FROM EVIDENCE TO UNDERDETERMINATION
ESSAYS IN THE WAY OF SCEPTICISM

Paolo Savini

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews

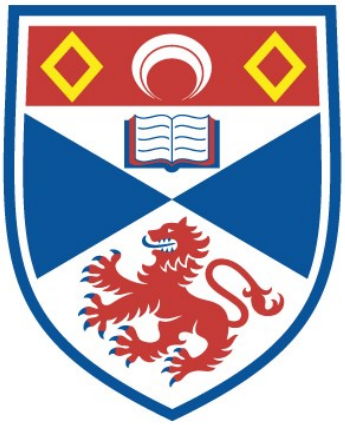

2023

From Evidence to Underdetermination

Essays in the Way of Scepticism

Paolo Savino

University of
St Andrews

This thesis is submitted in partial fulfilment for the degree of

Doctor of Philosophy (PhD)

at the University of St Andrews

December 2022

Candidate's declaration

I, Paolo Savino, do hereby certify that this thesis, submitted for the degree of PhD, which is approximately 60,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree. I confirm that any appendices included in my thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

I was admitted as a research student at the University of St Andrews in September 2018.

I received funding from an organisation or institution and have acknowledged the funder(s) in the full text of my thesis.

Date 13.04.2023 Signature of candidate

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree. I confirm that any appendices included in the thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

Date 13.04.2023 Signature of supervisor

Permission for publication

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Paolo Savino, confirm that my thesis does not contain any third-party material that requires copyright clearance.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

Date 13.04.2023 Signature of candidate

Date 13.04.2023 Signature of supervisor

Underpinning Research Data or Digital Outputs

Candidate's declaration

I, Paolo Savino, hereby certify that no requirements to deposit original research data or digital outputs apply to this thesis and that, where appropriate, secondary data used have been referenced in the full text of my thesis.

Date 13.04.2023 Signature of candidate

Abstract

Scepticism about justification is the view that justification is impossible. Underdetermination scepticism is scepticism that turns on the idea that our beliefs are underdetermined by the evidence relative to certain sceptical hypotheses. This thesis provides an elucidation and a defence of underdetermination scepticism on an evidentialist framework for justification and a mentalist conception of evidence. The thesis consists of five chapters and a conclusion. Chapter 1 introduces the Underdetermination Argument for scepticism and explains the core concepts of the thesis. Chapter 2 explores the relationship between closure and underdetermination scepticism. Chapter 3 responds to the Infallibility Objection, the idea that the Underdetermination Argument is a bad argument because the inference from sameness of evidence to underdetermination presupposes infallibilism. Chapter 4 responds to the charge that the Underdetermination Argument relies on excessive demands on the cognitive accessibility of evidence. Chapter 5 responds to attempts to resist scepticism on the ground that it is a Moorean fact that our beliefs are justified. The conclusion reviews and generalizes the results of the previous chapters. The upshot is that a significant set of objections against underdetermination scepticism fails. At the end of the day, we might have to take the possibility of living with scepticism seriously – or at least more seriously than we thought.

Acknowledgements

I would like to thank the following for their help and support.

My supervisors. First, my thanks go to Philip Ebert. Philip became my primary supervisor at a time of personal and professional confusion. Philip patiently helped me to find my way back into the PhD. Ever since, I have benefited regularly from Philip's kindness and intelligence and, from time to time, from his mountaineering prowess. It is cheesy, but true: I could not wish for a better supervisor.

Jessica Brown, Patrick Greenough, and Roy Sorensen have all acted as either my primary or secondary supervisor at some point. The present work is better for their advice than it would be otherwise.

My annual reviewers. For similar reasons, I would like to thank my annual reviewers: Derek Ball, Aaron Cotnoir, Adrian Haddock, Peter Milne, and Walter Pedriali. Special thanks to Adrian and Walter for their encouraging words at a moment when I needed to hear those words.

My examining committee. Crispin Wright and Kevin McCain have examined this thesis. I am thankful for their thorough engagement with my work, their sharp and helpful feedback, and for making my viva a memorable experience.

SASP & Arché. Overall, the St Andrews & Stirling Graduate Programme in Philosophy (SASP) has provided me with a friendly and stimulating research environment. The same is true of Arché, the philosophical research centre in St Andrews, where I learned a great deal of what I know about epistemology. Thanks to everyone who makes them function on a daily basis, from faculty members to professional services. Special thanks to Franz Berto.

The SASP PhD community. I have shared the joys and sorrows of being a PhD student with some amazing colleagues and friends. Katharina Bernhard, Jimena Clavel Vázquez, Andrés Hernández Villarreal, Lixiao Lin, Matteo Nizzardo, Andrea Oliani, Quentin Pharr, Stefano Pugnaghi, Ravi Thakral, and Xintong Wei stand out. They know why.

