



THE UNIVERSITY *of* EDINBURGH

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

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

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

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

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

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

Meta-epistemological Scepticism:
Criticisms and a Defence



Submitted for the degree of PhD by Research
The University of Edinburgh
2015

Declaration

I, Christopher Ranalli, hereby declare the following. The present thesis, submitted for examination in pursuit of a Ph.D., by Research in Philosophy, has been entirely composed by myself, and it has not been submitted in pursuit of any other academic degree, or professional qualification.

Signature:

Date: 30-01-2015

Acknowledgements

Descartes's *Meditations on First Philosophy* is the first piece of philosophy I've ever read. Since then I've been interested in scepticism about our knowledge of the external world and related issues. This thesis grew out of a debate about the nature of scepticism and its arguments. I owe a significant debt to some of the core participants of that debate.

In particular, I would like to thank Professor Duncan Pritchard for his helpful guidance and encouragement throughout. His work has been one of the major sources of inspiration for this thesis. Moreover, his unwavering support, advice, and kindness as a supervisor has been a huge motivation. I would also like to thank Allan Hazlett for his very helpful comments and suggestions. A special thanks also goes to Professor Barry Stroud for his helpful discussions, which have significantly helped me to understand the philosophical issues surrounding scepticism. Furthermore, I would also like to thank the School of Philosophy, Psychology, and Language Sciences for awarding me the Career Development Scholarship, which supported me during the PhD.

I am also indebted to the postgraduate community at the University of Edinburgh, and in particular the members of the Epistemology Research Cluster. I am especially grateful to Cameron Boulton, Lee Whittington, Alex Whalen, Lani Watson, Ju Wang, Lauren Ware, Natalie Ashton, Robin McKenna, Scott Kerzmann, Ben Sworn, Di Yang, Joey Pollock, and Tim Kunke. All of you have helped contribute to a wonderful atmosphere, excellent discussion, and to being amazing philosophical companions. I would also like to thank friends from the Edinburgh Philosophy Society, especially Anna Baker, Sam Carter, Jamie Sutherland, Andrea Blomqvist, and Richard Fitzpatrick. All of these people have been extremely helpful and kind, and have helped to make it all worthwhile and wonderful. I am especially thankful to Karoline von dem Berge, for her amazing personal support, kindness, companionship, and immense commitment. I am extremely grateful and blessed to have such amazingly wonderful people in my life.

Finally, I would like to extend a sincere thanks to the examiners and the chair for all of their time, effort, and patience. I am grateful to all of you.

Abstract

The epistemological problem of the external world asks:

- (1) “How is knowledge of the world possible given certain obstacles which make it look impossible?”

This is a “how-possible?” question: it asks how something is possible given certain obstacles which make it look impossible (cf. Cassam 2007; Nozick 1981; Stroud 1984). Now consider the following question, which asks:

- (2) “How is a philosophically satisfying answer to (1) possible?”

Scepticism is the thesis that knowledge of the world is impossible. It therefore represents a negative answer to the first question. Meta-epistemological scepticism is the thesis that a satisfying philosophical explanation of how our knowledge of the world is possible is itself not possible. It therefore represents a negative answer to the second question.

In this thesis, I explore the prospects of meta-epistemological scepticism. In particular, I structure the thesis around two master arguments from Stroud (1984, 2000, 2004, and 2009) for meta-epistemological scepticism. The first argument is what I call “Stroud’s puzzle”, and the second argument is “Stroud’s dilemma” (cf. Cassam 2009). I argue that Stroud’s puzzle fails to provide adequate support for meta-epistemological scepticism. However, I also argue that Stroud’s dilemma withstands serious objections (e.g., from Sosa 1994, Williams 1996, and Cassam 2009). In short, while Stroud’s puzzle fails to provide adequate support for meta-epistemological scepticism, Stroud’s dilemma does seem to provide adequate support for meta-epistemological scepticism. This thesis therefore represents a partial defence of meta-epistemological scepticism. Meta-epistemological scepticism is therefore a live option in epistemology.

In Chapter 1, I explain what meta-epistemological is, present Stroud's puzzle and Stroud's dilemma for meta-epistemological scepticism, and argue that meta-epistemological sceptics are not committed to first-order scepticism. In Chapter 2, I examine what I call the "anti-revisionist" premise of Stroud's puzzle and argue that it lacks adequate support. In Chapter 3, I examine the "conditional scepticism" premise of Stroud's puzzle and argue that it lacks adequate support. In Chapter 4, I look at Williams's (1996) master argument against Stroud's dilemma, and argue that it fails. In Chapter 5, I look at externalist responses to Stroud's dilemma, and in particular, Sosa (1994). I argue that Sosa's objection fails, and therefore Stroud's dilemma survives serious externalist objections. In Chapter 6, I explain Cassam's (2009) argument against Stroud's dilemma, and I argue that it fails. Chapter 7 concludes the thesis, summarising the main results.

Publications

This thesis includes versions of the following publications (additional author's contributions omitted):

1. Ranalli, Chris (2014). Luck, Propositional Perception, and the Entailment Thesis. *Synthese* 191 (6): 1223-1247.
2. Pritchard, Duncan and Chris Ranalli (*forthcoming*). "On Meta-epistemological Scepticism". *Traditional Epistemic Internalism*, (eds.) M. Bergmann & B. Coppenger, (Oxford University Press).

Preface

This thesis is structured around two arguments in support of what I call “meta-epistemological scepticism”. This is the thesis that a satisfying philosophical explanation of how knowledge of the world is possible is not possible.¹ Both of the arguments in support of meta-epistemological scepticism are extracted from Barry Stroud’s work.² The first argument is what I call “Stroud’s puzzle” and the second argument is what I call “Stroud’s dilemma”.³ I argue that while Stroud’s puzzle fails to provide adequate support for meta-epistemological scepticism, the arguments against Stroud’s dilemma are defective, so that meta-epistemological scepticism remains a live option in epistemology.

Although work on scepticism has been thriving since the late 1970’s, in part due to advancements in our understanding of the form and commitments of the best sceptical arguments, and the development of epistemic externalist, contextualist, and neo-Moorean views, not much attention has been given to the meta-epistemological disputes that arise out of one’s engagement with scepticism.^{4, 5, 6} This thesis can be read as at least one attempt to fill this gap.

¹ In Chapter 1, I explain in detail what meta-epistemological scepticism is.

² Stroud’s puzzle is extracted from Stroud (1984, 2009). Stroud’s dilemma is extracted from his (2004) paper “Perceptual Knowledge and Epistemological Satisfaction” in *Ernest Sosa and his Critics*, and his (2000) collection of papers titled *Understanding Human Knowledge*. The relevant papers include his “Scepticism and the Possibility of Knowledge” pp. 1-8, “Understanding Human Knowledge in General” pp. 99-121, “Epistemological Reflection on Knowledge of the External World” pp. 122-138, and “Scepticism, ‘Externalism’, and the Goal of Epistemology” pp. 139-154.

³ Cassam (2009) also refers to the argument as “Stroud’s dilemma”.

⁴ For example, the main epistemic externalist responses to scepticism come in at least two varieties: “neo-Moorean” views which argue, from externalist principles of knowledge or justification, that we can know the denials of sceptical hypotheses, and “closure denying” views, which argue that the conditional premise of the sceptical argument, that if you know anything about the world, then you know that the sceptical hypotheses are false, is false. For the former, see Pritchard (2005), Greco (2000), Hill (1999), and Sosa (2014, 2009, 2008). For the latter, see Dretske (2014, 2003, and 1970) and Nozick (1981). Internalist “neo-Moorean” views come in at least three varieties: *dogmatism*, *disjunctivism*, and *abductivism*. According to dogmatism, we can know the denials of sceptical hypotheses on the basis of sensory experience (given the right kind of phenomenology), even though our sensory experience only gives us non-entailing evidence for our beliefs about the world. See Pryor (2014, 2004, 2001, 2000), and Huemer (2013, 2007, 2001). According to epistemological disjunctivism, we can know the denials of sceptical hypotheses on the basis of sensory experience because they give us reflectively accessible entailing evidence for our beliefs about the world. See McDowell (2008, 1995), Pritchard (2012, 2011, 2008), and Pritchard and Ranalli (*forthcoming*). For a related type of anti-sceptical argument, see

Here is the structure of the thesis. Chapter 1 introduces the reader to “scepticism”, “higher-order scepticism”, “meta-epistemological scepticism”, Stroud’s puzzle and Stroud’s dilemma in support of meta-epistemological scepticism. I argue that meta-epistemological sceptics are not committed to first-order scepticism, even though first-order sceptics are plausibly committed to meta-epistemological scepticism.

Chapters 2 and 3 focus on Stroud’s puzzle. Stroud’s puzzle is a dilemma arising out of certain meta-epistemological features of “Descartes’s condition”. This is the requirement that, for any proposition about the world P , if one knows that P , then they know that they are not merely dreaming that P . Chapter 2 examines the second premise of Stroud’s puzzle. This is the “anti-revisionist” premise. I argue that the arguments from Stroud (1984) fail to provide adequate support for this premise. I then consider some additional arguments and argue that these arguments also fail to provide adequate support for this premise.

Williamson (2000). And according to abductivism, we can acquire abductive evidence to reject the sceptical hypotheses. See Vogel (2014, 2007, and 1990). *Contextualists* typically argue that in “philosophical contexts” we fail to know the denials of sceptical hypotheses, while in “ordinary contexts” our ordinary knowledge of the world implies that we can know the denial of sceptical hypotheses. See Cohen (2014, 2001, 1999) and DeRose (1995). There are also “Wittgensteinian” responses, some of which we say we have a default, non-evidential warrant to reject the sceptical hypotheses. See Wright (2004, 1991) and Wright and Davies (2004). Some other “Wittgensteinian” responses are Hazlett (2006) and Pritchard (*forthcoming*).

⁵ Although not much work has been done on meta-epistemological scepticism, there are exceptions. Nagel (1989) argues for meta-epistemological scepticism on the basis of his theory of the objective-self. There are connections between Nagel and Stroud’s arguments that I highlight in Chapter 1. As I read Fumerton (1996, 2006), he supports meta-epistemological scepticism because he thinks that any satisfying theory of how our knowledge of the world is possible would need to explain how we could have direct acquaintance with the truth-makers of our beliefs about the world, but he denies that it’s possible to have direct acquaintance with those types of truth-makers. I compare Fumerton and Stroud’s arguments in Chapter 5. Pritchard (2005) endorses meta-epistemological scepticism on the basis of his view that any satisfying theory of how knowledge of the world is possible would need to explain how we can eliminate reflective luck, but he doubts that this is possible. See Pritchard (2005, 208-213). However, Pritchard (2012) seems to have retracted this view. Unger (1984) also seems to be committed to meta-epistemological scepticism. He argues, on the basis of his hypothesis of philosophical relativity, that the problem of the external world lacks an objectively correct answer. And one might plausibly regard there being an objectively correct answer here as a necessary condition on a satisfying philosophical answer to the problem of the external world. For general scepticism about philosophy, see Brennan (2010).

⁶ For work on the nature of the sceptical argument, see Brueckner (1994), Cohen (1998), Fumerton (1995), Pritchard (2005, *forthcoming*), Pryor (2000), Vogel (2004), Wright (1991), and Yalçın (1992). For recent work on whether there is an underdetermination-based sceptical argument that is separate from a closure-based sceptical argument, see McCain (2013).

Chapter 3 examines the first premise of Stroud’s puzzle. This is the premise that if Descartes’s condition is true, then scepticism is true. It’s the “conditional scepticism” premise. I evaluate and adapt four arguments from Stroud (1984, 2000, and 2009) in support of this premise, and argue that all of them are defective.

Chapters 4, 5, and 6 focus on Stroud’s dilemma. Chapter 4 examines Williams (1996) argument that Stroud’s dilemma arises because it depends on a false presupposition that he calls “epistemological realism” (Williams 1996, *xx*).⁷ I argue that it’s not clear that it depends on any controversial presupposition, and that even if it were dependent on the presupposition that Williams identifies, it’s not clear that it’s false.

Chapter 5 examines Sosa’s (1994, 1997) externalist response to Stroud’s dilemma. I defend Stroud’s arguments from Sosa’s response, and provide an independent argument for a requirement on satisfying philosophical explanations of how knowledge of the world is possible.⁸

Chapter 6 examines Cassam’s (2009) response to Stroud’s dilemma. I argue that Cassam’s argument is defective. Chapter 7 summarises the main results of each chapter, and concludes the thesis.

⁷ Cf. Williams (1996*b*, 2001, and 2011).

⁸ I also give some attention to Fumerton’s (2006) comments.

Table of Contents

Chapter 1

On Meta-epistemological Scepticism.....	14
1. Introduction.....	14
2 What is Meta-epistemological Scepticism?	17
2.1 Definitions.....	17
2.2 Summary	24
3 Meta-epistemological Scepticism: three prima facie challenges	26
3.1 Meta-epistemological scepticism: Conclusion.....	31
4 Scepticism and Meta-epistemological scepticism.....	32
4.1 The Argument from Ignorance.....	33
5 Stroud’s Puzzle	38
5.1 Stroud’s Puzzle: the basic dilemma	38
5.2 Stroud’s Puzzle: the argument	41
5.3 Stroud’s Puzzle: “Obvious Truths” and “Platitudes”	45
6 Stroud’s Dilemma.....	51
6.1 Generality.....	52
6.2 From EPR to Scepticism	55
6.3 From the denial of EPR to Philosophical Dissatisfaction	57
6.4 Stroud’s Dilemma	58
7 Conclusion.....	59

Chapter 2

Stroud’s Puzzle: Descartes’s Condition and Revisionism	61
1 Introduction	61
2 Revisionism	63
2.1 What is Revisionism?	63
3 Stroud and Anti-Revisionism	66
3.1 Stroud’s Argument: The Strong Version	66
3.2 Stroud’s Argument: The Weak Version.....	69
3.3 Kaplan’s Condition	76
4 Arguments for Anti-Revisionism	78
4.1 Argument from Reflective Equilibrium	79
4.2 Argument from Lewis’s Maxim.....	80
4.3 Argument from the Principle of Charity	83
5 Interlude.....	85

5.1	Descartes’s Condition and Platitudes.....	88
5.2	How can Descartes’s Condition depend on “platitudes”?.....	88
5.3	Response to Stroud’s Argument.....	96
5.4	Hinge Propositions.....	100
5.4.1	Hinge Propositions and their Properties.....	100
5.5	Deep epistemic grammar.....	104
6	Conclusion.....	106

Chapter 3

Stroud’s Puzzle: Conditional Scepticism	108
1 Introduction	108
2 Stroud’s Arguments	109
2.1 The Regress Argument.....	111
2.2 Stroud’s Regress Argument: Evaluation.....	113
2.3 The No Test Argument.....	114
3 Stroud’s No Test Argument: Evaluation.....	118
3.1 Descartes’s Condition: The Necessary Condition Reading	118
3.2 Descartes’s Condition: The Necessary Pre-condition Reading.....	123
4 Stroud’s New No Test Argument	129
5 Stroud’s New No Test Argument: Evaluation	134
5.1 The First Premise	135
5.3 Assessment of the First Premise	138
5.3.1 The Knowledge Reading.....	138
5.3.2 The Revelation Reading.....	138
5.3.3 The Discrimination Reading	139
5.4 The Second Premise.....	139
5.4.1 Assessment of the Second Premise: Reliable Indicators and Knowledge.....	141
6 Conclusion.....	142

Chapter 4

Stroud’s Dilemma: Williams and Epistemological Realism.....	143
1 Introduction	143
2 Stroud’s Dilemma.....	144
2.1 Stroud’s Dilemma: EPR.....	144
2.2 Responding to Stroud’s dilemma.....	150
3 The Traditional Epistemological Project	154

4	Williams’s Master Argument	155
5	Williams’s Master Argument: Premise 1	159
5.1	Williams’s Argument for Premise 1c	162
5.2	Williams’s Argument for Premise 1b	163
5.3	Response to Williams’s Argument for Premise 1c	164
5.4	Williams’s Argument for Premise 1e	166
5.5	Response to Williams’s Argument for Premise 1e	167
5.5.1	The First Question	171
5.5.2	The Second Question	171
6	Williams’s Master Argument: Premise 2	172
6.1	Williams’s Argument for Premise 2	173
7	Conclusion	176

Chapter 5

Stroud’s Dilemma: Internalism and Externalism	178
1 Introduction	178
2 Externalist Responses to Stroud’s Dilemma	180
2.1 The First Externalist Response	181
2.2 The Second Externalist Response	185
3 Stroud’s Argument: Two Interpretations	193
3.1 The Internalist Interpretation	193
3.1.1 Internalism and the “Straightforward” Answer	195
3.2 The Circularity Interpretation	198
4 Stroud’s Anti-Externalist Argument	203
4.1 A Comparison: The Externalist Epistemologist and the Externalist Astrologer	204
5 Objections and Replies	206
6 Conclusion	213

Chapter 6

Stroud’s Dilemma: Propositional Perception, Simple Perception, and Cassam’s Response	214
1 Introduction	214
2 Cassam’s Response to EPR: Part I	216
2.1 The Argument from Generality	216
2.2 Response to Cassam	220

2.3	Respecting EPR.....	223
3	Propositional Perception and the Entailment Thesis	225
3.1	Motivating the Entailment Thesis	225
3.2	The Argument from Luck	230
3.3	Resisting the Argument from Luck.....	237
3.4	The Argument from Defeat.....	243
3.5	The Argument from Belief.....	248
3.6	Neutralizing the Argument from Belief	250
3.7	Dissipating the Argument from Belief.....	254
4	Cassam’s Response to EPR: Part II.....	260
5	Conclusion.....	266
Chapter 7 Meta-epistemological Scepticism: Afterword		268
1	Introduction	268
2	Review.....	268
3	Objections and Replies.....	271
3.1	Objection 1	271
3.2	Objection 2	273
4	Conclusion.....	275
Bibliography		276

Chapter 1

On Meta-epistemological Scepticism

1. Introduction

This chapter outlines the plan for the thesis: the core problems and the theses defended throughout. The focus of this thesis are the arguments from Barry Stroud (1984, 2000, 2009*a*, 2009*b*, 2011) for the view that I call “meta-epistemological scepticism”. This is the thesis that a satisfying, philosophical explanation of how our knowledge of the external world is possible is not possible.⁹

Here’s the plan for this chapter. In section 2, I explain what meta-epistemological scepticism is in more detail, and distinguish it from several other, related sceptical theses. In particular, I argue that even a positive and correct explanation of how our knowledge of the world is possible might not be sufficient for philosophical satisfaction.¹⁰

In section 3, I defend meta-epistemological scepticism from three prima facie objections. I do this in order to motivate meta-epistemological scepticism. In particular, I argue that meta-epistemological scepticism does not entail first-

⁹ In the next section (section 2), I explain in detail what “meta-epistemological scepticism” means, and I distinguish it from other kinds of scepticism. Arguably, one finds arguments for meta-epistemological scepticism in the work of Nagel (1986), Fumerton (1995, 2006), and Williams (1996). Later in the thesis (Chapters 4 and 5) I examine more closely Williams (1996) and Fumerton (2006).

¹⁰ I use the phrase “philosophical satisfaction” to pick out a pro-attitude towards philosophically satisfying explanations of how some type of knowledge is possible. In section 3, I discuss some possible necessary conditions on philosophical satisfaction. Throughout the thesis (chapters 5 and 6, in particular) I critically evaluate Stroud’s conditions for philosophical satisfaction.

order scepticism. In section 4, I outline first-order scepticism, and the argument from ignorance that is used to provide adequate support for that thesis. In section 5 and section 6, I provide a rough sketch of Stroud's two master arguments for meta-epistemological scepticism. That it is not obvious that these arguments are unsound provides motivation for the thesis that meta-epistemological scepticism does not require the truth of first-order scepticism. Instead, first-order scepticism is just one way among many ways of supporting meta-epistemological scepticism.

In the remainder of this section, I want to briefly outline Stroud's two master arguments for meta-epistemological scepticism. The first of these arguments presents the epistemologist who aims to answer "the problem of the external world", the problem of explaining how our knowledge of the external world is possible, with a certain meta-epistemological puzzle. The puzzle, I suggest, takes the form of a dilemma. Roughly, the dilemma is that either we accept scepticism about the external world, or else we accept that a "general procedure we recognize and insist on in making in assessing knowledge-claims in everyday and scientific life" is false (Stroud 1984, 30-31). But, intuitively, neither option is philosophically satisfying. On the one hand, scepticism is intuitively an unsatisfactory explanation of how knowledge of the world is possible, because it says that, despite appearances to the contrary, it's impossible. On the other hand, the falsity of an otherwise 'ordinary' or 'everyday' epistemic principle is intuitively unsatisfying as well. As Stroud argues: "we have no notion of knowledge other than what is embodied in those procedures and practices" (Stroud 1984, 31). So, on Stroud's view, neither option leaves the epistemologist with a satisfying explanation of how knowledge of the world is possible.¹¹

The second of these arguments also presents the epistemologist who seeks to answer the problem of the external world with a puzzle. The puzzle takes the

¹¹ Abstracting from the details of Stroud's presentation of the puzzle, Williams (1996) calls this the "epistemologist's dilemma". See Williams (1996), pp. 17-22, 45.

form of a dilemma. On the one hand, the kind of explanation that is needed to answer the problem satisfactorily is a fully general, philosophical explanation. According to Stroud, what epistemologists seek is a “completely general” philosophical explanation of how knowledge of the world is possible (Stroud 2000, 101):

We aspire in philosophy to see ourselves as knowing all of most of the things we think we know and to understand how all of that knowledge is possible. We want an explanation, not just of this or that item or piece of knowledge, but of knowledge, or knowledge of a certain kind, in general (Stroud 2000, 144).

The trouble that this kind of aim gets us in is that it seems to lead the following sort of requirement on a philosophical explanation of how our knowledge of the world is possible:

[...] we must explain it on the basis of another, prior kind of knowledge that does not imply or presuppose any of the knowledge we are trying to explain. Without that, we will not be explaining the knowledge in question in the proper, fully general way (Stroud 2000, 104).

But Stroud thinks that once this requirement is in place, scepticism follows. That’s the first horn of the dilemma.¹² On the other hand, if we reject this requirement, that too gets us in trouble. For according to Stroud:

[...] if in order to resist that [sceptical] conclusion, we no longer see ourselves in the traditional way, we will not have a satisfactorily general explanation of all of our knowledge (Stroud 2000, 106).

For Stroud then, this consequence should also “leave us dissatisfied”, since it means that we will have repudiated what looked like a worthwhile intellectual goal (Stroud 2000, 121). This is the second horn of the dilemma.

Stroud has furnished us with two master arguments for meta-epistemological scepticism. Both arguments take the form of a dilemma, where each horn has the consequence that a satisfying philosophical explanation of how our knowledge of the world is possible is not possible.

¹² Treating this problem as a dilemma, “Stroud’s Dilemma”, has a precedent in Cassam (2009), and, to a lesser extent, in Williams (1996), Chapter 3.

2 What is Meta-epistemological Scepticism?

In this section, I first explain the difference between the epistemological problem of the external world and the meta-epistemological problem of the external world. I then use this distinction to show how scepticism is a possible negative answer to the first problem, while meta-epistemological scepticism is a possible negative answer to the second problem. In the final part of this section, I expand on the commitments of the meta-epistemological sceptic.

2.1 Definitions

What is meta-epistemological scepticism? In order to better understand what meta-epistemological scepticism is, we need to get a handle on the first-order problem of scepticism about our knowledge of the external world. For ease of exposition, I will refer to this problem as the “problem of the external world”.

Following Nozick (1981), Stroud (1984, 2000), and Cassam (2007), it is best to think of the problem of the external world as an epistemological ‘how-possible?’ question. This kind of question asks how something is possible given something which makes it look impossible. What the problem of the external world asks is how knowledge of the external world is possible given certain obstacles which makes that kind of knowledge look impossible.¹³

Now the meta-epistemological problem of our knowledge of the external world can also be framed as a ‘how-possible?’ question. This problem asks how a philosophically satisfying answer to the problem of the external world is possible, given certain obstacles which make a philosophically satisfying answer to that problem look impossible.

In this thesis, I argue that even if there are good reasons to think that knowledge of the external world is possible, there are nevertheless good reasons to think that a satisfactory philosophical explanation of how that kind

¹³ See Stroud (1984, 140-144). Cf. Stroud (2000). See also Cassam (2007, 1-10), and Nozick (1981, 8-11).

of knowledge is possible is *not* possible. This later claim, that a satisfactory philosophical explanation of how our knowledge of the world is possible is itself not possible, is what I call “meta-epistemological scepticism”. As we will see in this chapter, there are two interesting arguments from Stroud (1984, 2000, 2009, 2011) in favour of meta-epistemological scepticism. What I will be arguing in Chapter 2 and Chapter 3 is that the first of these arguments from Stroud (1984, 2000, 2009) is not a good argument. However, in the second half of this thesis, I turn to the second argument from Stroud (2000, 2011) in support of meta-epistemological scepticism, and argue that the responses to this argument fail.¹⁴ In this fashion, I argue that Stroud’s second argument for meta-epistemological scepticism remains intact, surviving various objections that have been raised against it.

Now that I have provided a brief exposition of the problem of the external world, the meta-epistemological problem of the external world, and the core theses that I will be arguing for in this thesis, I want to turn our attention to some other terms that I would like to distinguish from meta-epistemological scepticism.

I will use the term “scepticism” to pick out the thesis that knowledge of the external world is impossible. When I use the term “the external world” (or “the world” for short), I mean the set of all actual true propositions about non-psychological reality. And by “non-psychological reality”, I mean the set of all actual true propositions which have the following property: they do not entail that anyone stands in any psychological relation to any of them.¹⁵ Here,

¹⁴ The responses include: Williams (2001, 1996a, 1996b, 1988), Sosa (1994, 1997), Cassam (2007, 2009), Fumerton (2006).

¹⁵ I am using G. E. Moore’s conception of physical objects here. Moore includes within his description of what a physical thing is that it is independent of anyone’s psychological states: “many ‘things’, e.g., after-images, double images, bodily pains, which might be said to be ‘presented in space’, are nevertheless not to be reckoned as ‘things that are to be met with in space’[.] [T]here is no contradiction in supposing that there have been and are ‘to be met with in space’ things which never have been, are not now, and never will be perceived, nor in supposing that among those of them which have at some time been perceived many existed at times at which they were not being perceived” (Moore 1959, 156). In his (1994; 2000) discussion of transcendental arguments, Stroud says that many transcendental

“psychological relation” picks out the relations that the following phrases express: “ x thinks that p ”, “ x believes that p ”, “ x has an experience as of p ”, “ x represents p ”, “ x knows that p ”, “ x understands p ”, and so forth. I don’t offer an analysis of these relations. I maintain only that, from any non-trivial true proposition about non-psychological reality, no proposition with a psychological-relation as a part logically follows.

Now facts about numbers, sets, line-segments, and other mathematical or abstract entities might satisfy the conditions for being part of non-psychological reality. So, I will restrict my attention to those non-psychological truths which are about physical entities. By “physical entity” I mean any entity that has spatio-temporal properties. This includes water, human beings, Saturn, my desk, stones, and so on. And I include propositions about events which involve physical entities, such as cricket matches, the Earth revolving around the Sun, spilling coffee, rocks falling, and so on.

Scepticism is thereby the thesis that knowledge of non-psychological truths about physical things or events is impossible. If scepticism is true, then no one can know that water exists, that they are watching a cricket match, that the Earth revolves around the sun, that there are books on their desk, that there are human beings, and so on.

Intellectually, scepticism is devastating. It is devastating because it says that, for many of the propositions we think we know, not only do we not know them, but we cannot know them. And this is why scepticism is “absurd”. It is absurd because it says that for many propositions we think or believe that we know, we can’t know. It contradicts what we all otherwise believe. It would be less troubling if it only said that there are some propositions about non-psychological reality that we couldn’t know. Perhaps there are propositions which are peculiarly difficult to know. But scepticism doesn’t say that there are propositions about non-psychological reality which are peculiarly difficult to

arguments aim to establish conclusions which are true “independently of all such broadly psychological facts”—that they “appear to state or imply nothing about anyone’s thinking of experience things in certain ways (Stroud 2000, 256).

know. Instead, it says that there aren't any propositions about non-psychological reality that we can know.

If "scepticism" picks out the thesis that knowledge of non-psychological, physical reality is impossible, then let "meta-scepticism" pick out the thesis that we cannot know whether knowledge of non-psychological, physical reality is possible. This expresses the thesis that no one can know whether the philosophical thesis picked out by "scepticism" is true.

Now, knowledge is factive: *S* knows that *p* only if *p* is true. So, scepticism implies that no one knows that knowledge of non-psychological, physical reality is possible because that kind of knowledge is impossible. In short, scepticism implies second-order scepticism (and *n*-order scepticism). Nevertheless, scepticism doesn't imply meta-scepticism. That scepticism is true doesn't imply that we can't know that it's true (that I can't know anything about non-psychological physical reality does not imply that I can't know that I can't know anything about non-psychological physical reality). And meta-scepticism doesn't imply scepticism. That we cannot know whether we can know anything about non-psychological, physical reality doesn't entail that we can't know anything about non-psychological, physical reality.

What I'm calling "meta-epistemological scepticism" is different from what I called "scepticism", "meta-scepticism", and "second-order" or "higher-order" scepticism. Meta-epistemological scepticism is a sceptical thesis about a certain kind of philosophical explanation of knowledge. In particular, it says that that a satisfying, philosophical explanation of how our putative knowledge of non-psychological, physical reality is possible is not possible.¹⁶ A core statement of the view is expressed in Stroud (2000). He says that:

[...] however much we came to learn about this or that aspect of human knowledge, thought, and perception, there might be still be

¹⁶ Cf. Williams (1994, 15). See also Williams (2012) discussion of Wright's (1985, 2004) anti-sceptical strategy, where he argues that Wright is a meta-epistemological sceptic as I have defined "meta-epistemological scepticism". See specifically (Williams 2012, 373).

nothing that could satisfy us as a philosophical understanding of how human knowledge is possible (Stroud 2000, 100).

I find the force and resilience of scepticism in the theory of knowledge to be so great, once the epistemological project is accepted, and I find its consequences so paradoxical, that I think the best thing to do now is to look much more closely and critically at the very enterprise of which scepticism or one of its rivals is the outcome: the task of the philosophical theory of knowledge itself. [...] I wonder whether there is a coherent point of view from which we could get a satisfactory understanding of ourselves of the kind we aspire to (Stroud 2000, 141).

I give reasons for thinking that, because of the way we apparently want to understand human knowledge is philosophy, there is and could be no satisfactory answer to the question [of how knowledge is possible] (Stroud 2000, xiv).

As I noted in the introduction to the thesis, the best explanation of Stroud's sceptical position is that his scepticism must reside at some higher-order level, since he disavows scepticism at the first-order.¹⁷ On the other hand, while Stroud does endorse a conditional form of scepticism, such that, if certain principles of knowledge are true, then scepticism is true, this is not the extent of his sceptical worries either. Rather, as Williams (1996) highlights, Stroud seems to think that "nothing could satisfy us as a fully general explanation of how we come to have knowledge of the world" (Williams 1996, 377).¹⁸

What Stroud thinks is impossible, then, is a certain kind of philosophical achievement. The particular philosophical achievement he has in mind is a satisfying philosophical understanding of how our knowledge of the external

¹⁷ See Stroud (1984, 2), and Stroud (2010), in (eds.) Edmonds and Warburton. Indeed, Stroud (2004) disavows scepticism at the *second-order* level as well, since he thinks that it's possible for us to know that other people see (and therefore know) that there are physical objects and properties. However, he maintains that the fact that we can know, by perceiving, that other people know, by perceiving, that there are physical objects and properties, does not generalize into a satisfying *philosophical explanation* of how our knowledge of any physical objects and properties is possible. I explore this element of Stroud's (2004) in relation to Fumerton and Sosa in Chapter 5. See also Pritchard and Ranalli (*forthcoming*) for a critical discussion of Fumerton and Stroud on this score.

¹⁸ cf. Williams (2011): "Since any satisfying philosophical view of the world and our place in it must accommodate our sense that we are capable of knowing about the world around us, and since skepticism suggests that no view will ever succeed in this, we can turn away from investigating philosophical skepticism only at the cost of *skepticism about philosophy*" (Williams 2011, 45). See also Bridges and Kolodney (2011, 11).

world is possible.¹⁹ Taking it out of the first-person, what Stroud thinks is impossible, then, is a satisfying philosophical explanation of how knowledge of the external world is possible.

Following Stroud and Williams, notice that I use the term “satisfying” to modify “philosophical explanation”, because no one denies that there are *philosophical explanations* of how knowledge of the external world is possible. That would be an uninteresting thesis, because it is obvious that it’s false: several prominent epistemologists have proposed philosophical explanations of how our knowledge of the world is possible. What Stroud denies is that these explanations can or should provide us with ‘philosophical satisfaction’. On his view, there are insurmountable obstacles to such explanations amounting to philosophically satisfying explanations.

Since there are explanations of how our knowledge of the world is possible, at least one thing at issue is the *correctness* of these explanations: are any of these explanations actually correct? But notice that mere correctness isn’t sufficient for philosophical satisfaction. After all, scepticism is a philosophical theory of knowledge.²⁰ It’s simply the *negative* theory which says that knowledge of non-psychological, physical reality is not possible. By implication, if scepticism is true, then none of the positive theories of how knowledge of the world is possible are true. And, *prima facie* at least, scepticism is *not* a satisfactory philosophical theory of knowledge, even if it’s correct.

Moreover, results in cognitive psychology imply that we can perceptually

¹⁹ A close statement of this view from Stroud can be put as follows: “once we really understand what we aspire to in the philosophical study of knowledge, [...] we will be forever unable to get the kind of understanding that would satisfy us” (Stroud 1989, 32), in (eds.) M. Clay and K. Lehrer (1989). Compare this with Stroud (2009, 569).

²⁰ cf. Stroud (2000, 141). Note that I am following Stroud here in his use of ‘theory of knowledge’ to pick out *explanations* of *how* knowledge of some type is *possible*. This should be contrasted with more orthodox uses of ‘theory of knowledge’, which pick out explanations of the individually necessary and jointly sufficient conditions for *S knows that p*.

know about non-psychological, physical reality.²¹ In this way, certain theories of cognitive psychology are in tension with (if not incompatible with) scepticism. But we might think that cognitive psychological theories of human perception are not satisfying philosophical theories of knowledge, and not because they're not correct. Instead, it's because they're not *philosophical*.²² Compare with physical theories of matter. These theories imply that there are composite material objects.²³ Yet, there are also philosophical theories of material objects, some of which say that there aren't any composite material objects.²⁴ Naturally, one might think that those philosophical theories are dissatisfying because they say that ordinary composite material objects like tables and chairs do not exist. Equally naturally, however, one might also think that the physical theory is dissatisfying if one presents it as an answer to the problem of material constitution. Why? One reason is that the physical theory is not aiming to resolve any specifically *philosophical problem* about material constitution, even if it has logical consequences that are incompatible with certain philosophical theories which *do* aim to resolve philosophical problems about material constitution. So, even if there are physical theories of objects that are correct, this doesn't imply that there are satisfying philosophical theories of physical objects. This would take the further step of showing that the correct physical theory of objects *is* a philosophically satisfying theory of physical objects. But this is not obvious.

So too, even if there are correct theories about how human beings perceptually know about non-psychological, physical reality, this doesn't imply that they are satisfactory philosophical explanations of how human beings

²¹ At least, certain mainstream theories in cognitive psychology are committed to the possibility of perceptually knowing about ordinary, macro-physical objects. For example, see Gibson (1950).

²² Cf. Stroud (2000, 99-101). A similar point is made in Nozick (1981, 12-13).

²³ Prima facie anyway, the atomic theory of matter seems to be in tension with mereological nihilism (the thesis that *there are no composite material objects*), even if, in the end, it's not incompatible with mereological nihilism. In fact, van Inwagen (1990, 72), a proponent of a restricted version of mereological nihilism, says that there is a tension between nihilism and, say, the existence of hydrogen atoms.

²⁴ For example, see Merricks (2003) and van Inwagen (1990).

perceptually know about non-psychological physical reality. It would take a further step to show that any correct theory of how human beings know about non-psychological, physical reality, is also a philosophically satisfying explanation of how human beings know about non-psychological, physical reality.²⁵

2.2 Summary

In the previous section ([section 2.1](#)), I presented some arguments in order to make clear what the meta-epistemological is and is not committed to. In particular, I argued for the thesis that the mere correctness of a theory of knowledge, a theory of how our knowledge of the world is possible, doesn't imply that it's philosophically satisfying. I provided two reasons in support of that thesis.

First, not every correct theory is necessarily a *positive* theory. Scepticism is a philosophical theory of knowledge, but it's not philosophically satisfying.²⁶ So, we might think that a *negative*, though *correct* philosophical theory of knowledge is not philosophically satisfying.

²⁵ There are two points I want to flag here in connection with the view that there are *philosophical* theories of how our knowledge of the world is possible, and the view that there are *correct* (though *non-philosophical*) theories of how our knowledge of the world is possible. The first point is that a naturalized epistemologist (or an "extreme" naturalized epistemologist) might want to argue that while there are philosophical theories of how our knowledge of the world is possible, there are nevertheless no *distinctively* philosophical theories of how our knowledge of the world is possible. Why? Because the *correct, philosophical explanation* of how our knowledge of the world is possible simply *is* (or is reduced to) the *correct, scientific explanation* of how our knowledge of the world is possible. This drains the philosophical explanation of its distinctiveness. The second point is that one might want to argue instead that while there are *correct* theories of how our knowledge of the world is possible, these theories are *scientific*, and not philosophical. So, on this view at least, it's open that there are any correct, *philosophical* theories of how our knowledge of the world is possible, even if it's *not* an open question whether there are any correct *scientific* theories of how our knowledge of the world is possible.

²⁶ As Pritchard (2002, 2005, 2012, *forthcoming*), Williams (1996, 2011), and Wright (1991) highlight, there is a distinction between scepticism *qua* paradox, and scepticism *qua* philosophical theory of knowledge. In the main text, I am using "scepticism" to refer to a philosophical theory of knowledge, and not as standing for a paradox. See also DeRose (1995), Sosa (1999), and Vogel (1999), and Williamson (2005) for framing the sceptical problem as a paradox.

Second, not every *positive* and *correct* theory is necessarily *philosophical*. Certain cognitive-psychological theories plausibly entail that human beings can know about non-psychological, physical reality. So, we might think that a positive, though *non-philosophical* theory of knowledge is not philosophically satisfying, even if it's correct.

Now *meta-epistemological scepticism* is the thesis that a *satisfying philosophical explanation* of how our putative knowledge of non-psychological, physical reality is possible *is not possible*. So, a meta-epistemological is *not* committed to the following theses:

- “There are no *philosophical theories* of how our putative knowledge of non-psychological, physical reality is possible”
- “There are no *positive philosophical theories* of how our putative knowledge of non-psychological, physical reality is possible”
- “There are no *correct philosophical theories* of how our putative knowledge of non-psychological, physical reality is possible”

Furthermore, it is not clear that a meta-epistemological sceptic is even committed to the conjunction of the last of those two theses: that there are no positive and correct philosophical theories of how our knowledge of the world is possible. For it might be that there are positive *and* correct philosophical theories of how our knowledge of non-psychological, physical reality is possible, but these two properties aren't sufficient for philosophical satisfaction (e.g. for the theory being a philosophically satisfying theory). That is, the mere positivity and correctness of an account of how our knowledge of the world is possible doesn't make it a philosophically satisfying account of how our knowledge of the world is possible. For it is at least not obvious that there is a straightforward implication from the positivity and correctness of the philosophical explanation that it is or ought to be philosophically satisfying.

To summarize, a meta-epistemological sceptic does not deny that there are

positive, philosophical theories of how our putative knowledge of non-psychological, physical reality is possible. Indeed, it's not even obvious that a meta-epistemological sceptic has to deny that there are *correct* philosophical theories of how our putative knowledge of non-psychological, physical reality is possible. Why? A meta-epistemological sceptic might think that *scepticism* is the correct account of our knowledge of non-psychological, physical reality. And this might be *why* she's a meta-epistemological sceptic. But what's not clear is that scepticism must be *the only reason* for being a meta-epistemological sceptic. In sections 5 and 6, I overview two reasons for being a meta-epistemological sceptic which do not require the truth of scepticism. For now, however, I want to attend to some *prima facie* considerations against meta-epistemological scepticism.

3 Meta-epistemological Scepticism: three *prima facie* challenges

In this chapter, I defend the following thesis: it is possible that we know about non-psychological physical reality, even if a satisfactory philosophical explanation of how we can know about non-psychological, physical reality is not possible. In short, I defend the thesis that meta-epistemological scepticism does not *entail* scepticism. For clarification, we should note that a commitment to this thesis does not entail that scepticism *is false*. Neither is it a commitment *to* meta-epistemological scepticism. Instead, it's just that one can combine meta-epistemological scepticism with the negation of scepticism and have a consistent view.

Notice that if meta-epistemological scepticism does not entail scepticism, that fact might serve to make meta-epistemological scepticism a more interesting thesis. It would be more interesting because its truth would not turn on the truth of scepticism, and so an independent debate over meta-epistemological scepticism can arise. In fact, it would leave room for the meta-

epistemological sceptic to be an anti-sceptic at the first-order level: a first-order anti-sceptic with respect to our knowledge of the world around us.

How does that make it an interesting view? Insofar as we think that certain philosophical puzzles cannot be resolved, it's an interesting view because it maintains that what looks like a core epistemological problem cannot be resolved to our satisfaction. But that it cannot be resolved to our satisfaction, on this view, would not mean that it is *because* scepticism is true. This feature of the view, then, would serve to make it more interesting.

However, we can imagine that an opponent of the view that meta-epistemological scepticism does not entail scepticism might maintain that the impossibility of a satisfying, philosophical explanation of how our putative knowledge of the world is possible raises serious questions about whether knowledge of the world is possible. In this section, I consider this objection and respond to it.

Prima facie, not all of the obstacles to the possibility of a satisfying philosophical explanation of how our knowledge of the world is possible are obstacles to the possibility of knowledge of the world. That is, it is plausible that not everything which might make a satisfying philosophical explanation of how something is possible necessarily makes that very thing look impossible. In general, the inference from “a philosophically satisfying explanation of how *F* is possible is not possible” to “*F* is not possible” just looks fallacious.

In order to foster this intuition, consider the following example from the debate on freedom of the will. One question we might raise here is whether we ought to insist, on *a priori* grounds, that the impossibility of a philosophically satisfying explanation of the possibility of freedom of the will entails that no one is free? Prima facie at least, there answer is “no”. I use the term “prima facie” here to highlight how this should strike us as intuitive. According to this intuition, it seems like there can be considerations which lead one to think that a satisfactory, philosophical explanation of how some phenomena is possible, like freedom of the will, without commitment to the impossibility of the

phenomena.

To return to our example, a philosopher might think that freedom of the will is possible, but maintain that a satisfying, philosophical explanation of how freedom of the will is possible is not possible. In fact, we have an example of some philosophers who maintain this kind of view.²⁷ “Mysterionism” is the name for the thesis that a satisfying, philosophical explanation of how freedom of the will is possible is impossible. Proponents of this view maintain that (i) the argument for the incompatibility of free-will with causal-determinism seems to be valid and sound; that (ii) the argument for the incompatibility of free-will with causal indeterminism seems to be valid and sound; that (iii) the intuitively compelling thought that freedom of the will is actual renders at least one of those arguments either invalid or unsound, but that (iv) it’s a mystery which argument is invalid or unsound, so that it’s a philosophical mystery how freedom of the will is possible.²⁸

The possibility of this kind of meta-metaphysical thesis in the debate over freedom of the will is useful for our purposes. It’s useful because it highlights at least one case where its proponents hold a negative thesis about the possibility of a satisfying philosophical explanation of some phenomena, whilst rejecting the thesis that that phenomena is impossible. On their view, the truth of “mysterionism” does not require that the argument for the incompatibility of free-will with causal determinism be valid and sound, nor does it require that the argument for the incompatibility of free-will with causal indeterminism be valid and sound. Instead, it requires the truth of the thesis that we cannot know which of the arguments is invalid or unsound.

²⁷ For example, see van Inwagen (2000).

²⁸ For a version of this argument, see van Inwagen (2000), and Ekstrom (2003, 153-155) for discussion. A similar sort of thesis has been argued for in the debate over the metaphysics of consciousness. According to proponents of the “cognitive closure” thesis, we simply can’t understand, due to necessary cognitive limitations, the nature of consciousness. We might think that this theory is correct, and philosophical, but nevertheless not a philosophically satisfying theory of the nature of conscious. However, what might be philosophically satisfying is a correct, philosophical explanation of *why* the cognitive closure thesis is true, even if the cognitive closure thesis is itself philosophically unsatisfying. See Flanagan (1991), McGinn (1993), and Nagel (1974).

I now want to consider an objection to this argument. The argument I've presented tried to bolster the intuition that there is no strictly *logical* reason to think that the impossibility of a satisfying philosophical explanation of some phenomena *F* implies that *F* is impossible. But according to the objection, it might be that there are *other* sorts of reasons which make it hard to see how *F* is possible, if a satisfying philosophical explanation of *F* is impossible.

A proponent of this objection might tell us that there have to be some reasons which prevent us from achieving a philosophically satisfying explanation of how knowledge of non-psychological, physical reality is possible. And what better reason could there be then that it is impossible? But notice that this objection can be accommodated. It might be true that the best explanation of why it's impossible to achieve a philosophically satisfying explanation of how knowledge of non-psychological physical reality is possible is that knowledge of non-psychological physical reality is impossible. But that some explanation is the best explanation of *p*'s being true (or false) need not be the actual reason that it is true (or false): it need not be the correct explanation because it's the best explanation. Of course, there is a trivial sense in which the correct explanation of *p*'s truth is "the best explanation" of its truth. It's de facto the best explanation because it's the correct one. None of this contradicts the thesis that meta-epistemological scepticism does not entail scepticism.

Let's consider a second objection. Suppose that an epistemologist has presented us with compelling reasons to think that their positive theory of knowledge is correct (that their positive theory of how knowledge of the world is possible is correct). Some philosophers might think that if a philosophical theory of knowledge is both positive and correct, then it ought to be philosophically satisfying as well. On this view, the positivity and correctness of a philosophical theory of how knowledge of non-psychological, physical reality is possible, is sufficient for the falsity of meta-epistemological scepticism. Call this thesis "Sufficiency":

Sufficiency: If a philosophical theory of non-psychological, physical reality is positive and correct, then a satisfying, philosophical explanation of how knowledge of non-psychological, physical reality is possible is possible.

Now the challenge that the proponent of Sufficiency might raise is: what more should one want for philosophical satisfaction other than the positivity and correctness of their explanation? This presents a compelling *prima facie* challenge to the proponent of the thesis that meta-epistemological scepticism does not entail scepticism, because if a positive theory of knowledge is true (and so scepticism is false), from Sufficiency it follows that that positive theory of knowledge is a philosophically satisfying explanation of how our knowledge of the world is possible.

The problem with Sufficiency, however, is that it ignores other plausible criteria for philosophical satisfaction. For example, one criterion might be that it's not rule-circular. A rule-circular explanation is one that exploits an epistemic principle which is itself part of the things we want to explain how it's possible for anyone to know.²⁹ For example, let's suppose that we wanted to explain how inductive knowledge is possible, and our explanation contained a proposition describing a principle which presupposed that we have inductive knowledge. While it's not obvious that this kind of explanation is dissatisfying, it's not obvious that it's not dissatisfying either. Proponents of Sufficiency would be committed to saying that it's not dissatisfying only if the explanation is positive and correct. On their view, then, rule-circular explanations which are positive and correct can be philosophically satisfying.

²⁹ Fumerton (1995) appeals to this kind of principle in order to rule out, not whether a theory or explanation of how knowledge of the world is possible is *philosophical*, but whether it is *philosophically satisfying*. This point is brought out more sharply in Fumerton (2006), where he argues that reliabilists, for example, can use their perceptual processes in order to non-inferentially justify the reliability of their perceptual processes, but that this type of justification won't be "philosophically satisfying noninferential justification" (Fumerton 2006, 190). See Pritchard and Ranalli (*forthcoming*) for a discussion of Fumerton and Stroud's disagreement about what's required for a philosophically satisfying explanation of knowledge.

We can generalize this case: if positivity and correctness are sufficient, then positivity and correctness plus any condition compatible with them is sufficient as well. (This is just a case of strengthening the antecedent with any proposition compatible with it). So, that it's rule-circular, positive, and correct is sufficient; that it's rule-circular, positive, correct, and not believed by anyone is sufficient; that it's rule-circular, positive, correct, not believed by anyone, and incompatible with otherwise plausible principles of knowledge is sufficient. In short, the proponent of Sufficiency is committed to positivity and correctness trumping any other possible condition that is compatible with them. But it's an interesting question whether *there are* any other possible conditions.

3.1 Meta-epistemological scepticism: Conclusion

In the previous section (section 3), I tried to answer three *prima facie* objections to meta-epistemological scepticism. The first objection raised the suspicion that meta-epistemological scepticism could not be divorced from scepticism. On this view, the truth of meta-epistemological scepticism raises serious questions about the possibility of knowledge of the world; so much so that we might wonder whether knowledge of the world is possible after all. Against this, I argued that there is no strictly logical reason to think that meta-epistemological could not be divorced from scepticism.

The second objection put a sharper point on the first objection. It said that the best explanation of the truth of meta-epistemological scepticism is scepticism, so that it is hard to see why one would be a meta-epistemological sceptic and not a first-order sceptic. Against this objection, I suggested that while scepticism might be one reason for being a meta-epistemological sceptic—*perhaps even the best explanation*—it's just not the only possible reason. There can be other reasons, reasons which are, as I will show below, *prima facie* plausible, even if incorrect.

The third objection presented a challenge to the proponent of the thesis that meta-epistemological scepticism is compatible with the denial of scepticism.

According to this challenge, the positivity and correctness of a theory of how our knowledge of non-psychological, physical reality is possible should be sufficient for philosophical satisfaction. That is, it ought to be a philosophically satisfying explanation. This is what I called “Sufficiency”. So, the worry here is how it could be that one is not a sceptic, but nevertheless a meta-epistemological sceptic. I argued that Sufficiency ignores other possible criteria for philosophical satisfaction, so that it’s less puzzling why one would deny scepticism but accept meta-epistemological scepticism. On this view, it’s an open question whether the positivity and correctness of the theory or explanation of how our knowledge of the world is possible is sufficient for a philosophically satisfying explanation, since one can very easily conceive of additional criteria over and above positivity and correctness that are plausibly required for philosophical satisfaction as well. In the next section, I show how Stroud (1984, 2000) presents additional criteria.

4 Scepticism and Meta-epistemological scepticism

Stroud’s (1984) meta-epistemological argument is structured around “the problem of the external world”, and the attendant sceptical argument. So, the first question we need to address before we can look at Stroud’s argument is this: what is “the problem of the external world?” What’s problematic is just which problem the phrase “the problem of the external world” picks out, since different philosophers use that phrase to talk about what seem, on the surface at least, to be different problems.³⁰ For example, most philosophers understand

³⁰ For example, Williams (2001, 77) says that “the problem of the external world” is an *underdetermination problem*. See also Yalcin (1992) for framing the problem this way. However, other philosophers, such as Pritchard (*forthcoming*, 2005) think there are at least two separate sceptical problems, one arising out of the closure principle, and one arising out of underdetermination worries. Still, some other philosophers, such as McDowell (2009, 2011), think that it’s a problem arising out of a “veil of perception”, and that the problem of the external world thereby turns on certain principles about the nature of perception. The issue surrounding the ‘sources’ of scepticism is discussed in Pritchard’s (*forthcoming*) “The

the problem along the following lines:

- There are apparently compelling arguments for the thesis that no one can know anything about non-psychological, physical reality. And “the problem of the external world”, the problem of explaining how any knowledge of non-psychological, physical reality is possible, is simply the problem of identifying and explaining which premises of those arguments are false and explaining why they are false (or more weakly: which premises of those arguments fail to adequately support the sceptical conclusion and why they fail to adequately support the sceptical conclusion).

Let’s suppose that this is correct. What we need to know now is what the apparently compelling arguments are. The most prominent way of expressing the argument is to frame it around an “argument from ignorance”. I will use the argument from ignorance because typically this is used as a template for the sceptical argument, where the mechanics of the problematic can be brought out more sharply in other ways (e.g. as an underdetermination problem; as a problem which centres on a closure principle and other epistemic principles; or as a problem which centres on principles about perceptual experience). In what follows, I’ll explain what the argument from ignorance is, and how Stroud’s (1984) meta-epistemological argument makes contact with the argument from ignorance.

4.1 The Argument from Ignorance

Let “SH” pick out what epistemologists call “sceptical hypotheses”. A “sceptical hypothesis” is simply an alternative ‘sceptical’ explanation of the causes of our sensory experiences and beliefs about non-psychological,

Sources of Scepticism”, *International Journal for the Study of Skepticism*, special issue on ‘Hinge Epistemology: Basic Beliefs After Moore and Wittgenstein’, (eds.) A. Coliva & D. Moyal-Sharrock. Neta (*forthcoming*) provides a recent discussion of the different ways the problem of the external world can arise.

physical reality, in which the causes are either systematically not what we believe them to be, or else, even if they are what we believe them to be, the causes are systematically unreliable, so that (intuitively) we're never in a position to know that any of our beliefs about non-psychological, physical reality are true.

Examples include the hypothesis that there is an evil-demon who has caused me (and continues to cause me) to have all of the sensory experiences that I have ever had, and has caused me to have all of the beliefs that I have about non-psychological, physical reality. However, almost everything that I believe about non-psychological, physical reality is false. An updated version of this hypothesis is that I am a bodiless brain in a vat, plugged into a simulation which causes me to have the sensory experiences that I have and the beliefs that I have, even though my beliefs are systematically false. Finally, there is the dreaming hypothesis, in which I have been kept in a permanent dream. It need not be the case that my beliefs about non-psychological, physical reality are systematically false, because it's compatible with dreaming that I have a body, that I have hands, that there are trees, and so forth. But the dreaming hypothesis highlights how, even if my beliefs weren't systematically false, my sensory experiences would nevertheless in no way put me in a position to know that they're not.³¹

What all of the sceptical hypotheses have in common, however, is that they vividly highlight the possibility of being systematically wrong, or systematically ignorant, about how non-psychological, physical reality is, including the proposition that there even is a non-psychological, physical reality.³² More broadly, they all highlight how it is logically compatible with

³¹ Cf. Williams (2000, 70).

³² Of course, the dreaming hypothesis doesn't entail that most of our beliefs about non-psychological, physical reality are false, because dreaming that p is compatible with p . But the hypothesis is still put into the service of showing that it's possible that everything we believe about non-psychological, physical reality is false. Our beliefs *might* be true, but possibly, none of them are. The problem then is how we know that they're not false. In any case, it's not clear that the actual falsity of belief matters here. Why? Consider the brain-in-

our having the sensory experiences and beliefs that we have that we are nevertheless not in a position to know that any of those beliefs are true or that any of those experiences are veridical.

So, now that we have an initial understanding of what the sceptical hypotheses are, we should ask whether we can know that they're not actual. In short, can we know that what SH says is true is not true?

On the one hand, it should strike us as obvious that we can. After all, I know that I have a body, and I know that my hands are in front of me. But then I need to remember what SH says: it's compatible with my having the sensory experiences as of having hands, as of having a body, and so on, that I'm merely a bodiless brain-in-a-vat, who falsely believes that I have hands, that I have a body, and so on. I naively suppose that I'm appealing to the fact that I have a body, the fact that I have hands, and so on, in order to know that SH is not the case. But the possibility of SH being true undermines that view. Why? Consider the following thought from Williams:

Sceptical hypotheses seem to show that there are endlessly many ways that the world might be, even though my experience of it remains unchanged. Accordingly, my experience fails to provide an adequate basis for favouring my actual system of beliefs over alternatives that seem logically just as coherent. But when it comes to knowing about the world, my experience is all that I will ever have to go on. I have no magical faculty for intuiting how things are in my surroundings. My only basis for my beliefs about the world, however unshakeable, are oddly groundless: mere beliefs rather than genuine knowledge (Williams 2000, 71).

This is supposed to motivate the premise of the argument from ignorance that no one can know that the sceptical hypotheses do not obtain.^{33, 34} We can

a-vat (BIV) hypothesis. It typically stipulated that being a BIV is incompatible with having hands. But as Roush (2010) points out, BIV's can have hands. And if we think that it's plausible that we don't know that we're not BIV's (and reason, by closure, that we don't know that we have hands), is it any less plausible that we don't know that we're not *handed* BIV's (and so on for many propositions about the world?). See Sosa (2008) for the worry that dreaming scepticism threatens perceptual knowledge by threatening the safety of our perceptually based beliefs about the world.

³³ See Williams (2001, 73), DeRose (1995, 1), and Pritchard (2002, 217) for the formulation of the argument from ignorance. Cf. Williamson (2005).

express this premise as follows:

- For any SH, we cannot know that \sim SH.

At this juncture, we might wonder whether we *have* to know \sim SH in order to know anything about non-psychological, physical reality. After all, that we can't know \sim SH might be surprising, but it doesn't threaten the possibility of knowledge of non-psychological, physical reality unless the following principle is true:

- For any proposition about non-psychological physical reality P , S knows that P only if S knows that \sim SH.

This principle says that it's a necessary condition of anyone's knowing that P that they know \sim SH. If that principle is false, then the premise that it's not possible to know \sim SH doesn't provide adequate support for scepticism.

However, consider now the fact that deduction is a means by which we can extend our knowledge to the known consequences of what we know. For example, if I know that 2 is a prime number, and I know that if 2 is a prime

³⁴ As I explained in the beginning of this section, the argument from ignorance can be presented as a template sceptical argument, where the mechanics of the sceptical argument can be brought out more sharply in other ways. For example, consider the premise that I don't know that \sim SH. Some epistemologists (e.g., Brueckner 1994, Williamson 2000, McDowell 1995, 1998) think that this premise is best supported with an evidential underdetermination argument of the following sort: I have exactly the same evidence in the good case, the case in which \sim SH, as I do in the bad case, the cases in which SH. So, my evidence doesn't favour \sim SH over SH. So, I don't know that \sim SH. Still, some epistemologists (e.g., Nozick 1981, DeRose 1995) think that this premise is best supported by showing that it fails to satisfy a sensitivity condition on knowledge, the condition that S knows that p only if, if p were false, S would not believe that p . Unger, for example, argued that S knows that \sim SH only if S could be absolutely certain that \sim SH, but that S could not be absolutely certain that \sim SH. Another formulation of the argument appeals first to an argument from illusion, or an argument from hallucination, and then tries to show that anything we know about non-psychological, physical reality is by inference from what we know about sensory appearances plus principles which link facts about our sensory experiences with facts about non-psychological reality. The second step is that there is no epistemologically adequate inference that we can make: either it is a circular inference or else it fails to bridge the gap between sensory appearance and non-psychological reality. The conclusion is that our knowledge is therefore strictly confined to our sensory appearances (and anything else which is logically consistent with the non-existence of non-psychological, physical reality). See Fumerton (1995) for this argument.

number then there is at least one even prime number, then I can deduce from these two propositions that there is at least one even prime number. I extended my knowledge to the known consequences of what I know by deduction. We can formulate this idea into the following deductive “closure” principle, which says that knowledge can be gained by competent deduction:

Deductive Closure: If S knows that p and S knows that (if p then q), and S deduces q from p and (if p then q), coming to believe that p and retaining their knowledge that p throughout, then S comes to know that q .

The Deductive Closure principle (or ‘closure principle’ for short) is not an epistemically demanding principle. It simply says that whenever S knows that p , and S knows that p implies q , S is thereby in a position to know, by deduction, that q .

With the closure principle in play, we can straightforwardly support the conditional premise of the argument from ignorance. First, suppose that we know some proposition about non-psychological physical reality P . Second, we know that P implies that $\sim SH$, given a suitable P and a suitable SH (for example, *that I have hands* and *that I’m not a handless BIV*). So, from the closure principle, the conditional premise that S knows that P only if S knows that $\sim SH$ follows.³⁵ But from modus tollens on those two premises, the sceptical conclusion follows:

³⁵ Stroud (1984, 2009) appeals to Descartes’s dreaming hypothesis as the target sceptical hypothesis. As the reader will notice, however, it’s hard to see how the closure principle could provide adequate support for the first premise of the sceptical argument (above), since dreaming that p is not incompatible with p . Instead, a stronger closure principle would need to be used (cf. Stroud 1984, 25). This caveat is not important for our current purposes, however, because Stroud’s argument that Descartes’s Condition (the principle that a necessary condition of knowing that p is knowing that I am not merely dreaming that p) implies that scepticism is true works just as well for the evil demon hypothesis. *Mutatis mutandis*, Stroud could replace the dreaming hypothesis with the evil demon hypothesis. In fact, Stroud (1977, 2000) appeals to the evil demon hypothesis in his argument for the thesis that Descartes’s Condition implies scepticism. See in particular Stroud (2000, 47-48). In Chapter 3, I provide a detailed exposition and evaluation of Stroud’s argument for this thesis.

- For any proposition about non-psychological, physical reality, we cannot know that P .

We can thereby represent the sceptical argument from ignorance as follows:

(S1) If S knows that P , then S knows that \sim SH.

(S2) S does not know that \sim SH.

Therefore,

(S3) S does not know that P .

