, memory, and testimony, self-knowledge – may simply have too little in common out of which to frame some more basic understanding of what it is for a source to yield knowledge that something is the case. This is not to deny they have anything significant in common, but the things they have in common do not give us the resources that we need.

Common factors tend to fall into two classes. On the one hand, they may be independently intelligible, but so general that when we abstract away from other features of the examples, we just aren't left with something that intelligibly adds up to knowledge. As previously mentioned it's very natural to think of reliability in this way. It's plausible to think that it is at least a necessary condition for a source to be a source of knowledge that it reliably produces true beliefs and that is obviously an independently intelligible ideal. However, it is not a sufficient condition; not all reliable sources are sources of

knowledge. So we have to try and add something to reliability to get us back up to knowledge.

In trying to flesh out those further conditions, however, we then face the opposite problem: there simply don't seem to be any further features that different sources of knowledge all have in common that mark them out as sources of knowledge. At least, none that are not trivial. Being sources of knowledge is something they all have in common, but that could hardly be reckoned a common factor in the present context.

Once we move away from trivialities such as these, though, problems set in. For instance, I have claimed that conscious awareness is crucial to perceptual knowledge. Yet there is no clear sense in which it applies to knowledge grounded in intellection. We can talk about 'seeing' by the natural light of understanding, but such talk is at best metaphorical and the metaphors themselves are of no real help. More metaphorically speaking these sources do have something in common. It is tempting to say that both make it 'manifest' to us that things are a certain way, both possess a certain kind of 'evidence' or perspicuity. But it is not as if we can use these notions to frame the required understanding of knowledge since these notions are not independently intelligible. The only way to get hold of the relevant notion of 'evidence' is by reference to these very examples.

This brings us back to the point with which we begun: the idea of knowledge or a source of knowledge may simply be too heterogeneous in the relevant respect to be grasped

other than via its instances – that is, by grasping particular ways in which we do in fact come to know things about the world around us. Knowing, we can then say, is having this sort of access to the facts, or this sort of connection (pointing now to perception, now to reason...). That may seem surprising until we stop and ask why it should be otherwise. Why should different cases of knowledge have anything more in common with each other than their simply being cases of knowledge? Not all concepts can be explained in more basic in terms; some must be primitive and there is no reason to suppose knowledge is a particularly bad candidate in this respect. While that may be frustrating from a philosophical perspective, it ought not to be surprising.

If this is right – if the concept of knowledge is essentially grasped through examples of which perception and reason are paradigms – that would explain why we cannot explain why or how these are sources of knowledge in more basic terms. Moreover, the fact that perception is, as a matter of fact, the most pervasive of these paradigms in the particular circumstances in which we find ourselves, and the fact that it plays a role in so many of the others, like memory and testimony, would do something to explain the continuing pull of the perceptual model and our tendency to try and explain other forms of knowledge along these lines (even when it comes to equally basic sources such as reason).

This is not at all implausible. We certainly do not have as firm a grasp on how it is that reason puts us in a position to know about the world – none, at least, that it is not shaken by the ‘slightest philosophy’. Nor do we have so clear a picture of how we know about

our own minds. This seems importantly different from the case of perception. Philosophy can try and persuade us that we never perceive the world around us, though even here it frequently struggles to convince. But even philosophy is rarely so bold as to tell us that seeing an object plainly is not, in principle, a way of coming to find out about what that object is like. In this respect, perceiving that something is the case is a peculiarly good way of coming to know that it is the case: a conclusive response to the question ‘how do you know?’ in a way in which testimony say, intuitively is not.¹³⁴ So while perception may not be our only model of what it is to come to know, it is the one we actually understand best.

6. Conclusion

I began this thesis by distinguishing between inferential and non-inferential justification. Most of the thesis has defended the claim that there is such a thing as non-inferential justification and that perceptual justification is an example. These claims are still widely rejected. In this chapter, though, I have argued that there may be other sources of non-inferential justification: memory and testimony may also furnish us with such grounds. So if all it takes to be a foundationalist is to think that there is such a thing as non-inferential justification, then the foundationalist has already won.

In this chapter I have claimed that is not enough. I have argued that a foundationalist is also essentially someone who thinks that among the different sources of knowledge and

¹³⁴ I am not committed to thinking that it’s necessarily true that perceiving is a way of knowing or even that it is rationally unrevisable. It may be knowable apriori but so is the truth that all bachelors are unmarried males. The latter is not, for that reason, reason rationally unrevisable. For a defence of this distinction see (Giaquinto 1996). Of course, it is an empirical question whether we can ever be said to see that anything is the case.

justification that we have access to, perception is special. What does that amount to? My first answer was that perception is a non-dependent generative source of knowledge and that that is what distinguishes it from other sources of empirical knowledge. By non-dependent I mean a source whose functioning as a source of knowledge does not essentially depend upon the existence and functioning of another. The same is not true of memory and testimony; they may be sources of knowledge, but to function as such they rely essentially upon the existence of other sources. I stand by what I said there: perception is a non-dependent source of knowledge in a way in which memory and testimony are not and this is something any foundationalist has got to think.