My Heidelberger *friends*. Fausto Carcassi and Richard Roth have been friends and intellectual role models since my undergraduate days. I do the philosophy I do the way I do because of them. Marco Bianchi, Antonin Dubois, and Stefania Girometti have

provided me with support and good times for many years now. *E più d'onore ancora assai mi fenno...*

My Italian friends. I have spent many pleasant and productive hours working online with Luca Morelli in the final stages of my PhD. And I have spent too many to count – and I hope to spend many more – with Alessandro Cherubini bagging hills and mountains.

Miscellanea. Zara Treharne has been with me for the last year of my PhD – her company has made it more pleasant than I could have imagined. I presented parts of this thesis at the Arché epistemology seminar (several times), at the SASP Friday Seminar (several times), at a SGSAH Philosophy Catalyst Workshop, at the European Epistemology Network Conference 2022 in Glasgow, and at the 11th congress of the Society for Analytic Philosophy (GAP.11) in Berlin. Thanks to members of the audiences for their feedback. Avertical World, Dundee's climbing centre, has been an important venue to unwind in the final months of the PhD – thanks to Miriam Bowen, Elisa Camminatore, and Paul Conlan for all the climbing and the chats.

My family. Finally, I cannot overstate the importance of my family's support: my parents Fulvia and Luigi, my sister Silvia, and my aunt Mariella. I love them very much. This thesis is for them.

Funding

This work was supported by a full fee waiver of the Department of Philosophy at the University of St Andrews.

Contents

Acknowledgements i
1. Introduction: Underdetermination Scepticism 1
1.1 Scepticism 2
1.2 Evidence and Justification 5
1.2.1 An Evidentialist Conception of Justification 6
1.2.2 A Mentalist Conception of Evidence 8
1.3 Underdetermination 11
1.4 The Underdetermination Principle 16
1.5 The Underdetermination Argument 19
1.6 Scepticism, Paradox, and Self-Refutation Charges 21
1.7 The Plan 25
2. Underdetermination Scepticism and Closure Scepticism 29
2.1 Preliminaries: Soundness and Cogency 31
2.1.1 The Soundness Question 32
2.1.2 The Cogency Question 34
2.2 Answering the Soundness Question 34
2.2.1 ~FAV and ~J 35
2.2.2 The Underdetermination Principle and the Closure Principle 37
2.3 Answering the Cogency Question 48
2.3.1 Motivating the Underdetermination Argument's Premises 50
2.3.2 Motivating the Closure Argument's Premises 53
2.4 Concluding Remarks 56
3. The Infallibility Objection 61
3.1 Fallibilism, Infallibilism, and Scepticism 62
3.2 The Infallibility Objection 65
3.2.1 Brueckner 66
3.2.2 Dodd 67
3.2.3 Smith 68
3.3 Confrontation and Sidestepping 72
3.4 Two Unsuccessful Attempts to Sidestep the Infallibility Objection 73
3.4.1 The Causal Explanation Principle 73

3.4.2 The Entailment Principle 75
3.5 Confrontation: A Bayesian Account of the Sceptical Inference 79
3.5.1 Establishing Premise 1 80
3.5.2 Establishing Premise 2 83
3.6 Concluding Remarks 88
4. Sameness of Evidence and Anti-Luminosity 91
4.1 The Anti-Luminosity Argument 92
4.1.1 Luminosity and Being in a Position to Know 92
4.1.2 The Argument 93
4.2 From Luminosity to Sameness of Evidence 95
4.2.1 The Sceptic's Argument 97
4.2.2 Williamson's Reply 99
4.3 A Summary and a Preview 101
4.4 What the Anti-Luminosity Argument Does Not Show 101
4.4.1 Adjacent Cases and the Good and the Bad Case 104
4.4.2 Small Differences in S's Evidence Do Not Matter 106
4.5 Sameness of Evidence Without Luminosity 110
4.5.1 Sameness of Evidence and Mentalism 111
4.5.2 Sameness of Evidence, Rationality, and the New Evil Demon 114
4.6 Concluding Remarks 116
5. Underdetermination Scepticism and Common Sense 122
5.1 Mooreanism, Scepticism, and Dogmatism 123
5.1.1 Mooreanism: What It Is and What It Is Not 123
5.1.2 Mooreanism and the Dogmatism Challenge 127
5.2 Philosophy, Science, and Common Sense 129
5.2.1 Rinard's Argument 129
5.2.2. A Significant Disanalogy Between Philosophy and Science? 130
5.3 Mooreanism and the Norms of Belief Change 133
5.3.1 Two Constraints 134
5.3.2 Plausibility 134
5.3.3 Certainty 138
5.3.4 Knowledge 140
5.4 Mooreanism and Philosophical Method 141