Since our concern is not with first-order responses to the argument, responses which target the soundness or validity of this argument, I will not discuss those responses. What I want to highlight in the next section (section 5) is how Stroud's (1984) puzzle can be framed around the argument from ignorance.

5 Stroud's Puzzle

5.1 Stroud's Puzzle: the basic dilemma

I began the introduction with a brief statement of Stroud's (1984) argument for meta-epistemological scepticism. I said that it provides the epistemologist with a puzzle. Stroud presents the puzzle, perhaps in its most general form, as follows:

We find ourselves with questions about knowledge that lead either to an unsatisfactory sceptical conclusion or to this or that 'theory' of knowledge which on reflection turns out to offer no more genuine satisfaction than the original sceptical conclusion it was meant to avoid. After several disappointments of this kind we can come to wonder whether there could ever be a general explanation of human knowledge that remained sufficiently non-sceptical [...] to satisfy us (Stroud 1984, 168).

The puzzle arises out of different sorts of desiderata for a satisfying, philosophical explanation of how knowledge of the world is possible. As Stroud tells us in the quotation, one intuitive desideratum is that the explanation

be a positive explanation: it must explain how we have the knowledge that we think we've got. Stroud expresses this desideratum as follows:

A satisfactory 'theory' or explanation of [knowledge of the world] must depict us as knowing all or most of the things we think we know. It must explain, given what it takes to be the facts of human perception, how we none the less know the sorts of things we think we know (Stroud 2000, 145).

This, of course, simply expresses the need for the explanation to be anti-sceptical. So the desideratum that the explanation be a positive explanation is the desideratum that it be an anti-sceptical explanation. But we might wonder whether there is anything else that is plausibly a desideratum for a philosophically satisfying explanation of how our knowledge of the world possible.

For Stroud, another desideratum that is that the explanation not be conceptually revisionary: it must not make substantive changes to our ordinary concept of knowledge. The problem, however, is that Stroud thinks that the epistemologist is faced with two unpalatable options: scepticism or revisionism.³⁶ As Williams's (1996) expresses the point, Stroud's (1984, 2009) arguments present the epistemologist who seeks to explain how our knowledge of the world is possible with a dilemma:

We can either accept scepticism, or make changes in our pre-theoretical thinking about knowledge that shrink the domain, or alter the status, of what we previously thought of as knowledge of objective fact (Williams 1996, 22).

Now let's take the first horn of the dilemma. Scepticism is itself an extreme form of revisionism, because we all believe that we know (or can know) about

³⁶ Cf. Fogelin: "scepticism (in its most challenging forms) is not based upon such prior theoretical commitments, but rather is the natural outcome of unrestricted exploitation of a feature already present in our everyday concept of knowledge" (Fogelin 1999, 159). See also Heil (1986), who says that: "The skeptic, despite appearances, does not deny common opinion. On the contrary, he abandons one platitude for the sake of others. He appeals not to arcane considerations, hidden details, but to familiar intuitions. These lead him perhaps to reject something important but, from his point of view, to do so only for the sake of the larger edifice" (Heil 1986, 1). For a discussion, see Leite (2010). Leite argues that there is no sense in which the sceptic, using the argument from ignorance, could be appealing to ordinary epistemic principles or common sense intuitions. See also Leite (2006) for the relationship between epistemic theorising and ordinary epistemic practice.

the world around us. I believe that I know that a computer is front of me; I believe that other people know that there are trees, houses, other people, books, and so on. I believe that I know about the world around me, that other people know about the world around us, and that those people share this belief about others. The number of beliefs we have about what we know about the world around us is perhaps as large as our beliefs about the world around us. Whatever its number, the point here is just that we have a substantive amount of higher-order, epistemological beliefs: beliefs about what we know (and don't know) about the world around us; beliefs about what we can come to know (and perhaps can't come to know) about the world us; and beliefs about how we can (and perhaps how we cannot) come to know about the world around us.

Scepticism implies that this belief is false (and perforce that most of the higher-order, epistemological beliefs are false as well). Of course, the thesis that revisionary explanations of how knowledge of the world is possible are philosophically unsatisfying is not so shallow a thesis that *no* form of error (or no degree of error) is permissible. We can be wrong about our beliefs about what we know, what we can come to know, and how we know come to know them, without philosophical dissatisfaction. This is just part and parcel with epistemic improvement: we seek to update our beliefs to fit with the facts (and therefore update our higher-order beliefs about what we know, etc., to fit with the facts).

The proponent of thesis that revisionary explanations of how our knowledge of the world is possible are philosophically unsatisfying explanations might think that scepticism is unpalatable here because scepticism revises not just one belief, but a substantive number of beliefs—an entire class of belief—in a domain that we care about, where getting things right rather than wrong matters to us (cf. Zagzebski 2009, 2004).³⁷

³⁷ I am presenting this as an *additional* reason to the *intuitively compelling* thought that scepticism is philosophically unsatisfying. At least one reason for motivating other reasons, besides its sheer intuitiveness, is that a sceptic might maintain that scepticism is *prima facie*

It's plausible that scepticism, then, in conjunction with the thesis that there is an entire class of higher-order epistemic beliefs for which substantive error seems non-negotiable, implies that if scepticism is true, then meta-epistemological scepticism is true. The argument in support of this conclusion is that scepticism is a form of revisionism because (i) it implies that an entire class of higher-order epistemic beliefs is in substantive error, but (ii) that this class of belief is not in substantive error strikes us as non-negotiable.

If it is intuitive that scepticism provides adequate support for meta-epistemological scepticism, we might wonder how intuitive the other horn of the dilemma is. Roughly, this horn of the dilemma is that revisionism about our ordinary conception of knowledge is philosophically dissatisfying. For Stroud, then, scepticism is one way to arrive at a philosophically unsatisfying explanation of how our knowledge of the world is possible, just not the only way. Another way is to substantially revise our ordinary conception of knowledge. In the next section (section [5.2](#)), I provide an outline of Stroud's argument, and then, in section [5.3](#), I explain how Stroud intends to adequately support the premises of this argument.

5.2 Stroud's Puzzle: the argument

Stroud's puzzle arises as follows. First, if Descartes's Condition is true, then scepticism is true.³⁸ As Stroud explains, "[i]f it is in general a necessary condition of our knowing something about the world around us that we know we are not dreaming, it follows that we can never know that we are not dreaming" (Stroud 1984, 43).³⁹ Second, if Descartes condition is false, then

philosophically unsatisfying, but not philosophically unsatisfying on sufficient reflection. I discuss this objection from proponents of scepticism in Chapter 2.

³⁸ Stroud registers this throughout, and refers to it as a kind of "conditional scepticism". See his (1984), pp. 19, 20, 23, 24, 30, 43, 48.

³⁹ This argument first appears in his (1977), reprinted in his (2000). The argument is developed in his (1984), and later updated in his (2009). For explicit statements of the thesis that Descartes's Condition implies that scepticism is true, see Stroud (1984), pp. 20, 21, 23, 48, and 54 (see also the abstract to his (1984) "The Problem of the External World", on *Oxford Scholarship Online*, 2003).

some “obvious truths” or “platitudes” are false.⁴⁰ Third, neither scepticism nor an explanation which implies that some “obvious truths” or “platitudes” are false is a philosophically satisfying explanation of how our knowledge of the world is possible.⁴¹ But since either Descartes’s Condition is true or it is false, it follows that a philosophically satisfying explanation of how our knowledge of the world is possible is not possible. That is, meta-epistemological scepticism is true.

A more formal and general presentation of the puzzle can be rendered as follows. First, suppose that the first premise of the sceptical argument is true. This premise tells us that, for any proposition about non-psychological, physical reality *P*, if S knows that *P*, then S knows that it’s not the case that they are not merely deceived about *P* (whether it is dreaming that *P*, hallucinating that *P*, fooled into falsely believing that *P*, or, as a counter-part to the perception of brains-in-vats, BIV-perceiving that *P*).

The first horn of Stroud’s puzzle is that if this first premise of the sceptical argument is true, then the second premise of the sceptical argument is true as well. This premise is that S doesn’t know that that they are not deceived about *P* (etc.). So the first horn of Stroud’s puzzle is that:

First horn of Stroud’s puzzle: If the first premise of the sceptical argument, (S1), is true, then the second premise of the sceptical argument, (S2), is true.

The first horn of Stroud’s puzzle leads to scepticism. After all, from the premise that, if S knows that *P*, then S knows that it’s not the case that they are not deceived about *P*, and the premise that S does not know that they are not deceived about *P*, it follows, from modus tollens, that S does not know that *P*. Now Stroud’s claim, recall, is that the first premise of the sceptical argument, the conditional premise, entails the second premise of the sceptical argument.

⁴⁰ See, for example, Stroud (1984), pp. 30-31, 76-77, 82.

⁴¹ See, for example, Stroud (1984), pp. ix, 81-82, 168-169,

Since the conjunction of the first premise and the second premise of the sceptical argument entail scepticism, it follows that, if Stroud is right that the first premise entails the second premise, then the first premise entails scepticism as well.⁴²

How does Stroud argue for the thesis that the first premise entails the second premise? In particular, how does Stroud argue for the thesis that Descartes's Condition implies that it cannot be satisfied? In brief, Stroud tells us that:

if Descartes is right to insist that in order to know something about the world around him he must know that he is not dreaming, then he is also right that he has no such knowledge, because the condition for knowledge that Descartes accepts can never be fulfilled: fulfilling it would require knowledge which itself would be possible only if the condition were fulfilled (Stroud 2003, 1).

This is just one of central argument that Stroud appeals to in order to provide adequate support for the thesis that Descartes's Condition implies that it cannot be satisfied. In Chapter 3, I consider the arguments in detail. For now, we just want to get an initial grasp of Stroud's argument in favour of this thesis.

A short version of Stroud's argument can be put like this. First, Stroud maintains that Descartes's Condition applies to every proposition about non-psychological physical reality. But if it applies to every proposition about non-psychological physical reality, then it also applies to the proposition that I am not merely dreaming that p, since this too is a proposition about non-psychological physical reality. So, applying Descartes's Condition to that proposition gives us the following principle: if I know that I am not merely dreaming that p, then I know that I am not merely dreaming that I am not

⁴² Here's one way of presenting the argument. First, consider the following two premises of the first-order sceptical argument (see section 4.1):

(S1) If S knows that *P*, S knows that $\sim SH$.

(S2) S does not know that $\sim SH$.

The first horn of Stroud's puzzle is that:

If (S1) is true, then (S2) is true.

Now, from *modus tollens*, (S1) and (S2) imply scepticism. So, if Stroud is right that (S1) implies (S2), then we can simplify the first horn of Stroud's puzzle as follows:

If (S1) is true, then scepticism is true.

dreaming that p. And so on, ad infinitum.

Now we might wonder whether this kind of regress argument applies to the sceptical hypotheses in general. Stroud seems to think so. That is, it is not some idiosyncratic fact about dreaming that renders the regress argument plausible for Stroud. For example, consider Stroud's (1977, 2000) version of the argument:

If, for everything I claim to know (or everything about an 'external world'), it is true that in order to know it I must at least know that there is no evil scientist artificially inducing in me the beliefs on the basis of which I claim to know it, then, in particular, in order to know that no evil scientist is doing that to me, I must know that there is no evil scientist artificially inducing in me the beliefs that lead me to believe that no evil scientist is doing that to me. And then there is nowhere to stop (Stroud 2000, 47).

So, Stroud seems to think that the regress argument for the thesis that Descartes's Condition implies that it cannot be satisfied (and therefore that scepticism is true) can be applied in general as follows: the first premise of the sceptical argument implies that the second premise of the sceptical argument is true (that is: if, for every P, a requirement of knowing that P is knowing that \sim SH, then I cannot know \sim SH).

The second horn of Stroud's puzzle results from the rejection of the first premise of the sceptical argument. After all, if Stroud is right that the first premise entails the second premise, and therefore scepticism, one might think that a good way to respond is to reject the first premise. But on Stroud's view this is philosophically unsatisfying. Why? His general reason is that it is revisionary: it will require, or commit one to, rejecting the "merest platitudes", and "what seem to be obvious truths" (Stroud 1984, 76). So, the second horn of Stroud's puzzle can be rendered as follows:

Second horn of Stroud's puzzle: If the first premise of the sceptical argument, (S1), is false, then an "obvious truth", "mere platitude", or a "fact" of our ordinary conception of knowledge" is false (Stroud *ibid*).

Now the final thought is that neither horn has a satisfying consequence. For if, as the first horn contends, the first premise of the sceptical argument entails the second premise, then scepticism is true. But intuitively scepticism is not a satisfying philosophical explanation of how our knowledge of the world is possible. On the other hand, if, as the second horn contends, the first premise is false, then a fact about our ordinary conception of knowledge is false. And, intuitively, this consequence is not philosophically satisfying either. For it allows that a positive, philosophical explanation of the possibility of our knowledge of the world is possible after all, but it would revise our ordinary conception of knowledge.⁴³

5.3 Stroud’s Puzzle: “Obvious Truths” and “Platitudes”

Stroud’s core (1984) argument is a meta-epistemological argument. It begins after one has constructed a first-order argument for scepticism. The kind of argument that Stroud focuses on is the kind of argument we reviewed in section 4. On this kind of reasoning, Stroud tells us that:

The idea of ourselves and of our relation to the world that lies behind the sceptical reasoning seems to me deeply powerful and not easily abandoned. As long as it is even an intelligible way of thinking the sceptical conclusion will be defensible from attack. In trying to give expression to the idea it is natural to resort to what seems like the merest platitudes. If that is so, trying to avoid scepticism by throwing over the old conception will not be easy—it will involve denying what seem to be obvious truths (Stroud 1984, 76).

Here Stroud highlights how the sceptical reasoning presents us with a dilemma: either we accept scepticism or else we reject some ‘obvious truths’. In other words, either scepticism is true or else some ‘obvious truths’ are false, but neither consequence is philosophically satisfying. So, what are the so-called ‘obvious truths’?

The explanation that Stroud provides is complex. I address the

⁴³ Compare with Pritchard (2012) and Williams (1996).

complexities in the next chapter (Chapter 2). For now, I want to side step this large question and just register the propositions he claims are (or are grounded in) the ‘obvious truths’ or the “platitudes we would all accept” (Stroud 1984, 82). This will help us to appreciate the kind of meta-epistemological puzzle he thinks presents an insurmountable obstacle to the epistemologist engaged in the project of explaining how knowledge of the world is possible.

First, Stroud thinks that there is a general principle of knowledge which provides support for what we will “Descartes’s Condition” on our knowledge of the non-psychological, physical reality. This condition says, roughly, that a necessary condition of knowing any proposition about non-psychological, physical reality is knowing that one is not merely dreaming that that proposition is true. More formally, we can express Descartes’s Condition as follows:

Descartes’s Condition: For any proposition P about non-psychological, physical reality, if S knows that P then S knows that they are not merely dreaming that P .⁴⁴

Now as Stroud highlights, it’s initially hard to see how this principle expresses or encodes an ordinary, uncontroversial principle of knowledge. For instance, he tells us that:

[...] it is obvious that we do not always insist that people know they are not dreaming before we allow that they knowing in everyday life,

⁴⁴ As formulated, Descartes’s Condition is structurally the same as the first premise of the sceptical argument: for any proposition about non-psychological, physical reality P , S knows that P only if S knows that $\sim SH$ (in this case, *not merely dreaming that P*). Many epistemologists think, however, that Descartes’s Condition is less plausible because the sceptical hypothesis it uses is *dreaming*, which is compatible with the truth P , rather than *being a BIV* or *being the subject of an evil demon*, which are stipulated as being incompatible with P . As I argue in Chapter 3, these details don’t matter to Stroud’s argument. See also section 5.2 for this point. Of course, as I noted earlier in this chapter, one could not provide support for Descartes’s Condition by way of the closure principle. One would instead need to use a stronger closure principle. However, some philosophers are sceptical that any closure principle is necessary for the sceptical arguments. For example, see Leite (2004) and Pryor (2000).

or even in science or a court of law, where the standards are presumably stricter. So it can easily look as if Descartes reaches his sceptical conclusion only by violating our ordinary standards and requirements for knowledge (Stroud 1984, 39-40).

On the other hand, Stroud also thinks that Descartes's Condition is supported by a general epistemic principle which seems to be part of our ordinary epistemic practices:

Reflecting even on the uncontroversial everyday examples alone can easily lead us to suppose that it is something like this: if someone knows something, p , he must know the falsity of all those things incompatible with his knowing that p (or perhaps all those things he knows to be incompatible with his knowing that p) (Stroud 1984, 29-30).

From this passage we can see that the kind of principle Stroud has in a mind is a "closure" principle:

Strong Closure Principle: If S knows that p , and S knows that their knowing that p implies q , then S knows that q .

This is what leads Stroud to maintain that Descartes's Condition is therefore a "fact" of our ordinary conception of knowledge" (Stroud 1984, 31). It is a fact about our ordinary conception knowledge because it is an instance of an epistemic principle which is true of our ordinary conception of knowledge.

So, at least one question Stroud needs to answer is how Descartes's Condition, a principle which looks like a controversial, epistemic principle, is related to our ordinary epistemic practices. As a first pass, let's consider, for example, how Pritchard (2014) explains the relationship between Descartes's Condition and our ordinary epistemic practices:

[Stroud] argues that the sceptic's system of epistemic evaluation is licensed by our ordinary system of epistemic evaluation on account of the fact that the former is simply a 'purified' version of the latter. That is, if we employ our ordinary practices of epistemic evaluation with due diligence and set aside all purely practical limitations, then what we end up with is the system of epistemic evaluation employed by the sceptic, one that requires [...] that agents must be able to rule-out radical sceptical hypotheses if they are to have the everyday

knowledge that they standardly attribute to themselves (Pritchard 2014, 217).⁴⁵

So, according to Pritchard, Stroud's guiding thought is that Descartes's Condition is a "purified" version of more colloquial epistemic principles, principles which we would insist on being satisfied in everyday life. For example, we can plausibly imagine insisting on the satisfaction of the following sort of principles in everyday life:

- If I know that the bird in front of me is goldfinch, then I know that it is not a canary.⁴⁶
- If I know that the fruit in front of me is in an apple, then I know that it is not a grape.

Now as Pritchard has explained, if we lift our practical limitations, the dreaming-scenario will become a relevant alternative to our everyday knowledge claims. Stroud's point is that, from the mere fact that we don't insist on satisfying Descartes's Condition in everyday life (or even raise the possibility that we are dreaming as a challenge to our everyday knowledge attributions) does not imply that knowing that you're not dreaming is not a necessary condition of knowing anything about non-psychological, physical reality. More specifically, Stroud's point is that:

- The facts about our ordinary epistemic practices, about our use of the term "knows" and the concept *S knows that p*, are compatible with Descartes's Condition being true of our ordinary conception of knowledge of non-psychological, physical reality.

Now that we have an initial understanding of what the so-called 'obvious

⁴⁵ Cf. Pritchard and Ranalli (2013).

⁴⁶ Austin (1946) appeals to this kind of example. However, unlike Stroud (1984), he thinks that knowing that one is not merely dreaming is not a necessary condition of knowledge. Stroud replies to Austin and relevant alternatives theories in his (1984), Chapter 2. For an excellent discussion of Austin's response to scepticism, see Kaplan (2008).

truths' are which Stroud says must be rejected in order to avoid accepting scepticism, we will be in a better position to appreciate his meta-epistemological argument. For ease of exposition, I will formulate Stroud's puzzle around the sceptical argument that I reviewed in section 4:

Stroud's puzzle

1. If (S1) is true, then a satisfactory philosophical explanation of how our knowledge of non-psychological, physical reality is possible is not possible.
2. If (S1) is false, then a satisfactory philosophical explanation of how our knowledge of non-psychological, physical reality is possible is not possible.

Stroud's Puzzle can be expressed in the form of a dilemma arising out of premises 1 and 2. The first horn supposes that (S1) is true: it's true that a necessary condition of knowing any proposition P about non-psychological, physical reality that we know $\sim SH$. For Stroud, this premise is Descartes's Condition: a necessary condition of knowing P is knowing that we are not merely dreaming that P .

Now we suppose that (S1) is either true or it is false. From this supposition and premise 1, it follows by disjunctive syllogism that meta-epistemological scepticism is true (it follows that a satisfying philosophical explanation of how our knowledge of the world is possible is not possible).

The second horn supposes that (S1) is false: it is false that a necessary condition of knowing any proposition P about non-psychological, physical reality that we know $\sim SH$. From the second premise, then, together with the supposition that either (S1) is true or it is false, it follows by disjunctive syllogism that meta-epistemological scepticism is true.

Since we can construct a valid argument from premises 1 and 2 for the conclusion that meta-epistemological scepticism is true, a resolution to Stroud's

puzzle would thereby be one that shows which of premises 1 or 2 is false, and to explain why they're false (or more weakly: to identify which of the premises is not adequately supported and to explain why they're not adequately supported).

Thus far, I have explained how one can formulate a valid argument for meta-epistemological scepticism in a way which doesn't presuppose that scepticism is true. What I haven't done is argue that that argument is sound. And I haven't yet explained in detail why Stroud thinks that the premises 1 and 2 are true. Instead, I have tried only to motivate, if only incompletely, the reasons that Stroud thinks adequately support premises 1 and 2, so that we can appreciate the puzzle that he raises.

In Chapter 2, I explain why Stroud thinks that premise 2 is true—that is, why the falsity of the principle that a necessary condition of knowing any proposition about non-psychological physical reality *P* is knowing that you are not merely dreaming that *P* is philosophically unsatisfying. I begin with the second premise because it is tantamount to one of the arguments of Chapter 3 that Descartes's Condition can be rejected without philosophical dissatisfaction.

In Chapter 3, I explain why Stroud thinks that premise 1 is true—that is, why if a necessary condition of knowing any proposition about non-psychological, physical reality *P* is knowing that you are not merely dreaming that *P*, then you cannot know that you are not merely dreaming that *P*. And I argue in Chapter 2 that his argument in support of the conclusion that premise 1 is true is not a good argument. Moreover, in Chapter 3, I argue that his argument in support of the conclusion that premise 2 is true is not a good argument.

What I have argued for thus far, then, is the thesis that:

- It is possible to construct an interesting valid argument for the conclusion that meta-epistemological scepticism is true without presupposing that scepticism is true.

And the argument that I presented in favour of that thesis is that:

- There is an interesting valid argument from the premises of Stroud's puzzle for the conclusion that meta-epistemological scepticism is true.
- None of the premises of Stroud's puzzle requires the truth of scepticism. Rather, one of the premises requires the possibility of scepticism being true, not that it is actually true.

In the next two chapters, then, I argue that Stroud's puzzle fails to provide adequate support for meta-epistemological scepticism. In the next section, I will explain Stroud's second meta-epistemological argument.

6 Stroud's Dilemma

In the previous section (section 5) I showed how Stroud (1984) argues for meta-epistemological scepticism. I suggested that Stroud's argument is best understood as a certain kind of meta-epistemological puzzle that presents the epistemologist with a dilemma. In this section, I show how Stroud (2000) argues for meta-epistemological scepticism. As we will see, this argument is not identical to his (1984) argument. What Stroud presents the epistemologist with, then, are two 'master' arguments for meta-epistemological scepticism, neither of which require the truth of scepticism in order to support meta-epistemological scepticism.

In his (2000) book, Stroud presents another dilemma for the epistemologist. The dilemma arises out of the constraints on fulfilling the "traditional epistemological project". This is the project of trying to achieve a satisfying, philosophical understanding of how knowledge in some problematic domain is possible (e.g. the external world, other minds, the future, the past, the a priori). As Stroud puts it:

What we seek in the philosophical theory of knowledge is an account that is completely general in several respects. We want to understand how any knowledge at all is possible—how anything we currently accept amounts to knowledge. Or less ambitiously, we want to understand with complete generality how we come to know anything in a certain specific domain (Stroud 2000, 101).

Stroud argues that there are limits on what should count as a satisfying philosophical understanding of how our knowledge in the problematic domains is possible. We are limited in terms of what counts as admissible explanans. In what follows (sections 6.1 and 6.2), I will explain two requirements which Stroud argues are necessary for a satisfying answer to the traditional epistemological project.

6.1 Generality

The goal of the traditional epistemic project is to explain how it's possible to know anything at all in a certain domain. The traditional epistemological project for our knowledge of the external world, then, is the goal of explaining how it's possible to know anything at all about non-psychological, physical reality.

Now let's consider an initial problem. We have already identified the domain (propositions about the external world) and what we want to explain about it (how it is possible to know that any members of the domain are true). But if this were the only way of singling out our subject matter, we might still be at a loss. For what if there were only different, unique “pathways” or “channels” to knowing propositions in that domain, and no general “pathways” or “channels”?⁴⁷ We would then have to approach the question in a piece-meal fashion, explaining on each occasion how knowing that type of proposition on that particular, unique basis, is possible. It's plausible that this kind of project could not be successful.

⁴⁷ For this idea, see Goldman (2002), and Goldman (2007).

So, in order to make progress on this question, it's plausible that we have to first identify some general source of our knowledge of non-psychological, physical reality.⁴⁸ And according to Stroud:

We want a general answer to the question. It should be expressed in terms of a general 'way of knowing'. And we find that general source in what we call 'the senses' or 'sense-perception'. The problem then is to explain how we can get any knowledge at all of the world around us on the basis of sense-perception (Stroud 2000, 5).⁴⁹

So the traditional epistemological project for our knowledge of non-psychological, physical reality is the problem of explaining how it's possible to know anything at all in that domain on the basis of the senses.

This is one dimension of generality that the traditional epistemological project seeks to accommodate. It seeks to explain how knowledge in the relevant domain is possible at all, so that nothing that we know is left unexplained. For ease of exposition, we will call this requirement the generality requirement.⁵⁰

As Stroud (2000) and Williams (1996) observe, however, the generality requirement seems to constrain even further the admissible realm of explanations we can give, so that explanations which entail or presuppose that have the sort of knowledge that we want to be explained are inadmissible.

For example, let's consider the problem of other minds as a test case. In the problem of other minds, the goal is to explain how it's possible to know any truths about the psychological states of others (e.g. what the person is thinking or feeling, or even that they have a mind at all). Now let's see how the generality requirement interacts with this goal. The explanandum is knowledge

⁴⁸ Cf. Cassam (2009).

⁴⁹ Cf. with Stroud (1984): "It takes very little reflection on the human organism to convince us of the importance of the senses—sight, hearing, touch, taste, and smell. [...] The importance of the senses as a source or channel of knowledge seems undeniable. It seems possible, then, to acknowledge their importance and to assess the reliability of that source" (Stroud 1984, 6-7).

⁵⁰ The argument (very roughly) for the generality requirement is that it is constitutive of the traditional epistemological project that a satisfying explanation respect the generality requirement. In support of this argument, Stroud appeals to our intuitions about cases (see Stroud 2000, 2-5, 104-107). I don't discuss the arguments here because it's not strictly necessary for my expository purposes. In Chapter 6, I discuss and evaluate these arguments.

of other minds, and the generality requirement tells us that we must explain how any knowledge of other minds is possible at all, so that none of our knowledge of other minds is left unexplained. Can our explanans then feature propositions about other minds? If they did, then there would be some unexplained knowledge. This knowledge would be our knowledge of those propositions—the propositions about other minds that make up our explanans. But the goal is to explain how it's possible to know any of those types of propositions. So, naturally, the next question to ask would be: how is it possible to know that those propositions (that compose our explanans) are true? After all, this question would be one that we ought to ask, if our goal is to explain how it's possible to know any propositions about other minds, since those propositions are themselves about other minds.

This kind of argument motivates an epistemic priority requirement on any satisfying answer to the traditional epistemological projects. Stroud puts the requirement as follows:

We want to explain to a certain kind of knowledge, and we feel we must explain it on the basis of another, prior kind of knowledge that does not imply or presuppose any of the knowledge we are trying to explain. Without that, we will not be explaining the knowledge in question in the proper, fully general way. This felt need is what so easily brings into the epistemological project some notion or other of what is usually called 'epistemic priority'—one kind of knowledge being prior to another (Stroud 2000, 104).

We can express the epistemic priority requirement (henceforth "EPR") more sharply as follows: any satisfying philosophical explanation of how knowledge in a domain D is possible must neither imply nor presuppose any knowledge that any members of D are true.⁵¹ Instead, our explanation should be composed of propositions that we know are true, but none of those propositions can be part of the target domain.

⁵¹ Cf. Cassam (2009, 577).

6.2 From EPR to Scepticism

What sort of explanations of how our knowledge of the world is possible does EPR permit, and which sorts of explanations does it exclude? Consider what Dretske (1969) calls “epistemic perception”. This kind of perception is the perception of facts (of states-of-affairs that obtain), rather than the perception of things (of objects and properties). For example, someone can see an open bank (can see an x which is F), without seeing that the bank is open (without seeing that x is F). To see an open bank is to see an object with properties (a bank that is open), while to see that the bank is open is to see that the object has that property (that the bank is open).

EPR seems to exclude explanations of how our knowledge of the world is possible which appeal to epistemic perception alone (cf. Cassam 2009, 571). Why? Because epistemic perception seems to be knowledge-entailing. A knowledge-entailing attitude ϕ to a proposition p is one that satisfies the following schema: if S ϕ s that p , then S knows that p .⁵² So, for example, if I see that the bank is open, then I know that the bank is open.

So, suppose I try to explain how it’s possible for anyone to know anything at all about non-psychological, physical reality by citing the fact that human beings can perceive that non-psychological, physical states-of-affairs obtain. This would clearly violate EPR because I would be trying to explain how it’s possible for anyone to know anything at all about the world “in terms of a form of perceiving that already amounts to knowing about the world” (Cassam 2009, 571). I would be explaining how it’s possible for anyone to know anything about the world on the basis of a state which implies knowledge of the world.

On the other hand, if EPR excludes knowledge-entailing states like epistemic perception, this does not entail that it excludes all forms of perception. Indeed, as Stroud highlighted (section 6.1), it’s hard to see how sense-perception, as a general source of knowledge, could be left out of the

⁵² Cf. Williamson (2000, 37).

explanation. This raises the question: if EPR prevents us from appealing to perceptual states like epistemic perception, which entail knowledge of the world, then what sort of perceptual states does EPR permit?

A natural answer is perceptual appearance states. These are states that are picked out by sentences such as that “I seem to see that p ”, “It perceptually appears to me that p ”, “I have a sensory experience as of p ”, “It is perceptually for me just as if p ”. As Williams (1996*b*) comments:

[...] what sort of explanation would satisfy us? Well surely, the argument now goes, if we are to understand how it is possible for us to know anything at all about the external world, we must trace that knowledge to knowledge that we should have even if we knew nothing about the world. In this way we are led straight to the traditional doctrine of the priority of experiential knowledge over knowledge of the world, for presumably experiential knowledge is what remains when knowledge of the world is set to one side (Williams 1996*b*, 363).

I can know that I seem to see that p without seeing that p (and so without knowing that p). If I can appeal to perceptual appearance states in order to positively explain how it's possible to know anything at all about non-psychological physical reality, then I will have satisfied the desiderata of the traditional epistemological project. The problem is that once our explanation satisfies EPR, it is difficult to see how we can avoid a non-sceptical explanation of our knowledge. Stroud (2000) presents a two-step argument for this conclusion as follows:

If we really are restricted in perception to ‘experiences’ or ‘sense-data’ or ‘stimulations’ which give us information that is prior to any knowledge of objects, how could we ever know anything about what goes on beyond such prior ‘data’? It would seem to be possible only if we somehow knew of some connection between what we are restricted to in observation and what is true in the wider domain we are interested in. But then knowing even that there was such a connection would be knowing something about the wider domain after all, not just about what we are restricted to in observation. And then we would be left with no satisfactorily general explanation of our knowledge (Stroud 2000, 105-106).

The first step of the argument tells us that our explanation of how our knowledge of the world is possible is one that respects EPR. So, insofar as our

explanation appeals to what EPR allows that we can know on the basis of the senses, it permits explanations in terms of our knowledge of propositions about perceptual appearances, none of which entail or presuppose facts about the world (about non-psychological physical reality).

The second step of the argument is that it's not enough to know facts about perceptual appearances. What we also need to know are how those facts are connected with facts about the world (the knowledge of which we want to explain).⁵³ For example, consider bridge principles such as that perceptual appearances as of p are reliable indicators that p is true. Stroud thinks that knowing that the bridge principle is true depends on or amounts to knowing a proposition about the world, violating EPR. So, the conclusion is that we cannot accommodate EPR without our explanation being a sceptical explanation.

6.3 From the denial of EPR to Philosophical Dissatisfaction

If an explanation which respects EPR implies that it is a sceptical explanation, then perhaps the right response is to reject EPR. After all, how plausible is EPR? Here, we need to distinguish between EPR as a requirement on explaining how we come to know various sorts of propositions, and as a requirement on explaining how it's possible for us to come to know various sorts of propositions. Stroud intends EPR to be a requirement on the latter rather than the former.

For example, when we go to explain how someone knows that it's raining outside, it can be enough to appeal to epistemic perceptual states. I can explain how someone knows that it's raining outside by citing the fact that they see or hear that it's raining outside. This is not problematic for non-philosophical purposes. As Stroud argues, it's problematic when it's put forward as an answer

⁵³ Cf. Fumerton (1995), who supports this step in an argument for *first-order* scepticism.

to the philosophical question; the question of how it's possible to know anything at all about non-psychological physical reality:

[...] in philosophy we want to understand how any knowledge of an independent world is gained on any of the occasions on which knowledge of the world is gained through sense-perception. So, unlike those everyday cases, when we understand the particular case in the way we must understand it for philosophical purposes, we cannot appeal to some piece of knowledge we think we already got about an independent world (Stroud 2000, 132).

So, when it comes to a satisfying answer to the traditional epistemological project, Stroud thinks that fulfilling EPR is non-negotiable.

6.4 Stroud's Dilemma

In the previous sections (6.1 and 6.2), we saw that Stroud thinks that there are certain desiderata for fulfilling the traditional epistemological project, the project of explaining how any knowledge of the world is possible at all. A philosophically satisfying explanation of how any knowledge of the world is possible has to fulfil the generality requirement, and this requirement constrains further the admissible types of explanations one can use. It constrains it so much that the only admissible explanations can neither imply nor presuppose any knowledge in the target domain (cf. section 6.1). That is, it must fulfil EPR.

And then we are led to a dilemma. The first horn is that our explanation must fulfil EPR. But if it fulfils EPR, then it won't be an anti-sceptical explanation (see section 6.2). The second horn of the dilemma is that if our explanation fails to fulfil EPR, then it won't be a satisfyingly general explanation (see section 6.3). In either case, the conclusion is that we won't have a satisfying, philosophical explanation of how our knowledge of the world is possible:

In short, it seems that if we really were in the position the traditional account in terms of epistemic priority describes us as being in, scepticism would be correct. We could not know the things we think we know. But if, in order to resist that conclusion, we no longer see ourselves in that traditional way, we will not have a satisfactorily

general explanation of all our knowledge in a certain domain (Stroud 2000, 106).

Call this “Stroud’s dilemma”. Stroud’s dilemma presents a challenge to the possibility of a satisfying philosophical explanation of how our knowledge of the world is possible.⁵⁴

Notice that like Stroud’s puzzle, Stroud’s dilemma is not an argument for scepticism. Instead, it is a dilemma, where one horn of the dilemma leads to scepticism, while the other is non-sceptical. The problem is that both options are, according to Stroud, philosophically unsatisfying. So, we have been presented with a prima facie valid argument for meta-epistemological scepticism which is not directly an argument for scepticism.

7 Conclusion

The goal of this chapter was to make the commitments of meta-epistemological scepticism clear, and how there are two arguments from Stroud (1984, 2000) which present an obstacle to the possibility of a satisfying philosophical explanation of how our knowledge of the world is possible.

In section 2, I explained what meta-epistemological scepticism is, and distinguished it from other forms of scepticism. In section 3, I assessed three prima facie challenges to meta-epistemological scepticism, and argued that these challenges can be met. The goal here was to further motivate meta-epistemological scepticism. In section 4, I reviewed one of the central arguments for first-order scepticism. The goal here was to provide an adequate backdrop from which to explain Stroud’s first meta-epistemological argument (Stroud’s puzzle). In section 5, I explained what Stroud’s puzzle is, how it motivates the thesis that meta-epistemological scepticism does not entail scepticism, and how it raises a challenge to the possibility of a philosophically satisfying explanation of how our knowledge of the world is possible. Finally,

⁵⁴ Cassam (2009) calls this dilemma “Stroud’s Dilemma”. Cf. Williams (1996, 15).

in section 6, I explained Stroud's second meta-epistemological argument (Stroud's dilemma), and how this argument raises a challenge to the possibility of a philosophically satisfying explanation of how our knowledge of the world is possible.

Chapter 2

Stroud's Puzzle: Descartes's Condition and Revisionism

1 Introduction

In the previous chapter, I framed the thesis around two meta-epistemological arguments from Stroud in support of meta-epistemological scepticism: Stroud's puzzle and Stroud's dilemma. In this chapter, I focus on the first argument: Stroud's puzzle.

Stroud's puzzle consists of the following claims. The first premise is that if Descartes's Condition is true, then it cannot be satisfied, so that scepticism is true. The problem is that scepticism cannot contribute to a satisfying philosophical explanation of how knowledge of the world is possible. (cf. Chapter 1, §5.1) The second premise is that if Descartes's Condition is false, then an ordinary epistemic principle is false. The problem is that the falsity of ordinary epistemic principle cannot contribute to a satisfying philosophical explanation of how our knowledge of the world is possible either. Stroud's puzzle, then, is that these claims provide adequate support for meta-epistemological scepticism (cf. Chapter 1, §5).

The plan for this chapter is as follows. In section 2, I explain Stroud's argument for the second premise of the puzzle.⁵⁵ As we will see, Stroud's

⁵⁵ The following disclaimer about Chapter 2 is in order: I respond to the arguments for the second premise of Stroud's puzzle before the first premise because the success of these

argument is an anti-revisionist argument: it is an argument for the conclusion that substantial revision to an ordinary epistemic principle is philosophically unsatisfying. In section 3, I respond to Stroud's anti-revisionist argument. However, in section 4, I illustrate some additional anti-revisionist arguments which might lend support to the second horn of Stroud's puzzle. But I argue that that these arguments fail to provide adequate support for second premise of Stroud's puzzle as well.

In the second part of the chapter, we switch gears and consider some alternative arguments in support of the second premise of Stroud's puzzle. Unlike the previous arguments, these arguments do not turn on anti-revisionist considerations. Considered in a more general form, Stroud has argued that there is some feature *F* of Descartes's Condition which renders its rejection a reason for philosophical dissatisfaction. I consider several different properties for *F* other than "being an instance of or grounded in an ordinary epistemic principle". The properties I consider are ones that certain passages from Stroud (1984) suggest.

In the first part of the chapter, I argue that if *F* is that it's a part of our ordinary epistemic practices, not only does this not count in favour of the principle, but it fails to render its rejection or falsity philosophically unsatisfying. However, in section 5, I argue that if instead *F* is that it's a hinge-proposition, then its rejection is plausibly philosophically unsatisfying. Nevertheless, I argue the principle doesn't have the properties that other, contender hinge propositions have. Finally, I argue that if *F* is that it's a principle within our deep epistemic grammar, then its rejection is plausibly philosophically unsatisfying because it's not plausibly a principle that we can reject. However, it's not clear that, as philosophers, we're equipped to know that it has that kind of feature. In short, it seems to be an empirical question. Section 6 concludes this chapter.

responses are helpful to our assessment to the first premise of Stroud's puzzle. This will become clearer in Chapter 3.

2 Revisionism

In this section, I introduce the distinction between “revisionism” and “anti-revisionism” with respect to some domain of inquiry. The distinction is useful for our purposes because it provides us with a framework for discussing the second premise of Stroud’s puzzle. In particular, we will see that Stroud’s argument in support of the second premise of Stroud’s puzzle turns on substantive anti-revisionist theses.

2.1 What is Revisionism?

What is revisionism? “Revisionism” picks out the combination of two theses. The first is that there is some commitment that most people have, or should have, or would have, upon sufficient reflection or training. The second is that this commitment is defective. It might be that it’s false, inaccurate, or that it doesn’t carve nature at its joints. The revisionist about some phenomena accepts both theses about that phenomena, while the “non-revisionist” accepts the first thesis, but denies the second. As I will be using the term, the “anti-revisionist” is a non-revisionist who maintains that revisionism ought to be avoided.⁵⁶

The central idea is that the disagreement over the first thesis is a “descriptive disagreement”. It’s a disagreement about whether a commitment to the phenomena in question is a commitment that most people have, or should have, or would have. The disagreement over the second thesis, however, is a “normative disagreement”, in the sense that disputants will argue that the commitment ought to be given up or that it ought not to be given up.

⁵⁶ One famous example of revisionism is Mackie’s “error-theory” of moral judgements. According to Mackie, our “ordinary moral judgements involve a claim to objectivity which both non-cognitive and naturalist analyses fail to capture. Moral scepticism must, therefore, take the form of an error-theory, admitting that a belief in objective values is built into ordinary moral thought and language, but holding that this ingrained belief is false” (Mackie 1977, 48-49). In this passage, we can see that Mackie is a revisionist as I am using the term, because Mackie thinks that ordinary moral discourse commits speakers to the objectivity of moral value, but that this commitment is false.

For example, consider the metaphysician who says that everyone is committed, by their “folk ontology”, to the existence of tables, chairs, lamps, and in general, composite material objects. To say that their “folk ontology” commits them to the existence of those entities is just to say that the ordinary body of beliefs that many people have (or would have, if prompted) commit them to the existence of those entities.⁵⁷

With this theory of what the folk ontology is in play, we can take at least two different kinds of attitudes to it. The first is to say that the folk ontology is accurate. Our theory about what the folk’s ontology tells us is that the folk are committed to the existence of composite, material objects, and there really are (there exists) composite material objects. The second is to say that the folk ontology is inaccurate (it’s not the case that there are composite material objects).⁵⁸ Our theory about what the folk’s ontology is tells us that the folk are committed to the existence of composite, material objects, but there are no such things.

The first theory is a meta-ontological theory, because it doesn’t simply say that there are composite material objects. Instead, it says that (i) the folk ontology is committed to the existence of composite material objects, and (ii) there are composite material objects. It is the fit between (i) and (ii) and the makes the first theory a non-revisionist theory. It does not revise the folk’s commitments.

On the other hand, the second theory is a revisionist theory, because it accepts that (i) the folk ontology is committed to the existence of composite

⁵⁷ For example, consider Hirsch’s (2005) definition of “common sense ontology”: “The composite objects we ordinarily talk about really exist; they typically persist through changes in their parts; they typically do not have sums; and they typically do not have temporal parts” (Hirsch 2005, 68).

⁵⁸ Peter Unger (1979) argues for mereological nihilism. Peter van Inwagen (1990) and Trenton Merricks (2003) are *quasi-nihilists*, arguing that there are no *non-living* composite material objects. Indeed, van Inwagen is a contextualist about mereological statements, whereas Merricks is an invariantist. For example, the contextualist might claim that the sentence “there is a chair in the room” expresses a false proposition in the context of metaphysics, but the same sentence expresses a true proposition in everyday contexts. The invariantist denies this.

material objects, but it denies (ii): it says that there aren't any composite material objects. It would not be revisionist if the theory only maintained (ii). It's the conjunction of (i) and (ii) that renders the theory a revisionist theory.

Now consider an epistemological case. Where does scepticism fall? We might say that scepticism is a revisionary thesis because it is incompatible with "folk epistemology", the body of ordinary epistemological beliefs and principles that we all share (or would share after suitable communication and reflection). Most people believe that not only can we know about the world around us, but that such knowledge is widespread. Scepticism falsifies a wide class of beliefs that many of us share (cf. Chapter 1, §2-3).

Now there are two kinds of objections that one can raise against Stroud's argument for the second premise of Stroud's puzzle. The first objection is that, from the falsity of the epistemic principle *P*, like Descartes's Condition, it does not follow that our explanation of how our knowledge of the world is possible is or ought to be philosophically unsatisfying. One might ask here whether Stroud is right to think that, because *P* is part of our ordinary epistemic practices, that its falsity renders our explanation philosophically unsatisfying.

The second objection is that it's implausible that the epistemic principle *P* is part of our ordinary epistemic practices. If this is right, then Stroud's thesis that Descartes's Condition implies that scepticism is true would tell only against Descartes's Condition, rather than a principle that is part of our ordinary epistemic practices as well.

Now we have to be careful here to note that the core reason Stroud thinks that the falsity of *P* renders unsatisfying any consequent philosophical explanation of how our knowledge of the world is possible is that *P* has the property of being part of our ordinary epistemic practices.⁵⁹ Yet one might think that while this property doesn't render its falsity philosophically unsatisfying, there can still be other properties such that, if *P* had them, then its

⁵⁹ Note, however, that Stroud will sometimes say that it's a "fact" of our ordinary conception of knowledge" which looks like a stronger claim (Stroud 1984, 30).

falsity would be philosophically unsatisfying. I take up this consideration in later sections of this chapter (section 5).

3 Stroud and Anti-Revisionism

In this section, I evaluate Stroud's argument for the second premise of Stroud's puzzle. In particular, I evaluate two anti-revisionist arguments in favour of this premise: a strong anti-revisionist argument and a weak anti-revisionist argument. I argue that both of these arguments fail to provide adequate support for the second premise of Stroud's puzzle.

3.1 Stroud's Argument: The Strong Version

Let's suppose that Stroud is right about the following: that Descartes's Condition cannot be satisfied, and therefore, that if Descartes's Condition is true, then scepticism is true. Perhaps the best thing to do is take this as a reductio of Descartes's Condition. But Stroud doesn't think that this move will work. As he puts the point:

[*Scepticism*] can be avoided, it seems to me, only if we can find some way to avoid the requirement that we must know we are not dreaming if we are to know anything about the world around us. But that requirement cannot be avoided if it is nothing more than an instance of a general procedure we recognize and insist on in making an assessing knowledge-claims in everyday and scientific life (Stroud 1984, 31).

According to Stroud, then, Descartes's Condition "cannot be avoided" because it is "nothing more than an instance of a general procedure we insist on" in ordinary contexts of epistemic evaluation (Stroud *ibid*). That is, the epistemic principle which encodes Descartes's Condition is part of our ordinary epistemic practices. Elsewhere, Stroud writes that we can be lead to endorse the principle by "following the same procedures we all ordinarily follow in assessing knowledge-claims in everyday life" (Stroud 2000, 57). But the question I want

to ask here is in what sense does that make Descartes's Condition "unavoidable"?

One answer to this question is that it makes it look plausible that Descartes's Condition is true. We cannot avoid accepting Descartes's Condition because it is true. Indeed, this might be what is behind Stroud's suggestion that Descartes's Condition gives expression to "platitudes we would all accept": that a principle's having this sort of feature, the feature of being platitudinous, counts in favour of the truth of the principle (Stroud 1984, 82). It is tempting, then, to construe Stroud's argument here as follows:

1. If a principle *P* is part of our ordinary epistemic practices, then *P* is true.
2. The strong closure principle is part of our ordinary epistemic practices.

Therefore,

3. The strong closure principle is true.

The first premise registers an extreme form of anti-revisionism. This is the thesis that an epistemic principles' being part of our ordinary epistemic practices entails that the principle is true. The second premise affirms the antecedent of the first premise, with the strong closure principle replacing '*P*'. From modus ponens on these two premises, it follows that the strong closure principle is true.

There are problems with this argument. The first problem is the fact that one can accept, along with Stroud, that the second premise is true, while rejecting the first premise. It's implausible that a principles' being part of our ordinary epistemic practices entails that the principle is true.⁶⁰ At best, that a

⁶⁰ For example, consider how the epistemic practices of healthcare workers has changed over time, with some practices refined, and others discarded. The same can be said for the epistemic practices of professional detectives and professional educators.

principle is part of our epistemic practices gives one prima facie justification to believe that it is true. Indeed, it might even have the status of being a “default” (cf. Pritchard 2012). But there is no question that this sort of justification can be undercut or overridden in light of new evidence.⁶¹

The second problem is that, while the argument might articulate one answer to the question (the question of what it means for Descartes’s Condition to be “unavoidable”), on closer inspection it’s not clear that it’s Stroud’s answer. Recall that Stroud maintains that:

We have no notion of knowledge other than what is embodied in those procedures and practices (i.e., our ordinary epistemic practices). So if that requirement is a ‘fact’ of our ordinary conception of knowledge we will have to accept the conclusion that no one knows anything about the world around us (Stroud 1984, 31 my addition).

The main premise of this argument expresses a thesis about the source or ground of our concept or understanding of knowledge. The premise seems to be that we “no notion of knowledge” except for whatever epistemological principles and procedures are a part of “folk epistemology” (Stroud *ibid*). What I am calling “folk epistemology” is our everyday shared body of epistemological beliefs, commitments, principles, and practices.

Stroud’s point, then, does not seem to be that, because Descartes’s Condition is an instance of an epistemic principle which is part of our ordinary epistemic practices, it follows that the principle is true. Instead, Stroud’s point seems to be much weaker than this. What’s difficult here is just how we ought to unpack this weaker claim.

⁶¹ That believing that p provides prima facie justification in favour of p is called “epistemic conservatism”. Even if a proponent of the argument in the text were an epistemic conservative, and even if it were true that the strong closure principle is part of “common sense”, or at least a *tacit belief* of most people, that is compatible with the justification one has in favour of the strong closure principle being overridden. So, even epistemic conservatism does not provide adequate support for anti-revisionism. For a defense of conservatism, see McCain (2008).

3.2 Stroud's Argument: The Weak Version

Perhaps we can unpack Stroud's weaker claim along the following lines: the fact that a principle is part of our ordinary epistemic practices makes it the case that we ought to believe that it is true because "[we] have no notion of knowledge other than what is embodied in those procedures and practices (i.e., our ordinary epistemic practices)" (Stroud *ibid*).

So, we can reformulate Stroud's argument as follows:

- 1*. If a principle *P* is part of our ordinary epistemic practices, then we ought to believe that it is true.
- 2.* The strong closure principle is part of our ordinary epistemic practices.

Therefore,

- 3*. We ought to believe that the strong closure principle is true.

Stroud's argument in favour of premise 1* is that we "have no notion of knowledge other than what is embodied in" our ordinary epistemic practices (Stroud *ibid*). Let's call this argument for premise 1* the "source argument". I call it the "source argument" because Stroud finds it hard to see what the source of one's conception of knowledge is except for what's part folk epistemology.⁶²

Is the source argument a good argument? We can suppose that Stroud is right that the strong closure principle is part of our folk epistemology. The question that needs to be addressed is whether that fact implies that we ought to believe that the strong closure principle is true.⁶³

⁶² Cf. Stroud (2000, 133).

⁶³ Now once might wonder whether Stroud is right even to think that that principle is part of our ordinary epistemic practices. I'm granting that to Stroud for the sake of argument, although I do think it is contentious. Here, at least, is one line of support Stroud offers in favour of this idea: "The force we feel in the sceptical argument when we first encounter it is itself evidence that the conception of knowledge employed is the very conception we have been operating with all along" (Stroud 1984, 71).

There are at least two issues with the source argument. I will consider these issues in turn. The first issue is that we can allow that a principle *P* being a part of folk epistemology gives us good grounds to believe that *P* is true, but this does not mean that its truth is non-negotiable. We can acquire reasons to reject that *P*, reasons which defeat the original reasons that we had from *P*'s being a part of folk epistemology. This does not entail that *P* is not a part of folk epistemology after all; just that folk epistemology is defective.

The second issue is that it's not obvious that this leads to an objectionable form of revisionism. In his discussion of our "pre-theoretical conception of knowledge", Haddock (2010) presents the following general idea in favour of epistemological anti-revisionism:

[...] there is a presumption in favour of our pre-theoretical conception of any object: we ought to hold on to as much of this conception as we can unless there are excellent reasons for its abandonment (Haddock 2010, 203).

Now let's suppose that Stroud is right about the first premise of Stroud's puzzle. This is the premise that if Descartes's Condition is true, then scepticism is true. It's not implausible that scepticism counts against a principle which implies that it's true. But there is a better response. This response tells us that anti-revisionists should be moderate rather than extreme. Scepticism modifies our pre-theoretical conception of ourselves as epistemic agents so much that there is no question that the kind of revision that scepticism implies is far greater than the kind of revision the rejection of Descartes's Condition implies (cf. Chapter 1, §5). Scepticism is therefore an "excellent reason" for the abandonment of Descartes's Condition (Haddock *ibid*).

Third, notice that Stroud includes as part of our ordinary epistemic practices the epistemic practices of scientists. Should this be permitted? At least it's not clear that it should be permitted if the epistemic practices of philosophers are not allowed to inform us of what our "notion" of knowledge

is.⁶⁴ So, we might wonder why our scientific practices, legal practices, social practices, and so forth get the privilege of being a source of our notion of knowledge while philosophy is prohibited from doing so. According to this objection, what Stroud includes as part of our ordinary epistemic practices is ad hoc.⁶⁵

I take it that Stroud's anti-revisionism develops out of his views about knowledge. The idea is that our grasp of what knowledge is somehow emerges out of our grasp of how it is embedded in certain epistemic practices. The issue is that Stroud permits various sorts of epistemic practices, such as our scientific and legal practices, to count as part of our "ordinary" or "everyday" epistemic practices, while discounting some others.⁶⁶ In particular, Stroud discounts the epistemic practices of philosophers. Does Stroud have to do this? Here's one reason for thinking that he does: if the epistemic practices of philosophers weren't prohibited, it's no longer clear that the rejection of the strong closure principle would be revisionary in any way at all.

Why? Because now the epistemic practices in philosophy would count as being part of our ordinary epistemic practices, and second, there are apparently independently plausible arguments against Descartes's Condition and the strong closure principle. These arguments against the principle would make it look less like the principle was part of our ordinary epistemic practices. So, it's hard to see how it's an obviously true principle. The mere existence of independently plausible arguments (independent of our evaluation of

⁶⁴ For example, Stroud insists that Descartes's Condition "cannot be avoided if it is nothing more than an instance of a general procedure we recognize and insist on in making and assessing knowledge-claims in everyday and scientific life. We have no notion of knowledge other than what is embodied in those procedures and practices" (Stroud 1984, 30-31). Notice that Stroud includes our assessments of knowledge-claims in "scientific life" as a reason for thinking that Descartes's Condition "cannot be avoided" (Stroud *ibid*). Compare with Stroud (1984, 53).

⁶⁵ Or more weakly: there is no principled criteria for what counts as part of our ordinary epistemic practices and what doesn't. If scientific epistemic practices count, why not philosophy?

⁶⁶ Cf. Stroud (2000, 131).

scepticism) for the rejection of the principle renders it not obviously true, even if it is true.

According to this objection, if we broaden our “ordinary” epistemic practices to encompass the epistemic practices of philosophers, then it’s not clear that the rejection of Descartes’s Condition would entail revisionism. On other hand, if we narrow our “ordinary” epistemic practices such that the epistemic practices of philosophers are excluded, but not the epistemic practices of natural scientists, medical doctors, and so forth, the narrowing looks ad hoc. It looks like it’s excluded in order to guarantee revisionism.

Stroud has a response to this argument. The argument goes like this. Let’s suppose that the strong closure principle is part of our ordinary epistemic practices, but that this fact is not sufficient to show that we ought to believe that it’s true. That is, its rejection is not non-negotiable. But don’t we thereby risk changing the subject from one of knowledge to something else—to some other epistemic property—if we don’t believe that the epistemic principle is true of our concept of knowledge, because what else could be the source of our information on what knowledge is if not the principles which seem to be part of our ordinary epistemic practices?⁶⁷ And if that’s right, doesn’t the rejection of any of those principles thereby change the subject from knowledge to some other epistemic property?

Call this the “changing the subject” objection. The point of the objection is to show that revisionists are changing the subject.⁶⁸ So, one of Stroud’s worries

⁶⁷ Lewis (1996) makes a similar point in his discussion of modal realism. He says the following: “I am trying to improve that theory, that is to change it. But I am trying to improve *that theory*, that is to leave it recognisably the same theory as before” (Lewis 134, 1986).

⁶⁸ McDowell (1995) and Pritchard (2012) also seem to be anti-revisionists. In particular, their defence of epistemological disjunctivism appeals largely to its apparently cohering with ordinary epistemic practice. For example, in his discussion of the motivations for epistemological disjunctivism, Pritchard (2012) says that “the naturalness of this conversational exchange implies that it is the epistemological disjunctivist who is working with the commonsense position, and her detractor who is offering the revisionistic view” (Pritchard 2012, 17). In a footnote, Pritchard says that this view even “occupies a kind of default position in our thinking about perceptual knowledge, on account of its being rooted in a commonsense picture of our rational support for such knowledge” (Pritchard 2012, 53).

is that, if we don't at least believe that the epistemic principle which underlies Descartes's Condition is true of our concept of knowledge, then we risk not talking about what we were interested in in the first place: the referent of our concept of knowledge.

How good of an objection is the "changing the subject" objection? The objection has at least worried some philosophers. For example, consider Jackson (1998):

[...] if we wish to address the concerns of our fellows when we discuss the matter—and if we don't we will not have much of an audience—we had better mean what the folk mean. We had better, that is, identify our subject via the folk theory (Jackson 1998, 118).

Jackson takes it that there are "commonplaces or platitudes or constitutive principles that make up the core we need to share in order to count as speaking a common moral language" (Jackson 1998, 132). Let's suppose that the same is true of our epistemic language: that there is a cluster of "commonplaces or platitudes or constitutive principles that make of the core" of our epistemic language. In general, there is a cluster of principles and practices which make up the core of our epistemic practices. In conjunction with the "source argument" from Stroud, this view from Jackson (1998) implies that revision to Descartes's Condition, and so the principle that underlies it (the principle for which Descartes's Condition is an instance), would be a revision to a constitutive principle of our core epistemic practices.

Prima facie, then, Stroud has a powerful argument. What fixes the referent of the concept of knowledge are the constitutive principles of our core epistemic practices. On Stroud's view, Descartes's Condition—or the general epistemic principle which underlies it—is among them. So, the rejection of Descartes's Condition would be the rejection of a constitutive principle of our core epistemic practices. Now the question here is this: doesn't that make it look plausible that the revisionist is changing the subject from knowledge to some fledgling concept?

It's not clear that this does make this look plausible. So, perhaps the main problem with the "changing the subject" objection is that it's not obvious that the subject is being changed. Not all conceptual revision is conceptual elimination. It does not follow from the elimination of some part *P* of a folk concept *C* that there is a distinct concept *C** which has *P* and it's this concept which is the folk's concept. Indeed, as Vargas (2009) highlights, "[s]cientific cases of concept revision are plentiful. Our concepts of water, light, and temperature have undergone substantial, even radical change" (Varga 2009, 45). In these cases, conceptual revision did not render a change in subject. But if this is so in scientific cases, is there some special reason to think this cannot be so in epistemological cases? No.

Let's turn to a similar objection. Williams (1996) accepts the "changing the subject" objection. According to Williams, if Stroud is right about Descartes's Condition, then "no solution to scepticism will amount to a defence of our pre-theoretical claims to knowledge as we have always intended them to be understood" (Williams 1996, 19).

However, Williams suggests that the "changing the subject" objection motivates a stronger claim. The motivation can be stated as follows. If, in rejecting the epistemic principle that is part of our ordinary epistemic practices, we change the subject as I have outlined above, so that what we show to be possible is not whatever satisfies the folk's concept of knowledge, it's hard to see how this is not "just another way of agreeing with the sceptic" (Williams *ibid*).

Now even if the "changing the subject" objection is a good objection, Williams's additional claim does not follow from the success of that objection. The sceptic maintains that knowledge of the external world is impossible. So, one would not be agreeing with the sceptic if the folk concept of knowledge turns out to be defective. To agree with the sceptic, one needs to have a certain view about how the concept is defective. For the sceptic, our concept of knowledge is defective in virtue of the fact that it's impossible for our

knowledge claims about the world to be true. It's defective in that it cannot be satisfied in that domain. But not all defects in our concept of knowledge are the defects that the sceptic thinks she has identified.

A final objection is what we might call the "equal weight" objection. According to the "equal weight" objection, for any propositions p and q , if the rejection of p and q both result in revision, then there's no more reason to prefer the rejection of p over the rejection of q . Both propositions have an equal epistemic weight.

The problem with the "equal weight" objection is that it assumes that revision does not come in degrees. The degree to which the elimination of p is a revision might be less than the degree to which the elimination of q is a revision. The content of the propositions in question, and our knowledge of the potential effects of their rejection, ought to make a difference to our assessment of how much revision the rejection of p or q requires. The proponent of the "equal weight" objection ignores this.

Consider again the rejection of Descartes's Condition and scepticism. These represent the two potential revisions that Stroud's puzzle forces us to choose between. Of course, one might think that meta-epistemological scepticism is another option—Stroud's option—but this would be a philosophical revision, a revision to what we thought we could achieve as philosophers of knowledge, rather than a revision to our epistemic principles and commitments. One might think that this observation bolsters Stroud's acceptance of meta-epistemological scepticism. If the choice is between a philosophical revision and a revision to our quotidian epistemic principles and commitments, it's the latter which trumps the former.⁶⁹

⁶⁹ This seems to be a driving idea behind "commonsense" philosophers. For example, consider Lycan (2001): "Common sense beliefs can be corrected, even trashed entirely, by careful empirical investigation and scientific theorizing. No purely philosophical premise can ever (legitimately) have as strong a claim to our allegiance as can a humble common sense proposition. [...] Science can correct common sense; metaphysics and philosophical "intuition" can only throw spitballs" (Lycan 2001, 41). Compare with Gupta (2006): "Any theory that would wage war against common sense had better come loaded with some

However, the correct option here is the rejection of Descartes's Condition. The reason it's the correct option is that the degree to which the rejection of Descartes's Condition revises our epistemic principles and commitments is less than the degree to which scepticism revises our epistemic principles and commitments. "Meta-epistemological scepticism" is the third choice but the wrong choice.