The problem we encountered is that this still doesn't capture all of what is most distinctive or most interesting about foundationalism. A reliabilist can think that perception is a non-dependent source of knowledge, but there is some interesting sense in which reliabilism isn't a form of foundationalism.

That lead me to my next claim: perception is not only a non-dependent source of knowledge, it is also an *explanatorily basic* source of knowledge. Unlike the claim about dependence, I have refrained from describing this claim as an essential commitment of foundationalism. It is not an issue most of them explicitly addressed and many do not have the strong intuition that reliabilism is not a form of foundationalism to begin with. Nonetheless, I hope to have made it plausible that this idea forms an important part of foundationalist thought and one that is genuinely absent in its rivals.

I think this is what foundationalists traditionally thought. They thought that perception is a way of coming to know about the world and that there is an important sense in which there is nothing more to be said about why that is so. When it comes to understanding what knowledge is, perception is as basic as it gets. This is not to say that perception is the only source of knowledge, even the only basic source; maybe reason is another. So perception may not be the only basic source of knowledge. Nonetheless, when it comes to *empirical* knowledge perception is unrivalled: it is the only non-dependent and explanatorily basic source of such knowledge. This is not only true I suggest, but what the foundationalist has been saying all along. For beings such as ourselves the very idea of empirical knowledge is the idea of perceptual knowledge.

BIBLIOGRAPHY

Alston, W. (1976) 'Two Types of Foundationalism', The Journal of Philosophy, 73: 165-185)

Aristotle (1993), Posterior Analytics, 2nd edition, translated with a commentary by Jonathan Barnes (Oxford: Oxford University Press).

Audi, R. (1997), 'The Place of Testimony in the Fabric of Knowledge and Justification', American Philosophical Quarterly 34: 405-22.

Audi, R. (2006), 'Testimony, Credulity, and Veracity', in Lackey and Sosa (2006): 25-49.

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

Austin, J. L. (1979), 'Other Minds', in J. L. Austin, Philosophical Papers, 3rd edition, ed. J. Urmsen and G. Warnock (Oxford: Oxford University Press): 76-116.

Ayer, A. J. (1969), The Foundations of Empirical Knowledge (London: Macmillan).

Ayer, A. J. (2000), 'Knowing as having the right to be sure', in Bernecker and Dretske (2000): 7-12.

Ayers, M. R. (1991), Locke: vol. 1: Epistemology (London: Routledge).

Bernecker, S and Dretske, F. (2000), Knowledge: Readings in Contemporary Epistemology (Oxford: Oxford University Press).

BonJour, L. (1985), The Structure of Empirical Knowledge (Cambridge, Mass.: Harvard University Press).

BonJour, L. (1998), In Defence of Pure Reason (Cambridge: Cambridge University Press).

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

Burge, T. (1993), 'Content Preservation', Philosophical Review, 102: 457-88.

- Burge, T. (1998), 'Computer Proof, Apriori Knowledge, and Other Minds', in J. Tomberlin (ed.), Philosophical Perspectives, 12 (Malden, Mass.: Blackwell): 1-37.
- Burge, T. (2003), 'Perceptual Entitlement', Philosophy and Phenomenological Research, 67: 503-48.
- Campbell, J. (2002), Reference and Consciousness (Oxford: Clarendon Press).
- Cassam, Q. (2007), The Possibility of Knowledge (Oxford: Oxford University Press).
- Cassam, Q. (forthcoming), 'Can the Concept of Knowledge be Analysed?', in D. Pritchard and P. Greenough (eds.), Williamson on Knowledge.
- Cassam, Q. (forthcoming), 'Contemporary Responses to Descartes Philosophy of Mind', in J. Broughton and J. Carriero (eds.) A Companion to Descartes (Oxford: Blackwell).
- Chisholm, R. M. (1966), Theory of Knowledge (Englewood Cliffs, N. J.: Prentice-Hall Inc.)
- Chisholm, R. M. (1982), The Foundations of Knowing (Minneapolis: University of Minnesota Press).
- Cohen, S. (2002), 'Basic Knowledge and the Problem of Easy Knowledge', Philosophy and Phenomenological Research, 65: 309-30.
- Dancy, J. (1985), An Introduction to Contemporary Epistemology (Oxford: Basil Blackwell).
- Davidson, D. (2000), 'A Coherence Theory of Truth and Knowledge', Bernecker and Dretske (2000): 413-30.
- DePaul, M. (2001), Resurrecting Old-Fashioned Foundationalism (Lanham: Rowman & Littlefield Publishers).
- Descartes, R. (1996), Meditations on First Philosophy, edited by J. Cottingham

(Cambridge: Cambridge University Press).