5.4.1 Particularism, Generalism, and Reflective Equilibrium 142
5.4.2 Kelly's Argument 143
5.4.3 Scepticism Without Hyper-Methodism 144
5.5 Concluding Remarks 147
6. Conclusion: Last Thoughts of a Provisional Sceptic 151
6.1 Summary 151
6.2 Open Questions 159
Bibliography 161

Spiritus durissima coquit

1

Introduction: Underdetermination Scepticism

Underdetermination scepticism is scepticism that turns on the idea that our beliefs are underdetermined by the evidence. First, the sceptic identifies a class of target beliefs (e.g. beliefs about the external world, beliefs about the past, beliefs about other minds) and the evidence that bears on their truth (e.g. sensory experiences, apparent memories, behavioural evidence). The sceptic claims that said evidence can be accounted for by an alternative hypothesis incompatible with the beliefs under attack. Sensory evidence can be accounted for by the deceiving demon hypothesis, apparent memories can be accounted for by the hypothesis that the world, our memories included, was brought about five minutes ago, behavioural evidence can be accounted for by the philosophical zombie hypothesis, and so on. This point is often made with reference to a good case and a bad case: a good case, in which our beliefs are true, corresponds to a bad case in which our beliefs are false although we have exactly the same evidence in the two cases. The sceptic argues that, because we have the same evidence in the good and the bad case, the evidence supports our beliefs and the incompatible hypothesis equally well. Invoking an underdetermination principle, the sceptic concludes that our beliefs are unjustified, because the evidence is neutral between the incompatible alternatives, i.e. because it does not favour one over the other.

This thesis provides a defence of underdetermination scepticism about the external world. It does so by responding to some important objections to the Underdetermination Argument for scepticism. The defence is conditional as well as partial. It is conditional as it assumes, as most responses do, several philosophically controversial claims, for which no fully satisfactory arguments can be given here. I will, however, flag these assumptions and ensure to flag my assumptions as well. It is partial because it responds only to a small if significant set of objections.

In this introductory chapter, I lay out the Underdetermination Argument for scepticism about the external world and the core concepts of this thesis. Section 1 clarifies the concept of scepticism; section 2 explains the concepts of evidence and justification, offers an evidentialist picture of justification, and motivates a mentalist conception of evidence; section 3 clarifies the concept of underdetermination; while section 4

explains and motivates the underdetermination principle. Section 5 puts together sections 1-4 and lays out the underdetermination argument for scepticism discussed in this thesis. Section 6 explains the paradoxical character of scepticism and its consequences for the anti-sceptical import of self-refutation charges to scepticism. Finally, section 7 offers a preview of the next chapters.

1.1 Scepticism

What is scepticism? It is hard to say. 'Scepticism' is a contested term inside and outside philosophy. A review of the meanings of 'scepticism' and its cognates reveals a multiplicity of somewhat related yet distinct meanings. In ordinary speech[1], the word 'scepticism' assumes several different meanings and connotations that centre around a cluster of loosely connected concepts: disbelief, doubt, uncertainty. Thus, 'scepticism' may come to indicate a state, a cognitive trait, or even an action or practice. Similarly, scepticism receives a positive connotation in some contexts, a negative one in others.

In philosophy, 'scepticism', though a contested term, has become a term of art, at least within certain debates.[2] Stroud (1984: vii) remarks that 'in modern, and especially recent, times scepticism in philosophy has come to be understood as the view that we know nothing, or that nothing is certain, or that everything is open to doubt.' He was writing with, among others, Russell (1912), Austin (1946), Moore (1959), Wittgenstein (1969), Dretske (1970), Unger (1975), Cavell (1979), and Nozick (1981) in mind. Almost forty years later, Stroud's observation remains largely correct, at least judging by a very large number of publications in epistemology.[3]

Although the paradigm of scepticism so understood is the claim that we can never know anything, scepticism may target justification, certainty, rationality or some other

[1] By 'ordinary speech', I mean language that is typically used by the philosophical layperson in her day-to-day dealings. I contrast it with philosophical jargon, which is typically used by philosophers among themselves, when they are talking and writing philosophy, or in a professional context, like at a conference, in classroom, or in specialistic print.

[2] I do not want to suggest any sharp divide between ordinary and philosophical speech. On the contrary, philosophers often philosophise in ordinary language, and the term 'scepticism' is no exception to this rule.