3.3 Kaplan's Condition

In his (2000) Mark Kaplan presents a criterion for when revisionism is permissible. According to Kaplan:

Anyone who would tell us, on philosophical grounds, that we know less than we thought would tell us that we should stop taking for granted things we have hitherto taken for granted in our ordinary inquiries and decision-making. If [*the revisionist*] is to sway us, she had better offer us good reason to think we will be better off for changing our ordinary practices in the way required (Kaplan 2000, 302 my addition).

In short, it's not that revisionism is in general impermissible. Instead, it is impermissible unless it provides us with a clear "improvement" in our epistemic practices (Kaplan 2000, 303). So, on Kaplan's view at least, the anti-revisionist can avoid the objection that anti-revisionism is opposed to the epistemic goal of seeking truth and the norm that our pre-philosophical beliefs should not be insulated from our epistemic or otherwise theoretical inquiries:

[...] that our epistemology must be true to our ordinary practices [...] is not to insulate our ordinary practices from change. Nor is it to say that our epistemological investigations cannot provide us with reason to change our ordinary practices. What it is to say is that our philosophical assessment of our epistemic condition must reflect our ordinary assessment of that condition (Kaplan 2000, 301).

With these two points in hand, we can construe a general criterion for when epistemic revision is permissible along the following lines:

powerful ammunition. Philosophy is incapable of providing such ammunition" (Gupta 2006, 178).

Kaplan's Condition: S ought not revise an ordinary epistemic principle, concept, or practice *P* unless the proposed revision would improve upon *P*.

For example, Kaplan considers Bayesianism, which he tells us is a “revisionary epistemology” because it provides a “revisionary way of evaluating the credibility of claims” (Kaplan 2000, 303). On Kaplan’s view, then, even though Bayesianism is at odds with our ordinary epistemic practices, Bayesianism is not ruled-out for anti-revisionist reasons because it “offers the improvement [to our ordinary epistemic practices] as advertised”—that is, it provides a “powerful [...] philosophical rationale for changing” our ordinary epistemic practices (Kaplan 2000, 303 my addition).

This kind of moderate anti-revisionism allows for permissible cases of revision. The first question I want to address is whether the rejection of Descartes’s Condition is a permissible case. Following Kaplan’s Condition, it would be permissible if rejecting Descartes’s Condition, and the principle that it is an instance of, provides us with an improvement of our epistemic practices. In short, if eliminating Descartes’s Condition corrects a defect in our epistemic practices, then we are not barred from rejecting it, and we therefore avoid the consequence that Stroud maintains, that rejecting it results in philosophical dissatisfaction.

What makes this question difficult to answer is that it might be that rejecting Descartes’s Condition provides some improvements to our epistemic practices along some axes, while it worsens those practices on some other axes.

In order to illustrate this problem, let’s consider the principle that Stroud thinks Descartes’s Condition instantiates. Recall that, according to Stroud, the principle which underlies Descartes’s Condition is the strong closure principle:

Strong closure principle: If S knows that *p*, and S knows that S’s knowing that *p* implies *q*, then S knows that *q*.

Stroud thinks this principle flows from reflection on “uncontroversial everyday examples” where the best explanation of our intuitions about those examples is our adherence to the strong closure principle (Stroud 1984, 29).

Now a suggested revision from Kaplan, *a la* Austin, is that we have a “special reason” to know the denials of propositions we know to be incompatible with our knowing. This would be an improvement on the principle, because it would still require us to know the denials of what we know to be incompatible with our knowing, along with what Stroud takes to be part of our ordinary epistemic practices, but the sceptical scenarios would be discounted, unless there were some special reason for taking them into consideration.⁷⁰ The thought is that, with the rejection of Descartes’s Condition, we avoid the consequence that scepticism is true, but not at significant cost to our ordinary epistemic practices, because the revision to the principle which underlies Descartes’s Condition is not enough to render it false *tout court*.⁷¹

On the other hand, if Stroud is right about the consequences of Descartes’s Condition, that it cannot be satisfied (and therefore implies scepticism), what we gain in terms of theoretical knowledge comes at an extreme price: we lose knowledge of the world. If moderate revisionism is a live-option, however, the choice is much clearer: the rejection of Descartes’s Condition is better than accepting it.

4 Arguments for Anti-Revisionism

In the previous section, I examined Stroud’s (1984) argument for anti-revisionism. I argued that, contra Stroud, revisionism is a live option. In this section, I consider some alternative arguments for anti-revisionism. I argue that none of these arguments is successful.

⁷⁰ This is what some proponents of the “relevant alternatives” theory say. See Dretske (1970), Greco (1993), and Stine (1976/1999). This view originates in Austin (1946).

⁷¹ There is an interesting question here as to whether this is a form of pragmatic encroachment. See Vahid (2014) for discussion.

4.1 Argument from Reflective Equilibrium

A general methodological consideration one can use to argue for anti-revisionism is the method of reflective equilibrium. The general method can be summarized as follows: first, we begin with a cluster of beliefs. Often, these beliefs will be connected by subject matter or method. Second, we construct principles which seem to provide adequate support for those beliefs. Third, we work back-and-forth, pinning the beliefs against the principles, and the principles against the beliefs, until we reach the desired balance or equilibrium between our theoretical principles and our pre-theoretical beliefs, in a process of continuous refinement.

Some philosophers maintain that the method of reflective equilibrium is the proper starting point for philosophical theorizing. For example, Lewis maintained that our philosophical views need to be brought into equilibrium with our pre-philosophical “intuitions” or “opinions”:

Our “intuitions” are simply opinions: our philosophical theories are the same. Some are commonsensical, some are sophisticated; some are particular; some general; some are more firmly held, some less. But they are all opinions, and a reasonable goal for a philosopher is to bring them into equilibrium. Our common task is to find out what equilibria there are that can withstand examination, but it remains for each of us to come to rest at one or another of them (Lewis 1983, x).

In outline, the method of reflective equilibrium suggests that we bring our pre-philosophical beliefs and our philosophical beliefs into equilibrium. Now our task is to show that revisionism is a live option. So let us suppose that Descartes’s Condition is a part of our ordinary epistemic practices. The question we need to ask is whether, given this supposition, the method of reflective equilibrium prohibits the rejection of Descartes’s Condition.

We might think that the method of reflective equilibrium would not lead the inquirer to adopt scepticism over the rejection of Descartes’s Condition. The strong anti-revisionist (e.g., Stroud 1984) would conclude from this that, since the rejection of Descartes’s Condition is no more negotiable than adopting scepticism, meta-epistemological scepticism follows. On the other

hand, the more modest anti-revisionist can be moved here to reject Descartes's Condition. For not all epistemological beliefs, principles, and concepts are on par. Considerations such as theoretical virtue and belief preservation can move us to adopt one over the other. In this case, it's not implausible that the theoretical virtues of rejecting Descartes's Condition, together with the massive preservation of higher-order epistemic beliefs that comes with avoiding scepticism, might move the inquirer to favour the rejection of Descartes's Condition over adopting scepticism.

Of course, the strong anti-revisionist (e.g., Stroud 1984) maintains that this is not philosophically satisfying, because neither option is non-revisionist. This just registers Stroud's anti-revisionist requirement on satisfying philosophical explanations of how our knowledge of the world is possible. Nevertheless, the argumentative aim here is not to please the strong anti-revisionist, but to please the modest anti-revisionist. From section 3, the argument for Stroud's anti-revisionist requirement did not provide adequate support for that requirement. So, from the perspective of modest anti-revisionism, the method of reflective equilibrium does not provide a good argument against rejecting Descartes's Condition. As it turns out, the method of reflective equilibrium suggests that the rejection of Descartes' Condition brings our pre-philosophical beliefs and our philosophical beliefs into equilibrium better than accepting Descartes's Condition.

4.2 Argument from Lewis's Maxim

In his (1986), David Lewis proposed the following maxim when deciding between philosophical theories:

[...] never put forward a [...] theory that you yourself cannot believe
in your least philosophical and most commonsensical moments
(Lewis 1986, 135)

A proponent of anti-revisionism might think that Lewis's maxim provides at least some epistemic support for their anti-revisionism. After all, if scepticism

is true, then there is a serious error in our epistemic discourse and epistemic thought. We attribute knowledge to ourselves and others in a diverse range of activities: science, courts of law, classrooms, hospitals, etc.⁷² In our most commonsensical moments, is it plausible that we could accept scepticism? No.

The anti-revisionist will want to make a similar claim about the strong closure principle which underwrites Descartes's Condition. After all, if one accepts that the principle is a "fact" of our ordinary conception of knowledge", how could we accept that it's false in our least philosophical but most commonsensical moments? (Stroud 1984, 31).

On the other hand, what good is philosophical analysis if not to guide belief on a difficult subject matter, where philosophy prima facie bears on that subject matter? Remember that defending revisionism here is not the same as a defence of scepticism. Instead, it is a defence of the potential rejection or falsity of Descartes's Condition (and the epistemic principle which it is an instance of) not having to be philosophically dissatisfying in virtue of being a part of our ordinary epistemic practices.

Still, there is a better response to the anti-revisionist's use of Lewis's maxim. For it seems to speak for revisionism in our case rather than anti-revisionism. After all, scepticism is not a thesis about our knowledge of the world that we could believe in our least philosophical moments. So if Stroud is right that Descartes's Condition has the consequence that scepticism is true, then Lewis's maxim dictates that we should not put that forward as true because we know that, if Stroud is right about Descartes's Condition, then scepticism is true as well. Lewis's maxim suggests that if the choice is between scepticism and Descartes's Condition, it's Descartes's Condition that we ought to reject.

⁷² Even if this is not explicit, it is implicit in many activities. If I ask someone to pass the salt, it's implicit in what I request that the person knows where the salt is. If I think that I would like someone to pass the salt, so too it is implicit in what I am thinking that someone potentially knows where the salt is. The term "knows" and its cognates don't need to show up in speech for it to be implicit that it's there.

A few comments on this argument are in order. In calling it a “revisionist” argument, I do not presuppose that all forms of revision are on par. I am presupposing that Stroud is right on the following point: accepting scepticism and rejecting Descartes’s Condition are revisions of folk epistemology. What I am questioning is the claim that these two cases of revision are on par.

The sort of revision that flows from accepting scepticism is different from the sort of revision that flows from the rejection of Descartes’s Condition (see Chapter 1, §5). The sort of damage that scepticism does is not the same kind of damage that rejecting Descartes’s Condition does. Scepticism has the first-order effect that knowledge of the world is impossible, and the higher-order effect that our higher-order epistemic beliefs are false (e.g., *S* knows that the milk is in the fridge; *S* believes that someone knows that interest rates are rising). Since the damage is different, it is not surprising that our evaluation of the damage should be different as well.

Moreover, it is not obvious that the revisionist cannot respect Lewis’s maxim. Consider an argument from Lewis (1986) on how the method of reflective equilibrium interacts with common sense:

Common sense has no absolute authority in philosophy. It’s not that the folk know in their blood what the highfalutin’ philosophers may forget. And it’s not that common sense speaks with the voice of some infallible faculty of ‘intuition’. It’s just that theoretical conservatism is the only sensible policy for theorists of limited powers, who are duly modest about what they could accomplish after a fresh start. Part of this conservatism is reluctance to accept theories that fly in the face of common sense. But it’s a matter of balance and judgment. Some common sense opinions are firmer than others, so the cost of denying common sense opinion differs from one case to the next. And the costs must be set against the gains. Sometimes common sense may properly be corrected, when the earned credence that is gained by making theory more systematic more than makes up for the inherited credence that is lost. It is not to be demanded that a philosophical theory should agree with anything that the man on the street would insist on offhand, uninformed and therefore uninfluenced by any theoretical gains to be had by changing his mind (Lewis 1986, 134-135).

Lewis tells us that some common sense opinions are firmer than others, and draws from this observation that not all cases should favour non-revision over

revision. Let's see what Lewis's observation can tell us about Stroud's puzzle. Scepticism revises common sense opinions, but it's not so clear how the rejection of Descartes's Condition revises common sense opinions or beliefs. Instead, the rejection of Descartes's Condition results in the rejection of an epistemic principle which "grounds" an epistemic practice. As Stroud highlights, it's "obvious that we do not always insist that people know they are not dreaming before we allow that they know something in everyday life, or even in science or a court of law" (Stroud 1984, 39). So, one might think that the revision it would require would have limited effects on our ordinary epistemic practices.

4.3 Argument from the Principle of Charity

According to Davidson (1973) and Jackson (1998), as philosophers we ought to "work on the general presumption that the folk are not badly confused" in their beliefs and practices (Jackson 1998, 103). On a weak formulation of the principle, being charitable interpreters of the folk's beliefs, concepts, and practices requires us to presume that their beliefs, concepts, and practices do not suffer from massive, systematic error.

For example, we can construe the sceptic about our knowledge of the world as someone who maintains that none of the knowledge-claims we make in that discourse are true. This, it might seem, violates the principle because it attributes massive error to the participants of the discourse. Moreover, Stroud's revisionist is someone who accepts that the epistemic principle which underlies Descartes's Condition is a part of our epistemic practices, but that that principle is false. This too, it might seem, violates the principle because it attributes error to the participants of the epistemic practice.

So, like Lewis's maxim, this principle is put forward as a principle which, prima facie at least, speaks against revisionism. But should we accept the principle; and if we should accept the principle, does it provide us with good

reasons to avoid revisionism? Here is a general statement and defence of the principle, where the argument seems to be anti-revisionist in character:

How should we respond to claims to find an inconsistency in an ordinary concept? We have some reason to be suspicious of claims to find inconsistencies in our ordinary concepts. For the very fact that these ordinary terms survive in our language suggests that they are at least by and large well-functioning, and interpretive constraints like the principle of charity give us reason to resist interpretations that attribute widespread inconsistencies to our thought, especially where it concerns features of our everyday practical life (Thomasson 2010, 598).

Augmenting the argument to our epistemic case, we can re-state the argument as follows. The fact that our concept of knowledge (of the external world) has survived within our epistemic discourse provides us with good reason to believe that the concept is well functioning. But the sceptic's view about our concept of knowledge (of the external world) implies that it is not well functioning, because it cannot be satisfied. So, the principle gives us good reason to believe that scepticism is not the correct account of our concept of knowledge (of the external world).

Now we might imagine augmenting the argument further to provide support for anti-revisionism with respect to Descartes's Condition and the strong closure principle. For example, we first suppose that the strong closure principle explains or grounds various epistemic practices.⁷³ The supposition here is that there is an epistemic practice of requiring one to know that various error-possibilities do not obtain when the person knows that those error-possibilities are incompatible with their knowing that *p*. The second claim is that this epistemic practice is well-functioning, and has survived or resisted substantive change from within the practice; if the practice was defective, then we maximize the practitioners irrationality, and the principle of charity prohibits that.

What's problematic about this sort of appeal to the principle of charity is that it seems to speak in favour of revisionism in this case. Recall Stroud's

⁷³ See Stroud (1984, 26-30).

puzzle. If Stroud is right, then there are widespread inconsistencies in our epistemic discourse and practices. For on his view, if Descartes's Condition is true, there is a massive defect in our epistemic discourse, because scepticism turns out to be the correct account of the target discourse. On the other hand, if Descartes's Condition is false, there is a massive defect in our epistemic practices, because a principle, which is part of those practices, turns out to be false. What, then, does the principle prescribe in this case?

On the one hand, as I augmented the argument before, it tells against scepticism. But that it tells against that instance of revisionism does not entail that it tells against the other instance of revisionism that is at issue in Stroud's puzzle. The question we need to address here is how damaging would it be to our epistemic practices if Descartes's Condition turned out to be false?

On Stroud's view at least, it would reveal that our concept of knowledge was defective.⁷⁴ Let's suppose for the sake of argument that Stroud is right about this. Does the principle then prohibit us from claiming that the concept is defective? It's not clear that it does. That it is defective does not entail that our epistemic practices are overall defective. It does not entail that our knowledge-ascriptions are systematically defective (e.g., false), or that our epistemic practices are systemically irrational or unreliable. That kind of error is what the principle tries to prohibit. So, it is not clear that this instance of revision is incompatible with the principle.

5 Interlude

In previous section, I argued that the arguments against revisionism do not provide adequate support for the second premise of Stroud's puzzle. In particular, the argument from reflective equilibrium, the argument from Lewis's maxim, and the argument from the principle of charity, failed to provide adequate support for the second premise of Stroud's puzzle if the

⁷⁴ See Stroud (1984, 31). Cf. Stroud (1984) pp. 55-55 and 71-72.

relevant property that is supposed to make the rejection of the strong closure principle unsatisfactory is that it's a part of ordinary epistemic practice.

The general template argument for the second premise of Stroud's puzzle can be formulated as follows:

1. Descartes's Condition (and the principle for which it is an instance) is a part of our ordinary epistemic practices.
2. If Descartes's Condition (and the principle for which it is an instance) is a part of our ordinary epistemic practices, then the falsity of Descartes's Condition ought to be philosophically unsatisfying.

Therefore,

3. The falsity of Descartes's Condition ought to be philosophically unsatisfying.

This argument is valid. I argued in section 2 and 3 that the second premise is false (that it doesn't follow from the fact that Descartes's Condition is a part of our ordinary epistemic practices that its falsity ought to be philosophically unsatisfying).

Now I want to consider the following objection. I argued that even if Stroud is right that Descartes's Condition, and the principle for which it is an instance, has the property expressed by "being a part of our ordinary epistemic practices", it doesn't follow that its falsity ought to be philosophically unsatisfying. However, suppose that it has another property, *F*, such that its having *F* would render its falsity philosophically unsatisfying.

This objection is not ad hoc, because there are passages from Stroud (1984) which suggest Descartes's Condition has other properties which render

its falsity philosophically unsatisfying.⁷⁵ The arguments from section 2 and 3, however, have not spoken to those concerns. The next section tries to remedy this problem. I consider several properties:

- That it's a platitude
- That it's a hinge-commitment
- That it is a part of our deep epistemic grammar

Although this list is not exhaustive, it exhausts the plausible candidates suggested in Stroud's (1984). So, what I will be assessing is whether any of these properties (replacing *F* below) render the following argument sound:

- 1*. Descartes's Condition (and the principle for which it is an instance) has feature *F*.
- 2*. If Descartes's Condition has *F*, then the falsity of Descartes's Condition ought to be philosophically unsatisfying.

Therefore,

3. The falsity of Descartes's Condition ought to be philosophically unsatisfying.

I argue that the second premise 2* is false when *F* is that it's a platitude; I argue that the first premise 1* is false when *F* is that it's a hinge-commitment. I argue that the first premise 1* lacks adequate support when *F* is that it's a part of our deep epistemic grammar, and that it's plausibly not possibly false.

⁷⁵ For example, see Stroud (1984) pp. 29-30 (e.g. for its being a "simple and obvious fact about knowledge"), and 76-82 (e.g. for it being grounded in "platitudes" or "obvious truths").

5.1 Descartes's Condition and Platitudes

One suggestion from Stroud is that the rejection of Descartes's Condition results in the rejection of "platitudes we would all accept" (Stroud 1984, 82). How is this possible? As Stroud himself observes, Descartes's Condition doesn't look like a platitude, so how can it be that it nevertheless gives expression to or otherwise depends on a platitude?^{76,77}

Stroud's answer is complex. For this reason we will need to draw several distinctions and make important clarifications to Stroud's argument. The main question that we need to address is in what sense Descartes's Condition could depend upon "platitudes we would all accept" (Stroud *ibid*).

5.2 How can Descartes's Condition depend on "platitudes"?

How can Descartes's Condition depend on "platitudes that we would all accept" (Stroud *ibid*)? This seems implausible on its face because platitudes are mundane truths and it's hard to see how a theoretical principle like Descartes's Condition could depend on mundane truths alone.

Stroud's response to this objection exploits a distinction between the philosophical context and the non-philosophical context of epistemic evaluation. The thesis is that our epistemic evaluative practices in those two contexts operate under different constraints, but that these constraints are not

⁷⁶ For example, Stroud writes: "It is obvious that we do not always insist that people know they are not dreaming before we allow that they know something in everyday life, or even in science or a court of law, where the standards are presumably stricter. So it can easily look as if Descartes reaches his sceptical conclusion only by violating our ordinary standards and requirements for knowledge" (Stroud 1984, 40-41).

⁷⁷ According to the Oxford English Dictionary, a platitude is "A dull and commonplace or trivial remark or statement". For example "everyone has the right to an opinion" is considered platitudinous. I don't think Stroud intends this meaning of "platitude". Instead, he seems to something like "a basic, fundamental commitment". He also refers to the platitudes as "obvious truths", which is closer to the definition (see Stroud 1984, 76). One might think that in virtue of being a basic, or fundamental commitment, this prohibits it from being banal. So, there is a tension here with Stroud's use of "platitude" and the lexical definition.

sufficient to render the contents of the knowledge-ascriptions incompatible with one another.

First, Stroud is an “invariantist” about knowledge-ascriptions. “Invariantism” is the thesis that the truth-conditions for knowledge-ascriptions (e.g., “S knows that *p*”) are not sensitive to context in the sense that the truth-conditions do not shift when the ascriber’s context of utterance or evaluation changes.⁷⁸

Second, Stroud’s answer also turns on his conception of the philosophical context. On Stroud’s view, it’s not the case that within the philosophical context of epistemic evaluation, the requirements for the truth of “S knows that *p*” are higher or stricter than within practical contexts. Instead, Stroud thinks that the philosophical context is “purer” and more “general”, but the facts about the context being “purer” and more “general” don’t make it that case that truth-conditions for “S knows that *p*” are different there than from within the practical context. For example, here is Stroud on the practical context:

The standards or procedures we follow in everyday life find their source in the exigencies of action and in the general conditions under which actions must be performed. In the case of action, unlike that of belief and knowledge, truth is not the only important consideration. [...] This holds just as much for that action of saying something, or saying that you know something, or ruling out certain possibilities before you say that you know something, as for other kinds of actions (Stroud 1984, 65-66).

In practical contexts, practical norms, customs, and contingencies of time, effort, and reflection guide when it is permissible to issue a knowledge-ascription or claim that one knows that *p*.⁷⁹ These practical contingencies are lifted in philosophical contexts.⁸⁰

⁷⁸ Although this might be of limited value, Stroud has expressed this opinion in conversation. However, there is clear textual evidence for this as well. Stroud employs a “warranted assertability manoeuvre” against a contextualist account of knowledge-ascriptions in his (1984), Chapter 2. See also Bridges and Kolodny (2011) Chapter 1 for discussion. For recent work on knowledge-ascriptions, see Brown and Gerken (2012).

⁷⁹ What are “practical norms”? Examples include when it is permissible to issue a knowledge-ascription (e.g. assert that one or someone else knows that *p*), to raise possibilities of error, or to challenge a knowledge-ascription—that is, when it is permissible

According to Stroud, in the philosophical context we seek a “certain kind of understanding of our state or relation to the facts—what might be called an objective understanding of our position” (Stroud 1984, 79). What is an “objective understanding” of our epistemic position? For Stroud, it requires treating knowledge-ascriptions like other factual-ascriptions.⁸¹ This means treating the truth-conditions for “S knows that *p*” as one would treat the truth-conditions for “*p*”: as independent of the subject’s mental-states, conversational norms, and cultural customs. So, the first aspect of an “objective understanding” of our epistemic position is invariantism about the truth-conditions for knowledge-ascriptions.

The second aspect of an “objective understanding” of our epistemic position is that it “aims at an assessment of all our knowledge of the world all at once”. So, according to this aspect, it’s concerned with the set of all of our claims about the world we take to be cases of knowledge, where the aim is to form a “judgement on that knowledge from what looks like a detached ‘external’ perspective” (Stroud 1984, 209). The sense in which it is an “external perspective” is that our assessment of that set—whether there are members that amount to knowledge or not—cannot be from within an “internal” epistemic framework, where the epistemic evaluations are made given that knowledge of the world is possible, and given that some propositions about the world are

given our linguistic customs and conversational norms or maxims. See Audi (1989). For a discussion of knowledge-norms, see Benton (2014).

⁸⁰ For example, compare the quotation in the text with the following: “The special philosophical investigation is purely ‘theoretical’; it is detached from the practical concerns of everyday life. But it remains sufficiently connected with everyday life so that the general conclusion it reaches about our knowledge is to be understood as in direct conflict with what is strictly speaking said or implied in those everyday knowledge-claims” (Stroud 1984, 156).

⁸¹ For example, Stroud explains that “[w]hether someone (even ourselves) knows a certain thing is in that respect as objective a matter of fact as whether there is a mountain of a certain height in Africa, and what we seek is knowledge of whether or not that objective matter obtains, and perhaps in addition some understanding of how the conditions necessary and sufficient for its obtaining have been fulfilled” (Stroud 1984, 79-80).

known.⁸² From the “external perspective”, we cannot pass our judgement on that set given that some members of the set are known. Instead:

If all of our knowledge of the world around us is in question all at once we cannot then help ourselves to some independently reliable information about the world, as we usually do, to settle the question of whether our present course of experience is or is not on this occasion a reliable guide to the way things are (Stroud 1984, 209).

Stroud explains an example as follows:

For instance, on this view, in asking ‘Am I awake now or dreaming?’, as a philosophical question I am standing back as an ‘observer’ and asking simply under which concept ‘awake’ or ‘dreaming’, my present experience does in fact fall. Of course, I might not be able to ascertain which it falls under, but it does fall under one or the other whether I or anyone else can ever know it or not. One or the other is ‘objectively’ the case. This conception of what it is like to possess concepts, and to philosophize, expresses the traditional philosopher’s goal of complete or absolute objectivity (Stroud 2000, 30).⁸³

The claim is that the epistemologist can take a detached, external perspective on our epistemic practices, where the goal is to arrive at the correct description of the epistemic facts—the facts about which epistemic principles are true or false, or which epistemic practices are truth-reliable or unreliable—*independent* of our actual epistemic practices.

With this goal as the target, Stroud seems to think that the epistemologist ought to recognize that the elimination of sceptical possibilities (e.g., the dreaming scenario) is a genuine requirement on acquiring knowledge of the world, and that people can nevertheless be blameless in their failure to recognise that the elimination of the dreaming scenario is a requirement on knowledge of the world. This explanation of how the folk are blameless while

⁸² For example, see Stroud (1984, 212). Note that, as we will see in Chapters 4, 5, and 6, that this is closely related to EPR.

⁸³ Cf. “It is a quest for an objective or detached understanding and explanation of the position we are objectively in. What is seen to be true from a detached ‘external’ standpoint might not correspond to what we take to be the truth about our position when we consider it ‘internally’, from within the practical contexts which give our words their social point. Philosophical scepticism says the two do not correspond; we never know anything about the world around us, although we say or imply that we do a hundred times a day” (Stroud 1986, 81). Compare this with Nagel (1989).

the epistemologist is nevertheless correct is what Stroud thinks is grounded in platitudes.

Stroud appeals to the following sort of examples to bring out his idea:

Plane Spotter 1: In their airplane identification training-manual, planes with features A, B, and C are planes of type *F*. So, the careful plane spotter sees that the plane above has features A, B, and C, consults the training manual, and infers that the plane is an *F*, claiming to know that the plane is an *F*. The careless plane spotter, however, sees that the plane has features A and B, consults the manual, and infers that the plane is an *F*, claiming to know that the plane is an *F*.

Plane Spotter 2: The plane spotter 1 case, but, unlike that case, there is another kind of plane, *F**, which the manual fails to report because these planes are rare and cannot be detected from the ground, because it has features A, B, and C as well.

Stroud infers the following from the Plane Spotter cases.

- (a) in the plane spotter 1 case, the careful plane spotter knows that the plane is an *F*, while the careless plane spotter does not know that the plane is an *F*.
- (b) in the plane spotter 2 case, *neither* the careful plane spotter nor the careless plane spotter knows that the plane is an *F*. Moreover, the careful plane spotter is *blameless* in taking themselves to know that it is an *F*, whereas the careless plane spotter is not.
- (c) if the existence of *F**'s were made known to the careful plane spotter, the careful plane spotter would believe that *neither* the careful plane spotter nor the careless plane spotter knew that the plane was an *F*.

Stroud maintains that (a), (b), and (c) ought to be our intuitions about the plane spotter 1 case and the plane spotter 2 case, and that our actual epistemic

position with respect to the dreaming scenario is analogous to the careful plane spotter's relation to the planes with feature F^* in the plane spotter 2 case.⁸⁴

In short, Stroud maintains that, just as knowing that the plane is not an F^* is a requirement on knowing that the plane is an F , so too knowing that we are not merely dreaming that P is a requirement on knowing that P , even though we are blameless in our failure to recognise this requirement in non-philosophical contexts (and can be warranted for all practical purposes in taking ourselves to know without having fulfilled it).⁸⁵

How does this argument make contact with Stroud's claim that Descartes's Condition expresses or is otherwise grounded in a platitude, and that this fact is what makes its rejection philosophically unsatisfying? The answer is Stroud's diagnosis of the source of our intuitions about the plane spotter cases.

First, Stroud tells us that:

If our own more general practices of gaining and assessing knowledge in everyday life also operated under a similar set of practical constraints or restrictions [e.g., *the kind described in the plane spotter 2 case*] it looks as if it would also be possible for no one to know anything about the world around us even though our ordinary procedures are followed to the letter and our claims to know things are often beyond criticism (Stroud 1984, 75 my addition).

So, the first claim is that:

A. It is possible for an epistemic principle P to ground our ordinary epistemic practices without P being manifest in those practices.

⁸⁴ See Stroud (1984, 87-81). Stroud's explanation is arduous, so I have omitted the main quotations in the text and summarised the plane-spotter example and the key points instead. Stroud modifies the plane-spotter example from Clarke's (1972). See Kaplan (2000) for why Stroud's plane-spotter examples aren't effective ways of showing how Descartes's Condition could be grounded in ordinary epistemic principles.

⁸⁵ As Stroud puts the point: "the fact that we say what we do on different grounds from those available to speakers within the restricted practice [e.g. *the plane spotter cases*] does not show that the notion of knowledge we use is different from theirs. When they in their (justified) ignorance think that all the conditions of knowledge are in fact satisfied in the careful spotter's case, they mean by 'He knows it is an F ' what we mean by it, but they are simply mistaken (through no fault of theirs)" (Stroud 1984, 73-72 my addition).

What does principle A mean? How could it be that an epistemic principle is not manifest in our epistemic practices but could somehow ground those epistemic practices? Consider the following example from Pritchard:

the sceptical thought is that if one steps back from everyday life and employs everyday epistemic principles in a thorough-going way, while setting aside the practical limitations of everyday contexts, then one is led to sceptical doubts. It is precisely in this sense that radical scepticism is held to ‘fall out’ of our ordinary epistemic practices even while involving doubts that simply do not arise in quotidian situations (Pritchard *forthcoming*).

This helps to explain how it’s possible for Descartes’s Condition to depend on “platitudes we would all accept” without that principle itself being manifest in our epistemic practices. The next two questions we need to ask here are: how does it “depend on” the so-called “platitudes we would all accept” and what are those “platitudes we would all accept” (Stroud *ibid*)?

Stroud argues that the reason there’s this intuition that Descartes’s Condition, and the strong closure principle, could not be a part of our epistemic practices is that people fail to see that principle A is true. It’s possible for Descartes’s Condition, and the strong epistemic closure principle, to ground certain epistemic practices without being manifest within them. According to Stroud, the epistemologist should see people’s knowledge-ascriptions as “restricted in certain ways relative to what, with detachment, we can all recognize to be the full conditions of their truth” (Stroud 1984, 76). This is what the plane-spotter case is supposed to bring out.

The second claim is complex. The “platitudes” in question are platitudes about an “objective world” (Stroud 1984, 77). The idea seems to be this: just as the question of whether “*p*” is true or false is independent of what human beings know, believe, or think, so too the question of whether “*S* knows that *p*” is true or false is independent of what human beings know, believe, or think:

[...] with respect to knowledge and other human institutions, we have the same conception of objectivity. We want to know whether it is objectively true that somebody knows or has good reason to believe that sitting in a draught contributes to catching cold. In saying that the causes of cancer are still unknown we take ourselves to be making a

statement about the present state of human knowledge, and we think human knowledge is in whatever state it is in with respect to the causes of cancer quite independently of our not knowing or being in a position to reasonably assert that it is (Stroud 1984, 78-79).

Now let's return to the plane-spotter cases. Stroud thinks that if we told the careful plane-spotter about the existence of F^* 's, he would realise that "he had not been in the best position for determining whether he knows or for explaining how his knowledge is possible" (Stroud 1984, 81). Stroud thinks that learning about F^* 's would "give him a more objective understanding of his position" (Stroud *ibid*).

So now consider the contextualist's response. On their view, the truth-conditions for knowledge-ascriptions are sensitive to shifts in the attributor's context of epistemic evaluation. Their diagnosis of the plane-spotter case might be that once F^* 's are raised, the careful plane-spotter goes from knowing the plane is an F to not knowing that it's an F because he doesn't know it's not an F^* . Stroud, on the other hand, seems to think that the contextualist is ignoring the significance of acquiring a "more objective understanding" of one's epistemic position (Stroud *ibid*). The idea seems to be that the careful plane-spotter, in being made aware of F^* 's, has now acquire are "more objective understanding" of his epistemic position than he had before. The careful plane-spotter can now recognize that his previous previous knowledge-ascriptions were defective (Stroud *ibid*; cf. Stroud 1984, 76). In short, the careful plane-spotter has reached a better epistemic position for assessing knowledge-ascriptions (cf. Stroud 1984, 82).

Stroud thinks that the epistemological case runs parallel to the plane-spotter case. The epistemologist wants to understand how knowledge of the world is possible in general, and the suggestion is that:

It is a quest for an objective or detached understanding and explanation of the position we are objectively in. What is seen to be true from a detached 'external' standpoint might not correspond to what we take to be the truth about our position when we consider it 'internally, from within the practical contexts which give our words their social point. [...] I think the source of the philosophical problem

of the external world lies somewhere within just such as conception of an objective world or in our desire, expressed in terms of that conception, to gain a certain kind of understanding of our relation to the world (Stroud 1984, 81-82).

So, the suggestion is that the epistemologist comes to see that Descartes's Condition is true, and that it's just a purer or "more objective" expression of the principles that we follow in quotidian contexts.⁸⁶

The sort of dependence here however is weak. It's not that the epistemologist's goal for a more general, objective understanding of our epistemic position with respect to the world around us implies that the rejection of Descartes's Condition and the strong closure principle requires the rejection of that conception. This is consistent with granting to Stroud that that concept does depend on "platitudes". What is in question is whether the rejection of the relevant epistemological requirements entail the denial of those apparent "platitudes". It's hard to see how that could be possible. In the next section, I argue that it's implausible that this picture of what explains the difference between the practical and philosophical context provides adequate support for the thought that the rejection of Descartes's Condition implies the rejection of "platitudes we would all accept" (Stroud *ibid*).

5.3 Response to Stroud's Argument

In this section, I consider Johnston's "minimalism" about the justification of certain epistemic practices. Johnston distinguishes between what justifies our practices and the philosophical pictures of the facts about what justifies our practices. The minimalist then claims that what justifies our practices need not depend on the truth of the philosophical pictures of the facts about what justifies our practices.

To see minimalism in action, consider freedom of the will. We participate in practices where we attribute freedom and responsibility to others and

⁸⁶ Stroud tells us that adopting the principle would "give him a more objective understanding of his position" (Stroud 1984, 82).

ourselves. So, it might be that participants, with philosophical prompting, will tend to think of free action as an action that has its cause in the agent alone, rather than events in the remote past leading up to what the agent does. Furthermore, it might be that certain facts could undermine the whole practice (e.g., if it turned out hidden machines were controlling all of our actions). But according to Johnston:

[...] in the absence of such discoveries, the mere observation that freedom does not consist in the agent's uncaused initiation of action is not itself a criticism of the practice (Johnston 1992, 591).

The reason it need not be a criticism of the practice is that it's not obvious that what legitimates the practice are the facts about agent causation. So, while it might be true that facts about freedom of the will are what legitimates the practice, that leaves it open that facts about agent causation are what legitimates the practice. Johnston's point is that we can't read off from the practitioners' philosophical prompting—that practitioners tend to think of free action as agent causation—that facts about agent causation are what legitimate the practice. As he puts the point:

That the practice of attributing responsibility depends for its justification on facts about free agency—and that ordinary practitioners, under the right sort of philosophical urging, picture those facts as involving uncaused causings—does not settle it that the practice of attributing responsibility depends for its justification on facts about uncaused causings. The picture of free willing as uncaused causing may have only a minimal role. It may yet be epiphenomenal to the practice, and be redundant when it comes to justifying the practice (Johnston 1992, 592).

Minimalism has ramifications for our assessment of Stroud's argument. To see this, consider how Stroud argues for the thesis that Descartes's Condition is part of our ordinary epistemic practices. In general, the minimalist maintains that:

The metaphysical pictures associated with [*ordinary*] concepts do not represent central beliefs of the users of those concepts which guide the users in applying the concepts. [*Minimalism*] has it that the metaphysical pictures are philosophical epiphenomena. Metaphysical pictures, although they emerge from the experience of ordinary

concept users, do not guide ordinary practitioners in their everyday applications of the concepts and so do not represent the sort of central beliefs whose falsity would deprive the concepts of everyday application (Johnston 1993, 110).

We can adapt a version of the minimalist argument to the epistemological case. First, recall that Stroud thinks that the strong closure principle grounds certain epistemic practices, and that Descartes's Condition is just an instance of this principle.

The reason Stroud thinks that the strong closure principle grounds some of our epistemic practices is that that principle seems to be the common denominator in various sorts of cases. That is, it seems to be the principle which is guiding our intuitive verdicts about those cases. We can adapt these cases from one of Stroud's cases:

Suppose looking out the window I announce causally that there is a goldfinch in the garden. If I am asked how I know it is a goldfinch and I reply that it is yellow, we all recognize that in the normal case that is not enough for knowledge. 'For all you've said so far,' it might be replied, 'the thing could be a canary, so how do you know it's a goldfinch?'. A certain possibility compatible with everything I have said so far has been raised, and if what I have said so far is all I have got to go on and I don't know the thing is not a canary, then I don't know that there is a goldfinch in the garden. I must be able to rule out the possibility that it is a canary if I am to know that it is a goldfinch (Stroud 1984, 24-25).

Call this case "goldfinch". As Stroud acknowledges, in this case the proposition in question, that the bird is a goldfinch, is incompatible with the bird being a canary. So, the hypothesis is that the principle guiding our intuitions here would be the following closure principle:

Closure: If S knows that p , and S knows that p implies $\sim q$, then S knows that $\sim q$.

One problem is that Stroud doesn't think that the closure principle could be responsible for our verdicts in the following similar case:

If I come to suspect that all the witnesses have conspired and made up a story about the man's being in Cleveland that night, for example, and their testimony is all I have to go on in believing that he

was in Cleveland, I might find that I no longer know whether he was there or not until I have some reason to rule out my suspicion. If their testimony were all invented, I would not know that the man was in Cleveland (Stroud 1984, 25).

In this case, call it “Cleveland”, Stroud thinks that the person’s knowing that the man is in Cleveland is undermined, but not because the proposition that the man is in Cleveland is incompatible with the testimony being true. Instead, it’s incompatible with the person knowing that the man is in Cleveland. For this reason, Stroud thinks that it’s the strong closure principle which explains why our (putative) intuition is that he doesn’t know that the man was in Cleveland unless he knows that the witnesses’ testimony wasn’t invented. Stroud’s general idea is that the strong closure principle seems to be guiding our epistemic judgements.

However, why think that an epistemological principle is guiding ordinary practitioners in their everyday applications of epistemic concepts, and are not simply “epistemological epiphenomena”? This is the epistemological version of Johnston’s minimalist argument. There might be “epistemological pictures” associated with our ordinary epistemic concepts like knowledge, having good reasons, and evidence, but these pictures do not guide the everyday user’s application of those concepts.

This weighs in against Stroud’s thesis that what grounds Descartes’s Condition are “platitudes we would all accept” (Stroud 1984, 82). The first thought from minimalism is that the strong closure principle need not be what is guiding epistemic agents in their assessments of knowledge-claims, even in the cases that Stroud raises in favour of that. The second thought from minimalism is that Stroud’s conception of the epistemologist’s goal of acquiring a “detached ‘external’ standpoint” from which to assess our knowledge-ascriptions does not show that it depends on “platitudes”. This is a certain philosophical picture of the relationship between the epistemologist’s goal and the so-called “platitudes”. Since those philosophical pictures do not entail the so-called “platitudes”, it’s hard to see how the rejection of Descartes’s

Condition and the strong closure principle would result in the rejection of the “platitudes” (Stroud *ibid*).

5.4 Hinge Propositions

Does Descartes’s Condition, and the general epistemic principle for which it is an instance, express a hinge proposition?⁸⁷ The answer to this question turns on the properties of hinge propositions. In this section, I will argue while it might be plausible that if Descartes’s Condition expressed or depended on a hinge proposition, then it would be plausible that its falsity is philosophically unsatisfying. However, I argue that it’s hard to see how it expresses or depends on a hinge proposition.

5.4.1 Hinge Propositions and their Properties

Since Wittgenstein’s remarks do not form a definitive proposal on the nature of hinge propositions, we will need to look at some of the interpretations of Wittgenstein on hinge propositions. The two most basic interpretations are the epistemic and the non-epistemic conception of hinge propositions.⁸⁸

The epistemic conception is the thesis that hinge propositions can be known, while the non-epistemic conception is the thesis that hinge propositions cannot be known. It does not matter for our purposes which interpretation of hinge propositions is true, because the argument works on either conception.

Consider the non-epistemic conception of hinge-propositions. Let’s suppose that the principle which grounds Descartes’s Condition is a hinge proposition. If the non-epistemic conception of hinge-propositions is right, would this help to explain why its falsity ought to be philosophically unsatisfying? On the non-epistemic conception, the reason that hinge propositions cannot be known is either that hinge propositions are in fact not

⁸⁷ Cassam suggested to me that perhaps Stroud thinks that the epistemic principle which grounds Descartes’s Condition is a hinge proposition.

⁸⁸ For example, see Coliva (2010) and Pritchard (2005).

propositions all, or hinge propositions are non-factual norms and therefore not in the market for being true or false.⁸⁹

On either construal, however, this implies that the relevant epistemic principle cannot be known to be true. This implies the following: on the non-epistemic conception of hinge propositions, if Descartes's Condition and the strong closure principle are or depend upon hinge propositions, then these principles cannot be known to be true. The suggestion now is that its possession of this feature is what explains the truth of the second premise of Stroud's puzzle.

The problem, however, is that it's not clear how their possession of this feature would provide adequate support for the second premise of Stroud's puzzle. The principle cannot be known to be true or false; but how does that fact make its falsity philosophically unsatisfying any more than its truth philosophically unsatisfying? What we want to explain is the former, not the latter. But their possession of this feature does not help to explain it.

Let's suppose, then, that the epistemic conception of hinge propositions is correct. On this view, if Descartes's Condition and the strong closure principle are or depend on hinge propositions, then these principles can be known to be true or false. Is it plausible that Descartes's Condition or the strong closure principle are hinge propositions? That depends on what hinge propositions are. Consider Wittgenstein on the properties of hinge propositions:

At the foundation of well-founded belief lies belief that is not founded (OC §253).

[...] the *questions* that we raise and our *doubts* depend on the fact that some propositions are exempt from doubt, are as it were like hinges on which those turn (OC §341).

That is to say, it belongs to the logic of our scientific investigations that certain things are *in deed* not doubted. (OC §341-342).

⁸⁹ Pritchard presents a third option in his *forthcoming* "Epistemic Angst", that the hinge- propositions can be believed but are not "knowledge-apt", where this means not in the market for knowledge. This third option, however, does not affect the argument in the main text.

A candidate hinge-proposition that Wittgenstein discusses is that *the Earth existed before we did*. Wittgenstein tells us that this sort of proposition “forms the basis of action” and of “thought” because it belongs to their “foundations” (OC §411).

At least three important points about hinge propositions can be extracted from Wittgenstein’s remarks. The first point is that hinge propositions are “exempt from doubt” (OC §341). The reason for this is that hinge propositions belong to the “logic” of our inquiries (OC §341-342). That is, the fact that hinge propositions are not doubted (or ought not to be doubted) is what makes our inquiries possible.⁹⁰

A second point is that hinge propositions lack evidential support and fail to provide evidential support. These propositions don’t confer or receive evidence. Their epistemological role in our inquiries is some-how more basic than conferring or receiving evidence from within them.⁹¹

A third point is that a proposition’s being a hinge proposition is a contingent matter. A proposition can go from functioning as a hinge proposition to not functioning as a hinge proposition over time. This aspect of hinge propositions is brought out in the following quotation:

The propositions describing this world picture might be part of a kind of mythology. And their role is like that of rules of a game; and the game can be learned purely practically, without learning any explicit rules. It might be imagined that some propositions, of the form of empirical propositions, were hardened and functioned as channels for such empirical propositions as were not hardened but fluid; and that this relation altered with time, in that fluid propositions hardened and hard ones became fluid. The mythology may change back into a state of flux, the river-bed of thoughts may shift. But I distinguish between the movements of the waters on the river-bed and the shift of the bed

⁹⁰ For example, Wittgenstein writes: “If you tried to doubt everything you would not get as far as doubting anything. The game of doubting itself presupposes certainty” (OC §114–15).

⁹¹ This, at least, is a widespread reading of Wittgenstein, and at least part of this reading is extract from quotations like the following: “My having two hands is, in normal circumstances, as certain as anything that I could produce in evidence for it. That is why I am not in a position to take the sight of my hand as evidence for it” (Wittgenstein 1969, §250). See Pritchard (2011) for discussion, especially pp. 527-528.

itself; though there is not a sharp division of the one from the other (OC §95-97).

There are two questions we need to address here. The first question is whether Descartes's Condition and the strong closure principle have the properties of hinge propositions. The second question is whether, even if Descartes's Condition and the strong closure principle were hinge propositions, this would be sufficient to render their falsity philosophically unsatisfying.

Here, one might respond to the first question along the following lines. Hinge propositions look like general categorical propositions about the world (e.g., *the Earth existed before our births; the senses are reliable; human beings have hands; we are not made of sawdust; etc.*). Wittgenstein's claim is that these propositions "form the foundation of all operating with thoughts" (OC §401). So, we might think that Descartes's Condition and the strong closure principle aren't hinge propositions because hinge propositions are categorical rather than conditional.

However, while Wittgenstein's examples of hinge propositions are all categorical propositions, this alone does not mean that conditional propositions could not be hinge propositions. So, we will set the former objection to one side and consider a different response.

Here's a different response. One might think that Descartes's Condition and the strong closure principles don't look like pieces of common sense, and that hinge propositions do look like pieces of common sense.⁹² The thought here is how could those principles be hinge propositions if those principles are contentious, theoretical principles?

The response here, however, would be that although Descartes's Condition and the strong closure principle are not pieces of common sense, these principles follow from or depend upon pieces of common sense (cf. Stroud on "platitudes"; see §5), and those are the hinge propositions. But as we saw this in

⁹² Cf. Pritchard (2011) "Hinges, recall, are more akin to 'solidified' chunks of common sense, rather than contentious philosophical and theoretical theses" (Pritchard 2011, 539).

§5, it's implausible that those principles "depend on" or get support from pieces of common sense.

Let's now turn to the second question. The response to the second question grants that Descartes's Condition and the strong closure principle are hinge propositions, but disputes whether this fact would render their falsity philosophically unsatisfying.

On the one hand, one might reply that hinge propositions function as hinges only contingently. After all, whether a proposition is a hinge proposition is a contingent matter. So, one might think that even if Descartes's Condition and the strong closure principle are hinge propositions, this wouldn't be sufficient to render their falsity philosophically unsatisfying, because if ever they failed to function as hinge propositions, this property would thereby no longer explain why their falsity is philosophically unsatisfying. Put another way, this property would explain why their falsity is philosophically unsatisfying only if they couldn't fail to be hinge propositions. But that's precisely what they can fail to be.⁹³ We are therefore left none the wiser as to why the principle's being a hinge proposition implies that its falsity is (or ought to be) philosophically unsatisfying.

5.5 Deep epistemic grammar

One reading of a hinge-proposition is something close to what the generative linguists call "the principles and parameters" of universal grammar. These are abstract linguistic principles that are part of the subject's "deep" grammar, the

⁹³ A related observation from Reed (2011) is that "it seems clear that the so-called hinge propositions are ones that we are psychologically incapable of calling into question. This is, of course, compatible with their being false" (Reed 2011). If hinge propositions are incapable psychologically of being doubted and Descartes's Condition and the strong closure principle are hinge propositions, then two problems arise. First, since those principles are plausibly doubttable, they counts against them being hinge propositions after all. Second, if a hinge proposition can be false, why would their falsity be a reason for philosophical dissatisfaction? It would seem to be a reason for philosophical dissatisfaction only if they couldn't be revised, eliminated, or improved upon. But there's no reason to think that this cannot be done.

grammar which guides the subject's construction of sentences, but is not accessible to them through introspection.⁹⁴

Perhaps the strong closure principles help set the parameters of epistemic evaluation. For Stroud (1984), it is a fact about our epistemic practices because we all can and often do adhere to the principle in our individual epistemic evaluations. On this reading of Stroud, the strong closure principle (the principle that underwrites Descartes's Condition) has a transcendental status (cf. Strawson 40, 2000). What does it mean for it to have a transcendental status?

For the principle to have a transcendental status is for its truth to be a necessary condition for the possibility of certain sorts of cognitive activities. For example, on a modest construal, it might be that one must believe that it is true in order to reason in accordance with the principle. However, on a more ambitious construal, it might be that the principle must be true in order to reason in accordance with the principle.⁹⁵

Now the idea here is that the relevant epistemic principles *P* (Descartes's Condition and the strong closure principle) are part of our deep epistemic grammar, the epistemic analogue of our deep linguistic grammar. If this is right, it's plausible that the second premise of Stroud's puzzle would receive adequate support. After all, if it's part of our deep epistemic grammar, it's hard to see how *P* could be false. And even if it were false, then there would be the issue that we are irrational. We would be irrational because *P* would be impossible to change, because *P* is a part of our deep epistemic grammar, but *P*

⁹⁴ For example, see Chomsky (1957, 1986, 1995). See Devitt (2006) for discussion of the philosophical implications of transformational grammar.

⁹⁵ Transcendental arguments are often used for anti-sceptical purposes. Perhaps the most famous transcendental argument is Kant's (1781/1997). See Cassam (2003), Guyer (2008), Strawson (1985/2008) and (1975). The most recent versions include Putnam (1981) and Davidson (2006), from content externalist principles. For a transcendental argument from disjunctivism about perceptual experience, see McDowell (1998). Stroud (1968) famously argued that anti-sceptical transcendental arguments cannot be successful. For a discussion of these transcendental arguments, see Bridges (2006), Brueckner (1999), Sosa (2003), Pritchard (2013), Pritchard and Ranalli (*forthcoming*), and Ranalli (2013). For discussion of transcendental arguments more generally, see Stern (1999, 2000). For a discussion of modest transcendental arguments, see Brueckner (1996), and Stern (1997).

would also be something that is in error. This would render revision or elimination of *P* impossible, even if we learned that *P* is erroneous. It's plausible that that would be a good reason for philosophical dissatisfaction.

However, one might wonder where epistemologists are in a position to know or discover that *P* is part of deep epistemic grammar. After all, this looks like a straightforward empirical question for cognitive scientists. On this view, we might wonder what use Stroud's examples (e.g. the goldfinch case and the Cleveland case) reveal about the strong closure principle. It's hard to see how those examples provide adequate support for thinking that its part of our deep epistemic grammar. What else would be needed here is an extensive "epistemic corpus" on the judgments of speakers about epistemic cases. We would also need an experiment which gauges the speaker's epistemic capacities and solicits the right kind of intuitions. This is not an argument against Stroud's examples illiciting the intuition that the strong closure principle is plausible. It is also not itself an argument against Stroud's examples providing some prima facie support for the strong closure principle being the principle which guides our intuitions about his cases. Instead, the argument here is that that fails to establish that it is part of our deep epistemic grammar.

6 Conclusion

In this chapter, I focused on the arguments for the second premise of Stroud's puzzle. Stroud's (1984) argument for this premise is what I called an "anti-revisionist" argument. It is an anti-revisionist argument because Stroud maintains that the rejection of Descartes's Condition and the epistemic principle which underlies it entails the rejection of a "simple and obvious fact about knowledge", a "general procedure we recognize and insist on in making and assessing knowledge-claims", and the rejection of "platitudes we would all accept" (Stroud 1984, 27, 30-31, 82).

I argued that Stroud's anti-revisionist arguments fail to provide adequate support for the second premise of Stroud's puzzle. I then considered some additional arguments that anti-revisionists might appeal to in order to provide adequate support for the second premise of Stroud's puzzle, and I argued that none of them were sufficient either.

I then considered the following general objection. According to this objection, while it might be right that Descartes's Condition is grounded in our ordinary epistemic practices, it's not the case that this property is the relevant property which renders its rejection philosophically unsatisfying. Instead, it has some other property which renders its rejection philosophical unsatisfying.

However, I argued in each case that either it's not plausible that it has the relevant property in question, or it's not plausible that its possession of that property renders its rejection philosophically unsatisfying.

Chapter 3

Stroud's Puzzle: Conditional Scepticism

1 Introduction

In Chapter 2, I examined the second premise of Stroud's puzzle. I argued that none of the arguments provided adequate support for this premise. If that's true, then Stroud's puzzle fails to provide adequate support for meta-epistemological scepticism, because one of its premises lacks adequate support.

However, one might think that even if the first premise of Stroud's puzzle is true, we still have a problem. The problem is that the truth of Descartes's Condition would guarantee that it cannot be satisfied, so that scepticism is true.⁹⁶ This would not be as bad as being faced with Stroud's puzzle, but it would be bad if Descartes's Condition had compelling philosophical arguments in its favour. Moreover, there's no reason to let the case with Stroud's puzzle rest on our evaluation of the second premise alone. If both premises can be shown to lack adequate argumentative support, then so much the worse for Stroud's puzzle.

In this chapter, I examine the first premise of Stroud's puzzle. According to this premise, if Descartes's Condition is true, then it cannot be satisfied, and therefore scepticism is true. But scepticism is not philosophically satisfying (cf. Chapter 1, §2). I argue that none of the arguments provide adequate support for the first premise of Stroud's puzzle.

⁹⁶ See Chapter 1, §5.2.

In section 2, I review two versions of Stroud’s argument for this premise of the puzzle. The first argument is a regress argument. According to this version of the argument, if Descartes’s Condition is true, it cannot be satisfied because it leads to a regress (§2.1). I argue that this kind of regress, however, is unthreatening. The second argument is an argument from the impossibility of there being a test or way of coming to know that Descartes’s Condition is satisfied, to the conclusion that it cannot be satisfied. I call this the “no test” argument (§2.3). I argue that the “no test” argument is invalid (§3.1).

The third argument takes it that Descartes’s Condition is not a necessary condition but a necessary pre-condition on knowledge of non-psychological physical reality. I argue that even if this renders the “no test” argument valid, it’s not plausible that it’s sound because it’s not plausible that this version of Descartes’s Condition should be accepted (§3.2).

The fourth argument is from Stroud’s recent (2009). Stroud presents another version of the “no test” argument which neither appears to be invalid nor appears to express or otherwise depend on Descartes’s Condition understood as a necessary pre-condition. In section 4, then, I examine Stroud’s (2009) version of the “no test” argument. I present several plausible interpretations of the argument, and I argue that none of them are successful (§5).

The conclusion, then, is that Stroud’s puzzle fails to provide adequate support for meta-epistemological scepticism. As I argued in Chapter 2, the falsity of Descartes’s Condition is not a philosophically unsatisfying result. On the other hand, if Descartes’s Condition is true, it would not follow that it could not be satisfied, preventing the consequence that scepticism is true.

2 Stroud’s Arguments

In his (1984), Stroud tells us that:

I ascribed to Descartes the quite general thesis that knowing that one is not dreaming is a condition of knowing something about the world

around us on the basis of the senses. [...] If that really is a condition of knowing something about the world, I think it can be shown that Descartes is right in holding that it can never be fulfilled (Stroud 1984, 20).

In effect, this expresses the first premise of Stroud's puzzle: if Descartes's Condition is true, it could not be satisfied, having scepticism as a logical consequence. In Chapter 1, I presented a rough expression of Stroud's argument for this premise of Stroud's puzzle (see Chapter 1, §5). What I want to do here is examine in more detail Stroud's arguments for this premise of the puzzle.

In order to understand Stroud's arguments, we first need to understand how Descartes's Condition ought to be formulated. Stroud formulates this condition as a material conditional:

Descartes's Condition: For all S , and propositions about non-psychological, physical reality p , if S knows that p , then S knows that they are not merely dreaming that p .⁹⁷

For ease of exposition in this section, let's replace "merely dreaming that p " with the following dream operator, D , which takes propositions about non-psychological, physical reality as its complement. Following Stroud (1984), the logic of the D operator is that, from $D(p)$ indexed to a time t , it does not follow that p . However, $D(p)$ is compatible with p . It just doesn't entail that p . Since knowledge is factive, $D(p)$ does not entail that S knows that p .⁹⁸

Another property of D is that "anything which can be going on or that one can experience in one's waking life can also be dreamt" (Stroud 1984, 18). I will express this property of D as follows: if S can perceive that p , then S can dream that p .

⁹⁷ For formulations of the principle as a material conditional, expressing a necessary condition on knowledge of the world by means of the senses, see Stroud (1984), pp. 20, 23, 43, 48, 63, and 71. As we will see later in this chapter, Stroud (2009) modifies this principle.

⁹⁸ See Stroud (1984, 15-16). See also Wright (1991, 91).

2.1 The Regress Argument

As Stroud tells us, Descartes's Condition is a completely general epistemic requirement on knowledge of non-psychological, physical reality. It applies to any proposition about non-psychological, physical reality. And it is the generality of the requirement which generates a regress:

Descartes's conclusion [that scepticism is true] rests on the general requirement that we must know that we are not dreaming if we are to know anything at all about the world around us. That requirement is what renders inadequate any tests or procedures for determining that one is not dreaming; one would have to know that one was not simply dreaming that one was performing the test, and not dreaming that one was performing any of the other tests used to determine that, and so on (Stroud 1984, 48 my addition).

Notice here that Stroud mentions that one could attempt to use a "test" or some "procedure" in order to fulfil Descartes's Condition. But the problem is not with the test or procedure per se, but the fact that Descartes's Condition would then need to be applied to the propositions describing those tests or procedures.

This point is missed in the literature on Stroud's conditional sceptical argument.⁹⁹ One diagnosis of this is the failure to distinguish "dreaming that...", as an epistemic operator, from the proposition I am not dreaming, as a particular proposition that one must know does not obtain. For example, Descartes's Condition, as Stroud intends to use it at least, is better represented as containing two epistemic operators: knowledge (or "K") and dreaming (or "D").

In order to generate the regress, we need to embed the dreaming operator *D* within another instance of the principle. For example:

If *S* knows that *p*, then *S* knows that $\neg D(\neg D(p))$

Is this move valid? Stroud appears to think so. After all, the principle applies to any proposition about non-psychological, physical reality. The proposition that

⁹⁹ See, for example, Koethe (2005, 19). See Heil (1986) for an exception, esp. pp. 3-5 for an analysis of the regress argument.

I am not dreaming that (...) is therefore within the scope of the principle. The regress argument, then, should take the following form:

- (1) For all S , and propositions about non-psychological, physical reality, p , if S knows that p , then S knows that $\neg D(p)$.
- (2) If p is a proposition about non-psychological physical reality, and S is not dreaming that p is a proposition about non-psychological, physical reality, then any iteration of S is not dreaming that (...) is itself a proposition about non-psychological, physical reality.

Now we ask the question: can S come to know that they are not merely dreaming that p ? From (1) and (2), it follows that:

- (3) If S knows that $\neg D(p)$, then S knows that $\neg D(\neg D(p))$.

So now we ask the question: can S fulfil this necessary condition, embedded in the consequent of (3)? That is, can S come to know that they are not merely dreaming that they are not merely dreaming that p ? From (1), (2), and (3), it follows that:

- (4) If S knows that $\neg D(\neg D(p))$, then S knows that $\neg D(\neg D(\neg D(p)))$.

And so on, ad infinitum. This is what leads Stroud to say that:

If it is in general a necessary condition of our knowing something about the world around us that we know that we are not dreaming, it follows that we can never know that we are not dreaming (Stroud 1984, 43).

It is important to note that this version of Stroud's argument does not turn on specific features of the dreaming hypothesis as opposed to the evil-demon hypothesis. Here is another version of the same kind of argument, but instead of using dreaming as the specific sceptical hypothesis, Stroud uses the evil-demon hypothesis:

If, for everything I claim to know (or everything about an ‘external world’), it is true that in order to know it I must at least know that there is no evil scientist artificially inducing in me the beliefs on the basis of which I claim to know it, then, in particular, in order to know that no evil scientist is doing that to me, I must know that there is no evil scientist artificially inducing in me the beliefs that lead me to believe that no evil scientist is doing that to me. And then there is nowhere to stop (Stroud 2000, 47).