Dretske, F. (1969), Seeing and Knowing (London: Routledge & Kegan Paul).

Dretske, F. (2000a), 'Simple Seeing', in F. Dretske, Perception, Knowledge and Belief (Cambridge: Cambridge University Press): 97-112.

Dretske, F. (2000b), 'Conclusive Reasons', Bernecker and Dretske (2000): 42-62.

Dummett, M. (1978), 'The Justification of Deduction', in Truth and Other Enigmas (London: Duckworth): 290-318.

Dummett, M. (1994) 'Testimony and Memory' in B. Matilal and A. Chakrabarti (eds.) Knowing From Words (Dordrecht: Kluwer Academic Publishers): 251-72.

Evans, G. (1982), The Varieties of Reference (Oxford: Clarendon Press).

Giaquinto, M. (1996), 'Non-Analytic Conceptual Knowledge', Mind, 105: 249-68.

Ginet, C. (2005), 'Infinetism is not the Solution to the Regress Problem', in Steup and Sosa (2005): 140-49.

Goldman, A. (1986), Epistemology and Cognition (Cambridge, Mass.: Harvard University Press).

Goldman, A. (1992), 'Discrimination and Perceptual Knowledge', in Liaisons: Philosophy meets the Cognitive and Social Sciences (Cambridge, Mass.: The MIT Press): 85-104.

Haack, S. (1993), Evidence and Inquiry: Towards Reconstruction in Epistemology (Oxford: Blackwell).

Harman, G. (1973), Thought (Princeton: Princeton University Press).

Hume, D. (1975), An Enquiry Concerning Human Understanding, in L. A. Selby-Bigge (ed.), Enquiries Concerning Human Understanding and Concerning the Principles of

- Morals, 3rd edition, revised by P. H. Nidditch (Oxford: Clarendon Press).
- Hume, D. (1978), A Treatise of Human Nature, 2nd edition, edited by P. H. Nidditch (Oxford: Clarendon Press).
- Klein, P. (2005), 'Infinetism is the Solution to the Regress Problem', in Steup and Sosa (2005) :131-40.
- Lackey, J. (1999), 'Testimonial Knowledge and Transmission', The Philosophical Quarterly, 49: 471-90.
- Lackey, J. (2005), 'Memory as a Generative Epistemic Source', Philosophy and Phenomenological Research, 70: 636-658.
- Lackey, J and Sosa, E. (2006), The Epistemology of Testimony (Oxford: Oxford University Press).
- Lehrer, K and Paxson, J. (2000), 'Knowledge: Undefeated Justified True Belief', Bernecker and Dretske (2000): 31-41.
- Lehrer, K. (2000), Theory of Knowledge, 2nd edition (Boulder: Westview Press).
- Lewis, C. I. (1929), Mind and the World Order: Outline of a Theory of Knowledge (New York: Dover Publications).
- Lewis, C. I. (1946), An Analysis of Knowledge and Valuation (La Salle, Ill.: Open Court).
- Locke, J. (1975), An Essay Concerning Human Understanding, ed. P. H. Nidditch (Oxford: Clarendon Press).
- Martin, M. G. F. (1992), 'Perception, Concepts and Memory', Philosophical Review, 107: 745-63.
- Martin, M. G. F. (1995), 'Perception' in A. C. Grayling (ed.) Philosophy: A Guide

through the subject (Oxford: Oxford University Press): 26-43.

Martin, M. G. F. (2001), 'Epistemic Openness and Perceptual Defeasibility', Philosophy and Phenomenological Research, 71: 441-48.

Martin, M. G. F. (2002), 'The Transparency of Experience', Mind and Language, 17: 376-425.

Martin, M. G. F. (forthcoming), Uncovering Appearances.

McDowell, J. (1994), Mind and World (Cambridge, Mass.: Harvard University Press).

McDowell, J. (1998), Meaning, Knowledge, and Reality (Cambridge, Mass.: Harvard University Press).

McDowell, J. (1998a), 'Knowledge and the Internal', in McDowell (1998): 395-413.

McDowell, J. (1998b), 'Knowledge by Hearsay', in McDowell (1998): 414-43.