Prima facie at least, this is a good argument. The idea behind the argument is simple. Descartes’s Condition is a general epistemic principle. It applies even to the proposition that one is dreaming that something is so. One application of the principle on the proposition that you know that you are not dreaming that p , then, yields the result that if you know that you are not dreaming that p , then you know that you are not dreaming that you are not dreaming that p . And repeated applications of the principle generates a regress. The regress is supposed to be a regress of so far unfulfilled instances of the principle, each of which must be fulfilled for the preceding instance to be fulfilled.

2.2 Stroud’s Regress Argument: Evaluation

It’s plausible that this argument does show that there is a genuine regress of instances of Descartes’s Condition. But not all regresses are vicious. Some are benign. What we need to ask is whether the possibility of a regress of instances of Descartes’s Condition shows that those instances are unfulfillable.

Consider the truth predicate. Most philosophers accept the following disquotation principle for truth: p is true if and only if p . Moreover, most philosophers accept that truth is reflexive: if p is true, then it’s true that p is true. So now suppose that p . From disquotation, it follows that p is true. And from reflexivity, it follows that it is true that p is true. Repeating reflexivity yields: it is true that it is true that p is true, and so on, ad infinitum. This is a regress, but it is benign. From the truth of p , there is an infinite series of true propositions, each of which embeds the previous instance of ‘*is true that* [...]’.

Now let’s turn to Stroud’s regress argument. The regress argument can be put like this. First, we apply the principle that if $K(p)$, then $K(\neg D(p))$ —

Descartes's Condition—to the proposition that $\neg D(p)$ represents. We do this because knowing that $\neg D(p)$ is a necessary condition for knowing that any proposition about non-psychological physical reality is true. For this reason, it applies to proposition that I know that I am not dreaming that p , where p is also any proposition about non-psychological physical reality. This application of Descartes's Condition then yields the following conditional: if $K(\neg D(p))$, then $K(\neg D(\neg D(p)))$. Now since $\neg D(\neg D(p))$ is a proposition that falls within the scope of Descartes's Condition as well, it follows that if $K(\neg D(\neg D(p)))$, then $K(\neg D(\neg D(\neg D(p))))$, and so on, ad infinitum. From the truth of the proposition that if $K(\neg D(p))$, then $K(\neg D(\neg D(p)))$, and an application of Descartes's Condition on the proposition that $\neg D(p)$, we can generate an infinite series of instances of Descartes's Condition. Each new instance will take the previous instance of $\neg D(\dots)$ and embed it within its antecedent. This is a regress, but it's no more vicious than the previous regress with the truth predicate.¹⁰⁰

2.3 The No Test Argument

In the previous section, I explained the regress argument and argued that it is benign. In this section, I consider a different version of the argument, one which resembles the regress argument but features a significant difference. The difference is that Stroud considers whether one could satisfy Descartes's Condition on the basis of a test (or in general some way of knowing that they are not merely dreaming).¹⁰¹ Like the regress argument, Stroud thinks that a

¹⁰⁰ A similar point can be made about the KK principle. On one version of the KK principle, it's that if S knows that p , then S knows that (S knows that p). On another version of the KK principle, it's that if S knows that p , then S *can* know that (S knows that p). The first version seems to lead to a vicious regress, because it is constitutive of knowing that p that one know something else, namely, that one knows that one knows that p . However, the second version arguably does not lead to a vicious regress, even if it leads to a regress of conditional propositions. For example, see Smithies (2014) for this observation.

¹⁰¹ Cf. Wright (1991). Wright's argument appeals to the following principle:

“Proper Execution Principle (PEP):

If the acquisition of warrant to believe a proposition depends on the proper execution of some procedure, then executing the

new test will need to be introduced to show that the previous test was reliable or otherwise knowledge-conducive, so that there is a regress of so far untested tests. I call this the “no test argument”. The argument is complex. I’ll begin with some stage setting from Stroud:

Suppose [...] that Descartes does indeed know (somehow) that there is a test or circumstance or state of affairs that unfailingly indicates that he is not dreaming. Still, there is an obstacle to his ever using that test or state of affairs to tell that he is not dreaming and thereby fulfilling the condition for knowledge of the world. The test would have to be something he could know he had performed successfully, the state of affairs would have to be something he could know obtains. If he completely unwittingly happened to perform the test, or if the state of affairs happened to obtain but he didn’t know that it did, he would be in no better position for telling whether he was dreaming than he would be if he had done nothing or did not even know that there was such a test (Stroud 1984, 22).

Stroud’s first step is to suppose that Descartes knows that there is a test, *T*, the successful performance of which guarantees that he knows that he is not merely dreaming that *P*. However, Stroud suggests that the successful performance of *T* cannot be accidental. In general, the successful performance of *T* cannot be the product of mere luck. Stroud infers from this that Descartes must therefore know that *T* was performed successfully. This prompts Stroud to ask the question:

But how is he to know that the test has been performed successfully or the state of affairs in question does in fact obtain? (Stroud *ibid*)

This is where the second step of Stroud’s no test argument comes in. The second step registers the fact about dreaming that:

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).

For example, if I try to use some test to determine that I’m not now dreaming, then by (PEP), the execution of that test cannot give me a stronger warrant to believe that I’m not now dreaming than I have *independently* for believing that I executed the test properly. Here, the argument directly resembles Stroud’s *no test argument*, since the problem with acquiring an independent warrant is that that “belief cannot be independently warranted at all unless I have independent warrant for its component, that I really did execute the procedure and did not *merely dream its execution*” (Wright 1991, 100).

Anything one can experience in one's waking life can also be dreamt about; it is possible to dream that one has performed a certain test or dream that one has established a certain state of affairs obtains (Stroud *ibid*).

From this, Stroud infers that:

In order to know that his test has been performed or that the state of affairs in question obtains Descartes would therefore have to establish that he is not merely dreaming that he performed the test successfully or that he established that the state of affairs obtains (Stroud 1984, 22).

This is the third step of the argument. The claim is that, because it is possible to merely dream that one has performed *T* successfully (step 2), it follows that if Descartes's knows that he performed *T* successfully, then he knows that he is not merely dreaming that he performed *T* successfully. But then Stroud argues that this leads to a regress:

How could that in turn be known? Obviously the particular test or state of affairs already in question cannot serve as a guarantee of its own authenticity, since it might have been merely dreamt, so some further test or state of affairs would be needed to indicate that the original test was actually performed and not merely dreamt[.] But this further test or state of affairs is subject to the same general condition in turn. Every piece of knowledge that goes beyond one's sensory experience requires that one knows one is not dreaming. This second test or state of affairs will therefore be of use only if Descartes knows that he is not merely dreaming that he is performing or ascertaining it, since merely to dream that he has established the authenticity of the first test is not to have established it. And so on (Stroud 1984, 22-23).

Notice that right before we arrive at the regress of so far unfulfilled instances of Descartes's Condition, Stroud argues for an anti-circularity requirement on knowing that *T* was performed successfully. In short, Descartes cannot use *T* to show that his previous exercise of *T* was successful.¹⁰² From this, Stroud infers that Descartes's would therefore need a new test, *T*₁, the successful performance of which would guarantee that his performance of *T* was successful. But this new test *T*₁ would be subject to the same reasoning, and so

¹⁰² This requirement echoes Fumerton's requirement in his argument against epistemic externalism. Fumerton says that "you cannot *use* perception to justify the reliability of perception! You cannot *use* memory to justify the reliability of memory! You cannot *use* induction to justify the reliability of induction! Such attempts to respond to the skeptic's concerns involve blatant, indeed pathetic, circularity" (Fumerton 2000, 410).

the same set of requirements, leading to the introduction of new tests, T_2, \dots, T_n , each of which Descartes would need to know were performed successfully, and so on, ad infinitum.

To summarize Stroud's no test argument for the conclusion that, if Descartes's Condition is true, then it cannot be satisfied, we have the following five steps:

1. Suppose that S knows that there is a test, T , the successful performance of which guarantees that S is not at the time dreaming. It cannot be a matter of luck that S 's use of T is successful. So, S must know that T was performed successfully.
2. It is possible for S to dream that he has performed T successfully. So:
3. If S knows that he performed T successfully, then he knows that he is not merely dreaming that he performed T successfully.
4. S cannot, on the basis of successfully performing T , come to know that T was performed successfully. So:
5. If S knows that he is not merely dreaming that he performed T successfully, then S knows that there is another, distinct test, T_1 , the successful performance of which guarantees that T was performed successfully.

The no test argument is a regress argument because the new test, T_1 , will be subject to the same reasoning and the same set of requirements as the previous test. These five steps can then be repeated for that new test, and so on, for each test that is introduced in order to authenticate the agent's previous test. If this is right, then it follows from Descartes's Condition that it cannot be satisfied, and therefore that scepticism is true.

3 Stroud's No Test Argument: Evaluation

In the previous section, I explained Stroud's no test argument for the first premise of Stroud's puzzle. In this section, I provide a critical evaluation of this argument. In section 3.1., I argue that if Descartes's Condition imposes only a necessary condition on knowledge of non-psychological, physical reality, then the no test argument doesn't show that it cannot be fulfilled. In section 3.2., I argue, following McDowell (1998), Pryor (2000), and Cassam (2007), that if Descartes's Condition poses a necessary pre-condition on knowledge of non-psychological, physical reality, it's plausible that that form of the condition cannot be fulfilled, but it's not plausible that it has to be fulfilled.

3.1 Descartes's Condition: The Necessary Condition Reading

Stroud's no test argument features one supposition, that the agent knows that there is a test, the successful performance of which guarantees that they are not merely dreaming, and three premises: Descartes's Condition; the premise that, for any proposition about non-psychological, physical reality P , it is possible for one to merely dream that P is true; and the premise that one cannot use the same test to show that one's performance of that test was successful.

Notice that it is not so obvious that Descartes's Condition cannot be satisfied, if Descartes's Condition is understood as imposing only that one's knowing that one is not merely dreaming that P is a necessary condition on knowing P . After all, if a necessary condition on knowing that P is that one know that they are not merely dreaming that P , and fulfilling that condition requires the satisfaction of some test, it does not follow that the test cannot be fulfilled. Instead, it implies the weaker conclusion that if Descartes's Condition is fulfilled, then the test is fulfilled. More premises would be needed to show that it nevertheless could not be fulfilled.

At this juncture, I want to turn to a response. The response concedes the point that only the weaker conclusion follows: that if Descartes's Condition is fulfilled, then the test is fulfilled, so that additional premises are required to supplement the argument for the conclusion that it cannot be fulfilled.

There are two crucial premises needed to show that the condition cannot be fulfilled. The first is the anti-circularity requirement at step 4, that one cannot, on the basis of successfully performing the test T , come to know that T was performed successfully. For example, consider the following reasoning:

- (1) If I know that I am not merely dreaming that P , then there is some way that I know that I am not merely dreaming that P (e.g., a test, T).
- (2) I know that I am not merely dreaming that P on the basis of successfully performing the test T only if I know that I successfully performed T .
- (3) I cannot come to know that I successfully performed the test T on the basis of successfully performing T .

This set of requirements makes it hard to see how one could come to know that one is not merely dreaming that they successfully performed T because one cannot use T to show that one has successfully performed it. One could of course introduce a new test, T_1 , but then Stroud's worries about a regress would arise.

Still, we can draw a distinction between "using" a test T in the sense of appealing to T as a premise in an argument or as a reason for believing that T was performed successfully, and "using" T in the sense that the execution of T is a way of knowing (or coming to know) that T was successfully executed.¹⁰³

¹⁰³ Here's an example. In the first case, I see that S sees that it's raining outside. I then use the fact that "I see that S sees that a raining outside" as a premise to argue that "S sees that it's raining outside". One might think that if seeing that something is so is in question, then I cannot appeal to that fact that I see that S sees that it's raining outside as a reason to

One might think that the former is objectionable while the latter is not. After all, just think of that fact that one ‘uses’ one’s eyes to check one’s eyes in the mirror (or someone else’s eyes). There is a sense in which this circular, but the circularity here is benign. Stroud’s argument seems to depend on the agent having to provide an argument or to reason their way to the conclusion that *T* was performed successfully. If that’s the case, it’s plausible that appeals to *T* would be prohibited. But why think that the proper execution of *T* cannot simply be a way of knowing or coming to know that one is not merely dreaming that *P* (or that one is not merely dreaming that one performed *T* successfully)?

Moreover, while it is true that not being in a position to know that one is not dreaming precludes one from being in a position to acquire perceptual knowledge of the world, this does not mean that acquiring perceptual knowledge of the world precludes one from knowing that they are not merely dreaming. Instead, the latter can be used to sustain one’s knowledge that they are not dreaming. As McDowell (1998) puts it:

One’s knowledge that one is not dreaming, in the relevant sort of situation, owes its credentials as knowledge to the fact that one’s senses are yielding one knowledge of the environment—something that does not happen when one is dreaming (McDowell 1998, 238).

In short, in so far as knowing that one is not dreaming is only necessary for knowing anything about the world around us, our knowledge of the world can sustain our knowledge that we are not merely dreaming. The requirement can be met.

believe that it’s raining outside. However, consider a second case, where I see that S sees that it’s raining outside. Here, it’s plausible that I now know that S sees that it’s raining outside. But my knowledge that S knows that it’s raining outside is sustained by an appeal to that fact that I see that S sees that it’s raining outside as a premise or a reason to believe that S sees that it’s raining outside. Rather, my seeing that S sees that it’s raining outside is a way of knowing that S sees that it’s raining outside. Even if seeing that something is so is in question, this doesn’t prevent it from being a way of knowing, even if it prohibits appeals to this type of seeing as a premise in an argument.

A second problem facing the argument lies with Stroud's epistemic connection between it not being a matter of luck that one performs the test successfully and knowing that the test was performed successfully:

There is an obstacle to his ever using that test or state of affairs to tell that he is not dreaming and thereby fulfilling the condition for knowledge of the world. The test would have to be something he could know he had performed successfully, the state of affairs would have to be something he could know obtains. If he completely unwittingly happened to perform the test, or if the state of affairs happened to obtain but he didn't know that it did, he would be in no position for telling whether he was dreaming than he would be if he had done nothing or did not even know that there was such a test (Stroud 1984, 21-22).

An anti-luck requirement is expressed in this passage. The requirement is that because it can be a matter of luck that S performs the test successfully, S must therefore know that it was performed successfully in order to exclude the possibility of luck.¹⁰⁴

If the performance of the test was in fact not due to luck, it's hard to see why S would also need to know that the test was not due to luck. Suppose that the successful performance of *T* guarantees that one is not dreaming that *p*.¹⁰⁵ It is possible that that performance of *T* was successful, but not because of the intervention of luck. At least, there is no a priori reason to think that this is impossible. If that is right, what makes it plausible that that performance of *T* does not put S in a position to know that they're not dreaming that *p*?

This just expresses an anti-luck requirement on knowledge. If S knows that *p* on the basis of *b*, then S's forming a true belief rather than a false belief cannot be the product of luck. What we need to address is whether this is connected with our having to know that it is not the product of luck. In order to discuss this question, let's consider an example from Unger (1968):

¹⁰⁴ For a detailed overview of the necessity of an anti-luck requirement on knowledge, see Pritchard (2005). Pritchard (2007, 2012) takes an anti-luck requirement to express a platitude of our concept of knowledge. However, see Madison (2011) for an argument against this claim.

¹⁰⁵ This is an assumption that Stroud concedes. See Stroud (1984, 21).

[A] man may overhear his employer say that he will be fired and he may do so quite by accident, not intending to be near his employer's office or to gain any information from his employer. Though it may be an accident that the man came to know that he will be fired, and it may be somewhat accidental that he knows this to be so, nevertheless, from the time that he hears and onwards, it may well be not at all accidental that the man is right about its being the case that he will be fired. Thus, he may know, whether by accident or not (Unger 1968, 159).

In this case, it is a matter of luck that the person acquired their evidence. Had the agent not been there to overhear their co-workers, he wouldn't have acquired that particular evidence. But it is not a matter of luck that, given their evidence, their belief is true rather than false. The former kind of luck is benign, in that its presence is not incompatible with knowledge. If, on the other hand, it was a matter of luck that, given their evidence, their belief is true rather than false, then that kind of luck would prevent them from acquiring knowledge. That kind of luck is incompatible with knowledge.

In the no test argument, Stroud appears to put significant weight on the fact that the agent's successful performance of the test not be due to luck. But it's not clear that the kind of luck at issue in the argument is incompatible with knowledge. To see this, consider again the following passage from Stroud:

If he completely unwittingly happened to perform the test, or if the state of affairs happened to obtain but he didn't know that it did, he would be in no position for telling whether he was dreaming than he would be if he had done nothing or did not even know that there was such a test (Stroud *ibid*).

This statement flags that it might be a matter of luck that the agent performs the test, not that it is a matter of luck that the agent's performance of the test is successful. But this kind of luck is compatible with knowledge. Stroud makes the fallacious inference that, from the fact that the agent's performance of the test might have been due to luck, that therefore the agent would be in no better position for knowing that he is not dreaming that p than if he had failed to perform the test. In fact, as Unger's case brought out, it might put the agent in a position to acquire knowledge.

To see this, let's consider an account of knowledge-preventing epistemic luck. This is what Engel (1992) and Pritchard (2005) call "veritic" epistemic luck: it is a matter of luck that the agent's belief is true rather than false. As Pritchard (2005) explains, an agent's true belief is due to veritic luck if it is:

[...] true in the actual world, but in a wide class of nearby possible worlds in which the relevant initial conditions are the same as in the actual world—and this will mean, in the basic case, that the agent at the very least forms the belief in the same way as in the actual world—the belief is false (Pritchard 2005, 146).

Now let's consider the no test argument again. Stroud concedes that the agent can know of a test, *T*, the successful performance of which guarantees that the agent is not dreaming that *p*. Now if the agent comes to believe that they are not dreaming that *p* on the basis of the successful performance of *T*, then the fact that *T* guarantees that they're not dreaming that *p* makes it the case that their belief is not true because of veritic luck. Why? Because if the agent had performed the same test in a nearby world, their belief would still come out true.

I conclude that Stroud's no test argument does not provide adequate support for the thesis that if Descartes's Condition is true, then it cannot be satisfied.

3.2 Descartes's Condition: The Necessary Pre-condition Reading

Let's consider another interpretation of Descartes's Condition. On this interpretation, the requirement that one know that one is not merely dreaming that *p* has to be known independently of anything that one knows about non-psychological physical reality. That is, Descartes's Condition expresses a necessary pre-condition on knowledge of non-psychological, physical reality. We can formulate Descartes's Condition as follows:

Descartes's Condition*: For any proposition about non-psychological physical reality *P*, if *S* knows that *P*, then *S* knows that they are not

merely dreaming that *P* independently of and antecedently to any knowledge *P*.¹⁰⁶

The terms “independently” and “antecedently” are not clear. So we need say something about these terms if we’re to make progress assessing Stroud’s no test argument. Pryor (2000) provides a useful definition of the notion:

Your justification for believing p_1 is antecedent to your justification for believing p_2 just in case your reasons for believing p_1 do not presuppose or rest on your reasons for believing p_2 . Your reasons for believing p_1 can not beg the question whether p_2 (Pryor 2000, 525).

For example, let’s suppose that *S* believes that *God exists*, and their justification for believing that God exists is that *the Bible is reliable* (and it says that God exists), and that *God authored the Bible*. In this case, *S*’s justification for believing that *God exists* presupposes or rests on *S*’s justification for believing that *the Bible is reliable*.

Now consider the proponent of Descartes’s Condition*. They will say that the principle prohibits one from knowing that one is not merely dreaming that *P* on the basis of *b* if *b* implies or presupposes any knowledge of non-psychological physical reality. In this fashion, one’s knowledge that one is not merely dreaming that *P* is independent of anything one knows about non-psychological physical reality because that knowledge doesn’t imply or presupposes any knowledge of that type.

¹⁰⁶ Cf. Pryor (2000): “If you’re ever to know anything about the external world on the basis of your perceptual experiences, then you have to be in a position to antecedently know you’re not then being deceived by an evil demon” (Pryor 2000, 528). Pryor takes this premise to be essential to the sceptical argument. Compare with Williams (1996, 85), McDowell (1998, 225), and Cassam (2007, 31). Cassam (2007), Pryor (2000), McDowell (1998), and Williams (1996) all reject the principle. Wright (1991), however, maintains that this version of Descartes’s Condition is true, but that it *can* be fulfilled by having a default, *a priori* entitlement to believe that we are not merely dreaming that *P* which is sufficient for knowing that we’re not. This is related to the “liberalism” vs. “conservatism” debate about the relationship between perceptual knowledge and sceptical hypotheses (or defeaters more generally). Pryor (2000) and Huemer (2013, 2007, 2000) are “liberals”. This means that they deny that perceptual knowledge requires one antecedently knows the denial of sceptical hypotheses. Wright (2004) and Davies & Wright (2004) defend conservatism, the denial of liberalism. For discussion, see Neta (2010). For different arguments against Pryor and Huemer’s form of liberalism, see Siegel (2012).

How would the no test argument run with Descartes's Condition* in play? First we suppose that if S knows that they are not merely dreaming that *P*, then S knows it on some epistemic basis *b*. Descartes's Condition prohibits *b* from resting on or presupposing any knowledge of non-psychological, physical reality.

But how are we going to get from this claim to the claim that one cannot know that they are not merely dreaming that *P*? Here's one way. Suppose that the only way *b* could provide a way for one to know that they are not merely dreaming that *P* is that *b* is an argument whose premises we know, but this knowledge, in accordance with Descartes's Condition*, neither implies nor presuppose any knowledge of non-psychological physical reality. If that were the only way of coming to know that we are not merely dreaming that *P*, Stroud would have a plausible case for why Descartes's Condition cannot be fulfilled.¹⁰⁷

Here's another way. Suppose we have a test *T*, the proper execution of which guarantees that we're not merely dreaming that *P*. Now Descartes's Condition then prohibits one from knowing that there is some such test *T* because this knowledge would presuppose knowledge of non-psychological physical reality. From this, it would follow there is no test that one can use in order to know that they are not merely dreaming that *P* because Descartes's Condition prohibits one from knowing that there is some such test.

Even if we ignore this issue, recall that Stroud makes the following assumption in his no test argument: the proper execution of *T* cannot yield knowledge unless one knows that they have properly executed *T*.¹⁰⁸ I've argued

¹⁰⁷ I use term "plausible" because almost all epistemologists think that from premises which cannot entail or presuppose any knowledge of non-psychological physical reality, it's hard see how any such premises would put the agent in a position to know, for example, that they are non-merely dreaming. Still, some epistemologist's think that there are arguments from premises about sensory experience and *a priori* constraints on good explanation which *do* put us in a position to know that, say, we're not merely dreaming everything. For example, see Vogel (1990) and (2014) for a defence of abductive arguments against scepticism.

¹⁰⁸ For example, see Stroud (1984, 22).

that this assumption is negotiable (see §3). The problem now is that, if we don't reject that requirement, Descartes's Condition makes the fulfilment of this requirement impossible, because knowing that one has properly executed *T* would imply or presuppose knowledge of non-psychological, physical reality. And this follows because its proper execution implies that one is not merely dreaming that *P*.

So, Stroud's no test argument appears, prima facie at least, to be both valid and sound if Descartes's Condition is understood as expressing a necessary precondition on knowledge of non-psychological physical reality (i.e., as Descartes's Condition*). However, we can dispute the condition. I argued in Chapter 2 that there are good reasons to think that it does not follow from the rejection or falsity of Descartes's Condition that a satisfying philosophical explanation of how our knowledge of non-psychological physical reality is possible is not possible. So, the second premise of Stroud's puzzle would no longer threaten this move. This brings us to one proposal from Cassam (2007):

In any serious investigation of the conditions of knowledge, we start off with the idea that there are certain things that we know, or certain kinds of knowledge that we actually have. We regard some of the knowledge that we take ourselves to have as negotiable and some as more or less non-negotiable. An unacceptable consequence of an epistemic principle would be the undermining of knowledge in the latter category. [...] Clearly, there is a difference between a piece of knowledge being more or less non-negotiable and its being absolutely non-negotiable. Presumptions to knowledge, even presumptions to basic perceptual knowledge, are, in principle, capable of being defeated, but we shouldn't underestimate what it would take to defeat such a presumption. What we would need is an epistemological principle that is at odds with the presumption and that has such overwhelming independent plausibility that even the fact that it threatens to undermine what we previously took to be non-negotiable knowledge doesn't warrant the principle's rejection (Cassam 2007, 32-33).

There are two core ideas from Cassam in this passage. The first idea is that there are degrees of negotiability in terms of revising or readjusting what we take ourselves to know. Imagine a sliding scale. On one end of the scale are propositions for which knowledge is absolutely non-negotiable. On the other

end of the scale are propositions for which knowledge is the most negotiable. Suppose that the possibility of knowledge of non-psychological physical reality, even if not on the furthest end towards non-negotiable knowledge, is at least within its penumbra. Cassam's second idea, then, is that what it takes for an epistemological principle like Descartes's Condition* to override our belief that we have this knowledge would be for it to have a lot going for it in terms of compelling arguments, explanatory power, intuitive plausibility, and reasonable consequences.

Descartes's Condition* in fact lacks compelling arguments in favour of it. The arguments in favour of Descartes's Condition* consist of augmenting Stroud's argument that it's part of our ordinary epistemic practices.¹⁰⁹ But as we saw in Chapter 2, the falsity of Descartes's Condition need not lead to philosophical dissatisfaction. The second argument is summarised by Pryor (2000) as follows:

The skeptic's scenarios are not ordinary run-of-the-mill alternatives to what we purport to know on the basis of perception. They have special features. It's these special features that account for the sense some philosophers have that no course of experience would enable us to know whether or not those scenarios obtain. The skeptic should argue that there's something especially bad about the scenarios he puts forward, and that for this reason, we have to know that his scenarios do not obtain antecedently to knowing anything on the basis of perception (Pryor 2000, 526-527).

The argument is that sceptical scenarios have special features which make it plausible that knowledge of their denials must not rest on or presuppose knowledge of non-psychological physical reality. For example, in the dreaming scenario, it is compatible with any (or all) of one's sensory experiences and beliefs that one fails to know anything about non-psychological physical reality. In short, our "knowledge of the world is 'underdetermined' by whatever it is that we get through" sensory experience (Stroud 2000, 6).

¹⁰⁹ Another argument for it would be that it follows from the thesis that dreaming that p at t is incompatible with know that p at t , and the strong closure principle, and that the strong closure principle is plausible. Still, Cassam's (2007) argument as I'm using it in the text would still work here as well, since the strong closure principle is *not* within the penumbra of non-negotiable knowledge.

Let's suppose that this underdetermination thesis is true. This is the thesis that our sensory experiences are compatible with the falsity of (almost) anything we believe about non-psychological physical reality.¹¹⁰ Of course, one might dispute whether the underdetermination thesis is true. Indeed, one might wonder whether the mere possibility of a sceptical scenario obtaining provides adequate support for the underdetermination thesis. Let's set these issues to one side.

The question we need to address is whether the underdetermination thesis has epistemological significance. After all, the underdetermination thesis is not an epistemological thesis, so it's prima facie hard to see how it has epistemological significance without supplementation. Stroud supplies the needed epistemological addition as follows:

[...] if we acknowledge that our sensory experiences are all we ever have to go on in gaining knowledge about the world, and we acknowledge, as we must, that given our experiences as they are we could nevertheless be simply dreaming of sitting by a fire, we must concede that we do not know that we are sitting by the fire. Of course, we are in no position to claim the opposite either. We cannot conclude that we are not sitting by the fire; we simply cannot tell which is the case. Our sensory experience gives us no basis for believing one thing about the world rather than its opposites, but our sensory experiences is all we ever have to go on (Stroud 1984, 31-32).

In this passage, it seems that Stroud commits himself to the following principle:

- (U) If the underdetermination thesis is true, then sensory experience does not favour p over $\neg p$.

This principle U connects the non-epistemological underdetermination thesis with the epistemological underdetermination thesis that sensory experience as of p does not favour p over $\neg p$.¹¹¹

¹¹⁰ I say "almost" because it's not plausible that it's compatible with the fact that the person having the sensory experiences doesn't exist, and it's not plausible that that fact isn't a fact about non-psychological physical reality. Cf. Silins (2014, 63).

¹¹¹ This principle is used in the literature in order to formulate the following underdetermination-based sceptical argument: 1. S sensory experience as of p does not favour p over the hypothesis that S is merely dreaming that p . 2. If S's sensory experience as of p does not favour p over the hypothesis that S is merely dreaming that p , then S does

Now one might think that the underdetermination thesis is true and that U provides adequate support for Descartes's Condition*, that S knows that p only if S knows they are not merely dreaming that p independently of anything S knows about non-psychological physical reality. The problem is that it's not clear that it does provide adequate support for Descartes's Condition*. First, it doesn't follow from these considerations that if one's sensory experiences as of p fails to favour p over $\neg p$, then one cannot know that one is not merely dreaming that p . An argument would be needed for this claim. Second, it doesn't follow from these considerations that if one sensory experiences fail to favour that p over *I'm merely dreaming that p* , then I can know that *I'm not dreaming that p* only if I have independent epistemic support in favour of p . Why can't p sometimes be my epistemic support in favour of *I'm not merely dreaming that p* ?

Moreover, the consequences of Descartes's Condition* might tell against it. Let's suppose that it cannot be satisfied. Then this consequence speaks against it. It plausibly has no other explanatory power besides its power to explain the validity and potential soundness of Stroud's no test argument. And it is not intuitively compelling. It simply doesn't strike many of us that one's knowing that they're not merely dreaming that P must be independent of anything one knows about non-psychological physical reality. This requirement is so demanding that it prohibits arguments to the best explanation whose premises are not *a priori*, conceptual, or logical truths.

4 Stroud's New No Test Argument

In the previous section, I exposed some of the challenges facing the proponents of Stroud's (1984) no test argument in support of the thesis that if Descartes's Condition is true, then it cannot be satisfied, so that scepticism is true.

not know that p on the basis of their sensory experience as of p . Therefore, 3. S does not know that p on the basis of their sensory experience as of p . See Brueckner (2014) and Pritchard (2012, 2005).

The first version of the argument relied on Descartes's Condition, understood as expressing a necessary condition on knowledge of non-psychological physical reality. I argued that, on this reading, it's plausible that (i) Descartes's Condition can be fulfilled, and otherwise (ii) it didn't follow from the fact that a test is required to fulfil that condition that the test is unfulfillable.

The second version of the argument relied on Descartes's Condition, understood as expressing a necessary pre-condition on knowledge of non-psychological physical reality. I argued that, on this reading, it's plausible that (iii) Descartes's Condition could not be fulfilled, but that (iv) it's not plausible that it has to be fulfilled.

In his (2009), Stroud rejects both versions of the argument. In particular, Stroud concedes that the necessary condition reading, were it correct, would be a problem for the no test argument, since it wouldn't follow from the fact that Descartes's Condition expresses a general necessary condition on knowing any proposition about non-psychological physical reality *P*, that one could not know that one is not merely dreaming that *P*.

How then should we understand the role of Descartes's Condition in the no test argument? According to Stroud, the proponent of the no test argument:

[...] accepts that knowing that you are not dreaming is a requirement on knowing that there are 'things located outside' you, but [...] also insists that that condition is to be fulfilled on the basis of sensory experience alone (Stroud 2009: 561).

In turn, we can reformulate Descartes's Condition as follows:

Descartes's Condition**: For all *S*, propositions about non-psychological physical reality *P*, if *S* knows that *P*, then *S* knows on the basis of sensory experience alone that *S* is not merely dreaming that *P*.

So, the no test argument trades on Descartes's Condition**. Descartes's Condition** is what:

[...] represents a threat to knowledge, given the conception of sensory experience that is at work in the argument. On that conception, the requirement of knowing that you are not dreaming, and knowing it on the basis of sensory experience alone, introduces a regress that precludes fulfilment of what it says is a necessary condition of knowing anything about ‘things located outside’ us. For each putative case of knowing such a thing, a necessary condition of knowing it can be fulfilled only by knowing something else, and so by fulfilling a different necessary condition of knowing that different thing. Each step in the attempt to fulfil a necessary condition of the truth of the original knowledge-claim introduces a different and so far unfulfilled requirement (Stroud 2009, 561-562).

An immediate question we should ask is what makes Descartes’s Condition** different from Descartes’s Condition? Notice that Stroud (2009) tells us that the agent has to know that the condition is fulfilled on the basis of “sensory experience alone” (Stroud *ibid*). So what we need to address is what it means for that condition to be fulfilled on the basis of sensory experience alone.

In order to get a better picture of what is going on, then, we should consider Stroud’s (2009) formulation of the no test argument, and examine how it works:

If each sensory experience I have, as far as I can tell from that experience alone at the time, could be a dream, I could know of ‘things located outside me’ by means of sensory experience alone only if I know or have reason to believe that the sensory experience I am having is a real perception and not a dream. Nothing in that sensory experience itself can reliably indicate to me that it is not a dream. So to know on the basis of sensory experience that it is a real perception and not a dream I would have to have some other sensory experience or perform some sensory test that would indicate to me that the experience in question is real and not a dream. But whatever that sensory experience or the sensory outcome of that putative test happened to be, it would not in itself reliably indicate to me that the first experience I was testing for was real and not a dream unless I knew or had reason to believe that the new sensory experience of testing was itself real and not a dream. No sensorily distinctive mark or feature of that new experience could indicate that. So some further experience or test would be needed to determine whether it was a dream or not. And so on (Stroud 2009, 561).

Let’s consider the first statement. It looks like Stroud embeds Descartes’s Condition** in a conditional. I will formulate it as the following premise:

- (1) For all sensory experiences E , if for all S “can tell from that experience alone at the time, $[E]$ could be a dream”, then if S knows that P “on the basis of sensory experience alone”, then S knows that they’re not merely dreaming that P (Stroud *ibid*).

Notice that the antecedent of the first premise is about what S can “tell” on the basis of the sensory experience. I’ve retained Stroud’s phrases about “telling” and “knowing on the basis of sense experience alone”, because I will argue that these phrases are problematic in the argument. For now, I will bracket this problem so as to move forward with the exposition of the argument.

The second sentence of the quotation gives expression to the second premise:

- (2) It's not the case that there is something "in that sensory experience [which] itself can reliably indicate to me that it is not a dream" (Stroud *ibid*).

I take this premise to register the thesis that one cannot come to know, on the basis of the sensory experience alone, that it is not a dream. But this is problematic because why should we allow that as a premise in the argument? Stroud gives us the following reasons:

Any sensory experience you might appeal to to show that a particular experience is not a dream could, as far as you could tell, be a dream. But the dream possibility has this implication at this point not because something other than ordinary, familiar dreaming is in question. It is because of the way the requirement of knowing that you are not dreaming on a particular occasion must be fulfilled, as Descartes here thinks of it. It is not that the requirement cannot be fulfilled because some ‘extraordinary’ kind of ‘dreaming’ is in question. It is rather that the requirement cannot be fulfilled when the only resources available are limited to what can be known on the basis of the senses alone (Stroud 2009, 562).

So, because of the facts about “what can be known on the on the basis of senses alone”, we are unable to know that we are not merely dreaming that p (Stroud *ibid*). What are those facts? In his (2009), Stroud cites a passage from Descartes

which exposes Descartes as committed to the thesis that our sensory experiences are logically compatible with the world being radically different from what we take it be:¹¹²

If we acknowledge that our sensory experiences are all we ever have to go in gaining knowledge of about the world, and we acknowledge, as we must, that given our experiences as they are we could nevertheless be simply dreaming of sitting by the fire, we must concede that we do not know that we are sitting by the fire. Of course, we are in no position to claim the opposite either. Our sensory experience gives us no basis for believing one thing about the world rather than its opposite, but our sensory experience is all we have got to go on (Stroud 1984, 31-32).

After all, it's plausible that if it's impossible for the sensory experience to be a reliable indicator that it is not a dream, then it's hard to see how one could come to know that it's not a dream by appealing only to that sensory experience. The problem here is that Stroud doesn't formulate the premise as a conditional, but as a categorical premise.

From premise (1) and (2), Stroud makes the following inference:

- (3) "So to know on the basis of sensory experience that it is a real perception and not a dream I would have to have some other sensory experience or perform some sensory test that would indicate to me that the experience in question is real and not a dream" (Stroud *ibid*).

With premise (3) in play, Stroud can now set up a regress argument. Premise (1) says if *S* knows that *P* on the basis of sensory experience alone, then *S* knows that they are not merely dreaming that *P*. Premise (2) says, roughly, that

¹¹² Here is the passage from Descartes, quoted in Stroud (2009): "[...] every sensory experience I have ever thought I was having while awake I can also think of myself as sometimes having while asleep; and since I do not believe that what I seem to perceive in sleep comes from things located outside me, I did not see why I should be any more inclined to believe this of what I think I perceive while awake" (Descartes 1984, 53). See also Wright (1991) for an endorsement of this thesis. In particular, he says that: "the content and quality of an episode of experience is never logically sufficient to classify it as dream or perception" (Wright 1991, 91).

it's hard to see how the sensory experience as of *P* alone could provide *S* with a way of coming to know that it's not merely a dream that *P*. Stroud infers from this that some test *T* is necessary for knowing that the sensory experience as of *P* is not merely a dream that *P*. Here, however, Stroud applies Descartes' Condition** to the proposition that the test *T* reliably indicates that the sensory experience as of *P* is not merely a dream that *P*. From this, we get:

- (4) If *S* knows that the test *T* reliably indicates that the sensory experience as of *P* is not merely a dream that *P*, then *S* knows that they're not merely dreaming that the test *T* reliably indicates that the sensory experience as of *P* is not merely a dream that *P*.

Now with premise (4), a regress gets started. *S*'s knowing that the test is a reliable indicator of the sensory experience not being a dream is of course sufficient for fulfilling the requirement that they're not merely dreaming that the test reliably indicates that the sensory experience as of *P* is not merely a dream that *P*. But how does *S* come to know this? Stroud here falls back on premise (2), that there's nothing about the sensory experience as of *P* which reliably indicates that it's not merely a dream that *P*. And then a test is introduced to show that the sensory experience as of *P* is a reliable indicator that it's not merely a dream that *P*. But the proposition that this test is a reliable indicator that the sensory experience as of *P* is not merely a dream that *P*, is subject to the same reasoning, and the same set of requirements, introducing the regress.

5 Stroud's New No Test Argument: Evaluation

In the previous section, I outlined Stroud's (2009) new no test argument. The argument is similar to his previous (1984) no test argument, but with at least

one significant difference. The difference is that rather than appealing to what I called “Descartes’s Condition”, Stroud refines that requirement so that Descartes’s Condition must be fulfilled on the basis of “the senses alone” (Stroud 2009, 562). I called this “Descartes’s Condition**”. In this section, I critically evaluate Stroud’s (2009) no test argument.

5.1 The First Premise

Let’s begin with the first premise. Stroud puts the first premise as follows:

If each sensory experience I have, as far as I can tell from that experience alone at the time, could be a dream, I could know of ‘things located outside me’ by means of sensory experience alone only if I know or have reason to believe that the sensory experience I am having is a real perception and not a dream (Stroud *ibid*).

I noted the problematic phrase “as far as I can tell” in section 4. The problem with this phrase is that it’s implausible that Stroud intends the modal operator “can” to express the point that it’s possible that one is dreaming. Instead, the modal operator here is epistemic: it is a dream, for all that one can tell. But being able to tell, while in some uses, an epistemic verb, it is not obvious which epistemic state it signifies.

In what follows, I will disambiguate between different epistemic uses of telling that p in order to draw out the most plausible version of the first premise.¹¹³

One epistemic use of the verb is that telling that p is reporting that p , as with:

(1a) “He told us that he would be there soon.”

(1b) “The manager told us to get here earlier next time.”

In these cases, the verb expresses a communicative act, an act which functions to report that p is the case. This is not the use that Stroud has in mind, because

¹¹³ See Millar (2010, 124-143) for a developed theory on which telling that p implies knowing that p .

it's not plausible that our sense experiences communicate to us in the literal sense of being communicative acts.

Another epistemic use of telling that p is that telling that p is revealing that p , as in:

(2a) "The data tells us something quite different."

(2b) "The look on her face told us what had been happening".

This sense of telling is non-verbal, but epistemic. Figures, facts, pictures, or looks, for example, can reveal that something is so. For example, it might be that a stop sign tells us not just to stop, but that we have gone too far: for we knew that once we reached the stop sign, we have passed our destination. It is this second sense of telling that the use of 'tell' is meant to express. It is an epistemic sense of 'tell' that can, but perhaps need not be, capturing a non-verbal presentation that something is so.

On this reading, if for all we can tell, an experience as of p could be a mere dream that p , it is not that we are failing to tell something, but that the experience is failing to tell us something. It is failing to tell us that it is not merely a dream that p . In other words, experiences in which we dream that p do not have properties, identifiable to the subject of the dream, which put them in a position to know that the experience is a dream that p .

Here, we might wonder what it is about the experience and one's relation to it that disallows the experience as of p to reveal the kind of state that it is (e.g., a dream that p). I take this to just register a good question, rather than express disbelief in the principle. For it just registers that if a sensory experience as of p is a mere dream that p , it fails to tell us something about it: that is not a perception that p . Taken on its own, it is harmless. It registers the intuition that, in the bad case, a case where we are dreaming that p , and so don't know that p , we are not in a position to know that it is not the good case where we are not misled, and we are not dreaming that p . It does not follow from this

alone that, in the good case, where we are not misled, we are not in a position to know that we are not misled.

Another epistemic use of telling that *p* is that telling that *p* is distinguishing that *p* or discriminating that *p* from a range of alternatives. For example:

(3a) “I can’t tell the difference between them”.

(3b) “You can tell that this tree is much older than that tree”.

On this reading, if as far as we can tell, we could be dreaming that *p*, then we cannot discriminate dreaming that *p* from some alternatives.

A final reading is that telling that *p* is knowing that *p*:

(4a) “I can just tell that he’ll be late.”

(4b) “No one will be able to tell that it was us.”

This is a robust epistemic use of the verb, and is comparable with the purely epistemic sense of the verb to see, as when one can see that the war has ended, or that someone can see that the argument is a proof of *p*.¹¹⁴

I submit that these are the most plausible disambiguations of telling that *p* in the first premise of Stroud’s argument. We are left with the following disambiguated versions of the first premise of Stroud’s (2009) no test argument:

Revelation Reading: if S’s sensory experience *E* as of *p* does not reveal to S that it’s not a merely a dream that *p*, then if S knows that *p* on the basis of *E* alone, S knows that *E* is not merely a dream that *p*.

Discrimination Reading: if S cannot discriminate their sensory experience *E* as of *p* from merely dreaming that *p*, then if S knows that *p* on the basis of *E* alone, S knows that *E* is not merely a dream that *p*.

¹¹⁴ For the purely epistemic sense of such verbs, see Gisborne (2010). See also French (2010) and (2013).

Knowledge Reading: if S does not know that their sensory experience E as of p is not merely a dream that p , then if S knows that p on the basis of E alone, S knows that E is not merely a dream that p .

In the next section, I assess each version of the first premise.

5.3 Assessment of the First Premise

5.3.1 The Knowledge Reading

The knowledge reading is problematic because it guarantees that Descartes's Condition* is not fulfilled. To see this, suppose that the antecedent of the first premise on the knowledge reading is true: S does not know that their sense experience as of p is not a mere dream that p . Now consider the conditional that is embedded in the premise: if S knows that p on the basis of their sense experience E as of p alone, then S knows that E is not a mere dream that p (this is Descartes's Condition**). From modus tollens, it follows that S does not know that p on the basis of their sense experience E as of p alone. We cannot accept that as a premise in the argument.

5.3.2 The Revelation Reading

The revelation reading is problematic because it too guarantees that Descartes's Condition* is not fulfilled. The argument here is simple. First, revealing that p seems to entail that S knows that p . If a state of mind reveals that p to me, it's plausible that I (now) know that p . But just as with the knowledge reading of telling that p in the first premise, this is problematic because it guarantees that Descartes's Condition is not fulfilled. We cannot accept that as a premise in the argument.

5.3.3 *The Discrimination Reading*

The discrimination reading is problematic because it's hard to see how the antecedent implies the consequent; that is, it's hard to see how the discrimination version of the first premise is true. In what follows, I will try provide adequate support for this claim.

Now the kind of discrimination at issue in the premise is not perceptual discrimination, as when one perceptually discriminates a tomato from a banana, but introspective discrimination. The question we need to address is whether our inability to introspectively discriminate a perceptual experience as of p from a mere dream as of p implies that a necessary condition of knowing that p is knowing that one is not merely dreaming that p .

At least one prima facie case against this is the thought that it makes our knowledge of non-psychological physical reality depend far too much on our introspective discriminatory capacities. And absent an argument for why there is such a dependence, we have very little reason to accept the conditional. So, although we might be open to the (perhaps plausible) idea that our knowledge of non-psychological, physical reality depends (in many cases, e.g., for demonstrative judgements) on our perceptual discriminatory capacities, it's far from clear that this dependence reaches as far as our introspective discriminatory capacities.¹¹⁵ This brings us to the second premise of Stroud's new no test argument, where the case against a dependence between our introspective discriminatory capacities and our knowledge of non-psychological physical reality can be brought out more sharply.

5.4 The Second Premise

Recall that this is the premise that “[n]othing in that sensory experience itself can reliably indicate to me that it is not a dream” (Stroud *ibid*). I noted in section 4 that it's not clear that we should allow this to be an unargued premise

¹¹⁵ For example, see Goldman (1976) and Millar (2010).

in Stroud's (2009) no test argument, because it comes far too close to not knowing that their sense experience as of p is not a mere dream that p .

In this section, I will address this concern. The concern is that we cannot allow the second premise, as it stands, to be an unargued premise in the no test argument. What we are looking for, then, is some argument from Stroud (or on Stroud's behalf) which provides adequate support for this premise.

The first task is to make clear what Stroud means by "knowledge on the basis of the senses alone". What Stroud calls "knowledge on the basis of the senses alone" is knowledge that is grounded in one's sensory experiences. This kind of knowledge might amount to knowledge of one's immediate environment (of non-psychological, physical reality), or it might amount only to mere appearances as of one's immediate environment (of non-psychological, physical reality), where this latter type of knowledge is a species of self-knowledge.

For example, knowledge on the basis of visual experience alone might amount to seeing that p , and therefore knowing that p , where " p " is a proposition about non-psychological, physical reality. On the other hand, knowledge on the basis of visual experience alone might amount to seeming to see that p , where what one knows is a fact about one's mind: that it visually seems to one that p .¹¹⁶

¹¹⁶ That Stroud is conceiving of it this way is brought out in the following quotation: "For one thing, the demand for non-sensory knowledge to supplement what we get from the senses rests on the assumption that we never actually *perceive* anything to be so in the world around us" (Stroud 2009, 563). The "non-sensory knowledge" that Stroud mentions would be knowledge of some kind of epistemic principle, such as that 'sensory experiences as of p are reliable indicators that p '. There are two versions of this view he considers: an internalist and an externalist version. The internalist version leads to scepticism, whereas the externalist version is philosophically unsatisfying (for a summary of his reasons, see pg. 564). On this basis, Stroud tells that "to understand properly how our knowledge of the world is possible we cannot settle for less than our being able sometimes simply to *see* or in general *to perceive that such-and-such is so* in the world we take ourselves to have knowledge about. That is the moral I would draw from Descartes' sceptical argument. We should get as far away as possible from the conception of perceptual knowledge to be found in Descartes' *First Meditation*. We should abandon the idea of purely sensory knowledge that is neutral with respect to how things are in an independent world" (Stroud 2009, 564-565).

In the argument, the reason that Stroud thinks that “[a]ny sensory experience you might appeal to to show that a particular experience is not a dream could, as far as you could tell, be a dream” is that he is conceiving of our knowledge on the basis of the senses alone along the lines of seeming to perceive that p rather than perceiving that p (Stroud *ibid*). This might be a good inference, then, because it’s not implausible that if we can only appeal to seeming to perceive that p (and anything we know *a priori* about seeming to perceive that p), this would not count against our merely dreaming that p .

5.4.1 Assessment of the Second Premise: Reliable Indicators and Knowledge

Can’t a sensory experience as of p simply be a reliable indicator that p , and thereby not merely a dream that p ? If there are the right sort of causal, counterfactual, or reliable connections between experiences of that type and facts of that type, it’s at least not obvious that sensory experiences of that type cannot be reliable indicators of facts of that type, and thereby not merely dreams as of facts of that type.

Consider what Dretske (1986, 2014) calls “natural signs”, like rings in a stump or tracks in the snow. The rings in the stump are reliable indicators of the age of the tree; the tracks in the snow are reliable indicators of an animal having tracked through the snow. Even if the agent looking at the rings in the stump is unaware of this fact (even if the agent doesn’t know that the rings in the stump are reliable indicators of the age of the tree), those rings in that stump still are reliable indicators of the age of the tree.

Now consider again a sensory experience as of p (say, as of one’s book on a nearby shelf). That sensory experience might be a reliable indicator as of one’s book on the nearby shelf, even if one doesn’t know that their sensory experience as of their book on the nearby shelf is a reliable indicator of their book being on the shelf. Now recall that Stroud tells us that “[n]othing in that sensory experience itself can reliably indicate to me that it is not a dream”

(Stroud *ibid*). But this isn't obviously true, since intuitively many things can reliably indicate that something is so without us thereby knowing that they reliably indicate that it's so. It's not clear why one would have to satisfy the higher-order epistemic requirement of knowing that a potentially reliable-indicator is a reliable-indicator in order for it to reliably indicate that something is so.

6 Conclusion

In this chapter, I examined the first premise of Stroud's puzzle. This is the premise that if Descartes's Condition is true, then scepticism is true. I argued that this premise of Stroud's puzzle lacks adequate support.

In section 2, I reviewed two versions of Stroud's (1984) argument for this premise. The first argument is the regress argument. I argued that the regress is benign. The second argument is the no test argument.

In section 3, I argued that the no test argument is fallacious. I then considered a version of the no test argument which traded on different version of Descartes's Condition. This is the necessary pre-condition version of Descartes' Condition. I argued that even if this renders the no test argument valid, it's not sound because Descartes's Condition so understood should be rejected.

In section 4, I examined Stroud's (2009) no test argument. I explained how there are several different versions of the first premise of this argument. I argued in each case that the premise is implausible. Since none of the arguments for the first premise of Stroud's puzzle are successful, and none of the arguments for the second premise of Stroud's puzzle are successful, it follows that Stroud's puzzle fails to provide adequate support for meta-epistemological scepticism.

In the next chapter, we turn our attention to the second argument from Stroud: Stroud's dilemma.

Chapter 4

Stroud's Dilemma: Williams and Epistemological Realism

1 Introduction

In Chapter 2 and Chapter 3, I explained Stroud's puzzle in support of meta-epistemological scepticism. I argued that Stroud's puzzle fails to provide adequate support for meta-epistemological scepticism.

Recall that in Chapter 1, I distinguished between two arguments in support of meta-epistemological scepticism: Stroud's puzzle and Stroud's dilemma. Even if Stroud's puzzle fails to provide adequate support for meta-epistemological scepticism, that still leaves Stroud's dilemma unphased, and so the case for meta-epistemological scepticism stands.

In this chapter, I recap Stroud's dilemma, how one might respond to it, and I assess a complex response to Stroud's dilemma from Michael Williams (1996). I argue that Williams's argument fails to show that Stroud's dilemma does not provide adequate support for meta-epistemological scepticism. In what follows, I will explain the structure of this chapter.

Section 2 outlines Stroud's dilemma for meta-epistemological scepticism, and provides a brief overview of the responses to Stroud's dilemma. This is done to remind the reader of the direction the thesis is going in. Section 4 sets out Williams's core criticism of Stroud's dilemma, what I call "Williams's master argument" against Stroud's dilemma. Section 5 evaluates the first premise of Williams's master argument, while section 6 evaluates the second premise of Williams's master argument. Section 7 concludes this chapter.

2 Stroud's Dilemma

In Chapter 1, I explained the form of Stroud's dilemma and the basic arguments in support of each premise of the dilemma. In this section, I return to Stroud's dilemma in order to provide an adequate backdrop for our assessment of Williams's master argument against Stroud's dilemma.

2.1 Stroud's Dilemma: EPR

Stroud's dilemma arises out of the constraints on fulfilling the "traditional epistemological project" (Stroud 2000, 109). As I explained in Chapter 1, this is the project of trying to achieve a satisfying, philosophical understanding of how knowledge in some problematic domain is possible (e.g., the external world, other minds, the future, the past, and the *a priori*). Following Stroud (2000), I will focus on the traditional epistemological project for our knowledge of the external world (for our knowledge of non-psychological, physical reality).

How can the traditional epistemological project lead one to accept meta-epistemological scepticism? To answer this question, we need to examine what Stroud thinks are the desiderata for fulfilling the traditional epistemological project.

The first desideratum is that the explanation be a general explanation.¹¹⁷ Here it must be remembered that the explanation has to answer how knowledge of the world around us is possible at all. This is what Cassam (2007) calls a "how possible?" question; a question which asks how something is possible given certain obstacles which make it look impossible. These kinds of questions need to be distinguished from "how?" questions (e.g., how did S come to know that *p*?, "what makes it possible?" questions (e.g., what makes it possible for S to acquire that knowledge that *p* on the basis of *b*?), and "what is

¹¹⁷ In Chapter 6, I explore what is meant by a "fully general" explanation of how knowledge of the world is possible, in examining Cassam's (2009) response to Stroud's dilemma.

it?” questions (e.g., what is it to know that p ?). What counts as a good answer to one of these questions need not be a good answer to the “how possible?” question.

Second, to explain that knowledge of the world around us is possible is not to explain how it is possible.¹¹⁸ For example, one need not dispute that knowledge of the world around us is possible. One can grant that knowledge of the world around us is actual, but that this doesn’t go towards fully explaining how it is possible. It just states that it’s actual, which implies that it’s not impossible, but that knowledge of the world is not impossible does not amount to satisfying philosophical explanation of how it is possible.¹¹⁹

This point can be brought into sharper relief if we consider a scientific explanation of knowledge. This kind of explanation will accept that people know things about the world around them on certain bases, and these bases can

¹¹⁸ Cf. Nozick (1981). See also Stroud (1984, 139-140). In particular, Stroud tells us that: “the question philosophy must answer while avoiding the traditional sceptical account is how our knowledge of the world is possible. It is simply not enough simply to demonstrate *that* it is possible, or even that it is actual. [...] What is wanted is an *explanation*, but just any (even apparently true explanation) explanation of how our knowledge of the world is possible will do (Stroud 1984, 143). I take Stroud here to be telling us two things. First, the fact that human beings have knowledge of the world is not sufficient for a satisfying philosophical understanding of how it’s possible. Second, even a positive and correct explanation isn’t sufficient either, since it might not be *philosophical* (cf. Chapter 1, §2).

¹¹⁹ For example, see Nozick (1981), who argues that: “So the task of showing how p is possible cannot be done once and for all by a proof that p . What a proof can do—show us that p is true—is not what we need, for we already believe this. Why isn’t it enough to know that p is true, why do we also need to understand how it can be true? To see how p can be true (given these apparent excluders) is to see how things fit together. This philosophical understanding, finding harmony in apparent tension and incompatibility, is, I think, intrinsically valuable” (Nozick 1981, 10). Moreover, Nozick adds that: “the task of explaining how p is possible is not exhausted by the reargued action of meeting arguments from its apparent excluders” because there “remains the question of what facts or principles might give rise to p ” (Nozick 1981, 11). I think that Nozick is right about both points, with a caveat on the first point. Nozick takes it that *the reason* that a proof of p doesn’t contribute to explaining how p is possible is the fact that the philosopher already believes that p . But this isn’t right. It isn’t right because the arguments which make p look impossible might have made the philosopher withdraw their belief that p . This does not entail that the philosopher now believes that $\sim p$. Instead, the philosopher might just be agnostic about whether or not p . Indeed, it’s plausible that an argument which makes p look impossible ought to defeat one’s justification for believing that p . So, a proof of p might contribute towards explaining how p is possible, but this does not entail that it’s sufficient for explaining how p is possible.

be studied using computational models and in general the resources of cognitive science. The trouble with the scientific explanations (or even more colloquial explanations, such as that human beings can perceive that there are objects in their environment) is that “it explains how we know some particular fact in the area we are interested in by appeal to knowledge of some other fact in the same domain” (Stroud 2000, 101).¹²⁰

Now it must be reinforced here that what makes this kind of explanation unsatisfying is not that it is not true. Again, it is allowed that we know about the world around us, and that the natural sciences can inform how such knowledge is possible. Second, what makes this kind of explanation unsatisfying is not that it can't be an adequate answer to a “how possible?” question. There might be good questions about how, given the facts about our memories and our perceptual systems, human beings can come to know about their environment. As Stroud remarks, it would be “as absurd to cast doubt on the prospects of scientific investigation of human knowledge and perception as it would be to declare limits to our understanding of human digestion” (Stroud 2000, 99). The problem, rather, is that these sorts of answers cannot be adequate answers to the philosophical question of how our knowledge of the world is possible.

In Chapter 1, I argued that it doesn't follow from the possibility of a satisfying scientific explanation of how knowledge of the world is possible that a satisfying philosophical explanation of how knowledge of the world is possible unless there is a sound argument for the thesis that the scientific explanation is, amounts to, or otherwise ought to replace the philosophical explanation.¹²¹ This argument is part in parcel with Stroud's remark about scientific explanations.

¹²⁰ Cf. Stroud (1981). See also Fumerton (1994). One of Fumerton's worries is that: “[g]iven such [scientific] views, it is not clear that a philosopher *qua* philosopher is even in a position to speculate intelligently on the question of whether or not we have noninferentially justified belief in any of the propositions under skeptical attack” Fumerton (1994, 326 my addition).

¹²¹ This kind of strong naturalised epistemology can be found in Quine (1969). For discussion, see Feldman (2001).

This raises the question of what makes the philosophical question so different from the scientific question. After all, both questions can take the same form of words. At least one answer is that the obstacles that make one ask the respective questions are different in the two cases. For example, Stroud (1984*b*) expresses the obstacle as follows:

The difficulty in understanding how sense-perception gives us knowledge of any such truths is that it seems at least possible to perceive what we do without thereby knowing something about the things around us. There have been many versions of that fundamental idea. But whether it is expressed in terms of “ideas” or “experiences” or “sense data” or “appearances” or “takings” or “sensory stimulations,” or whatever it might be, the basic idea could be put by saying our knowledge of the world is “underdetermined” by whatever it is that we get through that source of knowledge known as “the senses” or “experience.” Given the events or experiences or whatever they might be that serve as the sensory “basis” of our knowledge, it does not follow that something we believe about the world around us is true (Stroud 1984*b*, 549).

For Stroud, this obstacle is what leads us to embark on the traditional epistemological project:

The problem is then to explain how we nevertheless know that what we do believe about the world around us is in fact true. Given the apparent “obstacle,” how is our knowledge possible? (Stroud *ibid*).

So, the obstacle brings into question the possibility of any knowledge of the world at all. This kind of obstacle is what leads one to raise the question that gives expression to the traditional epistemological project. Moreover, it is this feature of the question that also renders inadequate using any particular instances of what we take to be knowledge of the world (or citing the general fact that knowledge of the world is possible) in order to answer it. Stroud puts the point like this:

If all our knowledge of the world around us is in question all at once we cannot then help ourselves to some independently reliable information about the world, as we usually do, to settle the question whether our present course of experience of experience is or is not on this occasion a reliable guide to the way things are (Stroud 1984, 208).

This encodes the idea that our explanation should tell us how any knowledge of the world is possible at all, so that no case of knowing is left potentially unexplained. Call this the generality requirement on a satisfying philosophical explanation of how our knowledge of the world is possible (cf. Chapter 1, §6.1). The problem then is that explanations which cite, presuppose, or otherwise entail that we have knowledge of the world look inadmissible. In other words, the problem that the generality requirement raises is that it seems to constrain even further the admissible types of explanations one can put forward in order to fulfil the traditional epistemological project. Stroud (2009*b*) summarises the argument for an additional epistemic priority requirement¹²² (henceforth “EPR”) as follows:

If *all* knowledge of the world on the basis of the senses is what is in question, and nothing that is already taken to be known about the external world can be appealed to, further requirements appear to follow. It looks as if knowledge of the world could be explained as coming somehow from perception only if what we get through the senses alone is always something less than knowledge of anything that is so in the external world. So even if perception gives us knowledge of something, it is tempting to say that the general epistemological problem, properly understood, requires ‘explaining knowledge of the world on the basis of another, prior kind of knowledge that does not imply or presuppose any of the knowledge we are trying to explain’. We could even say, more noncommittally, that the problem requires ‘showing how our knowledge of the world comes to be out of something that is not itself knowledge of the world’ (Stroud 2009*b*, 591).

The argument here seeks to provide adequate support for the thesis that one’s explanation of how our knowledge of the world is possible cannot fulfil the generality requirement without fulfilling EPR.

Should we accept that this is true? First I want to work through the argument for why fulfilling the traditional epistemological project seems to

¹²² Here’s how Cassam (2009) expresses EPR: “EPR [...] says that a philosophically satisfying explanation of our knowledge of the world must represent it as coming out of something that neither implies nor presupposes any knowledge of the world” (Cassam 2009, 576). Compare with: “In its most general form EPR says that when we want to explain knowledge of kind K we must do so on the basis of another, prior kind of knowledge that does not presuppose or imply any knowledge of kind K” (Cassam 2009, 577).

require fulfilment of the generality requirement. The argument here is meant to provide at least some prima facie support for that thesis. I do not take it that the argument is a “knock-down” argument for the thesis that failing to fulfil EPR implies that the corresponding explanation is less than general. Instead, I presume that it motivates the thesis that failing to do so makes it look less general than we expect a philosophical explanation to be.

The domain in question consists of propositions about non-psychological physical reality. What we want to explain is how it’s possible to know that any member of that domain is true. Let our explanation, *E*, imply or presuppose that some member of the domain is true. Call this particular member *P*. Is *E* a fully general explanation of how it’s possible to know that any member of the domain is true? It is only if it explains how it’s possible to know that *P*. *P*, after all, is a member of the domain the knowledge of which we want to explain how it’s possible to know.

Why does this matter? Recall that any putative instance of knowledge of non-psychological physical reality is in question. In explaining how it’s possible to know that some particular member is true, this particular member is taken as representative of the entire domain. The problem then arises as follows: the goal is to explain how it’s possible to know that any member of the domain is true. Let *R* be the representative proposition here. Now suppose that (i) our explanation *E* of how it’s possible to know that *R* is true implies or presupposes that some other member of the domain *P* is true (in violation of EPR).

This raises the question of how it’s possible to know that *P* is true, since *P* is also a member of the domain in question. If knowing that *P* turns out to be necessary for explaining how it’s possible to know that *R* is true, then we are saying that knowing at least one proposition about non-psychological physical reality is necessary for explaining how it’s possible for anyone to know anything at all about non-psychological physical reality. How it’s possible for anyone to know that *P*, then, would be left unexplained. So, some knowledge

of non-psychological physical reality would be necessary for explaining how it's possible to know that any proposition about non-psychological physical reality is true. This is not a fully general explanation of how it's possible to know anything at all about non-psychological physical reality.¹²³

Stroud's dilemma can be put as follows: if EPR is true, then scepticism is true. Our explanation of how knowledge of non-psychological physical reality is possible is that it's not possible. But scepticism is not a satisfying philosophical explanation of how that type of knowledge is possible. On the other hand, if EPR is false, then a fully general explanation of knowledge non-psychological physical reality is possible is not possible. And anything less than a fully general explanation of how that type of knowledge is possible is not a satisfying philosophical explanation either.

Williams accepts both horns of Stroud's dilemma.¹²⁴ What he denies is that we have good reasons to even pursue the traditional epistemological project. In what follow, I'll provide a brief overview of how one might respond to Stroud's dilemma, providing a backdrop for Williams's response and the responses from the externalist and Sosa (Chapter 5) and Cassam (Chapter 6).

2.2 Responding to Stroud's dilemma

Stroud's dilemma seems to provide adequate support for meta-epistemological scepticism. How might one go about avoiding this consequence? Here, one might dispute whether the argument is sound, or whether the arguments in support of the premises themselves provide adequate support for the premises.

The first premise moves from acceptance of (EPR) to scepticism. The second premise moves from the rejection of (EPR) to rejection of the traditional epistemological project: that it's not possible after all to give a satisfying philosophical explanation of how knowledge of non-psychological reality is possible.

¹²³ Cf. this kind of argument with Stroud (2009, 568-569).

¹²⁴ See Williams (1996, 89-93).

Stroud's main argument for the first premise is that:

If several different possibilities are all compatible with our perceiving what we do, the question is how we know that one of those possibilities involving the truth of our beliefs about the world does obtain and the others do not. That would seem to require an inference of some sort, some reasonable hypothesis or some form of reasoning that could take us from what we get in sense-perception to some proposition about the world around us. That hypothesis or principle of inference itself either will imply something about the world around us or it will not. If it does, it belongs among those propositions our knowledge of which has yet to be explained, so it cannot help explain that knowledge. If it does not, how can our acceptance of it lead to knowledge of the way things are around us? If it itself implies nothing about such things, and we could perceive what we do without knowing anything about such things, how is our knowledge to be explained? If we are in fact in that position, how is our knowledge of the world around us possible? (Stroud 1984*b*, 549-550).

So, our goal is to meet the desiderata of the traditional epistemological project. According to Stroud, this requires that our explanation respect EPR. This implies that our explanation cannot consist of facts such as that we can see or otherwise perceive that propositions about the world are true if this amounts to knowing propositions about the world. Explanations which involved knowledge-entailing mental states or processes are therefore in violation of EPR.

The problem here, however, is that if our explanation respects EPR, then our explanans will need to be facts that we know, but those facts cannot be ones which entail or presuppose facts about the world. We have to make a transition from knowledge of those facts to knowledge of facts about the world. How is this to be done? Stroud tells us that it "would seem to be possible only if we somehow knew of some connection between what we are restricted to in observation and what is true in the wider domain in question" (Stroud 2000, 105). For example, consider the following principle:

(P) Perceptual seemings as of p are reliable indicators that p .

Stroud thinks that in order to explain how knowledge of the world is possible on the basis of our knowledge of perceptual seemings (for example), we need to know that a principle like (P) is true. A crucial premise of Stroud's argument for the first horn of the dilemma, then, is the following:

- (1) If the epistemologist's knowledge of facts about perceptual seemings can explain how knowledge of non-psychological physical reality is possible, then the epistemologist knows that (P), that facts about perceptual seemings are reliable indicators of facts about non-psychological physical reality.

The problem with knowing that (P) is true, however, is that "knowing even that there was such a connection would be knowing something about the wider domain after all, not just about what we are restricted to in observation" (Stroud 2000, 105). To summarise, Stroud thinks that the following premise is true:

- (2) If S knows that (P) is true, then S knows that a proposition about non-psychological physical reality is true.

Suppose that the antecedent of premise (1) is true. This is the assumption that the epistemologist's knowledge of facts about perceptual seemings can explain how knowledge of non-psychological reality is possible. In short, the assumption is that fulfilling EPR is consistent with the resulting explanation being a positive explanation: one which accounts for how the knowledge in question is possible.

From modus ponens on premise (1) and its antecedent, it follows that the epistemologist knows that the principle (P) is true. This is the antecedent of premise (2). From this and premise (2), then, it follows that the epistemologist's explanation requires them to know a proposition in the domain in question is true, in violation of EPR.

This is problematic because fulfilling the traditional epistemological project requires the resulting explanation to respect EPR. However, Stroud's

argument for the first horn of the dilemma is that one cannot both respect EPR and the resulting explanation be positive (anti-sceptical).

One response is to reject or refine EPR. One might think that Stroud is right to the extent that there are some constraints on what makes a proposed answer to the traditional epistemological project admissible or not. But this does not mean that a satisfying answer to the traditional epistemological project needs to meet all of Stroud's desiderata. In Chapter 6, I look at an objection from Cassam (2007) along these lines.

Another response takes issue with the first premise of the argument in favour of the first premise of Stroud's dilemma. This is the premise that the epistemologist needs to know that a certain bridge-principle between perceptual seemings and facts about the world is true. Externalists might think that it's enough that there is a true linking-principle between perceptual seemings and facts about the world. On this view, the requirement that the epistemologist needs to know that this linking-principle is true in order for the explanation to be admissible is spurious. Moreover, an externalist might think that, insofar as we have to know that there is such a true linking-principle, this demand can be met along externalist lines. In Chapter 5, I evaluate this kind of externalist response from Sosa (1994).

Another kind of response tries to block Stroud's dilemma from arising. For example, one can grant that if the fulfilment of the traditional epistemological project is a legitimate intellectual goal, then Stroud would be right that the epistemologist runs into a problematic dilemma. But perhaps the traditional epistemological project does not give expression to a legitimate intellectual goal. In outline, this is the response of Williams (1996) to Stroud's dilemma. In the remainder of this chapter, I evaluate Williams's response to Stroud's dilemma.

3 The Traditional Epistemological Project

In the previous sections, I rehearsed the arguments in support of the premises of Stroud's dilemma. In this section, I outline Williams's (1996, 1996*b*) response to Stroud's dilemma.

The core target of Williams (1996) critique of Stroud's dilemma is what I will call the "legitimacy thesis": the traditional epistemological project, the project of trying to give a satisfying philosophical explanation of how knowledge of non-psychological physical reality is possible, is legitimate.

What does it mean for the traditional epistemological project to be "legitimate"? I will not offer a definition or a set of criteria for a project to be legitimate. Instead, I will work with one requirement on a project being legitimate. The requirement is from Williams's (1996):

What matters is whether "our knowledge of the world" picks out the kind of thing that might be expected to be susceptible to uniform theoretical analysis, so that failure to yield such analysis would reveal a serious gap in our understanding (Williams 1996, 102).

The requirement here is that our explanandum be susceptible to uniform theoretical treatment. Still, the phrase "uniform theoretical treatment" is itself unclear. I will proceed, then, with an example from Williams in order to make the requirement clearer:

Consider the set of all things denoted by English words beginning with the letter "H": would we expect to understand, with complete generality, what these are and how they come to be? Clearly not: this does not even feel like the pursuit of a perfectly intelligible project. The reason why is obvious enough: we expect general intelligibility only with respect to kinds that exhibit some kind of theoretical integrity: paradigmatically, the kinds we call "natural kinds". To the extent that the traditional quest for a general understanding of our knowledge of the world does feel like the pursuit of a perfectly intelligible goal, it must be because we are willing to attribute some kind of theoretical integrity to the beliefs in question (Williams 1996*b*, 368).

As Williams highlights, we expect a uniform treatment of some subject matter when there are plausible reasons to suspect that the subject matter has some common nature shared between each member. This is the case with gold, water,

oak trees, wolves, plankton, and other natural kinds.¹²⁵ Williams's case with the letter "H" is not of this kind. All that we expect the set of all the objects, properties, or events that, in English, are referred to with words that begin with the letter "H" to have in common is being the referent of an English word that begins with the letter "H".

The point that Williams is driving at with this example is that the goal of explaining what all of the referents of English words that begin with the letter "H" have in common is not a legitimate intellectual goal. It is not one that one ought to pursue because it is not the sort of thing that is amenable to uniform theoretical treatment.

To ask whether the traditional epistemological project is legitimate, then, is to ask whether it expresses an intellectual goal that we ought to pursue. This will of course turn on our interests as epistemologists, but it will also turn on whether the subject matter in question is susceptible to uniform theoretical treatment. On this score, Williams thinks that we ought to first ask whether "knowledge of the external world" is like water, in being susceptible to uniform theoretical treatment, or instead like the referents of English words beginning with the letter "H", before we decide whether one ought to be dissatisfied with the prospects of failing to achieve the target intellectual goal.

4 Williams's Master Argument

Williams (1996) provides a master argument against Stroud's dilemma. It is "against" Stroud's dilemma in the following sense: if the argument is valid and sound, then it prevents Stroud's dilemma from posing an actual dilemma for the epistemologist engaged with the traditional epistemological project. The reason it prevents Stroud's dilemma from posing an actual dilemma to the epistemologist is that, if the argument is valid and sound, there is no such thing (or no such kind of thing) as knowledge of the external world to theorise about.

¹²⁵ For knowledge as a natural kind, see Kornblith (2002, 2007). For a discussion and overview of the debates on natural kinds, see Bird and Tobin (2015).

Williams's (1996) master argument first tells us that the traditional epistemological project presupposes an apparently contentious meta-epistemological thesis called "epistemological realism", but that "epistemological realism" is not true. What I am calling Williams's "master argument" against Stroud's dilemma can be formulated as follows:

1. If the traditional epistemological project is legitimate, then epistemological realism is true.
2. Epistemological realism is not true.

Therefore,

3. The traditional epistemological project is not legitimate.

In the previous section (section 3) I explained what Williams means for the traditional epistemological project to be legitimate. In the remainder of this section, I will concentrate on the formulation of "epistemological realism".

According to Williams (1996), epistemological realism is the thesis that there are "natural epistemological kinds and natural epistemological relations" (Williams 1996, 120). On this view, there is a mind-independent realm of "epistemological facts" which are "context-invariant" rather than "interest-relative" (Williams 1996, 121). Epistemological realism is a realism about the "objects of epistemological reflection" (Williams 1996, xx). On this account, there are such things, states, processes, or properties as knowledge of the external world, epistemic justification, entitlement, warrant, evidence, and so on.¹²⁶

Epistemological realism is not a clear thesis. We need to address what the phrases that Williams uses to articulate the thesis mean. Let's consider "natural epistemological kinds and natural epistemological relations" (Williams *ibid*). Williams does not mean "natural" in the sense of not being autonomous from

¹²⁶ Cf. with *moral particularism*. See Dancy (2004).

the natural sciences, or as being opposed to supernatural entities and supernatural methods. “Natural” in Williams’s sense means context-invariant or independent of situational factors—that is, independent of the topic in question, the discipline one is engaged in and its methods, the direction of the conversation, and other situational, pragmatic, and methodological constraints. So “epistemological realism” picks out the thesis that there are context-invariant epistemological “kinds” and “relations”.¹²⁷

What are epistemological “kinds” and epistemological “relations”? An example will help make clear what Williams means. Consider foundationalism. Foundationalism is a thesis about the structure of epistemic support (e.g., epistemic justification, warrant, or evidential support). It is the thesis that there are some justified beliefs which do not require that one has other justified beliefs, that having these justified beliefs does not come from one’s justification to believe other propositions, and that these justified beliefs serve as the “foundational” or “basic” beliefs from which all other justification is acquired and transmitted.¹²⁸

Here “epistemic justification” picks out a putative epistemic kind. This is the epistemological property that a person’s belief can possess or lack. Indeed, foundationalists maintain that it can be transmitted through inference, as when

¹²⁷ Cf. Greco (2000, 46-47) for a discussion of epistemological realism. Greco argues that, contra Williams, epistemological realism seems to be true. See also Grundmann (2005) for a discussion of Williams’s “inferential contextualism” (which is opposed to epistemological realism). Grundmann argues that Williams is wrong that the best sceptical arguments depend on the truth of epistemological realism. For a discussion of the similarities and differences between Williams’s contextualism and other forms of contextualism, such as attributer contextualism, see Pritchard (2002).

¹²⁸ Cf. Pryor (2001, 100-101), and Poston (2010). Most of the time, foundationalism is expressed as the thesis that “all knowledge and justified belief rest ultimately on a foundation of noninferential knowledge or justified belief” (Fumerton 2010). See Pryor (2005) and Huemer (2001) for a defence of a modest form of foundationalism. See Fumerton (1995, 2001, and 2006) for a defence of an extreme form of foundationalism, according to which non-inferential justification requires that the agent is acquainted with the truth-maker of their belief and acquainted with the fact that the truth-maker makes true the proposition that the agent believes. The classic objection to foundationalism is Bonjour (1985). Note too that externalist versions of foundationalism have been developed. For a general overview of externalist foundationalism, see Fumerton (1995), Chapter 4, and Bonjour (1980). For development of externalist foundationalism, see Alston (1989) and Sosa (1980).

one deduces q from a set of propositions p_1, p_2, \dots, p_n . The transmission of epistemic justification from a belief that p to another belief that q is an example of an epistemic relation between two beliefs.

Now the epistemological realist maintains that the facts about epistemological properties, such as epistemic justification, and the facts about the epistemological relations between beliefs, such as when justification is or can be transmitted from one belief to another, is context-invariant. So, one might be an epistemological realist and a foundationalist in the following sense: the foundationalist structure of epistemic justification, where some beliefs are “basic” or “foundational”, and others are “non-basic” or “non-foundational”, and the non-basic belief’s justification terminates in the justification of the basic beliefs, is context-invariant. This is what leads Williams (1996) to claim that, on the “foundationalist’s view, a belief’s intrinsic epistemological status derives from the content of the proposition believed” rather than facts about the context in which one believes the proposition (Williams 1996, 116).

Opposed to epistemological realism is epistemological anti-realism or what Williams (1996) calls “contextualism”. According to the epistemological anti-realist, “a proposition has no epistemic status whatsoever” independent of contextual factors (Williams 1996, 119). For example, consider the proposition that I have hands. Williams’s epistemological anti-realist tells us that that proposition derives its justification from a context. In some contexts, the proposition lacks justification, while in some others it has justification. And what counts as a justifier for that proposition in one context might not count as a justifier for that proposition in another context.

Now we might wonder how epistemological realism bears on the traditional epistemological project. In particular, we might wonder how the truth of epistemological anti-realism shows that there is some defect in the traditional epistemological project.

In order to see how the truth of epistemological anti-realism might undermine the traditional epistemological project, consider again the main

target of Williams's criticism. The main target of Williams's criticism is what I called the "legitimacy thesis". This is the thesis that the traditional epistemological project is "legitimate" in the sense that there is such a thing (or kind of thing) as knowledge of the external world which admits of uniform theoretical treatment. In particular, knowledge of the external world admits of the kind of positive philosophical explanation that is at issue in Stroud's dilemma.

I grant that Williams's criticism is plausible to the extent that if the traditional epistemological project is not legitimate, then it is hard to see how Stroud's dilemma poses a serious problem for the epistemologist. In other words, it would be hard to see how Stroud's dilemma presents an interesting meta-epistemological argument, because meta-epistemological scepticism would turn out to be on par with scepticism about witches or scepticism about phlogiston, which are not interesting forms of scepticism.¹²⁹

This covers some of the core terms that Williams uses in his master argument against Stroud's dilemma. In the next two sections (sections 5 and 6) I will examine the sub-arguments for the premises of Williams's master argument against Stroud's dilemma.

5 Williams's Master Argument: Premise 1

In the previous section, I outlined Williams's master argument against Stroud's dilemma, and the core terms which figure in the premises of that argument.

In this section, I will examine Williams's argument for the first premise of his master argument. Recall that the first premise of Williams's master argument is that if the traditional epistemological project is legitimate, then

¹²⁹ The reason this follows is that, if Williams's reason for rejecting the legitimacy thesis is correct, that there is no such thing (or kind of thing) as knowledge of the external world, then the thesis that "a satisfying philosophical explanation of how knowledge of the external world is possible is impossible" would be uninteresting. It would be akin to the thesis that "a satisfying physical explanation of phlogiston is not possible" because there is no such thing as phlogiston.

epistemological realism is true. The argument for this premise is complex. I explain the details of each sub-premise of the argument below. For now, it will be helpful to see this argument in outline:

- 1a. If the traditional epistemological project is legitimate, then our beliefs about the external world form a theoretical kind.
- 1b. If our beliefs about the external world form a theoretical kind, then either it's a topical kind or an epistemological kind.
- 1c. It's not a topical kind.

Therefore,

- 1d. If our beliefs about the external world form a theoretical kind, then it's an epistemic kind.

However,

- 1e. If our beliefs about the external world form an epistemic kind, then epistemological realism is true.

Therefore, from hypothetical syllogism on 1d and 1e:

- 1f. If our beliefs about the external world form a theoretical kind, then epistemological realism is true.

Therefore, from hypothetical syllogism on 1a and 1f:

- 1. If the traditional epistemological project is legitimate, then epistemological realism is true.

The conclusion is that the traditional epistemological project, if it is legitimate, requires epistemological realism to be true. The support for this premise comes in part from the thesis that if the traditional epistemological project is a

legitimate project, then knowledge of the external world is a “theoretical kind”—the kind of thing that admits of uniform theoretical treatment.

In section 4, I explained Williams’s view on what it takes for the traditional epistemological project to be legitimate. In short, it requires that knowledge of the external world, or belief about the external, form a unified “kind”, from which general theoretical explication is possible.¹³⁰

A consequence of this observation is that knowledge of the external world, or belief about the external world, forms a “theoretical kind”, if the project of explaining how knowledge of the external world is possible expresses a legitimate intellectual project. Knowledge of the external world must constitute “a genuine kind of thing”, such as water, buffalo, birch, or quark. (Williams 1996, 108).

Two questions need to be addressed. First, what does it mean for our beliefs about the world to form a “topical kind”? It means that what unifies them as the same in kind, as being beliefs about the external world, are facts about their subject matter or their topic. Second, what does it mean for our beliefs about the world to form an “epistemological kind”? It means that what unifies them as the same in kind, as being beliefs about the external world, are facts about their common epistemic sources, or their common epistemological properties.

The argument for the first premise is valid. I accept premise 1a and premise 1b. We can therefore set those premises to the side. However, the argument is defective. In short, premise 1c, that our beliefs about the external world do not form a topical kind, lacks adequate support. And premise 1e, that if our beliefs about the external world form a theoretical kind, then

¹³⁰ Cf. Williams (1996*b*): “We expect general intelligibility only with respect to kinds that exhibit some kind of theoretical integrity: paradigmatically, the kinds we call “natural kinds.” To the extent that the traditional quest for a general understanding of our knowledge of the world does feel like the pursuit of a perfectly intelligible goal, it must be because we are willing to attribute some kind of theoretical integrity to the beliefs in question (Williams 1996*b*, 368).

epistemological realism is true, lacks adequate support as well. In what follows, I will argue for these two claims.

5.1 Williams's Argument for Premise 1c

Premise 1c is the premise that our beliefs about the external world do not form a genuine topical kind. In other words, there is no uniform subject matter for our beliefs about the external world to be about.

I want to grant Williams two points about this premise. First, from the facts about our usage of phrases such as “knowledge of the world” in both quotidian and theoretical contexts, it doesn't follow that there is such a thing, state, relation, or property as being knowledge of the external world. The same argument applies to our usage of phrases such as “belief about the world”. As Williams observes, the facts about our use of those phrases are compatible with there being nothing more than a “loose aggregate of more or less unrelated cases” (Williams 1996, 102).

Second, we ought to grant Williams the observation that “to understand all knowledge in the standard epistemic domains [e.g., beliefs about the external world, other minds, the *a priori*, etc.] is to suppose that the beliefs in those domains hang together in some important way” (Williams 1996, 103 my addition). I maintain that we ought to concede this observation because without some principle or method for individuating beliefs about the external world from other types of beliefs, we won't have a clear enough domain from which to explain how it's possible to know that our beliefs in that domain are true.

Now that we have granted Williams these observations about premise 1c, let's look at his main argument in support of this premise. Williams suggests that:

“Knowledge of the external world” covers not only all the natural sciences and all of history, it covers all everyday, unsystematic factual claims belonging to no particular investigative discipline. Since, even within a single subject, theories, problems and methods tend to proliferate with the progress of inquiry, so that even the most systematic disciplines tend to become less rather than more unified, it

is doubtful whether we can take a synoptic view of physics, never mind everything we believe about the external world. It is not obvious that it makes sense even to try (Williams 1996, 103).

The challenge is that it's at least not obvious that "knowledge of the external world" or "beliefs about the external world" pick out a unified domain from which theoretical explication is possible. The reason it is not obvious is that the domain "the external world" encompasses our historical beliefs, testimonial beliefs, (most of our) perceptual beliefs, scientific beliefs, sociological beliefs, many of our ordinary beliefs, and so forth. The thought here is how could these domains be integrated, so as to fall under the more general domain of being about the external world?

We can concede the challenge. It's not obvious that "knowledge of the external world" or "beliefs about the external world" pick out a unified domain from which theoretical explication is possible. But then again the opposite is not obvious either. This is true of interesting philosophical theses: it's neither obvious that the thesis is true nor obvious that the thesis is false. What we will need to address later is whether something can be said for our beliefs about the external world being about a uniform subject matter.

5.2 Williams's Argument for Premise 1b

Premise *1b* tells us that if our beliefs about the external world form a theoretical kind, then it's either a topical kind or an epistemic kind. According to this premise, what integrates our beliefs or knowledge of the external world is either its shared subject matter (e.g., the external world), or its shared epistemological nature (e.g., their common epistemic support or sources). Williams thinks that *1b* is true because:

If we are to make sense of the project of explaining how anything we believe about the world amounts to knowledge, we need a way of reducing our beliefs to order. We have to bring them under principles or show them as resting on commitments that we can survey. We must reveal kind of theoretical integrity in the class of beliefs we want to assess (Williams 1996, 103).

The premise is plausible. But one potential issue with Williams's argument for premise *1b* is that it comes close to begging-the-question against the proponent of the thesis that our beliefs about the external world can be integrated on the basis of either its topic or its epistemic sources. After all, recall that Williams thinks that this cannot be done without presupposing epistemological realism, but epistemological realism is false. So, on Williams's view, it cannot be done without a false presupposition. The problem here is that, if the epistemologist must bring beliefs about the external world "under principles or show them as resting on commitments we can survey"—as having some kind of "theoretical integrity"—it might look as if any principles or commitments are going to count as 'contentious theoretical principles' for Williams (not just epistemological realism). But if this is right, then it's hard to see how Williams is presenting a non-questioning-begging defence of the premise, since his argument would render impossible any adequate reply from his opponent.

5.3 Response to Williams's Argument for Premise 1c

Let's side step the previous objection. We can grant Williams the premise and still show that the argument is defective.

Williams argues that our beliefs about the external world can neither be integrated along topical lines nor along epistemological lines without presupposing epistemological realism, and this spells bad news for the traditional epistemological project. What are his arguments for these claims? His argument for premise *1c*, that our beliefs about the external world cannot be topically integrated, goes as follows:

The very nature of the traditional project demands that the principles in question be all-pervasive [e.g., principles about how our beliefs about the external world should be individuated from beliefs with a different subject matter]. For example, if we are to assess the totality of our beliefs about the world, there must be principles that inform all putative knowledge of the world *as such*. But what could they be? I take it to be obvious that, in one way, our beliefs do not show any kind of theoretical integrity. They do not, that is, add up to an ideally unified theory of everything. There is no way now, and none in

prospect, of integrating all the sciences, much less all of anyone's everyday factual beliefs, into a single coherent system: for example, a finitely axiomatized theory with specified rules of inference. [...] We have got a "view of reality" but indefinitely many. The idea [...] that we have a "system" of beliefs ought to be suspect. "Our beliefs", then, do not amount to a single, integrated "view of reality". They are not *topically integrated* (Williams 1996, 103-104).

Williams maintains that there are no "clear criteria for individuating beliefs" about the external world from beliefs with a different subject matter, and that even if there were clear criteria for individuating our beliefs about the external world from beliefs with a different subject matter, there would still be "no systematic way of enumerating all the things that we believe" about the world anyway (Williams 1996, 102).

It's not at all clear that our beliefs about the external world cannot be integrated along topical lines (cf. Chapter 1, section 2). Consider the fact that our beliefs about tables, planets, tigers, subatomic particles, plasma, trees, etc., do have something in common at a sufficient level of generality: being about the physical, or being about the spatio-temporal. Furthermore, they all have the property of being mind-independent, of being about what exists independently of anyone's (or everyone's) psychology or psychological states.¹³¹ In short, these sorts of beliefs are about a non-psychological physical reality. With this understanding of "the external world" in hand, we can construe the epistemologist engaged in the traditional epistemological project as asking how it's possible to know that our beliefs about non-psychological physical reality are true.

¹³¹ Cf. Miller (2014): "There are two general aspects to realism, illustrated by looking at realism about the everyday world of macroscopic objects and their properties. First, there is a claim about existence. Tables, rocks, the moon, and so on, all exist, as do the following facts: the table's being square, the rock's being made of granite, and the moon's being spherical and yellow. The second aspect of realism about the everyday world of macroscopic objects and their properties concerns independence. The fact that the moon exists and is spherical is independent of anything anyone happens to say or think about the matter" (Miller 2014), "Realism", *The Stanford Encyclopaedia of Philosophy* (Winter 2014 Edition). See also Haldane and Wright (1993, 3-4) for discussion of "realism" and mind-independence, and Stroud (2000*b*), Chapter 1, for a discussion of mind-independence.

Beliefs about “the external world” can therefore be unified as the set of beliefs at the intersection of being about mind-independent objects, properties, events, or states-of-affairs, and being about spatio-temporal objects, properties, or states-of-affairs. So it’s false that our beliefs about the external world “do not show any kind of theoretical integrity” (Williams 1996, 104). Premise 1c is therefore not a non-negotiable premise of the argument.

5.4 Williams’s Argument for Premise 1e

Premise 1e tells us that if our beliefs about the external world can be integrated along epistemological lines, then epistemological realism is true of our beliefs about the external world. According to Williams:

As a way of classifying beliefs, “beliefs about the external world” is only quasi-topical, bringing together beliefs belonging to any and every subject, or no-well defined subject. They are united only by their supposed common epistemological status. The essential contrast to “beliefs about the external world” is “experiential beliefs” and the basis for the contrast is the general epistemic priority of beliefs falling under the latter heading over those falling under the former. “External” means “without the mind”; and it is taken for granted that we have a firmer grasp of what is “in” in the mind than what is outside it (Williams 1996, 106).

On this account, our beliefs about the external world are integrated because these kinds of beliefs are subject to the “same fundamental, epistemological constraints” (Williams 1996, 104). What epistemological constraints are these? Williams answer is the “the senses”: our beliefs about the external world have the senses as their most basic source (Williams *ibid.*)¹³² Now one might think that this observation is sufficient to single out what integrates our beliefs about the external world. However, for Williams, the problem here is that ‘the senses’, as a source of belief about the external world, can be divided into a “causal truism and a contentious epistemological doctrine” (Williamson 1996, 104).

¹³² Cf. Cassam (2007).

On the first version (the “causal truism”), our beliefs about the external world all have the senses as a common source in the sense that our perceptual organs are the basic causal source of our beliefs about the external world. The problem with this, however, is that it does not provide adequate support for “thinking of the senses as a source or channel of knowledge”, rather than as merely playing a crucial causal role in the genesis of beliefs” (Williams 1996, 71). On Williams view, if we take “the senses” as being what unifies our beliefs about the external world along epistemological lines, then we are “conflating genetic and epistemic considerations” (Williams 1996, 71).

On the second version (the “contentious epistemological doctrine”), our beliefs about the external world have a “common evidential ground” (Williams 1996, 104).¹³³ According to Williams, if our beliefs about the external world “name a coherent kind”, then that’s because those beliefs have a common evidential ground (Williams *ibid*).

5.5 Response to Williams’s Argument for Premise 1e

Should we accept that if our beliefs about the external world have a common shared nature, then what constitutes their shared nature is a common set of epistemological properties? The contentious part of Williams’s argument here is that the common set of epistemological properties are what Williams takes their common epistemological properties to be: a “common evidential ground” in the form of sense experience.

Perhaps this first issue one should raise is that it’s not obvious that the facts about the most basic causes of our beliefs about the external world have no bearing on their epistemic status. For example, it’s plausible that whether our perceptual capacities are reliable makes a difference to whether their causes,

¹³³ That Williams rejects that our beliefs about the external world can be topically integrated is brought out more fully in the following quotation I draw from in the main text: “Only by tracing our beliefs about the world to a common “source, which is to say a common evidential ground”, can we make “beliefs about the world” the name of a coherent kind” (Williams 1996, 104).

our beliefs about the external world, are justified or unjustified. So, it's not true (or at least not obviously true) that the senses, considered as the most basic causal source of our beliefs about the world, cannot be integrated under the same kind without "conflating genetic and epistemic considerations" (Williams 1996, 71).

I suspect that Williams can accept this claim without it doing damage to his argument. It's plausible that we move the senses from the non-epistemic "genetic" domain into the epistemic domain as soon as we start to talk about them as being reliable. However, a related problem emerges here. I will bracket that problem for now, as the problem will be clearer once we answer the second version of what Williams thinks "the senses" might mean.

Recall that, on the second version (the "contentious epistemological doctrine"), our beliefs about the external world have a "common evidential ground." What is their (putative) common evidential ground? On the one hand, we could take it that their common evidential ground is specified in terms of perceiving that p (where " p " is a proposition about the external world). The problem here, however, is that "the external world" is too problematic to characterize as a genuine kind (see sections 5.3-5.4 for an argument against this claim).¹³⁴ So, on Williams's view at least, the proponent of the traditional epistemological project needs to find something epistemological to characterise what unifies beliefs about the external world into a genuine kind.

Williams maintains that what is supposed to make our beliefs about the external world into a genuine kind is that these beliefs have a common evidential ground. Their common evidential ground is "inner experience", considered as appearances or seemings:

"External" in the phrase "external world", does not mean "in one's surroundings" but "without the mind." The essential contrast to external reality is inner experience. But if we ask why this contrast

¹³⁴ For example, Williams tells us that the phrase "[a]bout the external world" seems to thin and abstract a characterisation of our beliefs to mark out a significant epistemological boundary" (Williams 1996*b*, 368).

is theoretically significant, and particularly why it is epistemologically significant, there is only one answer. The contrast is regarded as significant because beliefs falling on one or the other side of the divide are thought to share a certain ultimate epistemic status: what really unites beliefs about the external world is how they are (or can be) known. Thus beliefs about the external world are all, in the last analysis, “inferential,” in contrast to experiential beliefs (or in some versions, experiences themselves) which are (at least relatively) more basic. From the standpoint of justification, beliefs about the external world depend on experiential beliefs in a way that experiential beliefs do not depend on them. This is, of course, the traditional doctrine of the “epistemological priority” of experiential beliefs over beliefs about the world. And it is an idea that is characteristic of a foundational conception of our knowledge of the world.

There is a lot in this passage. I will outline the main idea, and then work through the details.

First, what integrates our beliefs about the external world into a (putative) theoretical kind is their shared epistemic status: being justified on the basis of sensory experience. Second, on Williams’s view, their shared epistemic status presupposes a fixed, context-invariant epistemic priority relation between our beliefs about our sensory experiences, on the one hand, and our beliefs about the external world, on the other. For instance, on this view, our beliefs about the external world are inferred from (at least in part) our beliefs about our sensory experiences, while our beliefs about our sensory experiences are non-inferential. This is traditional foundationalism. Third, that there is this kind of epistemic priority relation between our beliefs about our sensory experiences and our beliefs about the external world presupposes epistemological realism.

As Williams puts it:

We can characterize foundationalism as the view that our beliefs, simply in virtue of certain elements in their contents, stand in *natural epistemological relations* and thus fall into *natural epistemological kinds* (Williams 1996, 116).

For this reason:

[...] foundationalism presupposes the more general outlook that I have called “epistemological realism”. The foundationalist is an epistemological realist because his view involves more than openness to the idea of an epistemology hierarchy, cutting across ordinary

subject-matter divisions and operating independently of all contextual factors: it involves definite conception of the *status* of this hierarchy (Williams 1996, 120-121).

In short, if our beliefs about the external world are integrated along epistemological lines, then traditional foundationalism is true.¹³⁵ But if traditional foundationalism is true, then epistemological realism is true.

Notice that there is therefore a weak and strong reading of premise 1e from Williams:

1e Strong: If the traditional epistemological project is legitimate, then traditional foundationalism is true (a form of epistemological realism).

1e Weak: If the traditional epistemological project is legitimate, then (some form of) epistemological realism is true.

Even if traditional foundationalism entails that epistemological realism is true, that traditional foundationalism is false does not show that epistemological realism is false. There need to be independent arguments against epistemological realism. For it leaves it open that there are non-traditional forms of foundationalism that are epistemological realist.¹³⁶

There are several questions we could ask here. I will focus on two of them. The first question is whether the fact that our beliefs about the external world have to be integrated along epistemological lines, that this implies that traditional foundationalism is true. On the face of it, this claim is implausible.

¹³⁵ Cf. “To treat “our knowledge of the world” as designating a genuine totality, thus as a possible object of wholesale assessment, is to suppose that there are invariant epistemological constraints underlying the shifting standards of everyday justification, which [...] allow us to determine, in some general way, whether we are entitled to claim knowledge of the world. But if this is so, foundationalist presuppositions are buried very deeply in the Cartesian project. They do not just fall out of the totality condition’s exclusion of any appeal to knowledge of the world in the course of our attempt to gain a reflective understanding of that knowledge. They turn out to be involved in the very idea of there being something to assess” (Williams 1994, 113).

¹³⁶ For example, Cassam (2009) presents what he calls a non-traditional form of foundationalism. I examine this view in Chapter 6. See also Buchanan (2002, 71) for the observation that Williams sometimes treats epistemological realism as nothing more than traditional foundationalism.

The second question is whether, even if some form of foundationalism is required to integrate our beliefs about the external world, this is a bad thing for those beliefs. On the face of it, this claim is not implausible. It will turn on whether that form of foundationalism is plausible, and whether epistemological realism is worth the price of that form of foundationalism.

5.5.1 The First Question

Traditional foundationalism combines a strong thesis about what our non-inferential beliefs are, together with the formal foundationalist thesis about the structure of epistemic justification. The strong thesis about our non-inferential beliefs is that all of our non-inferential beliefs have appearances or seemings as their subject-matter (e.g., It appears to S that *p*), and that all of our justified beliefs about the external world get their justification by being inferred from those non-inferential beliefs.

Williams is right that traditional foundationalism implies an epistemic priority relation between beliefs about appearances or seemings and beliefs about the external world, such that justified instances of the latter must be inferred from justified instances of the former, but not vice-versa. If traditional foundationalism is true then epistemological realism is true. What is implausible is that it has to be traditional foundationalism which integrates our beliefs about the external world along epistemological lines. All that epistemological realism requires is that there be objective epistemic facts and relations, so that there is an epistemic kind to investigate. It would be superfluous to suppose that traditional foundationalism is the only way to guarantee this. This brings us to the second question

5.5.2 The Second Question

To the question “what unifies your beliefs “about the external world” into a genuine totality; into a theoretical kind?” I argued that there is good non-epistemological answer to this question (see §5.3). According to this

supposition, there must be some epistemological way of integrating our beliefs into a genuine totality. Here's one intuitive way of doing it: it's those beliefs which are typically justified (if they are) by sensory experience. It doesn't follow from the fact that our sensory experiences are compatible with the falsity of what we believe that our beliefs are thereby less epistemically secure than beliefs about sensory experiences (and that they are justified only by being inferred from propositions about sensory experience). After all, there are more modest forms of foundationalism. One form says that our sensory experiences provide immediate (non-inferential) justification to our beliefs about the external world in virtue of the facts about the sensory experiences and their relation to the contents of our beliefs. Another form says that our sensory experiences can provide immediate (non-inferential) justification to our beliefs about the external world in virtue of the reliability of sensory experience in general. The question is whether it's possible for there to be an epistemological way of integrating our beliefs about the external world which isn't the traditional foundationalist's way of integrating our beliefs about the external world. And the answer seems to be that there is.

6 Williams's Master Argument: Premise 2

The second premise of Williams's master argument is that epistemological realism is false:

- (1) If the traditional epistemological project is legitimate, then epistemological realism is true.
- (2) Epistemological realism is not true.

Therefore,

- (C) The traditional epistemological project is not legitimate.

The problem, however, is that the traditional epistemological project presupposes epistemological realism, and it presupposes epistemological realism because it presupposes traditional foundationalism.

I argued that it's implausible that it presupposes traditional foundationalism—that the traditional epistemological project is legitimate just if traditional foundationalism is true. In particular, I argued that Williams's sub-argument for the premise that, if our beliefs about the external world are integrated along epistemological lines, then it's integrated along traditional foundationalist lines, fails to provide adequate support for that premise (see §5.5). I also argued that it's not clear that our beliefs about the external world cannot be integrated along topical, non-epistemological lines (see §5.3).

Still, it's not obvious that non-traditional forms of foundationalism do not presuppose epistemological realism. This would spell bad news for the traditional epistemological project if epistemological realism is false and that project presupposed that it isn't. The question we need to address is whether it is false (or rather: whether the arguments for its falsity provide adequate support for believing that it's false).

In outline, here is the argument against Williams's case against epistemological realism. The argument is that, on the one hand, if Williams is right that epistemological realism is false, then he proves too much. He could opt for less and still succeed (by his own lights at least). On the other hand, it's far less plausible that he is right. Epistemological realism survives Williams's objections. In short, it's far less clear that epistemological realism is false than that the traditional epistemological project is illegitimate. If epistemological realism is the price one pays for the legitimacy of the traditional epistemological project, perhaps it's not a bad price.

6.1 Williams's Argument for Premise 2

Should we accept epistemological realism? The second premise of Williams's master argument takes it that epistemological realism is false. What are his

arguments for this premise? His first argument is just that it does not flow from our epistemic practices:

Epistemological realism is by no means forced on us by evident features of ordinary justification. It represents a definite choice of theoretical orientation and is, in the sense, optional (Williams 2001, 171).

Now this argument tells us that a reason to accept epistemological realism is going to be a philosophical reason—that is, a reason which is the outcome of some philosophical argument in favour of it rather than reflection on our epistemic practices.¹³⁷ This alone need not count against epistemological realism. Not all epistemological theories need or even ought to look to epistemic practice for adequate support.

However, I take it that what's implicit in Williams's argument is that there is no compelling philosophical reason to accept epistemological realism. In fact, Williams suggests that there are compelling reasons to reject it. The compelling reason is that epistemological realism “encourages scepticism” (Williams 2001, 196).

Lots of philosophers have said that realism encourages scepticism. “Realism” here is the thesis that the world exists independently of anyone's (or everyone's) psychological states. But almost none of them have taken this as a reason to think that realism is false. “Encouraging”, “inviting”, or “making room for” scepticism is not the same as entailment or providing adequate support for scepticism. Indeed, this suggestion is compatible with Williams's thesis that traditional foundationalism implies scepticism and traditional foundationalism is an epistemological realist view. For one can accept that traditional foundationalism implies scepticism but deny that all forms of

¹³⁷ Here's a similar suggestion: “Norms, including epistemic norms, are standards that we set, not standards imposed on us by ‘the nature of epistemic justification’. A belief is no more justified, wholly independently of human evaluative attitudes, than a certain kind of tackle in football is a foul, wholly independently of our practices of judging certain types of tackle to be against the rules” (Williams 2001, 170). The thought here seems to be that our ordinary epistemic practices don't seem to lend any support at all to epistemological realist theories.

epistemological realism do. It leaves it open that there are non-sceptical epistemological realist theories.

This leads us to the following concern. Williams ties epistemological realism to traditional foundationalism. The thought is that our beliefs about the external world have to be integrated but nothing topical can do it, so something epistemological has to do it. I argued something topical could do this, but I also argued that even if something epistemological had to do it, it wasn't clear that it had to be traditional foundationalism (see §5.3 and §5.5).

Let's suppose that Williams is right about this. If he is right, then the denial of epistemological realism is superfluous. After all, if traditional foundationalism entails that scepticism is true, this counts against traditional foundationalism. In effect, the denial of epistemological realism here would prove too much: all the theories which take it that knowledge of the external world is a real phenomenon, a real kind of thing for which theoretical explication is possible, Williams would be rejecting all of them in order to show that the traditional epistemological project is illegitimate.

Here is a related complaint. Let's suppose that epistemological anti-realism is true. So, suppose that "knowledge of the external world" is one of the methodological necessities of the traditional epistemological project. This means that it is presupposed in order to make the project legitimate. Now Williams think that given that presupposition, Stroud's dilemma provides adequate support for meta-epistemological scepticism: a satisfying philosophical explanation of how knowledge of the external world is possible is not possible.

But then it's hard to see how Williams's disagreement with Stroud is substantive. After all, what makes the traditional epistemological project legitimate turns out to be contextual. On this view, in a certain context the traditional epistemological project is legitimate. So, in that context, meta-epistemological scepticism is true. This result is bad, however, because it seems as if epistemologists would still want to dispute whether, in the context of the

traditional epistemological project, meta-epistemological scepticism is true. So, it's not obvious that Williams has a substantive disagreement with Stroud after all.

A third related issue is that epistemological realism might be true because knowledge is a natural kind. This is at least given some prima facie motivation by the fact that epistemologists take themselves to be investigating a state or feature of the mind, and that this is the same state or feature that is shared with humans and non-human animals, and that it's this state or feature which biologists, ethologists, and cognitive scientists appeal to in order to make sense of certain human and non-human animal behaviours.¹³⁸ At least, these are concerns that Williams should speak to.

A weaker idea here is that epistemic facts supervene on non-epistemic, natural facts. For example, facts about justification supervene on facts about the cognitive capacities and the reliability of these cognitive capacities. If this supervenience thesis is true, it doesn't entail that knowledge (of some kind) is a natural kind, but it would make it look plausible that there is a domain of epistemic facts for which theoretical explication is possible.

7 Conclusion

This chapter focused on Williams's master argument against Stroud's dilemma. In particular, Williams's master argument targets what I called the "legitimacy thesis". This is the thesis that the traditional epistemological project is legitimate (see §4). Williams's master argument said that the traditional epistemological project is legitimate only if epistemological realism is true but epistemological realism is not true. I argued against both premise of the

¹³⁸ Here, one can leave it an open question as to *how* best to investigate the state or feature. It might be that intuition and other *a priori* "armchair" methods are less useful than epistemologist might have hoped. On the other hand, it might turn out that these methods are not irrelevant to an understanding of the state or feature. This is not the issue. The issue is that it's not as obvious as Williams suggests that knowledge *could not be* a natural kind, much less something which admits of general theoretical explication.

argument (see §5). The arguments I presented attempted to show that William's arguments for the two premise of his master argument fail to provide adequate support for those premises. So, I did not argue that the traditional epistemological project is legitimate after all. Instead, I defended it from the objection that it is not legitimate. Insofar as the traditional epistemological is legitimate, then, Stroud's dilemma poses a challenge. The challenge is that if the argument from Stroud's dilemma is valid and sound, the meta-epistemological scepticism is true.

Chapter 5

Stroud's Dilemma: Internalism and Externalism¹³⁹

1 Introduction

In Chapter 4, I argued that Williams's (1996) master argument against Stroud's dilemma fails to provide adequate support for the thesis that the traditional epistemological project is illegitimate. This does not entail that it's not illegitimate. But absent some other reason to think that it's illegitimate, then perhaps the best thing to do is to show that Stroud's dilemma fails to provide adequate support for meta-epistemological scepticism.

In this chapter, I examine one attempt to show this from Sosa (1994, 1997) and per force the proponent of epistemological externalism (henceforth "externalism" for short). I argue that Sosa's arguments fail to show that Stroud's dilemma does not provide adequate support for meta-epistemological scepticism.

There are interpretive difficulties between Sosa (1994, 1997) and Stroud (1994, 2000, 2004). There are difficulties for at least two reasons. The first reason is that it's not clear what the source of Stroud's objection to Sosa's (1994, 1997) externalist response to Stroud's dilemma is supposed to be. The first interpretation is that the source of Stroud's objection to Sosa's (1994,

¹³⁹ Parts of this chapter were presented at the Orange Beach Epistemology Workshop (May 2014), and published in *Traditional Epistemic Internalism (forthcoming)*, Oxford University Press, co-authored with Duncan Pritchard (his contributions have been omitted). I would like to thank Michael Bergmann, Richard Fumerton, Kevin McCain, Ted Poston, and Chris Tucker for helpful suggestions.

1997) “externalist” response is that it’s not an “internalist” response. In section 3.1 I argue that this interpretation is mistaken. The second interpretation is that the source of Stroud’s objection to the externalist response is that it’s circular. In section 3.2, I argue that this interpretation is mistaken.

The second reason is that Stroud thinks that Sosa has missed the point of his argument (e.g., Stroud’s dilemma) for meta-epistemological scepticism.¹⁴⁰ This is right, but there are important insights from Sosa that should not be overlooked.¹⁴¹ Moreover, while I defend the thesis that the externalist response to Stroud’s dilemma is not successful, this turns on a more robust defence of Stroud’s dilemma that Stroud has given.

In this chapter, I do several things. The first thing that I do is straighten up the debate between Sosa and Stroud. This is worthwhile because it will advance the debate on meta-epistemological scepticism. Once this is done, we have a backdrop from which to assess Sosa’s arguments against Stroud’s dilemma.

The second thing that I do is argue for (what I will explain later in section 3.2 is) an “internalist” meta-epistemological requirement on satisfying philosophical explanations of how knowledge of the world is possible. As we will see, while Sosa thinks that Stroud maintains an “internalist” meta-epistemological requirement, it’s not clear that he does. So the argument I provide in favour of that requirement is not Stroud’s argument, but it does lend support to meta-epistemological scepticism.

In section 2, I adapt Sosa’s (1994, 1997) response to Stroud’s dilemma. In section 4, I explain Stroud’s response. In section 5, I extend Stroud’s response and defend it from objections. In particular, I argue for an “internalist” constraint on understanding how knowledge of non-psychological reality possible.

¹⁴⁰ For example, see Stroud (2000), pp. 144-145, and (2011), pp. 322-323.

¹⁴¹ For example, that any fully general theory of what knowledge is will be circular but this is not necessarily a defect in the theory. See Sosa (1994), pp. 286-290.

2 Externalist Responses to Stroud's Dilemma

Stroud's dilemma aims to provide adequate support for meta-epistemological scepticism. The first horn of the dilemma is that if EPR is true, then scepticism is true. In turn, one's philosophical explanation of how knowledge of the world is possible turns out to be negative. The second horn of the dilemma is that if EPR is false, then the traditional epistemological project cannot be satisfied, the outcome of which is that no fully general, philosophical explanation of how our knowledge of the world is possible is possible. In either case, the resulting explanation is not satisfying, and therefore meta-epistemological scepticism is true.

An externalist has at least two different places to take issue with Stroud's dilemma. First, the externalist can dispute the first horn of the dilemma, and argue that it doesn't follow from EPR that scepticism is true. Here, what will make their response an "externalist response" is that EPR can be fulfilled, but one does not need to know that the linking-principle between the prior base and the posterior base is true (for example, that sensory experiences as of p are reliable indicators that p). Instead, it's enough that a linking-principle is true.

Second, the externalist can accept the demand that the epistemologist must know that the linking-principle is true, and argue that that demand can be met. This is Sosa's (1994) response to Stroud's dilemma. Here, what makes the response an "externalist response" is the explanation of how the demand is to be met. The explanation of how the demand is to be met will turn on the fulfilment of externalist principles about epistemic justification or knowledge.

What we need to address is whether the externalist's responses to Stroud's dilemma are successful. I will argue that neither of the two externalist responses are successful. The first argument is complicated because we need to clear up the debate between Sosa and Stroud. So, that's what I will do next (sections 2.1-2.2). The second argument is for (what I explain later in section 2.2. is) an "internalist" meta-epistemic requirement.

2.1 The First Externalist Response

Stroud's (2000) argument for premise 1 of Stroud's dilemma is the two step argument that, first, the epistemologist's explanation must respect EPR, and second, that the epistemologist must know that a linking-principle between the prior base, the explanans, and the posterior base, the explanandum, is true.¹⁴² However, the epistemologist cannot jointly fulfil both requirements, because knowing that the linking-principle is true either amounts to knowing a proposition about non-psychological reality, or else it relies on one's knowledge of non-psychological reality.

Let's take a closer look at that argument:

If we really are restricted in perception to 'experiences' or 'sense-data' or 'stimulations' which give us information that is prior to any knowledge of objects, how could we ever know anything about what goes on beyond such prior 'data'? It would seem to be possible only if we somehow knew of some connection between what we are restricted to in observation and what is true in the wider domain we are interested in. But then knowing even that there was such a connection would be knowing something about the wider domain after all, not just about what we are restricted to in observation. And then we would be left with no satisfactorily general explanation of our knowledge (Stroud 2000, 105-106).

Stroud's basic idea here is that if one's explanation is to respect EPR, and one's explanation appeals to a general source of knowledge, such as sense-perception, then one should not explain how knowledge of the world is possible on the basis of knowledge-entailing mental-states, such as perceiving that *p* (where "*p*" is a proposition about non-psychological physical reality). Rather, one should conceive of their knowledge from sense-perception as not

¹⁴² In the quotation, Stroud expressed this requirement the second sentence. Compare with the following quotation, where Stroud responds to Sosa's (1994): "If there are truths of this kind [e.g., truths linking sensory experiences with facts about the world], although no one has discovered them yet, that fact alone obviously will do us no good as theorists who want to understand human knowledge in this philosophical way. At the very least we must believe such truths; their merely being true would not enough to gives us any illumination or satisfaction. But our merely happening to believe them would not be enough either. We seek understanding of certain aspects of the human condition, so we seek more than just a set of beliefs about it; we want to know or have good reason for thinking that what we believe about it is true" (Stroud 2000, 146 my addition).

amounting to knowledge of the world, but instead as amounting to knowledge of how things seem or appear.

The claim here is not that one cannot come to know that p on the basis of seeing or otherwise perceiving that p (where “ p ” is a proposition about non-psychological physical reality). Rather, the claim is conditional: if one’s explanation is to respect EPR, then their explanation of how knowledge of the world is possible cannot be in terms of perceiving that p (or any knowledge-entailing mental-states).

So, let’s suppose that our explanation fulfils EPR. If that’s right, Stroud thinks that it’s not enough for us to know only facts about seemings or appearances (e.g., that it seems to me that p), because what I want to explain is how knowledge of the world is possible at all, not just knowledge of how the world seems or appears to be. This is supposed to motivate the following requirement: the epistemologist engaged in the traditional epistemological project needs to know that certain epistemic principles linking seemings or appearances with facts about the world are true (e.g., that if it seems to S that p , that it is likely that p).

Stroud’s argument then goes as follows: if the epistemologist knows that the linking-principle P is true, either this entails or presupposes knowledge in the domain in question or does not. If knowing that P entails or presupposes knowledge in the domain in question, then EPR is violated. Now let’s suppose that EPR is not violated. It follows that S’s knowing that P is admissible, but Stroud thinks that there should be some explanation of how S knows that P . This explanation needs to respect EPR as well. But if it respects EPR, then it’s hard to see how it’ll help to explain how knowledge in the domain in question is possible. Stroud expresses this argument as follows:

If it had been shown that there is a certain “postulate” or “principle” which we have to have some reason to accept if we are to know anything about, say, the world around us, we would not thereby have come to understand how we do know anything about the world around us. We would have identified something we need, but its indispensability would not show that we do in fact have

good reason to accept it. We would be left with the further question whether we know that that “principle” is true, and if so how. And all the rest of the knowledge we wanted to explain would then be hanging in the balance, since it would have been shown to depend on that “principle.” Trying to answer the question of its justification would lead right back into the old dilemma. If the “principle” involved says or implies something richer than anything to be found in the prior evidential base—as it seems it must if it is going to be of any help—there will be nothing in that base alone that could give us reason to accept it. But if we assume from the outset that we do know or have some reason to accept that “principle,” we will be assuming that we already know something that goes beyond our prior evidential base, and that knowledge itself will not have been explained. We would therefore have no completely general explanation of how we get beyond that base to any knowledge of the kind in question (Stroud 2000, 106-107).

Prima facie at least, the externalist is in a position to avoid this consequence. The externalist can argue that insofar as there is a true linking-principle P between the prior base and the posterior base (e.g., between appearances as of p and the fact that makes p true), then the epistemologist will be in a position to understand how knowledge in the posterior base is possible without disrespecting EPR. There is no need for the additional requirement that the epistemologist know or have a reason to believe that the the linking-principle P is true.

For example, if the prior base consists of propositions about seemings or appearances (e.g., that it appears to S that p) and the posterior base consists of propositions about non-psychological physical reality, and what we want to explain is how it's possible to know that any members of the latter are true without disrespecting EPR, it is sufficient that the following principle be true: appearances as of p are reliable indicators that p .

Now it's not enough for the linking principle P to be true without the epistemologist knowing that P is true. The crucial point here is to see that this does not entail that one cannot know that p is true on the basis of an appearance as of p unless one knows that appearances as of p are reliable indicators that p . The claim is about what the epistemologist needs to know if their aim is to

fulfil the traditional epistemological project. So, the question that needs to be addressed is not whether it's consistent with externalism that an agent can come to know that *p* on the basis of appearances of *p* without knowing that a linking principle between appearances as of *p* and *p* is true.

The question is whether the agent—*qua* epistemologist—will be in a position to fulfil the traditional epistemological project, if the epistemologist knows that she has appearances as of *p*, and believes that appearances as of *p* are reliable indicators as of *p*, but unbeknownst to her, appearances as of *p* are reliable indicators that *p*. One might think that one's merely believing the correct and positive explanation of how knowledge of the world is possible is alone not sufficient for having a satisfying philosophical explanation of how that type of knowledge is possible (cf. Chapter 1, §3).

This takes us to the debate between Sosa and Stroud. As we will see, Sosa thinks that the epistemologist does need to know that some such linking principle is true. Here, then, is Stroud on why a linking principle needs to be known:

If many people knew that there are external things, but no one knew that anyone had that knowledge, then no one would know the answer to the epistemological question. There would *be* a positive answer—the existence of external things would be something known and not just an article of faith—but no one would know it. If many people after epistemological investigation came to believe that human beings know that there are external things, that would not be enough for a satisfactory outcome of their epistemological investigations, just as everyone's believing that there are external things was not enough to settle the question of their knowledge. Even if the investigators' belief were in fact true, that still would not be enough. The truth of the answer they accept would not give them the understanding they seek unless they could recognize that they know or have good reason to believe that answer. Rightly finding themselves with knowledge or good reason to believe that the answer they accept is true would give them a satisfactory understanding of human knowledge (Stroud 2004, 166).

The point here is that even if the epistemologist knows that *p* (and therefore has knowledge of the world), their knowing that *p* alone is not enough to explain

how it's possible. The epistemologist would also need to know something more, such as that appearances as of p are reliable indicators that p .

2.2 The Second Externalist Response

The first horn of Stroud's dilemma moves from the fulfilment of EPR to scepticism. As Cassam (2009) expresses the point:

On the one hand, the intellectual goal of understanding human knowledge in general [...] leads one to accept EPR[.] Yet as soon as we try to think of our knowledge as arranged in level of epistemic priority we find it impossible to explain how the epistemically prior knowledge (of, as it might be, the character of our experiences) 'could yield any richer knowledge lying beyond it' (Cassam 2009, 573).¹⁴³

However, the second horn of Stroud's dilemma moves the rejection of EPR to the failure to fulfil the traditional epistemological project. Consider Cassam (2009) again:

The second horn results from the *rejection* of EPR. [...] However, if this is the line we take then we will find ourselves to give an appropriately general explanation of how any knowledge of any objects at all is possible by means of the senses. So the epistemic priority requirement is one that we can neither live with nor live without (Cassam 2009, 573).

In either case, meta-epistemological scepticism follows. The dispute between Sosa and Stroud occurs at the first horn of Stroud's dilemma. Sosa (1994, 1997) maintains that an externalist explanation of how knowledge of the world is possible can fulfil EPR and not result in scepticism. Stroud, on the other hand, does not dispute that the externalist explanation does not result in scepticism, but for Stroud it nevertheless still results in meta-epistemological scepticism because it fails to fulfil additional requirements for being a satisfying philosophical explanation of how knowledge of the world is possible.

Recall that Stroud's argument for the first horn of the dilemma is that the epistemologist would need to know that the linking-principle between the prior domain and the posterior domain (e.g. facts about appearances and facts about

¹⁴³ The single quotations are Stroud (2000, 120).

the external world). However, Stroud argues that the epistemologist could not explain their knowledge of the linking-principle without violating EPR, rendering their explanation inadequate, or, failing to violate EPR, a sceptical explanation.

Now Sosa's externalist response accepts that the epistemologist needs to know that the linking-principle is true. In this fashion, Sosa disagrees with the first externalist response (see §2.1). However, Sosa thinks that Stroud's argument against the externalist's explanation of how linking-principle is known turns on his commitment to the following "meta-epistemic requirement":

In order to understand one's knowledge satisfactorily one must see oneself as having some reason to accept a theory that one can recognize would explain one's knowledge if it were true (Sosa 1994, 272).

Let's adapt this to the language we have been using:

- Meta-epistemic Requirement: If *S*'s understands how knowledge of the world is possible on the basis of a philosophical explanation *E*, then *S* must "see oneself" as having a reason to believe that *E* is true, and *S* must know that if *E* were true, then *E* would explain how knowledge of the world is possible.

Let's suppose that one element of one's explanation is that sensory experiences are reliable indicators that *P* is true. This would be an epistemic linking-principle between one's epistemically prior base (e.g., facts about sensory experiences) and one's epistemically posterior base (e.g., facts about the world). For example, if I have a visual experience as of a tree in front of me, that's a reliable indication that a tree is in front, and so that the belief that a tree is in front of me is true.¹⁴⁴ The meta-epistemic requirement requires us to "see" ourselves as having a reason to believe that our explanation is true.

¹⁴⁴ See Alston (2010) for a discussion of this type of reliabilism.