Owens D. (2000) Reason Without Freedom: The Problem of Epistemic Normativity (London and New York: Routledge).

Peacocke, C. (1992), 'Scenarios, Concepts and Perception', in T. Crane (ed.) The Contents of Experience: Essays on Perception (Cambridge: Cambridge University Press): 105-35.

Plantinga, A. (1993) Warrant and Proper Function (New York: Oxford University Press)

Pollock, J. and Cruz, J. (1999), Contemporary Theories of Knowledge, 2nd edition (Lanham: Rowman & Littlefield Publishers, Inc.).

Pryor, J. (2000), 'The Skeptic and the Dogmatist', Nous, 34: 517-49.

Pryor, J. (2004), 'What's Wrong with Moore's Argument?', Philosophical Issues, 14: 349-78.

Pryor, J. (2005), 'There is Immediate Justification', in Steup and Sosa (2005): 181-202.

- Quinton, A. (1973), The Nature of Things (London: Routledge & Kegan Paul).
- Sainsbury, M. (1997), 'Easy Possibilities', Philosophy and Phenomenological Research, 57: 907-19.
- Radford, C. (1966) 'Knowledge - by examples', Analysis 27: 1-11.
- Ryle, G. (1949), The Concept of Mind (Chicago: University of Chicago Press).
- Searle, J. (1983), Intentionality (Cambridge: Cambridge University Press).
- Sellars, W. (1956), 'Empiricism and the Philosophy of Mind', in H. Feigl and M. Scriven (eds.), Minnesota Studies in the Philosophy of Science, vol. 1 (Minneapolis: Minnesota University Press): 253-329.
- Sextus Empiricus (1967), Outlines of Pyrrhonism, translated by R. G. Bury (Cambridge, Mass.; Harvard University Press).
- Shope, R. K. (1983), The Analysis of Knowing: A Decade of Research (Princeton, Princeton University Press).
- Snowdon, P. F. (1980-1), 'Perception, Vision and Causation', Proceedings of the Aristotelian Society, 81: 175-92.
- Snowdon, P. F. (1998), 'Strawson on the Concept of Perception', in L. Hahn (ed.) The Philosophy of P. F. Strawson (Chicago and Lasalle: Open Court): 293-310.
- Snowdon, P. F. (2005), 'The Formulation of Disjunctivism: A Response to Fish', Proceedings of the Aristotelian Society, 105: 129-41.
- Sosa, E. (1980), 'The Raft and the Pyramid', in P. French, T. Uehling and H. Wettstein (eds.) Midwest Studies in Philosophy 5: Studies in Epistemology (Minneapolis: University of Minnesota Press): 3-26.
- Sosa, E. (1991), Knowledge in Perspective (Cambridge: Cambridge University Press).

- Soteriou, M. (2005), 'The Subjective View of Experience and its Objective Commitments', Proceedings of the Aristotelian Society, 105: 193-206.
- Steup, M. and Sosa, E. (2005), Contemporary Debates in Epistemology (Oxford: Blackwell Publishing).
- Strawson, P. F. (1988), 'Perception and its Objects', in J. Dancy (ed.) Perceptual Knowledge (Oxford: Oxford University Press): 92-112.
- Stroud, B. (2000), 'Epistemological Reflection on Knowledge of the External World', in Understanding Human Knowledge (Oxford: Oxford University Press): 122-38.
- Travis, C. (2004), 'The Silence of the Senses', Mind, 113: 57-94.
- Van Cleve, J. (2006), 'Reid on the Credit of Human Testimony', in Lackey and Sosa (2006): 50-74.
- Vogel, J. (1990), 'Are There Counterexamples to the Closure Principle?', in M. Roth and G. Ross (eds.) Doubting (Dordrecht: Kluwer): 13-27.
- Williams, B. (1978), Descartes: the Project of Pure Enquiry (Harmondsworth: Penguin).
- Williams, M. (1977), Groundless Belief (Oxford: Blackwell).
- Williams, M. (2005), 'Doing without Immediate Justification', in Steup and Sosa (2005): 202-16.
- Williamson, T. (1995), 'Is Knowing a State of Mind?', Mind, 104: 533-66.
- Williamson, T. (2000), Knowledge and its Limits (Oxford: Oxford University Press).
- Williamson, T. (forthcoming), 'On Being Justified in One's Head'.
- Wright, C. (1985), 'Facts and Certainty', Proceedings of the British Academy, 71: 429-72.
- Wright, C. (1991), 'Imploding the Demon', Mind 100: 87-116.

Zalabardo, J. (unpublished), 'Internalist Foundationalism and the Problem of Epistemic Regress'.