As I understand the phrase “see oneself as having a reason”, it picks out an internalist conception of reasons. It is “internalist” in the sense that it requires one to have reflective access to one’s reason.¹⁴⁵ On this conception, the demand is not just that the epistemologist have a reason to believe that their explanation is true, but that the epistemologist have reflective access to their reason (that they understand why it’s true). With the distinction between having a reason and having reflective access to one’s reason, we can distinguish between two versions of the meta-epistemic requirement as follows:

- Internalist meta-epistemic requirement: If S understands how knowledge of the world is possible on the basis of a philosophical

¹⁴⁵ Epistemological internalism comes in different forms. These should not be confused with internalism about mental content, where “mental content” is the name for what one’s state of mind is about or directed at (in short, one’s representational content). The internalist about mental content maintains that at least some mental content is “internal” in the following sense: if a mental state of a person *S* has a representational content *C*, then *C* supervenes only on *S*’s “intrinsic properties”. Here, “intrinsic properties” denotes those properties that would survive *S*’s transportation to twin-Earth or envatment over time. Epistemological internalism comes in at least two forms: “accessibilist” internalism and “mentalist” internalism. For discussion, see Pryor (2001). In both cases, what is “internal” is some epistemic property *F*, or some property which determines that one’s belief has *F* (one can find both kinds of internalism in the literature. For example, see Conee and Feldman (1985) for the “mentalist” version, and Fumerton (1995) for a more “accessibilist” version. Other internalists include Brewer (1996), Huemer (2013, 2007, 2000) and Pryor (2000) and Smithies (2014)). The “accessibilist” maintains that this epistemic property can be always potentially be known to the person. The person can come to know, through reflection alone (e.g., through introspection or *a priori* reasoning) that one’s justification for their belief that *p* is *F*. The “mentalist” maintains that the epistemic property *F* is “internal” in the sense that *F* is a mental state or supervenes only on the person’s mental states. Now, the mentalist does not want to include “wide” mental states, mental states whose representational content is determined by one’s environment. Instead, they only include “narrow” mental-states, states whose representational content is “internal” in the sense specified above (e.g., it does not depend on one’s environment). For the relationship between accessibilism and mentalism, see Madison (2010). It is also worth noting that epistemological “externalism” is the name for a broad range of theories, each of which have the denial of epistemic internalism in common. For example, the causal theory of knowledge (e.g., Armstrong 1973), process reliabilism (e.g., Goldman 2008, 1999, 1993), indicator reliabilism (e.g., Alston 1993), truth-tracking theories (Nozick 1981) information theories (e.g., Dretske 2000), some virtue theories (e.g., Sosa 2007, Greco 2000), some relevant alternatives theories (e.g., Dretske 1970), and some forms of contextualism (e.g., DeRose 1995, 2009). See Bergmann (2006, 2000) for a detailed discussion of externalist responses to scepticism. Pritchard (2012) represents epistemological disjunctivism as a form of accessibilist internalism.

explanation *E*, then *S* must have reflective access to their epistemic support in favour of *E*.

- Externalist meta-epistemic requirement: If *S* understands how knowledge of the world is possible on the basis of a philosophical explanation *E*, then *S* must have epistemic support in favour of *E*.

A few remarks are in order. First, I am using the term “philosophical understanding” as the cognitive analogue of “philosophical explanation”. So, when I use the phrase “if the epistemologist has a philosophical understanding of how knowledge of the world is possible, then...”, one can replace “philosophical understanding” with “grasps the philosophical explanation...”. In turn, whatever requirements on a satisfying philosophical explanation of how knowledge of the world is possible will be requirements on a satisfying philosophical understanding of how knowledge of the world is possible.

Second, I drop the phrase “having a reason to believe” because most epistemologists talk less about reasons and more about justification, warrant, evidence, or entitlement. I use the more general term “epistemic support” for some positive epistemic status that can be conferred onto one’s belief that *p*. A related point here is that “having a reason to believe that *p*” seems more like an internalist epistemic concept, so the hope is that “epistemic support in favour of *p*” is not just more general but neutral.

Third, I drop the second condition of the meta-epistemic requirement, that the epistemologist must know that if their explanation *E* were true, then *E* would explain how knowledge of the world is possible, because this condition is not disputed between Sosa and Stroud. It is the first condition which is in dispute, and for this reason I focus on this element of the meta-epistemic requirement.

Fourth, I’m not providing an analysis of “reflective access of one’s epistemic support in favour of *p*”. I am using that phrase as short hand for “one is in a position to know, on the basis of reflection alone, that such-and-such is

their epistemic support in favour of believing that *p*”, and this is the orthodox use of that phrase and similar phrases in the literature. In short, “reflective access to one’s epistemic support in favour of *p*” means “one is in a position to know, on the basis of *a priori* reasoning or introspective awareness alone, that such-and-such is one’s epistemic support in favour of believing that *p*”.¹⁴⁶

Should we adopt the internalist meta-epistemic requirement or the externalist meta-epistemic requirement? Sosa (1994) thinks that the externalist rather than the internalist version needs to be fulfilled:

[...] one must in fact *have* good reason or at least justification to accept some appropriate explanation. Why must one also *see oneself as having* such a reason? (Sosa 1994, 273).

However, there are *prima facie* considerations against the externalist meta-epistemic requirement. And this makes the internalist meta-epistemic requirement look more plausible. For example, consider the following case, an application of Bonjour’s (1985) clairvoyance case, which makes the internalist epistemic requirement look plausible:

A and B are both epistemologists. Both A and B are interested in answering the question: “how is justification possible at all?” The regress argument seems to show that it’s impossible, so both A and B seek to explain how justification in general is possible. A is a foundationalist about the architecture of justification. A thinks that we can have justification to believe that *p* which is non-inferential, and so does not depend for its justification on other propositions we are justified in believing. B, however, is a coherentist about the architecture

¹⁴⁶ Cf. Pryor (2001, 103-104). A further complication are the terms “*a priori*” and “*introspective awareness*”. I’m assuming that one knows *a priori* (or by introspective awareness) that *p* if and only if one’s way of knowing (or coming to know) that *p* is independent of perceptual experience of the world. And I’m assuming that “independent of perceptual experience of the world” is *epistemic* independence. So, for example, I might know *a priori* that if it’s raining outside, then it’s raining outside, even if I wouldn’t know that proposition unless I had some perceptual experience of it raining outside. Yet my perceptual experience of it raining outside is an enabling condition of my knowing *a priori* that if it’s raining outside, then it’s raining outside, rather than an epistemic basis for this knowledge.

of justification. B thinks that we cannot have non-inferential justification. All of our justification, according to B, is grounding in inferences from other propositions we have justification to believe. Suppose, however, that A is right and B is wrong: foundationalism is true, and so how it describes the architecture of justification reflects its nature. But suppose that A is incapable of reflecting any reasons he seems to have for endorsing foundationalism. An evil demon prevents A from engaging in reflective thought and introspection as it concerns matters pertaining to foundationalism and coherentism. So B thinks that A is a poor epistemologist, since whenever B tries to engage A with philosophical argument about the architecture of justification, A is unable to report his (putative) reasons for endorsing foundationalism, and his (putative) reasons against coherentism. A just draws ‘mental blanks’, unable to speak in favour of foundationalism by philosophical argument, examples, or thought-experiments. But what the demon does not do is prevent A from believing foundationalism. And as it turns out, A believes that foundationalism is true because it was the product of a reliable belief forming process, a process which, unbeknownst to A, manifests his strange skill at detecting true philosophical theories from false philosophical theories. Since the truth of interesting philosophical theories are difficult to decide, A has never been able to detect this strange skill.

Call this case the “unreflective epistemologist”. Here’s the intuition about the unreflective epistemologist: A does not have a satisfying philosophical understanding of how justification is possible. Yet A believes that foundationalism is true (a philosophical explanation of how justification is possible) as the result of a reliable belief forming process. A is also correct: foundationalism is true.

A few remarks are in order. Here, I am stipulating that, absent defeaters to believe otherwise, or to withhold believing, believing that p as the result of a reliable belief forming process can give one epistemic support in favour of p . This is therefore not in question, since it is being stipulated. The question is instead: if having epistemic support in favour of one's philosophical explanation of how such-and-such is possible is required for it to be in the market for philosophical satisfaction, is it enough that the epistemic support is not of the internalist sort at issue in the internalist meta-epistemic requirement?

The crucial thing to notice is that A's philosophical explanation of how justification is possible is believed, true, and A has some epistemic support to believe that their philosophical explanation is true. A's epistemic support is the fact that his belief that foundationalism is true is the product of a reliable belief forming process. Yet if having some epistemic support in favour of one's philosophical explanation is required, it's hard to see how A's fulfilment of the externalist meta-epistemic requirement puts A in a position to get it.

To help strengthen the case, imagine now that the demon stops preventing A from being able to access his epistemic support in favour of foundationalism. A is now in a position to report his reasons for believing foundationalism, and to produce philosophical arguments, examples, and thought-experiments in favour of foundationalism, all of which do provide adequate support for foundationalism. In this case, one could maintain that A went from not understanding to understanding how justification is possible; or A went from a deficient and therefore less than fully satisfying philosophical understanding to a satisfying philosophical understanding of how justification is possible; or one could maintain that there's no epistemological improvement whatsoever.

The first reaction seems to be the natural reaction. After all, it's just hard to see how A understands at all how justification is possible. That A's belief is true, and his belief was, unbeknownst to A, the product of his reliable clairvoyance for detecting true philosophical theories and avoiding false

philosophical theories, doesn't seem to put him in a position to understand how justification is possible.

This isn't a "knock down" argument against the externalist meta-epistemic requirement. But it seems to provide some *prima facie* support against the externalist meta-epistemic requirement and in favour of the internalist meta-epistemic requirement.

This also brings us back to the question of Sufficiency (see Chapter 1, §3). Put in general terms, Sufficiency is the thesis that if a philosophical explanation of how some epistemic property *F* is possible is a positive and correct explanation, then a satisfying philosophical explanation of how *F* is possible is possible. Sufficiency connects the positivity and correctness of a philosophical explanation with philosophical satisfaction. We can add to Sufficiency the following condition about epistemic support:

Sufficiency*: If a philosophical explanation of how some epistemic property *F* is possible is positive, correct (i.e., true), and believed on the basis of epistemic support which in fact supports that belief, then a satisfying philosophical explanation of how *F* is possible is possible.

The unreflective epistemologist case doesn't speak against Sufficiency*. Instead, it speaks against one formulation of Sufficiency*. It speaks against the formulation in which fulfilment of the epistemic support condition is externalist rather than internalist (that is, where the epistemic support relation doesn't require one to have reflective access to that fact it provides that support). Indeed, it's not clear whether fulfilment of the epistemic support condition along internalist lines would be sufficient either. Whether or not it would be sufficient turns on whether there are other plausible requirements that we are missing. This is why focusing on specific necessary conditions for the possibility of a satisfying philosophical explanation of such-and-such is far more fruitful than focusing on sufficient conditions. *Prima facie* at least, it's

plausible that it would be sufficient but, as I will argue in section 3.1, Stroud doesn't seem to think that it is.

3 Stroud's Argument: Two Interpretations

In the previous section, I argued that the unreflective epistemologist case is a prima facie case against the externalist meta-epistemic requirement. This suggests that the internalist meta-epistemic requirement is true.

There are two interpretations of Stroud's anti-externalist argument in the literature. The first interpretation—call it the “internalist interpretation”—maintains that Stroud argues for internalism about perceptual knowledge, and that his anti-externalist view flows from his internalism. The second interpretation—call it the “circularity interpretation”—maintains that Stroud argues that externalism permits an unsatisfactory circular explanation of how knowledge of non-psychological physical reality is possible, and for this reason an externalist explanation of how this type of knowledge is possible is not a satisfying explanation.

3.1 The Internalist Interpretation

Let's begin with the internalist interpretation. If people know that there are tables, then people know that there mind-independent, external objects. This would be a positive answer to G. E. Moore's (1939) question of whether the existence of external things is known, rather than just an article of faith. However, Stroud reminds us that:

[i]f many people knew that there are external things, but no one knew that anyone had that knowledge, then no one would know the answer to the epistemological question. There would be a positive answer—the existence of external things would be something known and not just an article of faith—but no one would know it. [...] The truth of the answer they accept would not give them the understanding they seek unless they could recognize that they know or have good reason to believe that answer. (Stroud 2004, 166).

This passage does not express a first-order thesis about knowledge. Instead, it expresses a meta-epistemological thesis about answering a certain kind of philosophical question. The philosophical question at issue is what I call the “traditional epistemological project”. This, recall, is the project of providing a satisfying philosophical explanation of how knowledge of non-psychological physical reality is possible.

Stroud’s point here is that in order for the epistemologist to have a satisfying philosophical understanding how knowledge of non-psychological physical reality is possible, it’s necessary that the epistemologist know or “recognize” that knowledge of non-psychological physical reality is possible (Stroud *ibid*). We can express this requirement as follows:

Stroud’s Requirement: If an epistemologist *S* has a satisfying philosophical understanding of how knowledge of non-psychological physical reality is possible, then *S* knows that knowledge of non-psychological physical reality is possible.

In order to appreciate the point that Stroud’s Requirement is not a first-order thesis about knowledge, consider the distinction between the thesis that *S* knows that *p* only if *S* knows that *S* knows that *p* and the thesis that *S* has a satisfying philosophical explanation of how knowing that *p* is possible only if *S* knows that knowing that *p* is possible.

The first thesis is a thesis about knowledge. It is a first-order thesis about the requirements for knowledge in general, which tells us that if *S* knows that *p*, then *S* knows that *S* knows that *p*. The second thesis, however, is a thesis about the requirements for a satisfying philosophical explanation of how knowledge is possible. And it’s not a thesis about the requirements for a satisfying philosophical explanation of how *S* knows that *p* either. So, while there might be an interesting relationship between first-order internalist theses and Stroud’s Requirement, failing to distinguish them encourages a misreading of Stroud.

Now let's consider another internalist reading of Stroud. Fumerton (2004) distinguishes between modest and immodest internalist epistemological goals. The goal of each type of internalist is to philosophically understand how knowledge of some sort is possible, and to do so by respecting internalist principles. According to Fumerton, Stroud is our paradigmatic immodest internalist, because he "seems to locate the epistemologist's target as second-level knowledge (or understanding)" (Fumerton 2004, 181). Fumerton summarises Stroud's (putative) internalist requirement as follows:

In trying to understand whether and how we know various propositions in a given field of knowledge, we cannot presuppose that we know or even reasonably believe any propositions the alleged knowledge of which we are investigating. [...] In trying to understand how we know various propositions in a given field of knowledge, we cannot presuppose the legitimacy of any of the methods we employ in coming to believe propositions of the sort in question, and therefore cannot use any of those methods in studying the knowledge in question. (Fumerton 2004, 181-82)

Prima facie, however, it's not obvious that these constraints are internalist at all. Instead, they express constraints on philosophically understanding how knowledge of certain types are possible. They say that we won't philosophically understand how (say) perceptual knowledge of the world is possible, if our explanation features arguments with premises about the external world, and if we presuppose the legitimacy of any of the rules, methods, or principles that are indispensable to the acquisition of knowledge of the world. But none of them imply or presuppose that S knows that *p* only if S is in a position to know, by reflection alone, that they know that *p*, or that S knows that *p* only if what fixes S's epistemic support for *p* are S's internal mental states.

3.1.1 Internalism and the "Straightforward" Answer

On Stroud's view, seeing or in general perceiving that *p* is a form of knowing that *p*, such that perceiving that *p* entails knowing that *p*, where "*p*" is a

proposition about non-psychological physical reality.¹⁴⁷ How does this kind of perception interact with the traditional epistemological project? Stroud tells us that an answer to the traditional epistemological project which tells us that it's possible to know that p on the basis of perceiving that p "is one answer to a question about our knowledge of external things; what might be called the most straightforward answer." (Stroud 2004, 166). Following Stroud, let's call this the "straightforward answer" to the traditional epistemological project.

While Stroud believes that the straightforward answer is true, in that it is possible to know that p on the basis of perceiving that p , on his view it cannot offer us a satisfying philosophical understanding of how knowledge of this kind is possible. This might seem odd, since Stroud thinks that the straightforward answer is true, and that its truth implies that we know about non-psychological physical reality. One puzzle, then, is what's wrong with the straightforward answer, if it's not that it's false?

Different answers to this question are available. Perhaps one problem is that the straightforward answer fails to explain how philosophers know that we can perceive that p (and therefore know that p). According to this objection, one can accept that if we can perceive that p , then we can know that p (again, where " p " expresses a proposition about non-psychological physical reality). The problem is how we know that we can perceive that p . In short, the problem is that the second-order knowledge is left unaccounted for.

Now recall Stroud's Requirement. Stroud's Requirement requires the philosopher to have the second-order knowledge that knowledge of non-psychological physical reality is possible. So, in this fashion, the internalist interpretation of Stroud (e.g., Fumerton 2006; see §3.1) might be on the right track, with their emphasis of second-order knowledge being an issue for Stroud. We can formulate their interpretation of Stroud's argument as follows:

¹⁴⁷ Fumerton (2014) doesn't think that it's possible that we perceive that p , where " p " is a fact about non-psychological physical reality.

- (1) If meta-epistemological scepticism is false, then it's possible to know that knowledge of non-psychological physical reality is possible.
- (2) It's not possible to know that non-psychological physical reality is possible.

Therefore,

- (C) Meta-epistemological scepticism is true.

This is not Stroud's argument.¹⁴⁸ It's not Stroud's argument because Stroud thinks that premise (2) is false:

If perception is indeed a way of coming to know something about external things, then I can also know by perception that that answer to that epistemological question is correct. I can often see that someone right in front of me sees that there is a table in the room and thereby comes to know that there is a table in the room. [...] The truth of the straightforward epistemological explanation is something I can know to be true by perceiving that it is true. (Stroud 2004, 166-67)

So, Stroud (2004) maintains that it is possible to know that we perceive that *p* (and therefore know that *p*) on the basis of perceiving that people perceive that *p*. The second-order knowledge is secured, and therefore, on Stroud's view, Stroud's Requirement can be fulfilled.

I conclude that the internalist interpretation of Stroud's anti-externalist argument is incorrect. On the one hand, Stroud does maintain a second-order epistemic requirement on satisfying philosophical explanations of how knowledge of non-psychological reality is possible—e.g., Stroud's Requirement—but I argued that this is not equivalent to an internalist requirement. Second, even if Stroud's Requirement were an internalist

¹⁴⁸ This is not Stroud's argument, but one can make a plausible case that Dretske (2010) thinks that this argument provides adequate support for meta-epistemological scepticism. Dretske tells us that: "Skepticism might be true; but then, again, it might not. I can live with this conclusion. I think, in fact, that there is much to be said for second-order skepticism—the view that we probably, at least we might, know a lot, but that we can't know whether or not we do. We at least have an answer for the first-order skeptic, the skeptic who thought he could demonstrate that we don't know" (Drestke 2010, 132).

requirement, Stroud (2004) maintains that it can be fulfilled. But if that's right, then it's not true that Stroud's anti-externalist argument flows from internalism, and therefore the internalist interpretation is defective.

3.2 The Circularity Interpretation

If Stroud's anti-externalist argument does not flow from a commitment to internalism, what does it flow from? One temptation is to think that it flows from the thesis that the proponent of the straightforward answer and the proponent of the externalist answer are giving circular explanations of how the philosopher knows that knowledge of the world is possible, and that this type of circularity is problematic.

I have argued that Stroud maintains a second-order epistemic requirement, Stroud's Requirement, but that this requirement can be met on the basis of propositional perception (e.g., perceiving that we perceive that there are external things).

Sosa (1994) maintains that Stroud's anti-externalist argument flows from concerns about epistemic circularity:

What we want is [...] *knowledgeable* understanding. And this we will never have until we are in a good position to accept our view of our own faculties (of perception or memory, for example), a view which properly underlies our trust in their reliability. But this view we will never be able to justify without relying in turn on already attained knowledge of the external [world]. And this precludes our ever attaining a philosophically satisfactory understanding of all our knowledge in that domain (Sosa 1994, 275).¹⁴⁹

Sosa envisages an externalist fulfilling the traditional epistemological project as follows: (i) if, in general, appearances as of *p* are reliable indicators that *p*, then knowledge of the world is possible. And (ii) appearances as of *p* are reliable indicators that *p*. Therefore, (iii) knowledge of the world is possible. And (iv) since (iii) is arrived at by competent deduction from (i) and (ii), I am therefore

¹⁴⁹ Sosa comments as follows: "The new demands [from Stroud] do not derive simply from such formal internalism. They derive rather from a distinctively epistemic circularity that came to philosophical consciousness long ago" (Sosa 1994, 276).

in a position to know that (iii). Therefore, I'm in a position to know that knowledge of the world is possible.

Premise (i) registers that if an externalist theory of knowledge of the world is true, then knowledge of the world is possible. In particular, if reliabilism is true, then knowledge of the world is possible. Premise (ii) tells that appearances as of p are reliable, and from modus ponens, (iii) follows, that knowledge of the world is possible.

Premise (iv) registers the sub-conclusion (iii), the conclusion that knowledge of the world is possible, was arrived at by competent deduction, so that if knowledge is closed under competent deduction, then the agent is in a position to know that (iii) is true. Ex hypothesi, (iii) can be competently deduced, and therefore one is in a position to know that knowledge of the world is possible.

The crucial question is how the epistemologist knows that premise (ii) is true, that perceptual appearances as of p are reliable indicators that p . Alston (1993) tells us that it must be by a track-record argument:

if sense-perception is reliable, a track-record argument will suffice to show that it is. Epistemic circularity does not in and of itself disqualify the argument. But even granting that point, the argument will not do its job unless we are justified in accepting its premises; and that is the case only if sense perception is in fact reliable (Alston 1993, 17).

Alston thinks that a track-record argument cannot be a way of discriminating the reliable indicators from the unreliable indicators.

Sosa associates this response to Stroud (1994, 2000). However, Stroud thinks that coming to know that there are external things on the basis of perceiving that there are is not circular in any problematic way. At least, if it's circular, it's not a vicious circle.¹⁵⁰ The transition from perceiving that p to knowing that p is not an inferential transition. Rather, seeing that p just is a form of knowing that p . As Stroud explains:

¹⁵⁰ See Cassam (2009) for this point as well.

Just as there is no circularity in coming to know in that way that there are external things, so there is no suspicious circularity in coming to know in that same way that this answer to the epistemological question is true. The straightforward answer does not say that one reaches the “conclusion” that people know by perception that there are external things, and so perception is reliable, by inferring it from something else one sees or knows. It simply says that I can see and thereby know that people see and know in that way there are external things. There is no inference or chain of reasoning, and so no room for circularity. (Stroud 2004, 167)

So, Stroud doesn't think that the straightforward answer is problematically circular. But does this show that he thinks that externalist's answer is not problematically circular? Consider Stroud's comment:

[the externalist's explanation] is not unsatisfactory because it suffers from some kind of circularity. There is no circularity involved in fulfilling the conditions Sosa says are sufficient for perceptual knowledge. Nor do I protest that those conditions are not sufficient for knowledge (Stroud 2004, 169 my addition).¹⁵¹

So the defect in the externalist's explanation is not that it's circular. At least, if that is a defect, that's not the defect that Stroud thinks he has identified.

Now one might suspect that Stroud's worry with the straightforward answer, and externalism, is that their explanation of how they know that knowledge of the world is possible is not a proof that knowledge of the world is possible. Instead, it just takes it for granted both that this kind of perception is a form of knowing about the external world, and that human beings manifest this kind of perception.

But perhaps a proof of these claims is what is required for a satisfying philosophical understanding of how knowledge of the world is possible. According to this objection, then, absent a proof of those propositions, we have

¹⁵¹ Cf. “If that is the answer [e.g., the straightforward answer] Sosa would give to the question about our knowledge of external things from which he begins, there is no disagreement between us, either about circularity or about our knowledge of external things. We can “legitimately and with rational justification” arrive at a belief that knowledge of external things is acquired by perception, and we can “attain thereby a general understanding” of how we know many of the things we know” (Stroud 2004, 167 my addition).

not achieved a satisfying philosophical understanding of how our knowledge of the world is possible.

Here too, however, Stroud thinks that this objection is mistaken. On his view, the straightforward answer does not:

[...] assume or take it for granted that people can see that there is a table in the room or that other external things exist. It says that is something that almost anyone can see to be true right before his eyes; it is not something that is or must be assumed. So it can be established that people can see that there are external things. (Stroud 2004, 167)

What Stroud is getting at here is that we can just see that others see that there are external things. Of course, this does not mean we can perceive ourselves seeing that that there are external things. But how problematic should this be? After all, if we can see that other people see that there are external things, that implies that we know that the answer to the traditional epistemological project is that it is possible for us to know that knowledge of external world is possible.

On the other hand, Stroud seems to make a mistake here. The question is whether Stroud's anti-externalist argument flows from the externalist's explanation being circular (if it is). The suggestion is that that's not the case because Stroud doesn't think that it is circular.

We need to distinguish between premise-circularity and rule-circularity, or more broadly explanatory-circularity. Roughly, premise-circularity is a logical defect of arguments whereby a conclusion of an argument is equivalent to (or entails) one of the premises that is put forth as adequate support for that conclusion. However, rule-circularity is an explanatory defect, whereby one's explanation of the epistemic legitimacy of the rule, principle, or method R, the possibility of some kind of knowledge the acquisition of which the legitimacy of R is necessary, itself uses R in order to order to establish the legitimacy of R (or to establish the possibility of that type of knowledge).¹⁵²

¹⁵² I am grateful to Chris Tucker for this suggestion at the Orange Beach Epistemology Workshop (May 2014). The notion of explanatory circularity can be extended as follows. It's the idea that the possibility of some phenomena is in question, and one's explanation of how that phenomena is possible trades on some aspect of that phenomena. For example, someone might be wondering whether it's possible for anything to be non-psychical. If

Now consider the relationship between an explanation being rule-circular and it violating EPR. If an explanation is rule circular then it violates EPR. For example, consider the straightforward explanation. The straightforward explanation tells us that the epistemologist knows that knowledge of the world is possible on the basis of perceiving that other people perceive that there are external things. This explanation is rule circular, because the proposition in question is whether or not people can know that there are external things, and the proposition that it is possible because we can see that people see (and therefore know) that there are external things appeals to knowledge of external things.

Stroud is right that coming to know that the straightforward answer is true by perceiving that it is true isn't circular insofar as he means premise-circular, but it does seem to be rule-circular. If we think that rule-circularity is a defect of an epistemologist's explanation of how they know that knowledge of non-psychological reality is possible, then the straightforward answer is defective because it's rule-circular. This does not mean that rule-circularity is its only defect. It might be just one of the defects, though not the only one.

This raises questions about the relationship between Stroud's Requirement and EPR. Stroud's Requirement tells us that if an epistemologist *S* has a satisfying philosophical understanding of how knowledge of non-psychological reality is possible, then *S* knows that knowledge of non-psychological reality is possible. EPR tells us that a satisfying philosophical explanation of how knowledge of non-psychological reality is possible neither implies nor presupposes any knowledge of non-psychological reality.

The two requirements are not equivalent. This is obvious because the two requirements impose different demands on the epistemologist. But what might

one's explanation of how it's possible for anything to be non-physical trades on the fact that God can create anything, even non-physical things, it's not hard to see that this kind of explanation will be dissatisfying. For if their explanation simply takes it for granted that God *is* a non-physical thing, it doesn't really explain how *anything at all* could be a non-physical thing unless it also explains how *God* could be a non-physical thing.

be problematic is whether it possible for the epistemologist to meet both of those demands. Prima facie, it's hard to see how the epistemologist can fulfil both requirements.

Here's the argument. Suppose that Stroud's Requirement is fulfilled. This implies that the epistemologist knows that knowledge of the world is possible. Now assume that "knowledge of the world is possible" expresses a proposition which implies or presupposes knowledge of the world. From this assumption and Stroud's Requirement, it follows that EPR is violated.

The problem with the argument is that the assumption is false. The proposition that knowledge of the world is possible is a modal proposition, not a categorical proposition. That knowledge of the world is possible is true, where the modal operator is logical or metaphysical possibility, does not entail that knowledge of the world is possible implies or presupposes knowledge of the world. It's consistent with the truth of that proposition that there are no human beings. That the two are consistent shows that the modal proposition does not entail the categorical proposition. But the truth of the categorical proposition is what violates EPR, and the truth of the modal proposition does not commit one to the categorical proposition. So, the fulfilment of Stroud's Requirement and EPR are not exclusive after all.

So, Stroud's Requirement and EPR impose different demands on the epistemologist. Moreover, fulfilment of Stroud's Requirement and EPR are not obviously mutually exclusive.

4 Stroud's Anti-Externalist Argument

In the previous section (§3), I explored two interpretations of Stroud's argument against the externalist response to Stroud's dilemma. The first was that Stroud's dissatisfaction with the externalist response is grounded in internalist principles that the externalist rejects (§3.1-3.2). I argued that this interpretation is defective. The second was that Stroud's dissatisfaction with the

externalist response to Stroud's dilemma is grounded in the externalist's response being circular (§3.2). I argued that this interpretation is defective.

In this section, I will present a different interpretation of Stroud's argument against the externalist response to Stroud's dilemma.¹⁵³ I maintain that Stroud's argument against the externalist response to Stroud's dilemma represents a powerful case against the externalist response. Later, in section 5, I will defend this response from objections.

4.1 A Comparison: The Externalist Epistemologist and the Externalist Astrologer

Stroud's dilemma presents a powerful case for meta-epistemological scepticism. The first horn of Stroud's dilemma is the externalist's target. EPR must be fulfilled, but it's hard to see how it can be fulfilled without scepticism. In short, if one's explanation fulfils EPR, then it's a sceptical explanation.

Part of Stroud's argument for the first horn of Stroud's dilemma is that in order for one's explanation to respect EPR, it can neither entail nor presuppose knowledge of the world, and that it's hard to see how it can do that without commitment to scepticism (cf. Chapter 1, §6.2). So, suppose we appeal to our sense experiences as of p (where " p " represents any proposition about the world). Stroud's thought is that, contra the first externalist response (§2.1), it's not enough that a linking-principle is true, such as that sense experiences as of p are reliable indicators that p . Instead, the epistemologist needs to know that p is true.

Sosa (1994) thinks that this demand can be met.¹⁵⁴ After all, insofar as the theorist's belief that the linking-principle is true is true, and the product of a reliable-belief forming process, then the theorist can know that the linking-principle is true. There is no obvious logical-defect in this kind of response.

¹⁵³ This interpretation is from Pritchard and Ranalli (*forthcoming*), where Fumerton's (1995, 2004) reasons for meta-epistemological scepticism are compared with Stroud's (2000, 2004).

¹⁵⁴ For example, see Sosa (1994, 285).

Moreover, since the externalist ought to appeal to their own account of knowledge (what else should the externalist appeal to here?), it's hard to see what's so objectionable about appealing to their account. In turn, the externalist has a prima facie good case for how to get around Stroud's worries about the first horn of Stroud's dilemma.

Stroud's criticism of this externalist response appeals to a comparison between two kinds of theorists: an externalist epistemologist about perceptual knowledge of the world and an externalist astrologer about astrological knowledge of the world.

In rough outline, Stroud's argument against the externalists' explanation of how the linking-principle is known is no better than the externalist astrologer's explanation of how their corresponding linking-principle is known. In order to appreciate the argument in more detail, consider these two theorists, and the following two cases:

The Externalist Astrologer: I know that I am looking into the crystal ball, and I believe that if there are reliable connections between what happens in the ball and the world I form beliefs about, then I can know about the world I am forming beliefs about. I also believe that there are reliable connections between what happens in the ball and the world I form beliefs about.

The Externalist Perceiver: I know that I am having perceptual experiences as of an external world, and I believe that if there are reliable connections between what happens in (or is represented by) these kinds of perceptual experiences and the world I form beliefs about, then I can know about the world I form beliefs about. I also believe that there are reliable connections between what happens in our perceptual experiences and the world we form beliefs about.

Stroud thinks that their explanation of how their linking-principles are known (second conjunct of the first sentence of each), is no better or worse than the other. The intuition is that, so far as a philosophical explanation of how each theorist knows that their respective linking principle is true, the externalist about perceptual knowledge of the world is no better or worse than the externalist about astrological knowledge of the world. In other words, so far as achieving a philosophical understanding of how their knowledge of the world is possible, the externalist epistemologist about perceptual knowledge is no better off than the externalist astrologer.

5 Objections and Replies

The first objection to Stroud's argument is that it seems to beg the question against externalism. After all, one might ask whether these theorists are right about there being reliable connections between their epistemic bases and their target beliefs (whether their explanation of how their linking-principles are known is true). To this question, one might argue that:

The difference between the positions of the two theorists lies only in the believed-in connections between the relevant experiences and the wider world. The theory says in each case that if such connections hold, that theorist knows. Each theorist, confidently sticking to his own story, believes that they hold in his case and not the other. Each might even try to settle the matter by consulting his own experience and his own theory, and find himself content with the discovered result. In that respect, the two positions are equally satisfactory, or unsatisfactory. (Stroud 2004, 171)

However this response will strike most externalists as begging the question against externalism. After all, one might suggest that the reliabilist, for example, explain the epistemological difference between the externalist astrologer and the externalist epistemologist about perceptual knowledge as follows:

- The astrologer who relies on the images in the crystal ball does not have reliable beliefs, beliefs in the market for knowledge, because

crystal ball-gazing is unreliable, whereas the externalist about perceptual knowledge does have reliable beliefs, beliefs in the market for knowledge, because perception *is* a reliable belief forming basis.¹⁵⁵ Perhaps one will think that the externalist astrologer and the externalist epistemologist are epistemic equals, but to do that would be to reveal one's internalist intuitions, if crystal ball-gazing is unreliable, while perception is reliable.

However, what's not in question is whether each theorist, *qua* epistemic agent, are epistemic equals. Instead, what's in question is whether each theorist, *qua* explainer, are epistemic equals.

It might be that the externalist epistemologist, who relies on their perceptual experiences in forming beliefs about the world, has beliefs which are in the market for knowledge, whereas the externalist astrologer, who relies on the images in the ball, does not have beliefs which are in the market for knowledge, because the former has true beliefs while the latter has false beliefs. In this sense, it might just be that the two theorist's aren't epistemic equals in terms of their possession of knowledge of the world. But that's not what's at issue. What's at issue is whether one explanation is better than the other in the sense that one makes an epistemic difference to one theorist but not the other.

This takes us to the second objection. One might ask whether externalist astrologer and the externalist about perceptual knowledge are in a position to know that their beliefs are reliable, and so in the market for knowledge. In short, the question is whether the theorists can achieve higher-order knowledge of their respective epistemic sources: ball-gazing and perception.

Here the externalist's objection is that, because ball-gazing is unreliable, the externalist astrologer who relies on that epistemic basis is not in a position to know that their ball-gazing beliefs are reliable, and so in the market for knowledge, whereas the externalist about perceptual knowledge, who relies on

¹⁵⁵ see Sosa (1997) for this kind of objection. In the main text, I'm summarising this kind of objection. Cf. Sosa (1994, 275-276).

their sense experiences, is in a position to know that their perceptual beliefs are reliable, and so in the market for knowledge.

According to this objection, if externalism is true, then insofar as the theorist's belief that their beliefs formed on the basis of perception are reliable is in fact the product of a reliable belief forming process, then their higher-order belief will be in the market for knowledge as well.¹⁵⁶ That is, insofar as their belief that the linking-principle between their epistemic basis (sense experiences) and the world is in fact true, then their higher-order belief is in the mark for knowledge.

Stroud's response to this objection focuses on the following points:

- both the theorists are externalists about their putative knowledge, and;
- both of their epistemic resources are impoverished in the same way.

Stroud's thought is that if these two points are true, there is a serious question as to whether one theorist is in a position to understand how their knowledge of the world is possible while the other is not. In other words, can the externalist epistemologist about perceptual knowledge be in a better position with respect to knowing that their linking-principle is true than the externalist astrologer, if both of the above points are correct?

What is crucial to Stroud's proposal, then, is that neither is better off than the other, because of how impoverished both theorists' epistemic resources are supposed to be. For in both the perceiver's case and the crystal-ball gazer's

¹⁵⁶ One might instead argue that even if the perceiver's beliefs about the world are in the market for knowledge, the application of their externalist proposal the next level up is problematic, because we do not achieve what Fumerton calls "philosophical assurance" of our belief forming processes or methods. See Fumerton (2004, 2006). Fumerton's idea is that if one can justifiably ask "I believe that my belief that *p* is non-inferentially justified if it was produced by a reliable process, but *is it* produced by a reliable process?", then the justifiability of asking that question for oneself is prima facie reason to think that one's belief that *p* isn't really non-inferentially justified. For Fumerton, the kind of non-inferential justification that the epistemologist should want is the kind that gives one "philosophical assurance" that *p* is true, so that their "philosophical curiosity" is satisfied (Fumerton 2006, 184). For discussion, see Ballantyne (2011) and Poston (2010).

case, their epistemic basis for forming beliefs about the world is supposed to be limited to what Stroud calls a “restricted domain”: mere perceptual experiences as of an external world, and mere ball gazing images as of an external world.

The second claim is the objectionable claim. According to Stroud, if neither theorist is in a position to perceive that there are external things, then the externalist epistemologist about perceptual knowledge is in no better of a position for understanding how their putative perceptual knowledge of the world is possible than the externalist astrologer is for understanding how their putative ball-gazing knowledge of the world is possible, because “[w]hat we get in sense-perception bears the same relation to the world we think we know” about through perception, as the ball-gazer’s ball-gazing bears to the world the ball-gazer thinks she knows about through ball-gazing.” (Stroud 2004, 172)

Prima facie, this is right. If the externalist epistemologist about perceptual knowledge can appeal to no more than *a priori* reasoning and their sense experiences, this doesn’t look like an improvement over the externalist astrologer, who can appeal to *a priori* reasoning and images in the ball.

On the other hand, this response might seem obscure, because it’s not clear what makes the externalist committed to the thesis that we cannot perceive that there are external things, as opposed to just having perceptual experiences as of there being external things.

Stroud doesn’t think that the externalist is committed to the thesis that no one can perceive that there are external things if perceiving that *p* implies knowing that *p*. Instead, Stroud thinks that the following weaker claim is true: if the externalist were to maintain that it is possible to perceive that *p*, then it would just amount to nothing more than the straightforward response to the traditional epistemological project, that we can perceive that *p* (where “*p*” represents propositions describing facts about the world), such perception is knowledge-entailing, and we can perceive that people perceive that there are external things, so that we can know that knowledge of the world is possible.

Now remember that the externalist's response here is supposed to be in line with EPR. The first horn of Stroud's dilemma is that if one's explanation respects EPR, it's a sceptical explanation. Part of the argument for this horn of the dilemma, recall, is that the theorist would need to know that the linking-principle is true. The internalist, it seems, cannot account for this without falling into regress or lapsing into scepticism. The externalist, on the other hand, wants to tell us that Stroud's demand can be met: we can know that the linking-principle is true, insofar as the theorist's belief that the linking-principle in fact is the product of a reliable belief-forming process.

However, if all that "we are aware of in perception" is "restricted to features of our perceptual experiences", such that the "external facts we know as a result of those experiences are nothing we ever perceive to be so", it's hard to see how the externalist epistemologist about perceptual knowledge is better off than the externalist astrologer in terms of understanding how their knowledge of the world is possible (Stroud 2004, 171).

In order to motivate Stroud's case, consider the following example. First, imagine the externalist astrologist, who believes that *a book is there* on the basis of seeing an image of a book in the crystal-ball. Second, imagine the externalist epistemologist about perceptual knowledge, who believes the same proposition, but on the basis of their perceptual experience as of a book being there. Both of these epistemic sources are restricted in the sense that it does not follow from *S* ball-gazed that *p*, that *S* knows that *p* (or even that *p*), and it does not follow from *S* has a perceptual experience as of *p*, that *S* knows that *p* (or even that *p*). Moreover, what the perceiver and the ball-gazer appeal to in forming their beliefs about their target epistemic sources is restricted to the mere deliverances of those epistemic sources alone, and what one could reason from what can be known on the basis of those sources alone.

Now the crucial question here is how far each theorist's track-record of their epistemic sources would take them. After all, when the perceiver forms beliefs about how reliable their perceptual experience as of the external world

are, their comparison will be between previous perceptual experiences as of p being true, and their believing that p is true, their perceptual experiences as of q being true, and their believing that q is true, and so on.

For example, the perceiver will believe that their perceptual experiences as of the cat being on the mat are reliable experiences, because on most of the occasions she had that kind of experience, a cat was on the mat—or rather, it was for the perceiver just as if a cat was on the mat. What the perceiver can consult is more perceptual experiences, and whatever they can infer from what can be known *a priori*. Likewise, when the ball-gazer forms a track-record about the constellations and peoples' personalities, she too will be restricted to what can be gazed in the ball, and known on the basis of crystal-ball gazing alone. For example, when she goes to check whether the crystal-ball gazing experiences as of the Moon being full are reliable indicators of their friend Smith being sad, the crystal-ball gazer will be consulting the images in the ball and whatever can be inferred what she knows on that basis.

Which theorist is better off as far as understanding how their putative knowledge is possible? On the one hand, it might be that one theorist has more true beliefs than the other. In this sense, one is better off because one has more true beliefs than the other. But when it comes to them knowing that their beliefs are true, how could one of the theorists be better off than the other? As Stroud noted, the relation the perceiver stands to their perceptual experiences is the same kind of relation that the ball-gazer stands to their crystal-ball gazing experiences. Again, it might be the case that the perceiver's first-order beliefs about the world are true, and their second-order beliefs about the epistemic merit of their perceptual experiences are true, but when it comes to knowing that this is so, and so being in a position to understand how their putative knowledge is possible, it's hard to see how the perceiver could be better off than the crystal-ball gazer, because the perceiver will be consulting the

restricted deliverances of their perceptual experiences, just as the crystal-ball gazer will be consulting the restricted deliverances of their crystal-ball.¹⁵⁷

Of course, the externalist still might not think there's a real problem here, because perceptual experiences as of *p*—even if “restricted” along the lines that Stroud has in mind—are still reliable, whereas crystal-ball gazing is not. But Stroud's rejoinder here is that the externalist philosopher—in providing this kind of answer—must have relied upon their perceptual experiences as epistemic support for that claim. In this fashion, the externalist epistemologist about perceptual knowledge who maintains that we can know that the linking-principle is true without violating EPR provides a good test case.

The thought here is this: what else is the externalist epistemologist about perceptual knowledge going to appeal to in order to support their belief that their own and other people's perceptual experiences are reliable, if not that he has had perceptual experiences as of himself and others making reliable judgements about the world which have been correct in the past, and that they've had experiences as of other peoples' judgements, assertions, and expressions of belief being accurate representations of the world? So, the problem here is that what the externalist epistemologist must appeal to is no better at the higher-order level than it was at the first-order level: either it's as good or it's as bad. So, the problem with the track-record argument is that the externalist:

¹⁵⁷ Here's a simple thought experiment which suggests this. Imagine having been born in a room with a single television screen. Everything you learn is from whatever shows up on the screen. Now imagine that you've never left this room. You see what you call “tree” on the screen. You come to believe that, outside the room, there is tree. Now suppose that the following is true: whenever whatever you call “tree” shows up on the screen, that thing is present outside the room. Now, how could the person, inside the room, come to know that that conditional is true? They could reason that, whenever whatever they call “tree” shows up on the screen, then a tree is outside the room, and build a track-record of cases. But remember this track-record will be built up from only whatever this person within the room can know a priori and know on the basis of looking at the images on the screen. Even with the stipulation that the images on the screen reliably co-vary with what goes on outside the room, it's very hard to see how the person is going to be able to know that there is some such reliable co-variation.

[...] must acknowledge that he does not [...] ever see or otherwise perceive that those human beings and other external things that he is interested in are there. Nor does he ever perceive the reliability of the connections that he believes hold between them. The most he is perceptually aware of or presented with in experience are the qualities or character of his perceptual experiences. (Stroud 2004, 170)

On Stroud's view, then, it seems like the problem with the externalist using a track-record argument in order to provide support for the belief that perception is reliable, is that all that the externalist can appeal to falls short of the facts about non-psychological physical reality. All we would have are perceptual experiences as of *p*, rather than perceptions that *p*, to appeal to here, together with whatever could be known *a priori*. In short, the track-record does not improve upon the so-called impoverished epistemic sources, such as perception (so understood) and crystal-ball gazing.

6 Conclusion

In this chapter, I examined the externalist responses to Stroud's dilemma and responded to their arguments. In section 2, I considered two externalist responses to Stroud's dilemma. Both of these responses targeted the first horn of Stroud's dilemma: that if EPR is true, then scepticism is true.

The first response said that insofar as a linking-principle between one's prior epistemic base, such as one's sense experiences as of *p*, is in fact true, then it would be possible for EPR to be fulfilled without scepticism. The second response accepted the requirement that the epistemologist needs to know that the linking-principle is true, and argued that that requirement can be met. This is the response from Sosa (1994). I argued that both of these responses fail.

Chapter 6

Stroud's Dilemma: Propositional Perception, Simple Perception, and Cassam's Response

1. Introduction

The second part of the thesis has been focused on Stroud's dilemma. Chapter 4 examined Williams's response to Stroud's dilemma, and argued that Williams's response is defective. Chapter 5 examined Sosa's response to Stroud's dilemma and argued that it too is defective. Chapter 6 will examine a different sort of response from Cassam (2009). I will argue that Cassam's response fails, and therefore that Stroud's dilemma survives serious objections.

In the remainder of this section, I want to describe the structure of this chapter. In order to do this, it will be helpful to structure the discussion around the central arguments that will appear in this chapter. So, let's first re-examine Stroud's dilemma. The first horn of Stroud's dilemma is that our explanation must fulfil the epistemic priority requirement, EPR. However, if it fulfils EPR, then it won't be an anti-sceptical explanation. It will entail scepticism (cf. Chapter 1, §6.2). The second horn of Stroud's dilemma is that if our explanation fails to fulfil EPR, then it won't be a fully general explanation (cf. Chapter 1, §6.3). So, in either case, the conclusion is that the explanation won't be a satisfying philosophical explanation of how knowledge of the world is possible. Either it won't be anti-sceptical or it won't be fully general, and neither seem to be philosophically satisfying (cf. Chapter 1, §6.4).

Cassam's response focuses on the first horn of Stroud's dilemma. In particular, Cassam focuses on the argument for EPR and whether EPR can be fulfilled without commitment to scepticism. Notice that these are two different approaches to the first horn of Stroud's dilemma. The first horn of Stroud's dilemma tells us that if the explanation fulfils EPR, then it implies scepticism. So, there are several choices for how one might respond to the first horn of Stroud's dilemma. One can argue that either:

- The arguments intended to provide adequate support for EPR fail to provide adequate support for EPR.
- EPR is false.¹⁵⁸
- It's not the case that if one's explanation fulfils EPR, then it's a sceptical explanation.
- It's not the case that if one's explanation is a sceptical explanation, then it's not a satisfying philosophical explanation.

Cassam takes the first and the third option. He takes it that there is an argument for EPR—what he calls the “argument from generality”—but that this argument fails to provide adequate support for EPR (Cassam 2009, 572). This turns on his discussion of “full-blooded” versus “modest” explanations. On Cassam's view:

[...] the best way of motivating this requirement [EPR] is not to emphasize the extreme generality of what we seek in the philosophy of knowledge but to insist that a fully satisfying philosophical account of our knowledge of the world needs to be full-blooded rather than modest (Cassam 2009, 574 my addition).

The idea is that it's not the case that the “argument from generality” provides adequate support for EPR but instead the desire for one's explanation to be

¹⁵⁸ I don't intend these to be exclusive, since it might be that EPR is false *and* the arguments for EPR fail to adequately support it.

“full-blooded” (Cassam *ibid*). This leads us to the third option. Cassam argues that:

On this reading [e.g., the desire for the explanation being a “full-blooded” explanation], EPR has something going for it but not so much as to make it impossible to satisfy. For this reason, EPR doesn’t make scepticism inevitable (Cassam *ibid*, my addition).

So, Cassam argues that it’s not the case that if one’s explanation fulfils EPR, then it’s a sceptical explanation. As we will see, this turns in part on whether explanations in terms of simple perception, states in which one perceives an object *o* which is *F*, unlike propositional perception, where one perceives that *o* is *F*, can be non-sceptical and fulfil EPR.

Against Cassam’s first argument, I will argue that his argument for the “modest” explanations of how our knowledge of the world is possible being a viable option is not a good argument. In particular, I argue that “full-blooded” explanations seem to be preferable in this case. Against Cassam’s second argument, I will argue that an explanation in terms of simple perception, like explanations in terms of propositional perception, cannot be both non-sceptical and respect EPR.

2 Cassam’s Response to EPR: Part I

In this section, I will do three things. First, in §2.1, I will explain what Cassam (2009) takes to be Stroud’s (2000) main argument for EPR. Second, in §2.2, I will explain how Cassam argues against this argument, and then I will respond to Cassam’s argument.

2.1 The Argument from Generality

Cassam (2009) takes Stroud’s (2000) main argument for EPR to be the following “argument from generality”:

We take it that knowledge of objects [in the external world] comes to us somehow by means of the senses, but if we thought of sensory knowledge as itself knowledge of material objects around us we

would not get an appropriately general explanation of how any knowledge of any objects at all is possible by means of the senses. We would be explaining knowledge of some material objects only on the basis of some others (Stroud 2000, 105 my addition).

I take it that the thrust of Stroud's argument is this. A successful explanation of how our knowledge of the world is possible would not leave us wondering whether it is possible. It would settle the question. Yet if our explanation presupposed some knowledge of the world, then it would be leaving some knowledge unexplained. That is, some of it would be left in question. The knowledge that it presupposes would be the knowledge that is left in question, and therefore unexplained. The goal, however, is to explain how knowledge of that sort is possible at all, so presupposing that some of it is possible isn't going to explain how it's possible at all.

A few comments are in order about this argument. The first one is that EPR is put forward as a necessary condition on explaining how knowledge of the world is possible, not on whether we have knowledge of the world, or how people come to know about the world, or what knowledge is. So, even if explanations of how we have knowledge of the world, or what knowledge of the world is, can violate EPR and be satisfying philosophically, this does not entail that explanations of how knowledge of the world is possible can violate EPR and be satisfying philosophically.

Second, we might wonder how the argument from generality for EPR is related to Stroud's dilemma. Recall that the first horn of Stroud's dilemma is that if EPR is true, then scepticism is true. The thought is that a satisfying philosophical explanation has to be a fully general explanation and that a fully general explanation has to fulfil EPR. So, we can present the argument as follows:

1. If an explanation of how knowledge of the world is possible is a satisfying philosophical explanation, then it's a fully general explanation.

2. If an explanation of knowledge of the world is possible is a fully general explanation, then it respects EPR.

Therefore,

3. If an explanation of how knowledge of the world is possible is a satisfying philosophical explanation, then it respects EPR.

The argument is valid. Premise 1 is at least intuitive, and Stroud maintains that the argument from generality provides adequate support for premise 2. This raises questions about just how “general” the traditional epistemic project, the project of explaining how knowledge of the world is possible, is and to what extent that feature of the project makes explanations which violate EPR inadmissible.

Cassam maintains that while it might be that a good answer to the traditional epistemological project needs to be “appropriately general”, the sense in which it must be “appropriately general” is ambiguous between the following three senses:

- (I) We want to understand how any knowledge at all is possible—how anything we currently accept amounts to knowledge. Or, less ambitiously, how we come to know anything at all in a certain specific domain (Stroud 2000, 101).
- (II) We want an account of our knowledge that would make all of it intelligible to us all at once (Stroud 2000, 8).
- (III) We want to find a single, or very small number of very general 'ways of knowing' (Stroud 2000, 4).

Cassam tells us that “(I) is a less demanding and more plausible requirement than (II)” and that “we might accept (II) without accepting (III)” (Cassam 2009, 578-579). This puts pressure the idea that the desire for an “appropriately general” explanation of how our knowledge of the world is possible needs to respect EPR. His argument for these claims goes as follows:

If one of the things we currently accept is the proposition P, and we succeed in explaining how we know that P, then it cannot be said that we have not explained how *anything* we currently accept amounts to knowledge. But explaining our knowledge of P obviously does not

amount to making all of our knowledge of the world intelligible to us all at once. Again, it is not obvious that making *all* of our knowledge intelligible all at once is necessarily a matter of identifying a single way, or small number of very general ways, of knowing. Why not accept we have lots of different ways of knowing about the world around us, and concentrate on doing justice to the multiplicity of pathways to knowledge? In accepting the sheer heterogeneity of our ways of knowing what we lose is not generality but a certain kind of simplicity. But the more we insist on finding a small number of ways of knowing the harder we are likely to find it to give a plausible account of the full range of human knowledge (Cassam 2009, 578-579).

I accept Cassam's point that we can accept (II) without accepting (III), but it's not clear that (II) is not as plausible as (I). Cassam holds, for good reason, that on his reading of (I), (I) is less demanding and more plausible than (II), because it's clear that the project of answering the question:

1. "How is it the case that, for any proposition we accept about the external world P, P is known?"

is less demanding than the project of answering the question:

2. "How is it the case that, for every proposition we accept about the external world P, Q, ..., they are known (or knowable)?"

But notice that (I) demands more than just answering 1 to the philosopher's satisfaction. (I) also demands that "we understand how any knowledge [of the external world] is possible at all" (Stroud 2000, 101 my addition). Cassam's argument doesn't address this aspect of (I).

In order to see this, notice that one could explain how P is known, or why P is known, without explaining how knowing that P is possible. For example, it might be that we can explain how P is knowledge rather than mere true belief, but this wouldn't explain how the obstacles which make knowing that P look impossible are to be removed. Yet explaining how knowing that P (or propositions of that sort) is possible is the other demand that (I) imposes.

Now we might wonder whether (I) is therefore less demanding and more plausible than (II) after all. (II) demands that our explanation of how our knowledge of the world is possible “make all of it intelligible to us all at once” (Stroud 2000, 8). What does this mean? On the one hand, it might mean just what 2 means: how, for all of the propositions about the external world that we accept, are they knowable? On the other hand, it might mean: how is it possible to know that those sorts of propositions are true without our explanation taking it for granted that some of them are known? This sounds a lot like an explanation which fulfils EPR. After all, an explanation which respects EPR would make “all of it intelligible to us all at once” in the sense that it would not be leaving the possibility of knowing any proposition about the external world an open question. To explain how any of it is possible “all at once” would be to explain how it is possible without taking it for granted that any of them are known (Stroud *ibid*).

2.2 Response to Cassam

In the previous section, I raised some doubts about whether the argument from generality fails to provide adequate support for EPR. But let’s suppose that Cassam is right that it doesn’t. Where does that leave us? Cassam argues that EPR is compelling insofar as we think that satisfying philosophical explanations of how our knowledge of the world is possible need to be “full-blooded” explanations (Cassam 2009, 280). He argues for this as follows:

The point about explanations of knowing in terms of epistemic seeing or other knowledge-entailing mental states is not that they are not explanations at all but that they are not what might be called full-blooded explanations. Full-blooded explanations are informative in a particular way. They purport to explain the transition from not knowing to knowing by showing how our knowledge of the world could come to be out of something that is not knowledge of the world. This is the kind of explanation that, as epistemologists, we should be looking for (Cassam 2009, 580).

EPR tells us that an explanation of how our knowledge of the world is possible neither implies nor presupposes knowledge of the world. One of the main

reasons Stroud thinks that, in order to fulfil the traditional epistemological project, the project of explaining how our knowledge of the world is possible at all, one's explanation must fulfil EPR, is that without doing so it won't be a "general" philosophical explanation.

What we should ask here, then, is what the relationship is between the following three properties: being a general philosophical explanation in Stroud's sense, fulfilment of EPR, and being a full-blooded explanation in Cassam's sense. Cassam tells us the following:

We have arrived at a reading of EPR according to which the idea to which it gives expression is that philosophical explanations of our knowledge need to be full-blooded. While full-blooded explanations needn't be fully general in Stroud's sense, fully general explanations must be full-blooded. A full-blooded explanation of knowledge of kind K is one that plausibly represents it as originating in something that is not already knowledge of kind K. On this account, EPR is the demand for a full-blooded explanation of our knowledge, and a full-blooded explanation of our knowledge must at least be one whose explanans and explanandum are logically distinct. That is why, on the assumption that 'S sees that P' logically entails 'S knows that P', it is not a full-blooded explanation of S's knowledge that P to say that S sees that P. To explain S's knowledge in this way is to give an entailing-explanation of it whereas full-blooded explanations are non-entailing; unlike entailing explanations they do not logically entail the knowledge that they purport to explain (Cassam 2009, 580).

We might think that the motivation for EPR comes from one's goal of giving a full-blooded explanation of how our knowledge of the world is possible. The idea is that full-blooded explanations are unsatisfying because they aren't informative in the right way, and they aren't informative in the right way because they entail that the agent has the type of knowledge they are trying to explain how it's possible to have at all. As Cassam registers, then, one of the main motivations for EPR depends on the "assumption that explanations of our knowledge are philosophically satisfying only to the extent that they are full-blooded" (Cassam 2009, 581).

Cassam's argument against this requirement is what I'll call the "argument from modesty". The first premise of this argument is that entailing-explanations can be informative. If entailing-explanations are informative, that puts pressure

on the idea that entailing-explanations cannot contribute to an explanation of how our knowledge of the world is possible. For example, if I explain how it's possible for *S* to know that *p* on the basis of seeing that *p*, this is different from appealing to deducing *p* from premises about sense-experiences. So, it can be informative to the extent that it explains how knowledge of the world is possible in more specific terms.

The second premise is that “modest” explanations, explanations which are not full-blooded, and therefore entail or presuppose the kind of knowledge in question, are a prima facie live option. Cassam questions the assumption that only full-blooded explanations of how our knowledge of the world is possible are in the market for philosophical satisfaction as follows:

Why accept this assumption? Perhaps we would have to accept it if we thought that entailing explanations are totally uninformative but we have already seen that this is not the case. Explanations of knowing in terms of knowledge-entailing mental states are modest rather than full-blooded but modesty is a serious option in epistemology as in other areas of philosophy. Moreover, once we acknowledge that even philosophical explanations can be modest there is no longer any need to worry about EPR. The particular conception of philosophical explanation to which this requirement gives expression is not one that we have to accept (Cassam 2009, 581).

However, it's not clear that this argument is a compelling argument. That there are modest options does not entail that we ought to select those options. This, after all, is what should be in dispute: whether we should select the full-blooded options over the modest options.

Moreover, that modest explanations can be informative doesn't entail that being informative is sufficient for the philosophical explanation being in the market for philosophical satisfaction. Of course, it wouldn't be fair to credit one with thinking that informativeness alone is sufficient either. So, we might instead consider the thesis that being informative, true, and believed is sufficient for one's explanation being a satisfying philosophical explanation of how knowledge of the world is possible. This thesis is a live option in the sense that it's not obvious that it's false. But this alone gives us little in terms of

reasons which favour a modest explanation over a full-blooded explanation, since it's also not obvious that we shouldn't favour a full-blooded explanation.

This issue is connected with the sense in which Cassam tells us that modest explanations should be a live-option here just as it is elsewhere in philosophy. The sense in which modest explanations are live-options elsewhere is the following: philosophers have argued for theories of some concept C in which the explanans t_1, t_2, \dots, t_n , of the explanandum of C are not reductively analysed. Yet one can give a full-blooded explanation of how knowledge of the world is possible which is “modest” in that sense. This would be an explanation of how knowledge of the world is possible on the basis of some explanans t_1, t_2, \dots, t_n , which neither entail nor presuppose knowledge of the world but aren't reductively analysed. So, we can grant that modesty is a live-option in that sense but that it's compatible with demanding a full-blooded explanation.

A final issue is that we might just find it less plausible that a modest explanation should be preferred over a full-blooded explanation. Consider a modest explanation of how our knowledge of the world is possible. This might be formulated along the following lines: “knowledge of the world is possible by knowing about it in particular ways. Here are the paradigmatic ways: one can see that p , hear that p , deduce that p , etc.” The obvious objection to modest explanations is that the obstacles which make knowing that p look impossible are going to make the paradigmatic ways of knowing look impossible as well. And if that's right then it's hard to see how a modest explanation could amount to a satisfying explanation of how knowledge of the world is possible.

2.3 Respecting EPR

In sections 2.1 and 2.2, I argued that EPR gets at least some support from the argument from generality, and that even if Cassam is right that it's the desire for a full-blooded explanation which motivates EPR, it's not clear that a full-blooded explanation isn't preferable to a modest explanation.

This leaves with the idea that EPR should be fulfilled. But how is this to be done? In Chapter 1, I argued that it's hard to see how an explanation of how our knowledge of the world is possible in terms of propositional perception is going to be admissible since it seems to violate EPR. The reason it seems to violate EPR is that the following entailment thesis appears to be true: if *S* perceives that *p*, then *S* knows that *p*.

The more general challenge here is that, on the one hand, our explanation of how our knowledge of the world is possible ought to appeal to our perceptual capacities, since these capacities seem to be our more basic means of engaging with the world around us.¹⁵⁹ On the other hand, we have the following worries with appealing to our perceptual capacities. I describe these worries below.

First, appealing to propositional or epistemic perception appears to violate EPR because of the entailment thesis. As Cassam (2009) stresses the issue, perceiving that *P* seems:

[...] too close to knowing that *P*. Once we get as far as perceiving that *P* we have already reached our final destination. We were supposed to be explaining how we reach it and we need to do better than say that we reach it by reaching it in a particular way (Cassam 2009, 579-580).

Second, appealing to non-propositional “objectual” or “simple” perception seems far too non-epistemic to help us here. Simple perception (e.g. *S* sees an *o* which is *F*) doesn't entail knowledge (it doesn't entail that *S* knows that *o* is *F*), but it is consistent with knowledge (it is consistent with *S* knowing that *o* is *F*). Yet if we appeal to simple perception then we will need to explain how it's possible for simple perception of the world to put us in a position to know about the world. After all, simply seeing an *o* which is *F* seems hardly a compelling explanation of how it's possible to know that *o* is *F*.

Third, appealing to sensory appearances or seemings (e.g. *S* has an appearance as of *p*), in line with EPR, faces the first horn of Stroud's dilemma:

¹⁵⁹ Cassam (2007) seems to accept this point. See in particular pg. 5 and pg. 25.

that if the resulting explanation fulfils EPR, then it's a negative, sceptical explanation.

Cassam grants Stroud the third issue. If one's explanation of how knowledge of the world is possible appeals to only sensory appearances or seemings and *a priori* knowledge, in line with EPR, then that explanation will end up being a negative, sceptical explanation. So our focus, then, will be on the first and second issues: whether a satisfying philosophical explanation of how our knowledge of the world is possible can be given in terms of either propositional perception or simple perception.

The main challenge for proponents of appealing to propositional perception is that it's too epistemic: the entailment thesis seems to be true. The main challenge for proponent of appealing to simple is perception is that it's too non-epistemic: how could simply seeing objects and properties explain how knowledge of the world is possible?

There are at least two responses to the first challenge. The first is to argue that an explanation in terms of propositional perception doesn't violate EPR after all because the entailment thesis is false. The second is to argue that the explanation doesn't have fulfil EPR. In the next section, I argue that the first response fails because the entailment seems to be true. Thus far, I've provided some motivation for thinking that the second challenge fails because fulfilment of EPR seems non-negotiable.

3. Propositional Perception and the Entailment Thesis

3.1 Motivating the Entailment Thesis

I look out the window, and I see that it's raining outside. Do I now know that it's raining outside? According to proponents of the Entailment Thesis, I do (Williamson 2000; Stroud 2009, 2011; Cassam 2007*a*, 2007*b*, 2008; Peacocke 2005; Dretske 1969). If I see that *p*, I know that *p*. In general, the Entailment

Thesis is the thesis that if S perceives that p , then S knows that p . Of course, one could be a proponent of specific entailment theses for specific perceptual modalities without accepting the general Entailment Thesis, but I'm going to presume throughout that arguments we consider for and against it are applicable *mutatis mutandis* to more specific versions of the Entailment Thesis.

So, what might lead one to accept the Entailment Thesis? Perhaps what makes the Entailment Thesis look plausible is just that it's intuitive. After all, it sounds absurd to claim "I see that it's raining outside, but I don't know that it is", and perhaps what explains it sounding absurd is that it is absurd. The guiding intuition here is that what best explains the seeming contradiction is that we are expressing a genuine contradiction.¹⁶⁰

Moreover, there is a plausible semantics of perception-verbs in propositional contexts (e.g., case in which a perception-verb is followed by an embedded-clause), which seems to have the consequence that the Entailment Thesis is true. If this is right, then the Entailment Thesis just falls out of a plausible semantics for perception-verbs in those contexts. For example, Gisborne (2010) and French (2012) argue that there are certain senses of 'see' in ascriptions of 'S sees that p ' which are knowledge-ascribing.¹⁶¹ If this is right, then it is hard to see how an entailment thesis is not true of those senses of 'see' when it is embedded in sentences of the form 'S sees that p '.

¹⁶⁰ Cf. Unger (1972, 304-305) for this suggestion. But there is a problem for this suggestion. Consider, objectual-perception reports, as in the first conjunct of (1) 'S sees the bus stopping, but doesn't know that the bus is stopping'. (1) sounds just as absurd as (2) 'S sees that the bus is stopping, but doesn't know that the bus is stopping'. But, plausibly, there is no entailment from 'S sees *an x which is F*' to 'S knows *that x is F*'. This result seems to undercut the idea that the Entailment Thesis is what best explains the seeming inconsistency of (2). Dretske (1969) provides a *pragmatic explanation* for the seeming inconsistency of reports like (1), rather than a *semantic explanation* (which would appeal to considerations like the Entailment Thesis), and one might think that this kind of explanation can be extended to reports like (2), although it's not clear that it will work. Thanks to Craig French for pointing this out to me.

¹⁶¹ For example, to say that one can see that S is quickly moving up the ranks seems to entail that the speaker knows that S is quickly moving up the ranks, but not because they can *visually* see that S is quickly moving up the ranks. French (2013) calls this the purely epistemic sense of "sees that" ascriptions.

We might also appeal to the pragmatics of our explanations of knowledge. For example, a good explanation of how I know that it's raining outside in certain contexts is that I just see that it's raining outside. According to Cassam, these sorts of explanations of how I know that something is so “have a kind of finality that many other explanations lack” and “[...] are not open to the same [kinds of] challenges” that other explanations of how I know something might be open to” (2007a, 334 my addition).

For example, if I ask 'how does S know that *p*?', and I'm told 'S read that *p*', it makes sense to ask 'but does S really know that *p*?'. However, if the response had been 'S saw that *p*', it doesn't make sense to ask 'but does S really know that *p*?'. Why? According to Cassam, the reason is that, while the former statement is open to further challenge, “there is no room for further challenge” of the latter statement ““unless this is a way of questioning whether you really see that *p*” (my emphasis; 2007a, 334).¹⁶² Extending this idea, we might think that

¹⁶² Consider two views about visual propositional perception: (WK) 'seeing that *p* is a way of knowing that *p*' and (EWK) 'seeing that *p* is a particularly *exemplary* way of knowing that *p*'. Notice that it seems possible for (WK) to be true *without* (EWK) being true. How? First consider this question: what makes a way of knowing that *p* a particularly *exemplary* way of knowing that *p*? One condition might be that that way of knowing that *p* actually *implies* that the agent knows that *p*. So, for example, we could explain the truth of (EWK) in terms of the Entailment Thesis, even though the truth of the Entailment Thesis might *not* explain the truth of (WK). For example, consider (RWK): 'reading that *p* is a way of knowing that *p*'. Plausibly, we may think that one can *read*, in a reliable newspaper, say, that a building was just demolished, and thereby come to *know* that the building was just demolished. But *reading that p* doesn't seem to imply *knowing that p*. After all, we can *read that p* even if *p* is false. I can read that a certain election was rigged even if this isn't true; perhaps the journalist was misled, and given false information. Insofar as knowledge is *factive*, then, it is at least possible for reading that *p* to be a way of knowing that *p* without entailing that whenever S reads that *p*, S knows that *p*. How does this consideration relate to the passage from Cassam? First, we might think that (EWK) provides support for the Entailment Thesis, even if (WK) doesn't. On this score, I am appealing to the relevant passage from Cassam (2007a) pp. 343-344 in order to highlight a *prima facie* motivation for endorsing the Entailment Thesis. But whether this reason can be sustained under scrutiny is open to dispute, and we can see this by looking to some of the other things that Cassam says. For example, it's not clear that Cassam would endorse the *prima facie* motivation that I presented in the main body of the text. Here's why: according to Cassam, ϕ -ing that *p* can be a way of knowing that *p* even if 'S ϕ -ed that *p*' doesn't entail 'S knows that *p*'. Of course, Cassam could maintain that for ϕ -ing that *p* to be an *exemplary* way of knowing that *p* it must entail knowledge. But it is not obvious that Cassam does think this. See his (2007a) pp. 345-347 for some discussion on this issue. Moreover, Cassam also seems to hold that explanations of how someone knows that *p* in terms of *objectual*

what best explains this pragmatic phenomena is the truth of the Entailment Thesis.

But what turns on the truth of the Entailment Thesis? On the one hand, we might think that there is an intimate relationship between perception and knowledge which has the Entailment Thesis as a consequence. One relationship might be that we cannot understand the former without reference to the latter: we cannot understand what it is to perceive that p without reference to what it is to know that p .¹⁶³ As Cassam (2007a) highlights “[i]f perceiving that P is a form of knowing that P then it is easy to see why we could not explain all the features of the concept of perception without reference to the concept of knowledge” (Cassam 2007a, 38). So on this kind of account, the Entailment Thesis would be a consequence of our conceptual analysis of propositional perception.

A related idea is expressed in Williamson's (2000). One of his central aims there is to provide a non-reductive account of propositional knowledge. In particular, he maintains that “knowing is the most general stative propositional attitude such that, for all propositions where p , necessarily if one has it to p then

perception (e.g., seeing an object o which is F) can, in certain sorts of cases, be exemplary explanations of how someone knows that p in at least something like the way in which explanations of how someone knows that p in terms of seeing that p are exemplary. See Cassam (2008), pp. 39-41 in particular. So Cassam does not appear to maintain that just because certain explanations of how someone knows that p might be particularly exemplary, such as explanations in terms of propositional perception or objectual perception, the exemplariness of the explanation needs to be explained by appeal to the Entailment Thesis. Furthermore, if it is true that explanations of how someone knows that p in terms of objectual perception can be exemplary explanations, one might then want to reassess the motivating consideration I highlighted in the main body of the text. After all, if they can be exemplary without being knowledge-entailing, this will be a problem for appealing to the exemplariness of explanations of how someone knows that p in terms of propositional perception *being grounded in the Entailment Thesis*. This all requires further discussion, but for my purposes I can set it aside, since I only want to highlight some *prima facie* considerations which motivate the Entailment Thesis. I have used Cassam (2007a) and (2008) to serve that purpose. But how one ought to read Cassam's own commitments, and whether those motivating considerations provide adequate support for the Entailment Thesis, are further matters. For theoretical considerations in favour of the Entailment Thesis, see Cassam (2007b), Williamson (2000), Stroud (2002, 2009, 2011), Unger (1972), and Dretske (1969).

¹⁶³ For example, Strawson (1992, 1995) suggests that we cannot understand the concept of *perception* without reference to the concept of *knowledge*, and vice-versa.

p is true” (Williamson 2000: 39). So if seeing that *p* is a particular factive stative attitude, Williamson's account has the consequence that seeing that *p* entails knowing that *p*. So in order to maintain Williamson's non-reductive account of knowledge, it seems as if the Entailment Thesis must be true.¹⁶⁴

Despite its intuitive and theoretical appeal, there have been a number of powerful arguments against the entailment thesis (see McDowell 2002, Turri 2010, and Pritchard 2011, 2012). I will group them into three general kinds of argument. One argument proceeds from the premise, *a la* Gettier, that knowledge excludes luck, and the premise that propositional perception does not exclude luck, to the conclusion that one can see that *p* but not know that *p*. Call this the Argument from Luck. In §2, I argue that the Argument from Luck is not sound and therefore poses no threat to the Entailment Thesis. The second argument proceeds from the premise that we can see that *p* even if there are known defeaters to *p*; and if there are known defeaters to *p*, then we ought not believe that *p*. But if we know that *p*, then we ought to believe that *p*. Thus, we can see that *p* but not know that *p*. Call this the Argument from Defeat. In §3, I argue that this argument is either unsound or else relies on a crucial assumption which forms the basis of our third argument against the Entailment Thesis. The third and perhaps more powerful argument against the Entailment Thesis moves from the premise that knowledge entails belief, and the premise that we

¹⁶⁴ A final source of interest comes from concerns with scepticism. According to Stroud (2009, 2011), the possibility of propositional perception is central to avoiding scepticism. On his view, if we want to understand “how our knowledge of the world is possible we cannot settle for less than our being able sometimes simply to see or in general to perceive that such-and-such is so in the world we take ourselves to know about” (Stroud 2009, 564-565). But in order to make room for this view, we have to “abandon the idea of purely sensory knowledge that is neutral with respect to how things are in an independent world”—that is, we have to make room for the idea that “[propositional perception is] a way of knowing” about the world around us (Stroud 2009, 565-566 my addition). Stroud (2011) registers this idea in the first three paragraphs of that paper, building on from his (2009) work on scepticism about the perceptual knowledge of the world. Moreover, although Stroud's reasons are different from Dretske's, Dretske expresses a similar idea in his (2010) contribution to *A Companion to Epistemology* (2nd ed). For a discussion on how theories of propositional perception make contact with the arguments for scepticism, see Cassam's (2007b), Chapter 1, McDowell (2008), and Stroud (2011).

can see that p without believing that p , to the conclusion that we can see that p without knowing that p . Call this the Argument from Belief. In §4, I argue that the proponent of the Entailment Thesis can provide at least three diagnoses of where this argument fails.

3.2 The Argument from Luck

It is a platitude about propositional knowledge that it excludes epistemic luck. Often, this claim is put forward as an intuition which explains other, perhaps even more basic, epistemological intuitions. The intuition can be expressed along the following lines. Knowledge is a kind of achievement. But if one gets things right because of some imposition of luck—an imposition which luckily gave her a true, rather than a false belief, say—then it cannot also be the case that it was an achievement, something for which the agent can take credit. The intuition here is just that genuine achievements exclude success by luck, and knowledge is some such achievement.¹⁶⁵ Likewise, in the Gettier counterexamples (1963), it seems like although the agent formed a true justified belief, the agent nevertheless did not have knowledge. And what seems to explain her lack of knowledge is that she got things right rather than wrong because of luck. Indeed, we might just think that this was the point of the Gettier counterexamples: to expose that what was wrong with the tripartite definition of knowledge is that it did not exclude luck.¹⁶⁶

Now consider the Entailment Thesis. We can ask whether or not seeing or otherwise perceiving that p entails that we haven't luckily gotten a true belief. If knowledge excludes luck, and a state ϕ entails that S knows that p , ϕ excludes

¹⁶⁵ Several epistemologists (Sosa 2007; Greco 2003; Pritchard 2005) think that the anti-luck intuition on knowledge is grounded in some such intuition about how knowledge is a kind of achievement. Cf. Pritchard (2007, 277) for a detailed expression of this intuition.

¹⁶⁶ Cf. Dancy (1985): “[J]ustification and knowledge must somehow not depend on coincidence or luck. This was just the point of the Gettier counterexamples; nothing in the tripartite definition excluded knowledge by luck” (Dancy 1985, 84). See also Pritchard *forthcoming*, where he argues against some recent attempts to show that knowledge is compatible with epistemic luck.

luck as well. For if one can see that p but nevertheless get a true belief because of luck, then seeing that p could not entail that one knows that p . So, let's consider a case where it looks like we can see that p and believe truly that p because of some imposition of luck:

Out in a field, Smith and Jones are walking around what look like a hundred or so barns. Smith, some twenty-meters in front of Jones, stops in front of a particular barn and looks at it. He thinks to himself 'look at how bright this barn is'. Jones catches up, but a bug has become lodged in his eye. He looks at the barn, but doesn't recognize it as such. He asks 'is that a barn or what?'. Smith replies, 'yes, I can see that it is'. Unbeknownst to Smith and Jones, however, that is the one barn among hundreds of fake barns made to look just like barns.

Call this case BARN after Goldman's (1976) barn-façade case. What should we draw from BARN? Consider the following argument:

- (1) Smith sees that it is a barn.

This premise is intuitively plausible. After all, Smith is looking at a real barn, he applies the concept <barn> to the barn in front of him, the weather is clear, and his perceptual faculties are working fine. We have no prima facie reason not to think that he sees that it's a barn. Together with the Entailment Thesis, (1) implies:

- (2) Smith knows that it is a barn.

But this doesn't seem right. For this case looks like nothing more than a standard barn-case, where the intuition is that the subject does not know that it is a barn. Indeed, had Smith come upon one of those fake barns—something that could very easily have happened, to be sure—he would have mistaken it for a barn. So, it seems like Smith's getting a true belief rather than a false belief was due in part to luck. But knowledge excludes luck. So:

- (3) Smith doesn't know that it's a barn.

Contradiction. Now, how should we resolve this puzzle? We have at least three options. The first option would be to maintain the Entailment Thesis and the

anti-luck condition on knowledge, but deny that Smith sees that a barn is there. Call this the Straightforward Response. The second option would be to maintain the anti-luck condition, but also maintain that Smith sees that a barn is there. This leaves us with two further options. The first option is negative. If BARN is a case where luck prevents the agent from knowing that a barn is there, then we should reject the Entailment Thesis. Call this the Negative Response. However, if BARN is not a case where luck prevent the agent from knowing that a barn is there, then it is still open for one to maintain the anti-luck condition on knowledge, and just insist that, in this case and others like it, because the agent sees that a barn is there, they know that a barn is there as well. Call this the Positive Response. Of course, the burden is on the proponent of the Positive Response to explain how what looks like a standard barn-case is in fact not a standard-barn case. In what follows, I'll present an argument for the Positive Response, and explain how the verdict in this case doesn't conflict with our intuitions about barn-cases.

Consider first the Straightforward Response.¹⁶⁷ This might look like the most natural response for the proponent of the Entailment Thesis. After all, if the Entailment Thesis is true, then if Smith sees that a barn is there, he knows that a barn is there. But since he does not know that a barn is there, it follows that he does not see that it is there either.¹⁶⁸ While this response might look

¹⁶⁷ Both Millar (2011) and Cassam (2008) advance this kind of response. Their motivation is not to preserve the Entailment Thesis *per se*, however, since their motivation is just to describe the relationship between visual perception and knowledge, where a consequence is the truth of the Entailment Thesis. For Millar at least, environments with multiple ringers prevents one from exercising their relevant “recognitional capacities” and it's this which prevents the agent in barn-type cases from seeing that a barn is in front of them. Cassam, on the other hand, just seems to maintain that it's because S can't know in barn-type cases which prevents S from seeing that a barn is there as well, given that seeing that *p* is a way of knowing that *p*. See Millar (2011) pp. 333-337, and Cassam (2008) pp. 38-41, in particular.

¹⁶⁸ For an excellent discussion on different theories of propositional perception and their relation to perceptual experience, see French (2013). Moreover, French (2012) argues against Turri's (2010) arguments against the Entailment Thesis. But French appeals to broadly linguistic considerations, while in this paper I am appealing to broadly epistemological considerations. In fact, it would be interesting to see how one's theory of the nature of propositional perception informed how one ought to respond to the three general

plausible, the proponent of the Entailment Thesis is not confined to this response. On this score, the proponent of the Entailment Thesis might think that the Straightforward response trades on the idea that what explains Smith's not seeing that barn is there is the truth of the Entailment Thesis plus his not knowing that a barn is there. But now the following question arises: how are we to explain his not seeing that a barn is there together with the fact that he sees a barn, applied the concept <barn> to a barn, is competent with that concept, and his perceptual faculties are working fine, in otherwise good conditions? The point here is just that it might look hard to reconcile the fact that Smith saw a barn in close range, applied the concept <barn> to the barn, the barn looked to him to be a barn, with the alleged fact that Smith nevertheless failed to see that it was a barn. On the other hand, if one thinks that seeing that *p* is an epistemic state which amounts to knowing that *p*, then one is free to think that, since one of the epistemic conditions of seeing that *p* is not satisfied—knowing that *p*—it follows that Smith didn't see that *p* either. The thought here is that the considerations I adduced above won't count against the Straightforward Response, because the distinctively epistemic conditions on seeing that *p* weren't met.

The proponent of the Entailment Thesis could also mount a different objection. I will return to this point in a moment, where I present an argument for the Positive Response. But for now, let us turn to the second option: the Negative Response. The Negative Response looks plausible as well. After all, BARN at least looks like a Gettier case. And if it is a Gettier case, Smith doesn't know that a barn is there. But this is compatible, so the thought goes,

arguments I discuss here. For example, an anonymous reviewer pointed out that if we think of the nature of propositional perception as a propositional-perceptual state, where there need be no necessary epistemic conditions on its obtaining, then the worry with the Straightforward Response seems to be a good one. On the other hand, if we think of the nature of propositional perception as an epistemic-perceptual state—that is, as having epistemic conditions, such as putting the perceiver in a position to know that *p* for its obtaining—then the Straightforward Response looks more plausible for the proponent of this view of the nature of propositional perception.

with him still seeing that a barn is there. Indeed, as Pritchard has pointed out, wouldn't it make sense for Smith, in retrospect, to claim that while he saw that a barn was there, he just didn't know that it was a barn?¹⁶⁹ If this is right, then it looks as if one can see that p , but not know that p , because seeing that p does not guarantee that we have fulfilled the anti-luck condition on knowledge. Call this the Argument from Luck.

I want to pause here to examine this argument in more detail. First, we need to make more clear what epistemic luck is supposed to be. What we want to understand is how luck prevents the agent who sees that a barn is there from knowing that a barn is there. It is no longer the case that epistemologists have no clear and developed conception of epistemic luck. Instead, there are at least two accounts of the epistemic luck condition on propositional knowledge, expressed as modal conditions on belief. What we should examine in more detail, then, is whether Smith fails to meet these conditions. If he doesn't, then we will have good reason for thinking that the Argument from Luck stands.

How do we articulate the anti-luck condition on knowledge? Several epistemologists have framed the anti-luck condition in terms of certain modal conditions. The two most prominent modal conditions are the sensitivity condition (Nozick 1981; Dretske 1970) and the safety condition (Sosa 1999; Williamson 2000; Pritchard 2005). According to the sensitivity condition:

SENSITIVITY: S 's belief that p is sensitive iff in all nearby possible worlds in which $\sim p$, S doesn't believe that p .¹⁷⁰

¹⁶⁹ Pritchard (2011) pp. 442-443 and McDowell (2002) p. 277 makes the same point (indeed, the case McDowell gives is structurally analogous to Turri's (2010, 199) "Rabbit" case) where the agent in McDowell's case believed that the object merely *appeared* green because he justifiably thought it was under a colour distorting light, even though he later discovers there was no such light, and he actually *saw that it was green*.

¹⁷⁰ Nozick's expression of this condition was in terms of the subjunctive conditional: $\sim p \square \rightarrow \sim B(p)$. That is, if p were false, I wouldn't believe that p . He later refines the principle in his (1981) Chapter 3, §III, relativizing it to an epistemic basis. On this model, where b is a basis for belief, the condition would read: $\sim p \square \rightarrow \sim B(p)$ on basis b .

For example, to borrow a famous example from Nozick, if I were a brain in a vat I wouldn't believe that I was. That is, if 'I am not a brain in a vat' were false, I would nevertheless continue to believe that it's true. And this makes the belief that 'I am not a brain in a vat' not a sensitive belief, even if it is true that I'm not a brain in a vat.

Now, it seems like in all nearby possible worlds in which it's not the case that a barn is in front of Smith, he still believes that there is. After all, those fake barns look just like real barns. The proponent of the Argument from Luck can therefore appeal to the sensitivity condition in order to explain Smith's lack of knowledge.

Of course, this is just one popular articulation of the anti-luck condition on knowledge. But in order to avoid some of the complications with this principle, some epistemologists have adopted a similar, though quite distinct, modal condition on knowledge:¹⁷¹

SAFETY: *S*'s belief that *p* is safe iff in all nearby possible worlds in which *S* believes that *p*, *p* is true.¹⁷²

For example, I believe right now that a computer is in front of me because of what I see. And if I continued to believe that a computer is in front of me on the basis of what I see to be so, my belief that it is in front of me would continue to

¹⁷¹ The complications with sensitivity are not relevant for our purposes, so I'll only briefly register them here. The complications include the possibility of inductive knowledge with Sosa's (1999) 'Garbage Chute' case, and the denial of the closure principle for knowledge. Also, note that I am avoiding interesting issues about whether or not either of sensitivity or safety, however formulated, really does articulate the anti-luck condition on knowledge. There are arguments on offer for the claim that, while propositional knowledge excludes epistemic luck, neither safety nor sensitivity in any of their formulations properly expresses the anti-luck condition. See Avram Hiller and Neta (2007), Vogel (2007), and Lackey (2008).

¹⁷² I want to register that there are *at least* three different *expressions* of the safety condition on offer. Here's Williamson's (2000, 147) version: "if one knows, one could not easily have been wrong in a similar case". Pritchard (2007, 281): "S's belief is safe iff in most near-by possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world the belief continues to be true". And Sosa (1999, 146): "If S were to believe that *p*, *p* would be true", though Sosa relativizes safety to a basis later in his (2007).

be true. In general, if our belief forming mechanisms, or the epistemic bases from which we believe a target proposition *p* are good enough to make sure that we would continue to have a true belief in counterfactual situations, or in nearby possible worlds, then the belief is safe.

Now, it at least seems like Smith's belief is unsafe. In all those nearby possible worlds in which he believes that a barn is in front of him, he would be wrong. After all, too easily could he have come upon a fake barn which looks just like a barn.

The proponent of the Argument from Luck therefore has a powerful case. She can argue that Smith's belief that a barn is in front of him is neither safe nor sensitive. But if either of those principles properly articulates the anti-luck condition on knowledge, then Smith's belief fails to count as knowledge.

How should the proponent of the Entailment Thesis respond to this argument? I will argue that she should claim—contra initial appearances—that Smith meets both of those anti-luck conditions on knowledge. I will pursue this argument as follows. First, we need to distinguish between getting things right—e.g., forming a true belief rather than a false belief—as a matter of luck because of our circumstances, and getting things right as a matter of luck because of the basis from which we believe the target proposition.¹⁷³ Second, I will show how Smith's getting a true belief rather than a false belief is a matter of luck in the first sense but not the second sense. But the second sense is what matters for epistemic assessment and appraisal. The conclusion is that the Argument from Luck is not sound—Smith meets various anti-luck conditions on knowledge. If one retains the intuition that Smith nevertheless does not know that a barn is in front of him, something else must be explaining that

¹⁷³ This distinction can be found in Pritchard (2005), and first in Engel (1992) as the distinction between *evidential luck* and *veritic luck*. See also Unger's (1968) for three other distinctions between types of luck, none of which, he argues, are incompatible with knowledge. I appeal to these considerations against Stroud's puzzle, Chapter 3.

intuition if it is indeed accurate. If not, the intuition has otherwise been rendered harmless.

3.3 Resisting the Argument from Luck

The proponent of the Argument from Luck maintains that Smith sees that a barn is there, but does not know that a barn is there because he got things right rather than wrong because of some knowledge-excluding imposition of luck. It is then argued that because knowledge excludes luck, Smith does not know that a barn is there, even though he sees that it is there, which is incompatible with the Entailment Thesis.

But the proponent of the Entailment Thesis should question the second step. Did Smith get things right rather than wrong because of some imposition of luck? And if he did, was this imposition of luck the kind of luck that prevents him from knowing that a barn is there?

It should be granted that it is a matter of luck that what Smith came upon was a barn and not a fake-barn. And his coming upon a barn of course enabled him to see what he saw: a barn. But that his belief is true rather than false given his epistemic basis which supports his belief is not a matter of luck. Were he to believe that a barn is there on the basis of seeing that it is there, his belief would be true. Likewise, were it false that a barn was in front of him, it would also be false that he would believe the same proposition on the same basis.

Consider Sosa's account of the safety condition. Sosa relativizes safety to a basis, where a "belief cannot constitute knowledge if the believer might too easily have so believed on the same basis while his belief was false" (Sosa 2008:124). Now, Smith meets this safety condition. His perceptual basis is propositional perception. Indeed, both the proponent of the Entailment Thesis and the proponent of the Argument from Belief agree on this point. Both can agree that the agent sees that a barn is there and not just sees a barn. But then it follows that, from Smith's perceptual basis—propositional perception—not very easily would it have been false if based on that basis. In other words, given

Smith's basis, not very easily would his belief have been false because not possibly could his belief have been false—again, given his basis. Consider, then, this revised version of the safety principle:

SAFE BASIS: S's belief that p is basis relative safe iff in all nearby possible worlds in which S believes that p on basis B, S's belief that p is true.

Smith meets this condition as follows:

SAFE BASIS (Barn): In all nearby possible worlds in which Smith believes that a barn is there on the basis of propositional perception—for example, on the basis of seeing that a barn is there—his belief is true.

Smith also meets the basis-relative version of sensitivity as follows.¹⁷⁴

SENSITIVE BASIS (Barn): In all nearby possible worlds in which it is false that a barn is there, in front of Smith, he does not continue to believe that it is there on the basis of propositional perception—for example, on the basis of seeing that it is there.

When we index the agent's epistemic basis, the crucial premise of the Argument from Luck becomes neutralized. It should not be maintained that what the agent sees is that a barn is there, while he nevertheless does not know that a barn is there because of some imposition of knowledge-excluding luck. As I argued here, the kind of luck that does afflict the agent in BARN is not the kind that prevents him from knowing that a barn is there. The agent only luckily came upon a barn rather than a fake barn. This only luckily put him in a position to have the evidence that he has. But given the evidence he got through

¹⁷⁴ Nozick relativizes sensitivity to an evidential basis in his (1981, 179).

the good grace of good luck, his belief is true, and it would continue to be true insofar as he continued to believe the same proposition on that same basis.¹⁷⁵

I want to close this section with some important points. The first point is this: perhaps the sensitive and safe basis conditions are not the luck-excluding conditions required for knowledge. But if that is right, then so much the worse for the proponents of the Argument from Luck. The thought here is that the proponents of this argument have some conception of luck that they're appealing to in assessing whether or not the agent knows in BARN, and it is this conception which propositional perception can fail to exclude. But what conception is that? If it is not the conception embodied in either of the safety or sensitivity conditions on knowledge and their analogues, then it is not clear which conception they had in mind, if any at all. The second point is that it's unclear whether or not what I have argued here is incompatible with the verdict in Goldman's (1976) barn-case. First, notice that in Goldman's case, he was not specific about the agent's epistemic belief forming basis for the target proposition—at least, not as specific as the proponents of the Argument from Luck are with their barn case. I have argued that if the agent's belief forming basis is propositional perception, then given that propositional perceptions satisfies various anti-luck conditions, it follows that the agent knows that a barn is there. The proponent of the Argument from Luck accepts the antecedent, and I have urged here that proponents of the Entailment Thesis should as well. But one can grant this without also conceding that Goldman's verdict in the original barn case was wrong. The two verdicts are compatible. But how?

We should return to the distinction between good luck placing us in a position to have the epistemic basis that we have for a target proposition and

¹⁷⁵ Perhaps one will agree with me that in BARN, S satisfies the anti-luck conditions on knowledge, but maintain that S still fails to know that a barn is in front of them for other reasons independently of the anti-luck conditions. At this juncture, the proponent of the Entailment Thesis could then mount the Straightforward Response. After all, this is compatible with the argument against the Argument from Luck that I provided being successful. Thanks to an anonymous reviewer of *Synthese* for this helpful suggestion.

good luck making our belief true rather than false, irrespective of our epistemic basis for that belief. For example, if I flip through pages in a book, and just guess that I will land on page 323, then even if I formed a true belief, guesswork is not the kind of epistemic support, or evidential basis, or in general epistemic basis, which would have allowed me to continue to have true belief rather than a false belief in different circumstances. In these circumstances, guessing got me a true belief rather than a false belief, but had things been different, guessing could have given me a false belief. But this is not true of some epistemic bases. Propositional perception is such that whenever I believe that p on that basis of perceiving that p , the belief will be true. In this fashion, propositional perception is not susceptible to what Engels (1992) called “veritic luck”:

VERITIC LUCK: An agent S 's belief that p in circumstance C is veritically lucky iff given S 's evidence for p , it is just a matter of luck that S 's belief that p is true in C .

Veritic luck is the kind of luck that looks incompatible with knowledge. And notice that Smith's belief that a barn is there is not veritically lucky. For given that his basis in BARN is propositional perception, it is not just a matter of luck that his belief that a barn is there is true rather than false in those circumstances. But now here is the crucial point for the proponents of Goldman's verdict in his original barn case. Barn cases can be set up so that the agent does not have propositional perception as their evidence. We can imagine that Smith's epistemic basis is instead a visual experience or impression as of a barn being there. In this case, if Smith's epistemic basis for believing that a barn is there is his having a visual experience as of barn being there, this basis does not eliminate knowledge-excluding luck. In all nearby worlds in which Smith has an experience as of a barn being there, a barn isn't there. After all, there's plenty of nearby worlds where Smith has the same epistemic basis as he does in the actual world, since sufficient for having an experience as of a barn being there

is simply seeing a fake-barn. So, if we index Smith's epistemic basis to a visual experience as of a barn being there, he won't satisfy SAFE BASIS, since too easily could his basis have given him a false belief. So too with SENSITIVE BASIS, since had a fake-barn been in front of him, he still would have believed that it was a barn on the basis of having an experience as of a barn being there. In general, then, if the agent's epistemic basis is construed to be something like having a visual experience as of a barn being there, then even if in the actual world he is standing in front of the barn, and believes that a barn is there on the basis of having an experience as of a barn being there, this belief will be true, but veritically luckily true, and so not knowledge.

So, perhaps what makes the difference in whether we should assess the agent in a barn case as knowing or not knowing is their epistemic basis in those cases. If the basis is propositional perception, as the proponent of the Argument from Luck maintains, we have good reason for rejecting the claim that the agent does not know that a barn is there. But, again, this is compatible with different verdicts in other barn cases where the epistemic basis has changed.

But now, what if someone maintains that Smith fulfills those anti-luck conditions on knowledge, but nevertheless does not fulfill all of the relevant anti-luck conditions on knowledge? How should the proponent of the Entailment Thesis respond?

Of course, it is open to the proponent of the Entailment Thesis to fall back on the Straightforward Response, if that position can be sustained. But one can also inquire into the nature of the kind of luck that is, according to this opponent, knowledge-excluding in BARN. At this juncture, the kind of luck that is supposed to be knowledge-excluding is some kind of environmental-luck, if the argument I have given against the Argument from Luck is sound. On this view, Smith's coming upon the one barn and seeing that it is there rather than a fake barn is lucky in a way that is knowledge-excluding, despite the fulfillment of the other anti-luck conditions on knowledge that I

distinguished. Two kinds of environmental-luck are relevant to our discussion here:

- (1) The luckiness of 'there is a barn here' being true, in the circumstances in which there are more fake-barns around.
- (2) The luckiness of 'S came upon and saw that a barn was there' being true, in the circumstances in which there are more fake-barns around.

A proposition can be known even if it is a contingent truth. If one maintains that Smith does not know that a barn is there because of the kind of luck present in (1), it is hard to see how we avoid the implausible consequence that no one can know contingent truths. What explains the luckiness of (2), on the other hand, is that there is some sense in which Smith could very easily have come upon one of the fake-barns, given their prevalence in the area. If Smith had done so, he might have formed the false demonstrative belief that that (the barn in front of him) is a barn, or the true general belief that a barn is in the field. But I argued that once we index the anti-luck conditions to the relevant epistemic basis—propositional perception—it's not a matter of luck that the agent formed a true rather than a false belief. Otherwise, it offends against the orthodox on the nature of epistemic luck that the kind of luck present in (2) is knowledge-excluding, unless one's epistemic basis fails to have a non-accidental connection to the facts, which I argued earlier is not the case.¹⁷⁶

A final point on this score is that the argument I've presented here against the Argument from Luck is not an argument for the truth of the Entailment Thesis, nor is it an argument for the view that agents can know that a barn is in front of them in barn-cases. The conclusion of the argument is supposed to be modest. It's just that the proponent of the Entailment Thesis can grant the

¹⁷⁶ Cf. Unger (1968), Engel (1992), and Pritchard (2004, 2005, 2013). For discussion, see Engel (2011).

premise of the Argument from Luck that the agent sees that a barn is in front of them and consistently maintain that the agent thereby knows that a barn is there without offending against the core anti-luck intuitions.

3.4 The Argument from Defeat

One might be tempted at this juncture to maintain that what can go wrong for the agent in BARN is that there is a known defeater present which prevents the agent from knowing that a barn is there. For example, that fake-barns are present defeats, not his seeing that a barn is there, but his knowing that a barn is there. Consider the following reasoning adapted from Pritchard (2011, 442-443) in support of this view:

(D1) Smith sees that a barn is there.

(D2) Smith ought not believe that a barn is there. His reliable friend Jones told him that fake barns are abound, and that these fake barns look just like barns.

From (D1), it follows that:

(D3) If the Entailment Thesis is true, Smith knows that a barn is there.

Now suppose that the Entailment Thesis is true. From (D1) and (D3) it follows that:

(D4) Smith knows that a barn is there.

But from (D4) and (D2) it follows that:

(D5) Smith ought not believe that a barn is there and Smith knows that a barn is there.

(D5) is taken to be a contradiction. If Smith knows that p , then it is not the case that he ought not believe that p . So (D1) – (D5) is taken to be a reductio of the Entailment Thesis. Call this the Argument from Defeat.

The Argument from Defeat appears to be valid. But is it sound? Consider the following support for (D1) and (D2) from Pritchard (2011):

Suppose, for example, that one is in a situation in which one is genuinely visually presented with a barn and circumstances are in fact epistemically good (there's no deception in play, one's faculties are functioning correctly, and so on). But now suppose further that one has been told, by an otherwise reliable informant, that one is presently being deceived (that one is in barn facade county, say), even though this is in fact not the case. Clearly, in such a case one ought not to believe the target proposition, and hence one cannot possibly know this proposition either. (Indeed, if one did continue to believe the target proposition even despite the presence of this undefeated defeater, then one would still lack knowledge). Still, does it follow that one does not see that the target proposition obtains? I think not. [...] For suppose that one were to discover subsequently that the testimony one received was false, but that everything else one knows about the circumstances in which one was presented with this (apparent) barn remained the same. Wouldn't one now retrospectively treat oneself as having earlier seen that there was a barn? Think, for example, about how one would describe one's situation in this regard were one to be asked about it. Wouldn't it be most natural to say that one did see that there was a barn in the field, rather than to 'hedge' one's assertion by saying, for example, that one merely thought that one saw a barn? (Pritchard 2011, 442-443).

Of course, notice that in order for (D5) to be a genuine contradiction, it has to also be true that if S knows that p , S ought to believe that p . Contraposed, if it is not the case that S ought to believe that p , then it is not the case that S knows that p . This principle is exploited in Pritchard's reasoning in the third sentence: “Clearly, in such a case one ought not to believe the target proposition, and hence one cannot possibly know this proposition either” (Pritchard 2011, 442 - 443). With this in mind, we are now in a position to make more explicit Pritchard's reasoning and the principles that reasoning relies on:

PROPOSITIONAL PERCEPTION: S sees that a barn is there.

DEFEATER CONDITION: If a reliable informant tells *S* that fake barns are abound, then *S* ought not believe that a barn is there.¹⁷⁷

DEFEATER: A reliable informant tells *S* that fake barns are abound.

ENTAILMENT THESIS: If *S* sees that a barn is there, *S* knows that a barn is there.

EPISTEMIC OUGHT: If it is not the case that *S* ought to believe that a barn is there, then it is not the case that *S* knows that a barn is there.¹⁷⁸

From (PROPOSITIONAL PERCEPTION) and (ENTAILMENT THESIS), it follows that *S* knows that a barn is there. But from (DEFEATER CONDITION) and (DEFEATER) it follows that *S* ought not believe that a barn is there. However, from (EPISTEMIC OUGHT) it now follows that *S* does not know that a barn is there—contradiction. In order to resolve the contradiction, Pritchard eliminates (ENTAILMENT THESIS). But I think that this is the wrong move. Instead, we focus on (PROPOSITIONAL PERCEPTION), (DEFEATER), (DEFEATER CONDITION), and (EPISTEMIC OUGHT).

Consider first (DEFEATER). One might want to argue that if Jones is reliable and Jones tells Smith that fake barns are abound, then Smith ought not believe that a barn is there, in front of him. This looks plausible. But the proponent of the Entailment Thesis can be dogmatic, and insist that Smith's seeing that a barn is there defeats that defeater. We can imagine Smith telling Jones the following: "Perhaps there are fake barns around, but I can just see

¹⁷⁷ This premise is drawn out from the following quotation from Pritchard's (2011): "But now suppose further that one has been told, by an otherwise reliable informant, that one is presently being deceived (that one is in barn facade county, say), even though this is in fact not the case. Clearly, in such a case one ought not to believe the target proposition [...]" (ibid).

¹⁷⁸ I draw out EPISTEMIC OUGHT from the remaining quotation in footnote 12 of Pritchard's (2011): "[...] and hence one cannot possibly know this proposition either. (Indeed, if one did continue to believe the target proposition even despite the presence of this undefeated defeater, then one would still lack knowledge)." In any case, the Argument from Defeat would be fallacious without this principle in play.

that this is a barn.” Of course, the response to this jumps off the page. Jones might respond that Smith still ought not believe that a barn is there, even if it is true that he sees that a barn is there. This response maintains that rejecting (DEFEATER) is too dogmatic.

Moreover, notice that one could also mount a response to the Argument from Defeat which mirrors the 'Straightforward Response' from BARN. According to this response, if *S* doesn't know that a barn is there, because she ought not believe that it is there, then so too she doesn't see that it is there. The core thought here is that seeing that *p* is an epistemic state—a state where the satisfaction of '*S* sees that *p*' includes the satisfaction of certain epistemological conditions, the core condition being knowing that *p*. On this score, what is relevant to whether or not *S* knows in these circumstances will also be relevant to whether or not *S* sees that such-and-such is so in these circumstances. If *S* isn't in a position to know that a barn is there, because she ought not to believe that one is here, then she's also not in a position to see that it is there either.

Against this, the proponent of Pritchard's reasoning would need to give us further considerations in favour of the idea that *S* really does see that the barn is there, perhaps considerations that fall out of one's view of propositional perception. Instead, Pritchard's considerations appeal to the idea that it is simply less natural to hedge our claim from 'I saw that a barn was there' to 'I seemed to see that a barn was there', and that this is what provides support for *S*'s having seen that a barn is there. The debate at this juncture, however, should turn on how we should understand propositional perception, and not only on our intuitions about ascriptions of propositional perceptual states.

Now consider (DEFEATER CONDITION). This condition is problematic as well, but not because it is false. It is problematic because if we grant that Smith sees that a barn is there, and so he believes that it is there, how can he both see that a barn is there and withhold belief that it is there? Consider the following intuitive principles:

(+ EPISTEMIC OUGHT IMPLIES CAN) If S ought to believe that p , S can believe that p .

(– EPISTEMIC OUGHT IMPLIES CAN) If S ought to withhold belief that p , S can withhold belief that p .¹⁷⁹

The Argument from Defeat implies that Smith sees that a barn is there and that he ought not believe that a barn is there. But if Smith ought not believe that a barn is there, it must also be true that he can withhold believing that a barn is there. Of course, this is no obstacle in its own right. One can withhold believing that a barn is there and it can be true that one ought not believe that a barn is there. This, in itself, poses no threat. But it poses a threat when we conjoin it with the fact that Smith sees that a barn is there. The proponent of the Argument from Defeat relies on the premise that Smith sees that a barn is there. But how can the proponent of the Argument from Defeat consistently maintain that S sees that a barn is there, and that S ought to withhold belief that a barn is there? From (- EPISTEMIC OUGHT IMPLIES CAN), it follows that S can withhold belief that a barn is there. But that looks incompatible with seeing that

¹⁷⁹ Let's make clear a certain ambiguity in this principle. Let O be deontic operator for 'ought', and B the belief predicate. From $O(S, \sim B[p])$ it does not follow that $O(S, B[\sim p])$, nor does it follow from $\sim O(S, B[p])$ that $O(S, \sim B[p])$. In $\sim O(S, B[p])$, the negation operator is ranging over the ought operator, whereas in $O(S, \sim B[p])$ the negation operator is ranging over the belief predicate. In the first case, what ought not be the case is S believing that p . In the second case, what ought to be the case is S not believing that p . The (- EPISTEMIC OUGHT IMPLIES CAN) principles' antecedent says that what ought to be the case is S withholding belief that p , that is, not believing that p . So, the (- EPISTEMIC OUGHT IMPLIES CAN) principle says that $O(S, \sim B[p]) \rightarrow \diamond \sim B(S, p)$. The intuitiveness of this principle, I take it, derives from the analogous ethical principle that, if what ought to be the case is S not ϕ -ing, then it has to at least be logically possible for S not to ϕ . After all, if one's ethical principle entailed that everyone ought to successfully divide by zero, even though no one could ever successfully divide by zero, that would just be a bad ethical principle: it would entail that we should do what we logically cannot do. Here, I am just taking it that the same is true for normative epistemological principles.

a barn is there, since seeing that p seems to require believing that p . So, if the Argument from Defeat is sound, then the following principle is false:

PROPOSITIONAL PERCEPTION ENTAILS BELIEF: If S sees that p , S believes that p .

This principle seems to be true. How can we make sense of the fact that S sees that a barn is there rather than just seeing a barn, if it is not also the case that S believes that what he sees is a barn? Indeed, the most sustained theories of propositional perception (Dretske 1969; Cassam 2007; Stroud 2002, 2011) take it that what distinguishes objectual perception from propositional perception—seeing an object o which is F versus seeing that o is F —is that the latter, but not the former, requires acquiring the belief that o is F as a result of seeing an o which is F . The claim here, then, is that the Argument from Defeat is not conclusive, because it requires that S can see that p and not believe that p . But that is still an un-argued assumption. For all the proponent of the Entailment Thesis is concerned, we have good reasons for thinking that one cannot see that p unless we believe that p . The proponent of the Argument from Defeat will therefore have to mount an objection which shows that it is possible to both see that p and nevertheless not believe that p . Unless that claim can be sustained, the Argument from Defeat will be ineffective.

3.5 The Argument from Belief

The final argument against the Entailment Thesis that we will be considering is what I call the Argument from Belief. This argument moves from the premise that propositional knowledge entails belief, and the premise that one can see that p even if one does not believe that p , to the conclusion that one can see that p even if one does not know that p . The crucial premise is the second premise: S can see that p even if S does not believe that p .

In his (2010), Turri provides two cases which aim to support this premise. The first case involves an agent who looks at what appears to be Müller-Lyer

lines. Confidently, the agent asserts that the two lines are unequal in length. But in fact the lines aren't Müller-Lyer lines, and instead really are as they visually appear to be: equal in length. So, Turri argues that while the agent saw that the lines were equal in length, he just didn't believe that they were. Turri calls this case (LINES). In the second case, the agent thinks that he is involved in an experiment with hallucinogenic drugs. Later, a mechanical rabbit appears to enter the room. Yet, our agent doesn't believe that he sees a mechanical rabbit, because he is under the impression that he is hallucinating. However, he later discovers that a mechanical rabbit did enter the room, and so he saw that it was a mechanical rabbit. Turri calls this case (RABBIT).¹⁸⁰

Each of these examples follows the same pattern. The first step can be expressed as follows:

(STEP 1) *S* introduces a misleading defeater.

For example, in Pritchard's case it is the reliable testifier who tells *S* that all kinds of fake barns are around, even though there aren't. In Turri's (LINES) case, it is what *S* believes to be a visual illusion but isn't, and in his (RABBIT) case, it is what *S* believes to be a hallucinogen, but isn't. The second step is:

(STEP 2) *S* characterizes themselves as having seen that *p*.

In Pritchard's case, he relies on the idea that it makes sense, in retrospect, to treat *S* as having seen that a barn was in front of him, and that it also makes sense for *S* to not hedge their claim of having seen that a barn was there. In the (LINES) case, Turri relies on the idea that it makes sense, in retrospect, to treat *S* as having seen that the bottom line is the same length as the top line. Likewise, in the (RABBIT) case, Turri relies on the idea that it makes sense to treat *S* as having seen that the mechanical rabbit was in front of him. Note that the reason it seems to make sense to treat all of these agents as having seen that

¹⁸⁰ For the full examples, see Turri (2010, 198-199).

p is that there seems to be nothing built into these cases that makes it look like that agent could not have seen the relevant objects and properties and so could not have seen that certain propositions are true of those objects and properties. But then we also get:

(STEP 3) S characterizes herself as not having believed that p .

STEP 1 makes STEP 3 look plausible. Each agent believes that something abnormal is taking place, e.g., one believes that lots of fake barns which look like barn are around, one believes that they are undergoing a visual illusion, and one believes that they are hallucinating. But if S believes that lots of fake barns indistinguishable from barns are around, how can S also believe that what she sees is now a barn? Likewise, if S believes that she is undergoing a visual illusion, how can she also believe that the bottom line and the top line are the same length? And so on. Of course, notice that none of this implies that the agents could not believe that what seems to be the case is the case. But if it is intelligible for each agent to treat themselves, in retrospect, as having seen that p even though at the time none of the agents believed that p , this is supposed to be a counter-example to the thesis that if S saw that p , S believed that p .

3.6 Neutralizing the Argument from Belief

How persuasive is the Argument from Belief? As with the Argument from Luck and the Argument from Defeat, it is open to the defender of the Entailment Thesis to maintain a Straightforward Response. In this case, one can concede that because we do not believe that p , we don't know that p . And then if the Entailment Thesis is true, it follows that we did not see that p either. So we have no counter-example to the Entailment Thesis unless one provides good reasons for thinking that we do believe that p and that knowledge entails belief.

Are there other responses available to the proponent of the Entailment Thesis? We might think that the Argument from Belief puts just as much pressure on its premises as it does on the proponent of the Entailment Thesis.

For example, the proponent of the Entailment Thesis can grant the proponent of the Argument from Belief the central premise that:

Belief Premise: *S* sees that *p* and does not believe that *p*.

but can take issue with whether or not this alone is sufficient to generate a good argument against the Entailment Thesis. The proponent of the Argument from Belief, recall, argues that:

Knowledge Entails Belief: If *S* knows that *p*, *S* believes that *p*.

Belief Premise: *S* sees that *p* and does not believe that *p*.

Therefore,

Belief Conclusion: *S* sees that *p* and does not know that *p*.

This argument relies on the so far un-argued assumption that knowledge entails belief. Perhaps this view is non-negotiable. But the assumption that propositional seeing or perceiving entails knowing can take on the same status. This too might be an intuition that we are unwilling to concede.¹⁸¹ The point here, then, is just that we can make sense of the idea that the Entailment Thesis

¹⁸¹ Why might one think the Entailment Thesis is non-negotiable? Well, one could maintain that it is part of the concept of propositional perception that it has certain necessary epistemological conditions, such as *knowing that p*. Cf. Cassam (2008, 38) where he discusses this idea. See also Chisholm (1989, 40-41) and Moore (1953, 92) footnote 1. Note that their discussion does not bear directly on the present concern in the text, but one does find expression of the idea that propositional knowledge is a necessary condition of propositional perception. The core thought in the text, however, is that, insofar as we think we *can* and often *do* see that *p*, then because knowing that *p* is so is a necessary condition of that kind of perceptual state, the Entailment Thesis is, in consequence, non-negotiable. A similar thought for the Knowledge Entails Belief principle might run like this: it is *because* we can and often do know that *p* and because it is a *necessary condition* of this state being realised that we *believe* that *p*, so that the Knowledge Entails Belief principle is non-negotiable. What I am doing in the main body of the text is just highlighting how the reasoning on this score might run along the same lines, and so *simply appealing* to either of these principles alone, plus the relevant premise in the Argument from Belief, will be no more effective than (or at least as effective as) the other principle. The thought here is that both moves are available because one can, at least on the surface, consistently maintain that either principle is non-negotiable, so a simple appeal to one principle over the other will remain an ineffective dialectical move.

is non-negotiable. If that is right, then the central premise of the Argument from Belief can be used to argue in the reverse direction as follows:

Entailment Thesis: If S sees that p , S knows that p .

Belief Premise: S sees that p and does not believe that p .

Therefore,

\sim Knowledge Entails Belief: S knows that p and does not believe that p .

In other words, the premise that the proponent of the Argument from Belief uses to support the conclusion that the Entailment Thesis is false is the same premise that the proponent of the Entailment Thesis can use to support the conclusion that the Knowledge Entails Belief Thesis is false.¹⁸² The idea that I want to reinforce here is that the Argument from Belief as it stands gives us at least as much grounds for thinking that knowledge does not entail belief as it does for thinking that we can see that something is so but not know that it is so. A successful argument from belief against the Entailment Thesis will explain what makes this move unavailable. Otherwise, the Argument from Belief is just an appeal to intuitions. But we wanted the argument to decide the fate of the Entailment Thesis rather than leave us with a stalemate.

The second problem is that the example proves too much. It is less clear that the example shows that it is possible to see that p and not believe that p , rather than just showing that sentences of the following form are intelligible 'S saw that p , but S did not believe that p '. But that the sentence is intelligible need not be indefeasible evidence that it is not a contradiction. After all, the sentence (A) 'John found the mask under his desk, but it wasn't there' is intelligible, but

¹⁸² To be sure, *most* epistemologists do take the *Knowledge Entails Belief* principle to be non-negotiable. Yet it's not obvious that most epistemologists take it that knowledge entails *occurrent* belief that p . Indeed, there is also a long epistemological tradition which questions whether knowledge entails any type of belief. See Plato's *Republic* 476-479, Pritchard (1950), Woozley (1953), and Radford (1966).

false and false because it is a contradiction, just as (B) 'Jones remembered the car crash, but the police knew that it was not a crash' is intelligible and false, and false because it is a contradiction. So too, we can imagine that 'S saw that the lines were the same length, he just failed to believe it' is intelligible and false, and false because it is a contradiction. This might indeed not be true, but the example does not show that it is not true, just as the intelligibility of sentences (A) and (B) doesn't settle the matter over their contradictoriness or non-contradictoriness. To be clear, this argument is not intended to show that the case in favor of the second premise of the Argument from Belief is a not a counter-example to the thesis that *S* can see that *p* and not believe that *p*, but is instead intended to be an argument against its conclusiveness. What I am urging here is that there is gap between showing that a certain sentence or utterance is intelligible and showing that a certain sentence or utterance is not a contradiction. There are cases of sentences or utterances which are both intelligible and contradictions. So their intelligibility cannot be indefeasible evidence of their non-contradictoriness. What else is needed is some explanation which closes this gap; some explanation which tells us that what explains its intelligibility is its non-contradictoriness. And we might think that one way to close this gap is to provide an analysis of both '*S* sees that *p*' and '*S* believes that *p*' which reveals that the content of the latter is not contained in or otherwise presupposed by the content of the former.

We should not let the Entailment Thesis hang on the success of these arguments alone. There are other, perhaps more fundamental worries that we might have with the Argument from Belief. One such problem is that the case does not respect the distinction between occurrent belief and dispositional belief. Turri takes his case to show that an agent can see that *p* but not believe that *p*. But all that it seems to show is that the agent can see that *p* without having the occurrent belief that *p*; it is still an open question as to whether or not the agent can see that *p* and lack even the dispositional belief that *p*. In what follows, I will defend the thesis that the agent lacks the occurrent belief that the

lines are of unequal length, but possesses the dispositional belief that the lines are of unequal length. The argument can be applied, *mutatis mutandis*, to the other examples as well. The upshot of this argument is that insofar as propositional knowledge entails dispositional belief and not occurrent belief, the Argument from Belief will be, at best, inconclusive.

3.7 Dissipating the Argument from Belief

We have seen that there are two general problems facing the Argument from Belief. The first problem was that there seems to be a gap in showing that the sentence or utterance 'S saw that p , but did not believe that p ' is intelligible is not indefeasible evidence of it not being a contradiction. The second problem was the Argument from Belief puts at least as much pressure on the thesis that propositional knowledge entails propositional belief as it does on the thesis that propositional perception entails propositional knowledge. The crucial premise of the Argument from Belief can be used to argue against the Entailment Thesis just as well as it can be used to argued against the thesis that propositional knowledge entails propositional belief. In this fashion, the Argument from Belief is not an effective argument against the Entailment Thesis. But as I mentioned in the previous section, the proponent of the Entailment Thesis should not let the Entailment Thesis hinge on these two considerations alone. There is another, perhaps more fundamental concerns that the proponent of the Entailment Thesis can raise against the Argument from Belief. In this section, I explore these concerns.

Let me be explicit about where the Argument from Belief goes wrong. First, Turri, our token proponent of the Argument from Belief, does not respect the putative distinction between occurrent belief and dispositional belief.¹⁸³ Having an occurrent belief that p is the mental analogue of assertion—what is

¹⁸³ Cf. Campbell (1967), Audi (1994), and Schwitzgebel (2010) for this distinction. For the literature on whether or not knowledge entails occurrent belief or acceptance, see Armstrong (1969), Black (1971), and Radford (1970; 1990).

sometimes called judgment—or otherwise an occurrent endorsement of the content of a sentence or assertion. So to have the occurrent belief that p is to actively endorse the content p ; it is, roughly, to bring the proposition that p before one's mind and endorse it. But dispositional belief is not like this. A dispositional belief is, first and foremost, a disposition. If we dispositionally believe that p , and our disposition is triggered, we will manifest our disposition by occurrently believing that p . Dispositional beliefs have two conditions, a 'triggering condition' and a 'manifestation condition'. For example, a fragile glass has the disposition to shatter when struck (the triggering condition), but it can fail to shatter—it can fail to manifest that disposition, thereby failing to meet the manifestation condition—because of a “mask”. What masks the disposition is the soft packaging (Johnston 1992, 223). Likewise, I can have the disposition to actively endorse the content that the lines are of unequal length if I see a pair of lines which look to be of unequal length (a triggering condition), and nevertheless fail to manifest that disposition, because I know that most pairs of lines with that distinctive character are illusory. If we have the amount of knowledge that we think that we have, then most of our beliefs are dispositional and not occurrent. A person can attribute to me a number of beliefs that I have, and it can be true that I have those beliefs, even though I am not occurrently believing those attributed propositions. Mathematical propositions are a case in point. I believe that $1+1=2$, that $1+2=3$, that $1+3=4$, that $1+4=5$, and so on. If someone attributes to me the belief—indeed, the knowledge—that $1+9999=10000$, it doesn't follow that I then and there possess the occurrent belief that $1+9999=10000$.

Now that we have at least a working distinction between occurrent and dispositional belief, we should ask whether or not Turri and Pritchard's examples provide adequate support for the thesis that S can see that p and nevertheless not have the dispositional belief that p . After all, what seems to be the case is that if S knows that p , then S has the dispositional belief that p , and

need not have the occurrent belief that p . So, while the Argument from Belief can be expressed as:

Knowledge Entails Belief: If S knows that p , S believes that p .

Belief Premise: S can see that p and not believe that p .

Therefore,

~Entailment Thesis: S can see that p and not believe that p .

We can justifiably question the first premise. It is ambiguous between:

Knowledge Entails Occurrent Belief: If S knows that p , S possesses the occurrent belief that p .

And the second principle that:

Knowledge Entails Dispositional Belief: If S knows that p , S has the dispositional belief that p .

Let us consider the first disambiguation of that premise in the following modified Argument from Belief:

Knowledge Entails Occurrent Belief: If S knows that p , S has the occurrent belief that p .

Occurrent Belief Premise: S can see that p and not have the occurrent belief that p .

Therefore,

~Entailment Thesis: S can see that p and not know that p .

This argument is valid, but it is not sound. The first premise is false or else we lack most of our propositional knowledge. The second premise is true, but uninteresting. S can see that p (the lines are of unequal length; a barn is in front

of him; a giant mechanical rabbit is in front of him), and not at that time possess the occurrent belief that p . On the other hand, we can construe the argument so that it reads:

Knowledge Entails Dispositional Belief: If S knows that p , S has the dispositional belief that p .

Dispositional Belief Premise: S can see that p and not have the dispositional belief that p .

Therefore,

~Entailment Thesis: S can see that p and not know that p .

This argument is valid, but it too is not sound. The first premise is true, but the second premise is false. On the one hand, no proponent of the Argument from Belief has provided adequate support for the second premise (recall Turri and Pritchard); their arguments just do not speak against our seeing that p and lacking the dispositional belief that p . On the other hand, it is hard to even make sense of the idea that one can see that p and not have the dispositional belief that p . What would it be for me to see that o is F and at the same time not have a disposition to form the occurrent belief that o is F , given certain triggering conditions? At a minimum, the distinction between objectual perception and propositional perception, seeing an o which is F versus seeing that o is F , is grounded in the fact that the latter, but not the former, requires the possession of the concepts $\langle o \rangle$ and $\langle F \rangle$ and the predicational capacities required to predicate $\langle F \rangle$ of $\langle o \rangle$. We can wonder how it is even possible for one to see that o is F , and so predicate of $\langle o \rangle$ that it is $\langle F \rangle$, while at the same time lacking the dispositional belief that o is F , given certain triggering conditions. It would be not just interesting, but shocking, to learn that while we see that such and such is so, we can do so even if we lack dispositional beliefs about what we see to be

so. Thus far, no such argument has been mounted, nor does it look like an argument could successfully be mounted.

The question I want to ask now is whether or not Turri has shown that the agent lacks occurrent belief and dispositional belief, or just lacks occurrent belief. To see how the distinction between occurrent and dispositional belief can help us to dissipate the Argument from Belief, consider what the proponent of the Argument from Belief might think about propositional perception. They might think that even if seeing that p does not entail knowing that p , seeing that p can still have epistemological implications that mere objectual perception does not have. For example, even if we do not accept the Entailment Thesis, we can consistently accept the following thesis:

Weak Entailment Thesis: If S sees that p , then S is in a position to know that p .

What might lead epistemologists to accept the weak entailment thesis are cases which make it look like S can see that p and fail to believe that p , or believe that $\sim p$. If the distinction between propositional perception and objectual perception is a good one, we might still think that propositional perception has epistemic implications that mere objectual perception does not have. For example, seeing that p puts one in a position to know that p , provided that the agent's actual epistemic basis in support of their belief that p is their having seen that p . But coming to believe that p is up to the agent; they might do what they ought to do and refrain from believing that p because they think that things are not as they see them to be. So, in such cases what seems to be lacking is not belief simpliciter, but occurrent belief.

The proponent of the Entailment Thesis can concede that in such cases that the agent does not have the occurrent belief that p , and instead maintain that she has the dispositional belief that p . But does S have the dispositional belief that the lines are of unequal length? It seems so. After all, the capacities that allow S to see that the lines are of unequal length are the same capacities which would

allow him to get the occurrent belief that the lines are of unequal length. In seeing that the lines are of unequal length, *S* manifests those capacities. It is because *S* knows that lines which have that distinctive character are often lines of equal length, despite contrary appearances, that explains why he refrains from manifesting his capacity to know that the lines are unequal in length by refraining from believing, in this special case, that it is one in which how things look are how things are. We might put the point as follows. The triggering condition for believing that *o* is *F* (the lines are unequal) is present, since it looks to him as if the two lines are of unequal length. But the manifestation condition is not fulfilled because it is 'masked'. What masks the manifestation of the occurrent belief is his knowledge that two lines with that distinctive character are lines which, although they appear to be unequal in length, are nevertheless lines of the same length; this is what explains the ease with which he seems to recognize them as a familiar illusion. So, it is not the case, or at least not obviously the case, that he does not believe that the lines are of unequal length. Instead, it is plausible that he fails to occurrently believe that the lines are of unequal length, although he dispositionally believes that they are. He, therefore, knew right then and there that the lines were of unequal length, due in part to the fortunate circumstances he happened to be in, rather than the unfortunate circumstances he occurrently believed himself to be in.

One potential issue facing this diagnosis of the Argument from Belief is that it seems to have the consequence that *S* knew that the lines were of unequal length because he saw that the lines were of unequal length and had the dispositional belief that the lines were of unequal length. I want to concede that this is so. We can make sense of the idea that *S* knew that the lines are unequal in length. But it does not follow from that that it would be appropriate for *S* to assert that he knows that the lines are of unequal length. What explains the impropriety of the assertion that 'I know that the lines are of unequal length' is his occurrent belief that the lines are not of unequal length. Even if the case is modified to show that *S* withholds occurrent belief that the lines are of equal

length (as the lines appear to be), this is still compatible with *S* nevertheless having the dispositional belief that the lines are of unequal length. Indeed, the possession of the dispositional belief is what seems to explain the ease with which *S* recognizes that, although the lines appear to be unequal, it is just an illusion.

I have argued that three kinds of argument against the Entailment Thesis are unsuccessful. The Argument from Luck is unsound; the Argument from Defeat is either unsound or its soundness depends on the soundness of the Argument from Belief; and the Argument from Belief is at best inconclusive and at worst fallacious. There might be other considerations against the Entailment Thesis, but the three types of epistemological considerations I reviewed here don't seem to render the Entailment Thesis indefensible.

4 Cassam's Response: Part II

In the previous section, I argued that the entailment thesis can be defended from serious objections. Where does that leave us with Stroud's dilemma? If EPR is non-negotiable, and explanations in terms of propositional perception violate EPR, then explanations of how knowledge of the world is possible in terms of propositional perception are not philosophically satisfying.

This leaves us with explanations of how knowledge of the world is possible in terms of "simple perception". First, let's draw the distinction that Dretske (1969) draws between "simple perception" and propositional perception. Simple seeing is the seeing of an object which has properties. It's to see an object *o* which is *F* rather than to see that *o* is *F*. As Dretske puts it, "simple seeing concerns our seeing of objects and things—not facts about these things" (Dretske 2000, 98).

One of the crucial differences between simple seeing and propositional seeing is that the entailment thesis seems to be true of propositional seeing but not simple seeing. So, I can see a barn in front of me without knowing (without

evening believing) that there is a barn in front of me. In general, simple perception lacks the kind of epistemic consequences that propositional perception seems to have.

We might wonder how one can see an *o* which is *F* without seeing (and therefore knowing) that *o* is *F*. Cassam presents three possibilities:

- Conceptual lack (e.g. I see a barn but lack the concept *barn*).
- Failure to believe (e.g. I see a barn but don't believe that it's a barn).
- Environmental luck (e.g. I see a barn but I'm in barn-county).

As I argued in sections 3.2-3.7, environment luck doesn't prevent one from knowing that *p* if they see that *p*, and if one sees that *p*, then it's not clear that they could fail to believe that *p*, let alone fail to grasp that *p*. So, I've argued that those possibilities don't show that one can see that *p* but not know that *p*.

However, contrast the case of propositional perception with simple perception. It's far more intuitive that one can see a barn *and* lack the concept of a barn; that one can see a barn *and* not believe that it's a barn; that one can see a barn *and* be in barn-county.

How does this bear on Stroud's dilemma? We might think that:

- Explanations of how our knowledge of the world is possible in terms of propositional perception are unsatisfying because they violate EPR.
- Explanations of how our knowledge of the world is possible in terms of sensory appearances are unsatisfying because they lead to scepticism.

Yet this this leaves it open that:

- Explanations of how our knowledge of the world is possible in terms of simple perception *neither* violate EPR *nor* lead to

scepticism. So, these explanations are in the market for philosophical satisfaction.

Here's how Cassam expresses the point:

The challenge was to provide a full-blooded explanation of our knowledge of the world. How does the distinction between epistemic and simple seeing help us to meet this challenge? It would be a step in the right direction if it can be explained how any of our knowledge arises out of non-epistemic seeing or, more generally, out of non-epistemic perceiving. Such an explanation would be in line with EPR in the following sense: it would show how (some of) our knowledge of the world could come to be out of something that is not knowledge of the world (Cassam 2009, 583).

An explanation in terms of simple perception can respect EPR if simple perception neither implies nor presuppose knowledge of the world. This can be done if simple perception is a kind of knowledge, just not knowledge of the world, or if it's not a kind of knowledge at all. Cassam takes it that one can give a successful explanation of how our knowledge of the world is possible in terms of simple perception:

How is such a thing possible? How can something non-epistemic intelligibly be regarded as the source of something epistemic? To make this question a bit more concrete consider the barn example once again. How does my knowledge that there is a barn in front of me come to be? To explain how this particular piece of human knowledge comes to be is to explain how I know that there is a barn in front of me. Suppose, then, that my answer to the question 'How do you know there is a barn in front of you?' is 'I can see it'. To see the barn is to perceive an object rather than a fact. On the assumption that this kind of seeing is genuinely non-epistemic, and that my knowledge that there is a barn in front of me can be satisfactorily explained by saying 'I can see it', it looks like we now have a straightforward case of something epistemic coming to be out of something non-epistemic (Cassam 2009, 583).

This is the first step in Cassam's argument for the thesis that a successful explanation of how our knowledge of the world is possible can be given in terms of simple perception without violating EPR. The two central claims in the passage are that:

- A. Explanations in terms of simple perception can be explanatory of how one can know that *p*.

B. Explanations in terms of simple perception don't violate EPR.

The second step in Cassam's argument is to defend these two claims from objections, which I summarise as follows:

- i. Non-entailment objection: explanations in terms of simple perception are unsatisfying because they are not knowledge-entailing.
- ii. Incompleteness objection: if explanations in terms of simple perception are in the market for philosophical satisfaction, then they are elliptical for the corresponding explanations in terms of propositional perception.
- iii. Epistemic objection: explanations in terms of simple perception are epistemic.

I will consider Cassam's responses to these objections in turn.

Against (i), consider the following response: given EPR, one's explanation cannot be knowledge-entailing and be in the market for philosophical satisfaction, if the kind of knowledge it entails is knowledge of the world.

Moreover, as Cassam observes, it's not clear that explanations in terms of knowledge-entailing states are necessarily unsatisfactory, as these are often satisfying answers to questions like "how do you know that p ?" However, as Cassam fails to observe, it's hard to see how they could be satisfying answers to questions like "how is it possible to know that p ?", if we think EPR is true. After all, such explanations will entail knowledge of the kind in question, violating EPR. Still, this need not be the only reason to think that they are unsatisfying in this case. Another reason is that explanations in terms of knowledge-entailing states, such as that one can perceive facts about the external world, don't address the obstacles which make knowledge of the world look impossible. After all, if one's perception of facts about the external world

entails knowledge of those facts, then the obstacles which make the latter look impossible trivially make the former look impossible as well.

Cassam has an explanation of how simple perception of an object *o* which is *F* (e.g. a barn that's in front of one) can amount to knowing that *o* is *F* (knowing that a barn is in front of one) without presupposing knowledge of the world, and so without violating EPR:

If I see a barn in front of me and satisfy all the relevant subjective and objective conditions then I know that there is a barn in front of me. To say that I see that there is a barn in front of me is to imply that the subjective and objective conditions on epistemic seeing have been fulfilled. Seeing that there is a barn in front of me is, as it were, the result of seeing a barn in front of me in the right subjective and objective circumstances. Does it follow that, as (ii) claims, 'I can see it' is an incomplete answer to 'How do you know?', or only complete insofar as it is elliptical for 'I can see that . . .'?' On the first of these issues a natural thought is that what counts as a complete explanation of a person's knowledge that *P* is highly context-dependent. In a context in which there is no question that the relevant subjective and objective conditions are fulfilled 'I can see it' says everything that needs to be said to explain one's knowledge (Cassam 2009, 584-585).

This last point seems right. That *S* saw the barn can be a good explanation of how *S* knows that a barn is there, but to think that this could be a satisfying explanation of how it's possible to know that a barn is there equivocates what's a satisfying answer to an epistemological "how-possible?" question and what's a satisfying answer to an epistemic "how?" question. The former asks how it's possible to know that *p* while the latter asks how one knows or can come to know that *p*.

Now consider Cassam's response to (ii), the incompleteness objection:

On the issue of whether 'I can see it' is elliptical for 'I can see that there is a barn in front of me' what can be said is this: while seeing the barn can add up to seeing that there is a barn there when certain further conditions are fulfilled one would be hard pushed to give a non-circular account of what these further conditions are. In addition, from the fact that seeing the barn can add up to seeing that there is a barn there it does not follow that the latter is the source of one's knowledge. It is the seeing of the barn that tells one that there is a barn there. It wouldn't tell one that there is a barn there if, say, one lacked the concept of a barn but possession of this concept is just an enabling condition for knowing by seeing that there is a barn there (Cassam 2009, 585).

For Cassam, what makes seeing a barn compatible with not knowing that a barn is there is either lacking the concept of a barn, failing to believe that it's a barn, or the presence of environmental-luck. This is what Cassam means in the claim that seeing a barn if the right subjective and objective conditions are fulfilled can add up to knowing that a barn is there. So, Cassam appears committed to the following entailment thesis:

Cassam's Entailment Thesis: If S sees an o which is F and satisfies the subjective and objective conditions, then S knows that o is F .¹⁸⁴

Cassam's entailment thesis presents a problem for the proponent of the view that explanations in terms of simple perception can explain how knowledge of the world is possible in line with EPR. The problem is whether it is in line with EPR.

The problem with Cassam's entailment thesis is that it's hard to see how the explanation in terms of simple perception could explain how one's knowledge of the world is possible unless the right subjective and objective conditions are met. After all, not all simple perceptions amount to knowledge. Yet the problem is that if the subjective and objective conditions are met, then simple perception adds up to propositional perception. So, the following seems to be true: if S see an o which is F in the right subjective and objective circumstances, that S sees that o is F . How is this an explanation of a "case of something epistemic coming to be out of something non-epistemic"? (Cassam 2009, 583).

We can put a sharper point on the problem as follows. It's the fulfilment of the subjective and objective conditions during simple perception which explains the knowing. If that's right, then if explanations in terms of propositional perception are inadmissible because of their violation EPR, something Cassam commits himself to (cf. Cassam 2009, 574), the same is true

¹⁸⁴ Cf. Cassam (2009, 586). Another formulation of this principle might read: if S sees an o which is F and satisfies the subjective and objective conditions *for knowing that* o is F , then S knows that o is F .

of explanations in terms of simple perception. The idea here is that, in cases where the subject fails to fulfil the subjective and objective conditions, the subject's seeing doesn't explain the knowing. In fact, it couldn't.¹⁸⁵

For example, if the subject lacks the concept of a barn, they can't know that it's a barn. If the person doesn't believe that it's a barn, then they're not in a position to know that it's a barn. Moreover, if one thinks that the presence of environmental-luck prevents knowledge, then they couldn't know that it's a barn in that case. So, fulfilment of the subjective and objective conditions is necessary for the simple seeing to be explanatory of the knowing. But if that's right then we are back to the same issue for the proponent of explanations in terms of propositional perception. The main issue here was that such explanations violate EPR. So, contra Cassam, it's not the case that one can explain how knowledge of the world is possible on the basis of simple perception without violating EPR.

5 Conclusion

In this chapter, I examined Cassam's (2009) response to Stroud's dilemma. In section 2, I argued that Cassam's first response is defective. I then suggested that in order to explain how knowledge of the world is possible in line with

¹⁸⁵ A similar issue arises for Brewer's (2011) comment on EPR. He takes it that acquaintance with mind-independent objects and properties in the world puts us in a position to acquire knowledge of them. I think Brewer might be right that his explanation shows how knowledge of such objects could "come to be out of something that does not imply or presuppose any knowledge" of those objects (Brewer 2011, 153). But explaining the epistemological source (or epistemic genealogy) of that type of knowledge isn't obviously the same as explaining how any such knowledge of that kind is possible at all. If certain obstacles makes any and every case of knowledge of the world look impossible, explaining how a person can go from not knowing that *p* to knowing that *p*, where "*p*" is a proposition about the world, doesn't speak to the obstacles that make knowing that *p* in general look impossible. Moreover, it would violate EPR. It violates EPR *unless* what one takes themselves to be talking about, in providing their explanation, is something that neither implies nor presupposes any knowledge of the world, including *how* human beings perceive, *what* human beings perceive, and so forth. It's hard to see how the epistemologist's explanation in terms of simple perception, where this is an immediate acquaintance with external, mind-independent objects and properties, can do this without violating EPR.

EPR, one should appeal to sense perception. Cassam and Stroud maintain that appealing to sense perception in the form of perceptual appearances or seemings leads to scepticism. This left us with propositional perceptual states and simple perceptual states. The former is knowledge-entailing (§3), and the latter is knowledge-entailing when we appeal to them to explain knowledge (§4). The problem, then, is that it looks as if it's not possible to explain how knowledge of the world is possible without violating EPR but EPR seems to be necessary.

Chapter 7

Meta-epistemological Scepticism: Afterword

1 Introduction

The final chapter is an overview of the main results of each chapter and so the results of the thesis. In section 2, I review the main results of each chapter. In section 3, I turn to some general objections that might arise for these results, and I respond to these objections. In section 4, I conclude the thesis.

2 Review

This thesis focused on two important arguments from Stroud (1984, 2000) for meta-epistemological scepticism.¹⁸⁶ Chapter 1 set out to explain what meta-epistemological scepticism is, to distinguish it from other sceptical theses, and to defend the thesis that meta-epistemological scepticism does not entail first-order scepticism.¹⁸⁷ I then introduced the reader to two arguments in favour of meta-epistemological scepticism: Stroud's puzzle and Stroud's dilemma. I showed how valid arguments could be formulated, using Stroud's puzzle and

¹⁸⁶ Additional papers include Stroud (2004) and (2009).

¹⁸⁷ As I highlighted in Chapter 1, this helps make meta-epistemological scepticism look plausible. After all, first-order scepticism is not plausible. If meta-epistemological scepticism entailed first-order scepticism, this would count against it. But it's not clear that knowledge of the world should be amenable to philosophical explanation. We might think that while we have knowledge of the world, it might be that we cannot give a satisfying philosophical explanation of how that knowledge is possible. Compare this with Williamson's (2000, 2013) claim that, although we have knowledge, an analysis of knowledge is not possible. Cf. Zagzebski (1994). Other pessimistic claims about the limits of philosophical explanation can be found in Grundmann (2013) and S. Nichols, S. Stich & J. Weinberg (2003).

Stroud's dilemma, in favour of meta-epistemological scepticism without presupposing or requiring the truth of first-order scepticism. I argued that this further motivates the thesis that meta-epistemological sceptics are not committed to first-order scepticism.

Chapters 2 and 3 focused on Stroud's puzzle. Chapter 2 focused on the second premise of Stroud's puzzle. I focused on the second premise of Stroud's puzzle first because an assessment of this premise was tantamount to the argument of Chapter 3. The second premise of Stroud's puzzle is the premise that if Descartes's Condition is false, then an ordinary epistemic principle is false (or: "a 'fact' of our ordinary conception of knowledge" is false; or "platitudes we would all accept" are false), and this is philosophically unsatisfying (Stroud 1984, 31, 82). I argue that none of Stroud's anti-revisionist arguments provide adequate support for this premise. I then turned to additional anti-revisionist arguments, and argued that these are all defective as well. Finally, I consider the possibility that there is some other property of Descartes's Condition, *F*, the possession of which would render its falsity philosophically unsatisfying. But I argued in each case that it was either implausible that it had the relevant property, or that it was implausible that, if it had the property, then its falsity is philosophically unsatisfying.

Chapter 3 focused on the first premise of Stroud's puzzle. This is the premise that, if Descartes's Condition is true, then scepticism is true. But scepticism is not philosophically satisfying. I accepted that scepticism is not a satisfying philosophical explanation of how knowledge of the world is possible (cf. Chapter 1, §2 and §5). I argued that the first argument, the regress argument, failed to establish that if Descartes's Condition is true, then scepticism is true. I also argued that the second argument, the no test argument, failed to establish this. I then considered a different version of the no test argument, which took Descartes's Condition to express a necessary precondition on knowledge of the world, and argued that even if the no test argument is valid given this version of Descartes's Condition, we should reject

this version of Descartes's Condition. Finally, I considered Stroud's (2009) no test argument, and argued that the no test argument is defective. In turn, Stroud's puzzle fails to provide adequate support for meta-epistemological scepticism.

Chapters 4, 5, and 6 focused on responses to Stroud's dilemma. In Chapter 1 (§6), I highlighted how Stroud's dilemma, if sound, provides adequate support for meta-epistemological scepticism. Chapter 4 focused on Williams's (1996) response to Stroud's dilemma. In particular, it focused on what I called Williams's "master argument" against Stroud's dilemma. This argument targeted the legitimacy thesis, the thesis that the traditional epistemological project, the project of explaining how knowledge of the world is possible at all, is legitimate. The first premise of Williams's master argument is that the traditional epistemological project is legitimate only if epistemological realism is true, while the second premise is that epistemological realism is false. I argued that Williams failed to provide adequate support for both premises of his master argument, so that his master argument fails to show that the legitimacy thesis is false.

Chapter 5 focused on externalist responses to Stroud's dilemma. In particular, I focused on the response from Sosa (1994). The externalist targets the first horn of Stroud's dilemma, that if EPR is true, then the resulting explanation of how our knowledge of the world is possible is a negative, sceptical explanation. Stroud's argument for this horn of the dilemma rested on the premise that one must know that the linking-principle between the explanans and the explanandum are true. That is, one would need to know that, in general, sensory experiences as of p are reliable indicators that p . I argued that Stroud is right about this requirement. Sosa (1994) agrees, but unlike Stroud (2000, 2004), Sosa argues that this requirement can be met along externalist lines. I argued that this requirement cannot be met unless one violates EPR. This leaves the second horn of Stroud's dilemma. This is the horn that if EPR is false, then we can understand how our knowledge of the world is

possible “only by assuming that we have got some knowledge in the domain in question. And that is not philosophically satisfying” (Stroud 2000, 121). The reason it is not satisfying is that we will “have lost the prospect of explaining and therefore understanding all of our knowledge with complete generality” (Stroud *ibid*).

Chapter 6 focused on Cassam’s (2009) response to Stroud’s dilemma. Cassam targeted the first horn of Stroud’s dilemma in two different ways. First, the argument from generality fails to provide adequate support for EPR. Second, that EPR can be fulfilled without the resulting explanation being a sceptical explanation. I argued that both of Cassam’s arguments are defective.

3 Objections and Replies

3.1 Objection 1

In Chapter 1, I defended meta-epistemological scepticism from three prima facie objections. The main objection here was that one could not be a meta-epistemological sceptic without being a first-order sceptic. I argued that the argument in favour of that thesis failed to provide adequate support for it. In Chapters 2 and 3, however, I argued that Stroud’s puzzle failed to provide adequate support for meta-epistemological scepticism. So, one might raise again the following objection: how can one be a meta-epistemological sceptic without commitment to first-order scepticism?

The thought here is this. If, as I argued in Chapters 2 and 3, Stroud’s puzzle fails to provide adequate support for meta-epistemological scepticism, but, as I argued in Chapter 1, first-order scepticism provides adequate support for meta-epistemological scepticism, then doesn’t that make it look like one cannot be a meta-epistemological sceptic without being a first order sceptic, contra the author’s thesis in Chapter 1? In the next two paragraphs, I will respond to this objection.

First, we should distinguish between being committed to first-order scepticism because of one's commitment to meta-epistemological scepticism, and one's only reason for being a meta-epistemological sceptic is one's commitment to first-order scepticism. In Chapter I, I argued that Stroud's dilemma is, prima facie at least, a valid argument which, if sound, establishes meta-epistemological scepticism. What's important here is that Stroud's dilemma would continue to provide adequate support for meta-epistemological even if first-order scepticism is false. The reason this is so is that none of the premises of Stroud's dilemma require the truth of first-order scepticism. Instead, one of the premises of Stroud's dilemma requires the truth of the conditional thesis that if one's philosophical explanation respects EPR, then it's a sceptical explanation. So, if the arguments of Chapters 1, 4, 5, and 6 are successful, then it's not the case the only reason to be a meta-epistemological sceptic is first-order scepticism.

Second, let's suppose that I'm wrong about Stroud's dilemma. It's not the case that Stroud's dilemma has withstood the objections of Chapters 4, 5, and 6. If that's right, then it looks as if we lack non-sceptical reasons to be meta-epistemological sceptics, contra what I argued above. However, even if it were true that we lacked non-sceptical reasons for meta-epistemological scepticism, this doesn't entail that meta-epistemological scepticism implies first-order scepticism. Sure, it might be hard to see why one would be a meta-epistemological sceptic without being a first-order sceptic, if both Stroud's puzzle and Stroud's dilemma are defective, but that the non-sceptical reasons in favour of meta-epistemological scepticism have been removed does not entail that meta-epistemological scepticism implies first-order scepticism.¹⁸⁸

¹⁸⁸ A related point is here is that even though Stroud's puzzle fails to provide adequate support for meta-epistemological scepticism, separating Stroud's puzzle from Stroud's dilemma helps to bring into clearer focus a stronger argument for meta-epistemological scepticism. If the results of this thesis are correct, Stroud's dilemma provides the more compelling case for meta-epistemological scepticism.

3.2 Objection 2

In Chapters 4, 5, and 6, I defended Stroud's dilemma against serious objections that one or more of the premises of Stroud's dilemma is false or otherwise not adequately supported. But showing that those objections are defective doesn't amount to showing that Stroud's dilemma provides adequate support for meta-epistemological scepticism. So, it's not obvious that I have defended meta-epistemological scepticism.

I've argued that Stroud's dilemma in favour of meta-epistemological scepticism survives the objections that it fails to show this. On the one hand, this is a negative claim. It's the claim that some philosophers maintain that Stroud's dilemma fails to provide adequate support for meta-epistemological scepticism, but their arguments are defective. This leaves it open that Stroud's argument is also defective.

On the other hand, the claim might be more positive than initial appearances suggest. I take it that Stroud's dilemma provides prima facie justification to believe that meta-epistemological scepticism is true. It is "prima facie" because others considerations, such as the considerations I evaluated in Chapters 4, 5, and 6, can defeat this justification. But as I argued in those chapters, the objections all failed to show Stroud's dilemma is unsound or that the arguments in favour of the premises of Stroud's puzzle are bad arguments.

So, what's the source of this prima facie justification? It's not that the premises of Stroud's dilemma are intuitive. It's hard to see how highly theoretical propositions, such as the kind that Stroud's dilemma expresses, could be intuitively plausible. Instead, the premises are attractive. They are attractive because Stroud's sub-arguments in favour of the premises of Stroud's dilemma make them look attractive, and the objections from Chapters 4, 5, and 6 failed to show that those arguments are fallacious or otherwise unsound.

We can put this point in slightly different terms. I haven't provided a "knock-down" argument for meta-epistemological scepticism. Stroud's dilemma is no such argument. What I have done in terms of supporting the

premises of Stroud's dilemma is strengthen them: I have argued that the objections to the premises of Stroud's dilemma or the sub-arguments for those premises fail.

One might think that providing a defence of Stroud's dilemma would count in favour of meta-epistemological scepticism if meta-epistemological somehow occupied the default position. However, what in part makes meta-epistemological scepticism an interesting thesis is that it's a certain kind of philosophical scepticism: it's a scepticism about satisfying philosophical explanations of how knowledge of the world is possible. This might not be the default position, since perhaps we expect there to be a satisfying philosophical explanation of how knowledge of the world is possible. But on the other hand meta-epistemological scepticism has great explanatory power: it would explain why there is such radical disagreement over the satisfactoriness of the proposed explanations in the literature.

Furthermore, the dialectic here is not to convince one that meta-epistemological scepticism is true. Instead, the dialectic is to show that there is not just an obstacle to the possibility of knowledge of the world, but of a satisfying philosophical explanation of how knowledge of the world is possible. Stroud's dilemma poses such an obstacle. The objections that I evaluated in Chapters 4, 5, and 6 tried to show how either the obstacle could be dissipated or overcome. What I argued is that none of them are successful. This implies that the obstacle remains intact, and therefore it ought to continue to look as if a satisfying philosophical explanation of how our knowledge of the world is possible is not possible. In short, it ought to look as if meta-epistemological scepticism is true, even if its truth has not been established.¹⁸⁹

¹⁸⁹ This is consistent with agnosticism: although meta-epistemological scepticism looks true, this does not entail that it is true, or that one ought to accept that it is true—one might remain agnostic. There is a sense in which this attitude towards the main claim of the thesis comes close to capturing "Pyrrhonian" intuitions. For example, compare with Pritchard's (2000) take on Pyrrhonian scepticism: "Consider the Pyrrhonian focus upon claims to know. Right at the beginning of *Outlines of Pyrrhonism*, Sextus Empiricus emphasises a crucial difference between, on the one hand, the Dogmatists—who "have claimed to have

4 Conclusion

In this thesis, I have explained what meta-epistemological is, distinguished it from other forms of scepticism, and explained how one could be a meta-epistemological sceptic without commitment to first-order scepticism. I then argued that while Stroud's puzzle fails to provide adequate support for meta-epistemological scepticism, the responses to Stroud's dilemma failed to show that Stroud's dilemma is defective. So, Stroud's dilemma has withstood serious criticism. Meta-epistemological scepticism is therefore a live option. Further criticism is required to show that it lacks adequate support or is otherwise false.

discovered the truth" (SE 3)—and the Academics—who “have asserted that it cannot be apprehended” (SE 3)—and, on the other, the Sceptics who simply “go on inquiring” (SE 3). A contrast is thus established between those who claim knowledge or the lack of it, and the Pyrrhonian strategy of censuring each and every claim to know” (Pritchard 2000, 196). So, while we might be in a position to claim knowledge of the world, we're not in a position to claim that we have a satisfying philosophical explanation of how that knowledge is possible. For a general review of Pyrrhonism, see Machuca (2011). One might read Stroud's (2011) monograph as a version of Pyrrhonism about metaphysical explanation. Stroud tells us that his goal is to give “attention to the dissatisfaction that metaphysical reflection of the kind we have been concerned with appears to leave us with. I think its source lies in our unavoidable immersion in whatever conception of the world we seek metaphysical understanding of”, such that we cannot “achieve the kind of metaphysical understanding we seek” (Stroud 2011, 145). I draw parallels between his meta-epistemological views and his meta-metaphysical views in Ranalli (2013).

Bibliography

Alston, William P. (1993). *The Reliability of Sense Perception*. Cornell University Press.

Armstrong, D. M. (1973). *Belief, Truth and Knowledge*. London, Cambridge University Press.

Armstrong, D. M. (1969). Does Knowledge Entail Belief? *Proceedings of the Aristotelian Society*, 70, 21–36.

Audi, Robert (1989). *Practical Reasoning*. Routledge.

Audi, Robert. (1994). Dispositional Beliefs and Dispositions to Believe. *Nous*, 28, 419–434.

Austin, J. L. (1946). Other Minds. *Aristotelian Society Supplementary Volume* 20:148-87.

Ballantyne, Nathan (2012). Acquaintance and assurance. *Philosophical Studies* 161 (3):421-431.

Benton, Matthew A. (2014). Knowledge Norms. *Internet Encyclopaedia of Philosophy*.

Bergmann, Michael (2013). Externalist Justification and the Role of Seemings. *Philosophical Studies* 166 (1):163-184.

Bergmann, Michael (2008). Externalist Responses to Skepticism. In John Greco (ed.), *The Oxford Handbook of Skepticism*. Oxford University Press.

- Bergmann, Michael (2006). *Justification Without Awareness: A Defense of Epistemic Externalism*. Oxford University Press.
- Bergmann, Michael (2004). Epistemic Circularity: Malignant and Benign. *Philosophy and Phenomenological Research* 69 (3):709–727.
- Bergmann, Michael (2000). Externalism and Skepticism. *Philosophical Review* 109 (2):159-194.
- Bergmann, Michael & Coppenger, Brett (Eds.) (Forthcoming). *Traditional Epistemic Internalism*. Oxford University Press.
- Bird, Alexander and Tobin, Emma, “Natural Kinds”, The Stanford Encyclopaedia of Philosophy (Spring 2015 Edition), Edward N. Zalta (ed.), forthcoming URL = <<http://plato.stanford.edu/archives/spr2015/entries/natural-kinds/>>.
- Black, C. (1971). Knowledge without Belief. *Analysis*, 31, 152–158
- Bonjour, L. (1985). *The Structure of Empirical Knowledge*, Cambridge, MA: Harvard University Press.
- Brennan, Jason (2010). Scepticism about Philosophy. *Ratio* 23 (1):1-16.
- Brewer, Bill (2011). *Perception and Its Objects*. Oxford University Press.
- Brewer, Bill (1996). Internalism and Perceptual Knowledge. *European Journal of Philosophy* 4 (3):259-275.
- Bridges, Jason (2006). Davidson’s Transcendental Externalism. *Philosophy and Phenomenological Research* 73 (2):290-315.

Brown, Jessica & Gerken, Mikkell (Eds.) (2012). *Knowledge Ascriptions*. Oxford University Press.

Brueckner, Anthony (2008). Experiential Justification. In Quentin Smith (Ed.), *Epistemology: New Essays*. Oxford University Press.

Brueckner, Anthony (2005). Fallibilism, Underdetermination, and Skepticism. *Philosophy and Phenomenological Research* 71 (2):384–391.

Brueckner, Anthony (1996). Modest Transcendental Arguments. *Philosophical Perspectives* 10 (Metaphysics):265-280.

Brueckner, Anthony (1994). The Structure of the Skeptical Argument. *Philosophy and Phenomenological Research* 54 (4):827-835.

Brueckner, Anthony (1999). Transcendental Arguments from Content Externalism. In Robert Stern (Ed.), *Transcendental Arguments: Problems And Prospects*. Oxford: Clarendon Press.

Buchanan, Reid (2002). Natural Doubts: Williams's Diagnosis of Scepticism. *Synthese* 131 (1):57 - 80.

Byrne, Alex (2004). How Hard are the Sceptical Paradoxes? *Noûs* 38 (2):299–325.

Campbell, C. A. (1967). Towards a Definition of Belief. *The Philosophical Quarterly*, 17, 204–220

Cassam, Quassim (2009). Knowing and Seeing: Responding to Stroud's Dilemma. *European Journal of Philosophy* 17 (4):571-589.

Cassam, Quassim (2007). Ways of Knowing. *Proceedings of the Aristotelian Society* 107 (1pt3):339 - 358.

Cassam, Quassim. (2007). *The Possibility of Knowledge*. Oxford: Oxford University Press.

Cassam, Quassim. (2008). Knowledge, Perception, and Analysis. *South African Journal of Philosophy*, 27(3), 36–48.

Cassam, Quassim (2003). Self-Directed Transcendental Arguments. In Robert Stern (Ed.), *Transcendental Arguments: Problems And Prospects*. Clarendon Press.

Chomsky, Noam (1995). *The Minimalist Program*. MIT Press.

Chomsky, Noam (1986). *Knowledge of Language*. Prager.

Chomsky, Noam (1957). *Syntactic Structures*. Mouton.

Cohen, Stewart (2002). Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research* 65 (2):309-329.

Cohen, Stewart (1998). Two Kinds of Skeptical Argument. *Philosophy and Phenomenological Research* 58 (1):143 - 159.

Coliva, Annalisa (2010). *Moore and Wittgenstein: Scepticism, Certainty, and Common Sense*. Palgrave Macmillan

Conee, Earl (2004). Externalism, Internalism, and Skepticism. *Philosophical Issues* 14 (1):78–90.

Craig, Edward (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford University Press.

Dancy, Jonathan, Sosa, Ernest & Steup, Matthias (Eds.) (2010). *A Companion to Epistemology*, Second Edition. Blackwell.

- Dancy, Jonathan (2004). *Ethics without Principles*, Oxford: Clarendon Press.
- Dancy, Jonathan. (1985). *Introduction to Contemporary Epistemology*. Oxford: Wiley.
- Davidson, Donald (2006). *The Essential Davidson*. Oxford University Press.
- DeRose, Keith (2009). *The Case for Contextualism*. Oxford University Press.
- DeRose, Keith (1995). Solving the Skeptical Problem. *Philosophical Review* 104 (1):1-52.
- DeRose, Keith & Warfield, Ted A. (Eds.) (1999). *Skepticism: A Contemporary Reader*. Oxford University Press.
- Descartes, René (1996). *Meditations on First Philosophy: With Selections From the Objections and Replies*. Cambridge University Press.
- Devitt, Michael (2006). *Ignorance of Language*. Oxford: Clarendon Press.
- Dretske, F. (2010). Self-Profile. In J. Dancy, E. Sosa, & M. Steup (Eds.), *A Companion to Epistemology* (2nd ed.). Wiley-Blackwell.
- Dretske, Fred (2003). Skepticism: What Perception Teaches. In *The Sceptics: Contemporary Essays*. Aldershot: Ashgate Publishing.
- Dretske, Fred (2000). *Perception, Knowledge and Belief: Selected Essays*. Cambridge University Press.
- Dretske, Fred (1993). Misrepresentation. In Alvin Goldman (Ed.), *Readings in Philosophy and Cognitive Science*. Cambridge: MIT Press. 297.
- Dretske, Fred (1973). Perception and Other Minds. *Noûs* 7 (March):34-44.

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

Dretske, Fred (1970). Epistemic Operators. *Journal of Philosophy* 67 (24):1007-1023.

Ekstrom, Laura W. (2003). Free will, chance, and mystery. *Philosophical Studies* 22 (2):153-80.

Engel, Jr., M. (1992). Is epistemic luck compatible with knowledge? *Southern Journal of Philosophy*,30, 59–75.

Engel, Jr., M. (2011). Epistemic luck. *Internet Encyclopaedia of Philosophy*. <http://www.iep.utm.edu/epi-luck/>

Feldman, Richard, “Naturalized Epistemology”, *The Stanford Encyclopaedia of Philosophy* (Summer 2012 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/sum2012/entries/epistemology-naturalized/>>.

Feldman, Richard & Conee, Earl (1985). Evidentialism. *Philosophical Studies* 48 (1):15 - 34.

Fogelin, Robert J. (1999). The Sceptic’s Burden. *International Journal of Philosophical Studies* 7 (2):159 – 172.

Flanagan, Owen J. (1991). Consciousness. *The Science of the Mind*. MIT Press.

French, Craig (2013). Perceptual Experience and Seeing That P. *Synthese* 190 (10):1735-1751.

French, Craig (2012). Does Propositional Seeing Entail Propositional Knowledge? *Theoria* 78 (2):115-127.

Fumerton, Richard (2014). The Challenge of Refuting Skepticism, *Contemporary Debates in Epistemology* (2nd ed.), (eds.) M. Steup, E. Sosa, & J. Turri, 120-32, Oxford: Wiley-Blackwell.

Fumerton, Richard (2006). Epistemic Internalism, Philosophical Assurance and the Skeptical Predicament, *Knowledge and Reality: Essays in Honor of Alvin Plantinga*, (eds.)

Fumerton, Richard (1996). Metaepistemology and Skepticism, Boston, NJ: Rowman & Littlefield.

Fumerton, Richard (1994). Skepticism and Naturalistic Epistemology. *Midwest Studies In Philosophy* 19 (1):321-340.

Gettier, Edmund. (1963). Is Justified True Belief Knowledge? *Analysis*, 23(6), 121–123

Gisborne, Nicholas. (2010). *The Event Structure of Perception Verbs*. Oxford: Oxford University Press

Goldman, Alvin (2008). Reliabilism. *Stanford Encyclopaedia of Philosophy*.

Goldman, Alvin (1976). Discrimination and Perceptual Knowledge. *Journal of Philosophy* 73 (November):771-791.

Goldman, Alvin (2008). Immediate Justification and Process Reliabilism. In Quentin Smith (Ed.), *Epistemology: New Essays*. Oxford University Press.

Goldman, Alvin (2002). *Pathways to Knowledge: Private and Public*. Oxford University Press.

Goldman, Alvin (1999). Internalism Exposed. *Journal of Philosophy* 96 (6):271-293.

- Goldman, Alvin (1993). Epistemic Folkways and Scientific Epistemology. *Philosophical Issues* 3:271-285.
- Greco, John (2008). Skepticism about the External World. *The Oxford Handbook of Skepticism*. Oxford University Press.
- Greco, John (2000). *Putting Sceptics in their Place: The Nature of Skeptical Arguments and their Role in Philosophical Inquiry*. Cambridge University Press.
- Greco, John (1993). How to Beat a Sceptic without Begging the Question. *Ratio* 6 (1):1-15.
- Grundmann, Thomas (2013). Doubts about Philosophy? The Alleged Challenge from Disagreement. In Tim Henning & David Schweikard (Eds.), *Knowledge, Virtue, and Action. Essays: On Putting Epistemic Virtues to Work*. Routledge. 72-98.
- Grundmann, Thomas (2004). Inferential Contextualism, Epistemological Realism and Scepticism: Comments on Williams. *Erkenntnis* 61 (2-3):345 - 352.
- Gupta, A. (2006). *Empiricism and Experience*. Harvard University Press.
- Guyer, Paul (2008). Kant's Transcendental Idealism and The Limits of Knowledge: Kant's Alternative to Locke's Physiology. In Daniel Garber & Béatrice Longuenesse (Eds.), *Kant and the Early Moderns*. Princeton University Press.
- Haddock, Adrian & Macpherson, Fiona (Eds.) (2008). *Disjunctivism: Perception, Action, Knowledge*. Oxford University Press.

- Hazlett, Allan (2012). Factive Presupposition and the Truth Condition on Knowledge. *Acta Analytica* 27 (4):461-478.
- Hazlett, Allan (2010). The Myth of Factive Verbs. *Philosophy and Phenomenological Research* 80 (3):497 - 522.
- Hazlett, Allan (2006). How to Defeat Belief in the External World. *Pacific Philosophical Quarterly* 87 (2):198–212.
- Heil, John (1986). Review “The Significance of Philosophical Scepticism” by Barry Stroud (Oxford University Press). *Philosophy and Phenomenological Research*. Vol. 47, No. 2, pp. 331-336
- Hiller, A., & Neta, Ram. (2007). Safety and Epistemic Luck. *Synthese*, 158, 303–313
- Hill, Christopher S. (1996). Process Reliabilism and Cartesian Scepticism. *Philosophy and Phenomenological Research* 56 (3):567-581.
- Hirsch, Eli (2005). Physical-object Ontology, Verbal Disputes, and Common sense. *Philosophy and Phenomenological Research* 70 (1):67–97.
- Huemer, Michael (2013). Phenomenal Conservatism. Internet. *Encyclopaedia of Philosophy*.
- Huemer, Michael (2007). Compassionate Phenomenal Conservatism. *Philosophy and Phenomenological Research* 74 (1):30–55.
- Huemer, Michael (2000). Direct Realism and the Brain-In-A-Vat Argument. *Philosophy and Phenomenological Research* 61 (2):397-413.
- Jackson, Frank (1998). *From Metaphysics to Ethics*. Oxford University Press.

- Johnsen, Bredo C. (2009). The Argument for Radical Skepticism Concerning the External World. *Journal of Philosophy* 106 (12):679-693.
- Johnston, Mark (1992). Reasons and Reductionism. *Philosophical Review* 3 (3):589-618.
- Johnston, Mark (1992). Constitution is not identity. *Mind* 101 (401):89-106.
- Kallestrup, Jesper (2005). Contextualism between Scepticism and Common-Sense. *Grazer Philosophische Studien* 69 (1):247-266.
- Kallestrup, Jesper & Pritchard, Duncan (2014). Virtue Epistemology and Epistemic Twin Earth. *European Journal of Philosophy* 22 (3):335-357.
- Kant, Immanuel (1997). *Kant: Critique of Practical Reason*. Cambridge University Press.
- Kaplan, Mark (2008). Austin's Way with Skepticism. In John Greco (Ed.), *The Oxford Handbook of Skepticism*. Oxford University Press.
- Kaplan, Mark (2000). To What must an Epistemology be True? *Philosophy and Phenomenological Research* 61 (2):279-304.
- Kornblith, Hilary (2002). *Knowledge and its Place in Nature*, Oxford: Clarendon Press.
- Koethe, John (2005). *Scepticism, Knowledge, and Forms of Reasoning*. Cornell University Press.
- Lackey, Jennifer. (2008). What Luck is Not. *Australasian Journal of Philosophy*, 86(2), 255–267.

- Leite, Adam (2010). How to Take Skepticism Seriously. *Philosophical Studies* 148 (1):39 - 60.
- Leite, Adam (2006). Epistemic Gradualism and Ordinary Epistemic Practice: Response to Hetherington. *Philosophia* 34 (3):311-324.
- Leite, Adam (2005). Epistemological Externalism and the Project of Traditional Epistemology. *Philosophy and Phenomenological Research* 70 (3):505–533.
- Leite, Adam (2004). Skepticism, Sensitivity, and Closure, or why the Closure Principle is Irrelevant to External World Skepticism. *Croatian Journal of Philosophy* 12:335-350.
- Lewis, David (1996). Elusive Knowledge. *Australasian Journal of Philosophy* 74 (4):549 – 567.
- Lewis, David K. (1986). *On The Plurality of Worlds*. Blackwell Publishers.
- Lewis, David (1983). *Philosophical Papers Vol. I*. Oxford University Press.
- Lycan, William G. (2001). Moore Against the New Sceptics. *Philosophical Studies* 103 (1):35 - 53.
- Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. Penguin.
- Machuca, Diego E. & Reed, Baron (Eds.) (Forthcoming). *Skepticism: From Antiquity to The Present*. Bloomsbury.
- Machuca, Diego E. (2011). Ancient Skepticism: Pyrrhonism. *Philosophy Compass* 6 (4):246-258.

Madison, B. J. C. (2011). Combating Anti Anti-Luck Epistemology. *Australasian Journal Of Philosophy* 89 (1):47-58.

Madison, B. J. C. (2010). Epistemic Internalism. *Philosophy Compass* 5 (10):840-853.

McCain, Kevin (2013). Two Skeptical Arguments or Only One? *Philosophical Studies* 164 (2):289-300.

McCain, Kevin (2008). The Virtues of Epistemic Conservatism. *Synthese* 164 (2):185-200.

McDowell, John (2008). The Disjunctive Conception of Experience as Material for a Transcendental Argument. In Fiona Macpherson & Adrian Haddock (Eds.), *Disjunctivism: Perception, Action, Knowledge*. Oxford University Press. 19-33.

McDowell, John. (1994). *Mind and World*. Cambridge, MA: Harvard University Press.

McDowell, John. (2002). Responses. In N. H. Smith (Ed.), *Reading McDowell on Mind and World* (pp. 269–305). London: Routledge.

McGinn, Colin (1993). *Problems in Philosophy*. Blackwell.

Merricks, Trenton (2003). *Objects and Persons*. Oxford University Press.

Millar, Allan. (2008). Perceptual-recognitional Abilities and Perceptual Knowledge. In A. Haddock & F. Macpherson (Eds.), *Disjunctivism: Perception, action, knowledge*. Oxford: Oxford University Press.

Miller, Alexander, "Realism", *The Stanford Encyclopaedia of Philosophy* (Winter 2014 Edition), Edward N. Zalta (ed.)

Moore, G. E. (1939). Proof of an External World, *Proceedings of the British Academy* 25, 273-300.

Nagel, Jennifer (2010). Epistemic Anxiety and Adaptive Invariantism. *Philosophical Perspectives* 24 (1):407-435.

Nagel, Thomas (1986). *The View From Nowhere*. Oxford University Press.

Nagel, Thomas (1974). What is it like to be a bat? *Philosophical Review* 83 (October):435-50

Neta, Ram (Ed.) (2014). *Current Controversies in Epistemology*. Routledge.

Neta, Ram (2010). Liberalism and Conservatism in the Epistemology of Perceptual Belief. *Australasian Journal of Philosophy* 88 (4):685-705.

Nichols, Shaun, Stich, Stephen & Weinberg, Jonathan M. (2003). *Metaskepticism: Meditations in Ethnoepistemology*. In S. Luper (Ed.), *The Sceptics*. Ashgate.

Peacocke, Christopher. (2005). Another “P”: Representing conscious states, perception and others. In J. Bermúdez (Ed.), *Thought, reference, and experience: Themes from the philosophy of Gareth Evans*. Oxford: Oxford University Press.

Poston, Ted (2010). Similarity and Acquaintance: A Dilemma. *Philosophical Studies* 147 (3):369 - 378.

Prichard, H. A. (1950). *Knowledge and Perception*. Oxford: Oxford University Press.

Pritchard, Duncan & Ranalli, Chris (Forthcoming). On Meta-epistemological Scepticism. *Traditional Epistemic Internalism*, (eds.) M. Bergmann & B. Coppenger, (Oxford University Press.

Pritchard, Duncan & Ranalli, Chris (Forthcoming). Putnam on Brains in Vats and Radical Skepticism. In Sanford Goldberg (Ed.), *Putnam on Brains in Vats*. Cambridge University Press.

Pritchard, Duncan & Ranalli, Chris (Forthcoming). Disjunctivism and Scepticism. In Baron Reed & Diego E. Machuca (Eds.), *Skepticism: From Antiquity to The Present*. Bloomsbury.

Pritchard, Duncan (Forthcoming). *Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing*. Princeton University Press.

Pritchard, Duncan & Ranalli, Chris (2013). Rorty, Williams, and Davidson: Skepticism and Metaepistemology. *Humanities* 2 (3):351-368.

Pritchard, Duncan. (2013). There Cannot be Lucky Knowledge. In M. Steup, J. Turri & E. Sosa (Eds.). *Contemporary Debates in Epistemology* (2nd ed.). Blackwell.

Pritchard, Duncan (2013). Davidson on Radical Skepticism, Blackwell Companion to Donald Davidson, (eds.) E. LePore & K. Ludwig. Blackwell.

Pritchard, Duncan (2012). Anti-Luck Virtue Epistemology. *Journal of Philosophy* 109 (3):247-279.

Pritchard, Duncan (2012). *Epistemological Disjunctivism*. Oxford University Press.

Pritchard, Duncan (2011). Wittgenstein on Scepticism. In Marie McGinn & Oskari Kuusela (Eds.), *The Oxford Handbook Of Wittgenstein*. Oxford University Press.

Pritchard, Duncan. (2011). Epistemological Disjunctivism and the Basis Problem. *Philosophical Issues (supplement to Nous)*, 21, 434–455.

Pritchard, Duncan (2010). Doubt Undogmatized: Pyrrhonian Scepticism, Epistemological Externalism and the ‘Metaepistemological’ Challenge. *Principia* 4 (2):187-214.

Pritchard, Duncan, Millar, Allan & Haddock, A (2010). *The Nature and Value of Knowledge: Three Investigations*. Oxford: Oxford University Press.

Pritchard, Duncan (2009). Scepticism and the Possibility of Knowledge. *Analysis* 69 (2):317-325.

Pritchard, Duncan (2008). Sensitivity, Safety, and Anti-Luck Epistemology. In John Greco (Ed.), *The Oxford Handbook of Skepticism*. Oxford University Press.

Pritchard, Duncan (2007). Anti-Luck Epistemology. *Synthese* 158 (3):277-297.

Pritchard, Duncan (2007). How to be a Neo-Moorean. In Sanford Goldberg (Ed.), *Internalism and Externalism in Semantics And Epistemology*. Oxford University Press. 68-99.

Pritchard, Duncan (2006). McDowellian Neo-Mooreanism. In Fiona Macpherson & Adrian Haddock (Eds.), *Disjunctivism: Perception, Action, Knowledge*. Oxford University Press. 283-310.

Pritchard, Duncan (2005). *Epistemic Luck*. Clarendon Press.

Pritchard, Duncan (2005). The Structure of Sceptical Arguments. *Philosophical Quarterly* 55 (218):37-52.

Pritchard, Duncan (2005). Neo-Mooreanism, Contextualism, and the Evidential Basis of Scepticism. *Acta Analytica* 20 (2):3-25.

Pritchard, Duncan (2005). Neo-Mooreanism versus Contextualism. *Grazer Philosophische Studien* 69 (1):1-24.

Pritchard, Duncan (2005). Scepticism, Epistemic Luck, and Epistemic Angst. *Australasian Journal Of Philosophy* 83 (2):185-205.

Pritchard, Duncan. (2004). Epistemic Luck. *Journal of Philosophical Research*, 29, 193–222.

Pritchard, Duncan (2002). Recent Work on Radical Skepticism. *American Philosophical Quarterly* 39 (3):215-257.

Pryor, James (2013). There is Immediate Justification. In Steup, Matthias, Sosa, Ernest, and Turri, John (Eds.), *Contemporary Debates In Epistemology*, 2nd Edition. Blackwell.

Pryor, James (2004). What's Wrong with Moore's Argument? *Philosophical Issues* 14 (1):349–378.

Pryor, James (2001). Highlights of Recent Epistemology. *British Journal For The Philosophy Of Science* 52 (1):95--124.

Pryor, James (2000). The Skeptic and the Dogmatist. *Noûs* 34 (4):517–549.

Putnam, Hilary (1981). *Reason, Truth, and History*. Cambridge University Press.

Quine, W.V.O (1969). *Ontological Relativity and Other Essays*, New York: Columbia University Press.

Radford, C. (1966). Knowledge—By Examples. *Analysis*, 27, 1–11.

Radford, C. (1970). Does Unwitting Knowledge Entail Unconscious Belief? *Analysis*, 30, 103–107

Ranalli, Chris (2014). Luck, Propositional Perception, and the Entailment Thesis. *Synthese* 191 (6):1223-1247.

Ranalli, Chris (2013). Skepticism, Invulnerability, and Epistemological Dissatisfaction. In C. Illies & C. Schaefer (Eds.), *Metaphysics or Modernity?* Bamberg University Press. 113-148.

Reed, Baron, “Certainty”, *The Stanford Encyclopaedia of Philosophy* (Winter 2011 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2011/entries/certainty/>>.

Reed, Baron (2007). The Long Road to Skepticism. *Journal of Philosophy* 104 (5):236-262.

Ribeiro, Brian (2002). Cartesian Skepticism and the Epistemic Priority Thesis. *Southern Journal of Philosophy* 40 (4):573-586.

Schwitzgebel, E. (2010). Acting Contrary to our Professed Beliefs, or the Gulf between Occurrent Judgment and Dispositional Belief. *Pacific Philosophical Quarterly*, 91, 531–553

Siegel, Susanna (2012). Cognitive Penetrability and Perceptual Justification. *Noûs* 46 (2):201 - 222.

Smithies, Declan (2014). Can Foundationalism Solve the Regress Problem? In Ram Neta (Ed.), *Current Controversies in Epistemology*. Routledge. 73-94.

Smithies, Declan (2014). The Phenomenal Basis of Epistemic Justification. In Jesper Kallestrup & Mark Sprevak (Eds.), *New Waves in Philosophy of Mind*. Palgrave Macmillan. 98-124.

Sosa, Ernest and Kim, J., (eds.) (2000). *Epistemology: An Anthology*. Blackwell.

Sosa, Ernest (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge*, Volume I. Oxford University Press.

Sosa, Ernest (2008). Skepticism and Perceptual Knowledge. In Quentin Smith (Ed.), *Epistemology: New Essays*. Oxford University Press. 121.

Sosa, Ernest (2003). Davidson's Epistemology. In Kirk Ludwig (Ed.), *Contemporary Philosophy In Focus: Donald Davidson*. Cambridge University Press.

Sosa, Ernest (2000). Skepticism and Contextualism. *Noûs* 34 (S1):1-18.

Sosa, Ernest (1999). How to Defeat Opposition to Moore. *Philosophical Perspectives* 13 (S13):137-49.

Sosa, Ernest (1997). Reflective Knowledge in the Best Circles. *Journal Of Philosophy* 94 (8):410-430.

Sosa, Ernest (1993). Skepticism and our Knowledge Circle. *Philosophical Issues* 4:211-226.

Sosa, Ernest (1988). Beyond Scepticism, to the Best of our Knowledge. *Mind* 97 (386):153-188.

Sosa, Ernest & Stroud, Barry (1994). Philosophical Scepticism. *Aristotelian Society Supplementary Volume* 68:263 - 307.

Starmans, Christina & Friedman, Ori (2012). The Folk Conception of Knowledge. *Cognition* 124 (3):272-283.

Stern, Robert (2007). Transcendental Arguments: A Plea for Modesty. *Grazer Philosophische Studien* 74 (1):143-161.

Stern, Robert (2000). *Transcendental Arguments and Scepticism: Answering the Question of Justification*. Oxford University Press.

Steup, Matthias & Turri, John (Eds.) (2013). *Contemporary Debates in Epistemology*. Blackwell.

Stine, Gail (1976). Skepticism, Relevant Alternatives, and Deductive Closure *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, Vol. 29, No. 4, pp. 249-261.

Strawson, P. F. (1992). *Analysis and Metaphysics: An introduction to philosophy*. Oxford: Oxford University Press.

Strawson, P. F. (1995). My Philosophy. In P. Sen & R. Verma (Eds.), *The Philosophy of P. F. Strawson*. New Delhi: Indian Council of Philosophical Research

Strawson, P. F. (1985). *Scepticism and Naturalism: Some Varieties*. Routledge.

Strawson, P. F. (1975). *The Bounds of Sense: An Essay on Kant's Critique of Pure Reason*. Distributed By Harper & Row, Barnes & Noble Import Division.

Stroud, Barry (2011). *Engagement and Metaphysical Dissatisfaction: Modality and Value*. Oxford University Press.

Stroud, Barry (2011). *Philosophers Past and Present: Selected Essays*. Clarendon Press.

Stroud, Barry (2009). Explaining Perceptual Knowledge: Reply to Quassim Cassam. *European Journal of Philosophy* 17 (4):590-596.

Stroud, Barry (2009). Scepticism and the Senses. *European Journal of Philosophy* 17 (4):559-570.

Stroud, Barry (2004). The Epistemological Promise of Externalism. In Richard Schantz (Ed.), *The Externalist Challenge*. De Gruyter.

Stroud, Barry (2004). Perceptual Knowledge and Epistemological Satisfaction. In John Greco (Ed.), *Ernest Sosa and His Critics*. Malden MA: Blackwell Publishing. 165-173.

Stroud, Barry (2003). The Goal of Transcendental Arguments. In Robert Stern (Ed.), *Transcendental Arguments: Problems And Prospects*. Clarendon Press.

Stroud, Barry (2002). Sense-Experience and the Grounding of Thought. In Nicholas Smith (Ed.), *Reading McDowell: On Mind and World*. New York: Routledge.

Stroud, Barry. (2002). Sense-experience and the Grounding of Thought. In *Reading McDowell: On Mind and World*. New York: Routledge.

Stroud, Barry (2000). *Understanding Human Knowledge: Philosophical Essays*. Oxford University Press.

Stroud, Barry (1994). Scepticism, 'Externalism', and the Goal of Epistemology. *Proceedings of the Aristotelian Society* 68 (Supp. Vol):290-307.

Stroud, Barry (1989). Understanding Human Knowledge in General. In Marjorie Clay & Keith Lehrer (Eds.), *Knowledge and Skepticism*. Westview Press.

Stroud, Barry (1984). *The Significance of Philosophical Scepticism*. Oxford University Press.

Stroud, Barry (1981). The Significance of Naturalized Epistemology. *Midwest Studies In Philosophy* 6 (1):455-472.

Stroud, Barry (1968). Transcendental Arguments. *Journal of Philosophy* 65 (9):241-256.

Stroud, Barry (1977). Review “Ignorance: A Case for Scepticism” by Peter Unger (Oxford University Press). *The Journal of Philosophy*, pp. 246-257. Reprinted (as “Taking Scepticism Seriously”) in *Understanding Human Knowledge: Philosophical Essays*, Oxford University Press 2000.

Tobin, Emma & Bird, Alexander (2009). Natural kinds. *Stanford Encyclopedia of Philosophy*.

Thomasson, Amie L. (2010). The Controversy Over the Existence of Ordinary Objects. *Philosophy Compass* 5 (7):591-601.

Turri, J. (2010). Does Perceiving Entail Knowing? *Theoria*, 76(3), 197–206.

Unger, Peter (1979). There are no Ordinary Things. *Synthese* 41 (2):117 - 154.

Unger, Peter (1972). Propositional Verbs and Knowledge. *Journal of Philosophy* 64 (11):301-312.

Unger, Peter (1971). A Defense of Skepticism. *Philosophical Review* 80 (2):198-219.

Unger, Peter (1968). An Analysis of Factual Knowledge. *Journal of Philosophy* 65 (6):157-170.

Unger, Peter (1984). *Philosophical Relativity*. Oxford University Press.

Unger, Peter (1979). I Do Not Exist. In Graham F. Macdonald (Ed.), *Perception and Identity*. Cornell University Press.

Unger, Peter (1975). *Ignorance: A Case for Scepticism*. Oxford University Press.

Vahid, Hamid (2014). Varieties Of Pragmatic Encroachment. *Acta Analytica* 29 (1):25-41.

Van Inwagen, Peter (2008). How to Think about the Problem of Free Will. *Journal Of Ethics* 12 (3/4):327 - 341.

Van Inwagen, Peter (2000). Free Will Remains a Mystery. *Philosophical Perspectives* 14:1-20.

Van Inwagen, Peter (1998). The Mystery Of Metaphysical Freedom. In Peter Van Inwagen & Dean W. Zimmerman (Eds.). *Metaphysics: The Big Questions*. Blackwell. 365-373.

Van Inwagen, Peter (1990). *Material Beings*. Cornell University Press.

Vogel, Jonathan (2008). Epistemic Bootstrapping. *Journal of Philosophy* 105 (9):518-539.

Vogel, Jonathan (2007). Subjunctivitis. *Philosophical Studies* 134 (1):73 - 88.

Vogel, Jonathan (2013). The Refutation of Skepticism. In Steup Matthias, Sosa Ernest, And John Turri (Eds.), *Contemporary Debates In Epistemology*. Blackwell.

Vogel, Jonathan (2004). Skeptical Arguments. *Philosophical Issues* 14 (1):426–455.

Vogel, Jonathan (2000). Reliabilism Leveled. *Journal of Philosophy* 97 (11):602-623.

Vogel, Jonathan (1997). Skepticism and Foundationalism. *Journal of Philosophical Research* 22:11-28.

Vogel, Jonathan (1990). Cartesian Skepticism and Inference to the Best Explanation. *Journal of Philosophy* 87 (11):658-666.

Williams, Michael (2007). Why (Wittgensteinian) Contextualism Is Not Relativism. *Episteme* 4 (1):93-114.

Williams, Michael (2004). Scepticism and the Context Of Philosophy. *Philosophical Issues* 14 (1):456–475.

Williams, Michael (2001). *Problems of Knowledge: A Critical Introduction to Epistemology*. Oxford University Press.

Williams, Michael (1996). Understanding Human Knowledge Philosophically. *Philosophy and Phenomenological Research* 56 (2):359 - 378.

Williams, Michael (1991/1996). *Unnatural Doubts: Epistemological Realism and The Basis Of Scepticism*. Blackwell.

Williams, Michael (1988). Epistemological Realism and the Basis of Scepticism. *Mind* 97 (387):415-439.

Williamson, Timothy (2013). Knowledge First. In Matthias Steup, Ernest Sosa, and John Turri (Eds.), *Contemporary Debates In Epistemology*. Blackwell.

Williamson, Timothy (2005). Knowledge and Scepticism. In Frank Jackson & Michael Smith (Eds.), *The Oxford Handbook Of Contemporary Philosophy*. Oxford University Press.

Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford University Press.

Wittgenstein, Ludwig (1969). *On Certainty*, (Eds.) Anscombe and Von Wright. Harper Torchbooks.

Wong, W., Kolodny, N. & Bridges, J. (eds.) (2011). *The Possibility of Philosophical Understanding: Essays for Barry Stroud*. Oxford University Press.

Woozley, A. D. (1953). Knowing and not Knowing. *Proceedings of the Aristotelian Society*, 53, 151–172.

Wright, Crispin & Davies, Martin (2004). On Epistemic Entitlement. *Aristotelian Society Supplementary Volume 78*:167 - 245.

Wright, Crispin (2004). Wittgensteinian Certainties. In Denis Mcmanus (Ed.), *Wittgenstein and Scepticism*. Routledge.

Wright, Crispin (2004). Warrant for Nothing (And Foundations for Free)? *Aristotelian Society Supplementary Volume 78* (1):167–212.

Wright, Crispin (1991). Scepticism and Dreaming: Imploding the Demon. *Mind, New Series*, Vol. 100, No. 1, pp. 87-116.

Yalçın, Ümit D. (1992). Skeptical Arguments from Underdetermination. *Philosophical Studies* 68 (1):1 - 34.

Zagzebski, Linda (2009). *On Epistemology*. Wadsworth.

Zagzebski, Linda (2004). Epistemic Value and the Primacy of What We Care About. *Philosophical Papers* 33 (3):353-377.

Zagzebski, Linda (1996). *Virtues of the Mind: An Inquiry Into The Nature Of Virtue And The Ethical Foundations of Knowledge*. Cambridge University Press.

Zagzebski, Linda (1994). The Inescapability of Gettier Problems. *Philosophical Quarterly* 44 (174):65-73.

Zalabardo, José L. (2012). Wright on Moore. In Annalisa Coliva (Ed.), *Mind, Meaning, and Knowledge: Themes from the Philosophy of Crispin Wright*. Oxford University Press. 304–322.