[3] The bibliography at the end of this thesis provides some evidence for this claim. Additional evidence can be found in two reviews of work on scepticism from the early 2000s: see Pryor (2001) and Pritchard (2002).

positive epistemic condition.[4] Moreover, scepticism can vary in modal strength. For instance, Frances (2018: 581) calls 'disagreement scepticism' the view that some of our controversial beliefs are epistemically defective because we are aware that they are controversial. Similarly, Feldman (2006: 217) writes that in many cases of recognised disagreement with informed and intelligent people we lack reasonable belief and knowledge, and calls this 'a kind of contingent real-world skepticism.'[5]

Like its modally stronger counterpart, the modally weaker claim makes a negative statement about our epistemic situation, which is the hallmark of scepticism.[6] The common factor is the idea that scepticism is a claim with a distinctive epistemological content, *that creatures like us are not or cannot be in some positive epistemic condition C with respect to some propositional domain D*[7]. Thus, we may call this way of thinking about scepticism '*epistemological dogmatic*', and the corresponding type of scepticism '*ED-scepticism*'.[8] It is

[4] Positive epistemic conditions are desirable ends in our investigations, things we, as epistemic agents, see as valuable, like knowledge, certainty, justification, rationality etc.

[5] See also Christensen (2009), Kelly (2010), Grundmann (2019).

[6] Still, in some contexts, it may be convenient to mark the difference between modally strong and modally weak forms of scepticism by reserving the label 'scepticism' for modally strong claims. For instance, Rosenkranz (2007: 58) distinguishes scepticism from True Agnosticism: 'unlike scepticism, True Agnosticism does not rule out the possibility of knowledge altogether: its truth does not preclude that we might know the truth-value of the statements in question by as yet unconceivable methods.'

[7] Propositional domains are sets of propositions, typically sharing a topic, like the set of propositions about the external world, the set of propositions about the past, the set of propositions about other minds etc.

[8] ED-scepticism can vary along several dimensions. It can vary along C by targeting different positive epistemic conditions (knowledge, justification, etc.) or their subspecies (perceptual knowledge/justification etc., inductive knowledge/justification etc., testimonial knowledge/justification etc., and so on), and it can vary along D by targeting different propositional domains (propositions about the external world, propositions about the past, propositions about other minds etc.). The strength and reach of ED-scepticism varies with variations in C and D. If knowledge requires justification, but not the other way around, justification ED-scepticism entails (ceteris paribus, i.e. assuming D and the modal strength are fixed) knowledge ED-scepticism, but not the other way around. And if perceptual justification is necessary for testimonial justification, but not the other way around, perceptual justification ED-scepticism is logically stronger than testimonial justification ED-scepticism as well as testimonial knowledge ED-scepticism. Similarly, ED-scepticism about other minds is logically stronger than ED-scepticism about non-human minds, and justification ED-scepticism about the past is logically stronger than justification ED-scepticism about the distant past (ceteris paribus, i.e. assuming C and the modal strength are fixed). Strictly speaking, scepticism can vary along an additional variable that I have not considered here. I have assumed that sceptical claims always target creatures like us, but that need not be the case. Accordingly, scepticism may target larger or smaller groups of agents.

dogmatic because it conceives of scepticism as a claim or a doctrine.[9] It is epistemological because it conceives that claim as having a distinctive epistemological content.[10] This is how I shall understand underdetermination scepticism, too.

Comesaña and Klein (2019) define scepticism with respect to a field of propositions F as 'the claim that the only justified attitude with respect to propositions in F is suspension of judgement.' Does their definition fit the epistemological-dogmatic conception? It does on the assumption that there are only three doxastic attitudes one can have towards a proposition: belief, disbelief, and suspended judgement. On that assumption, the claim that suspended judgement is the only justified attitude is equivalent to the ED-sceptical claim that one is never justified in believing anything.[11]

Some authors refer to scepticism as *the denial of certain truths or of the existence of certain entities*.[12] For instance, Levy (2015: 192) refers to the claim that agents are never morally responsible for their actions as 'scepticism about moral responsibility'. Similarly, Sinnott-Armstrong (2006: 11-2) calls 'scepticism about moral truth' the claim that no (substantive) moral belief is true, and 'scepticism about moral reality' the claim that there are no moral facts or properties, while Watson (1977: 316) classifies Socratism, the claim that weakness of will is impossible, as a form of scepticism. According to this way of thinking about scepticism, sceptical claims need not have an

[9] Here, I am using the term 'dogmatic' in the Ancient Greek sense. In this sense, a dogma is simply a doctrine. Nowadays, the word 'dogmatic' has assumed a pejorative sense that was lacking in Ancient Greek (Annas and Barnes 1985: 1-2).

[10] ED-scepticism is a family of epistemological views. These views should not be confused with ontological views like the view that there are no physical objects, no moral facts, no abstract entities etc. Although the latter views count as sceptical according to another way of thinking about scepticism (see further), they are not forms of ED-scepticism.

[11] One might object that the claim that suspended judgement is the only justified attitude and the claim that one is never justified in believing or disbelieving any proposition are not equivalent, because the latter does not entail the former. After all, it might be true that one is never justified in believing or disbelieving any proposition although no one is ever justified in suspending judgement either. Perhaps, no doxastic attitude is ever justified. Although this is right, it need not mean that Comesaña and Klein's definition does not fit ED-scepticism. Instead, it might mean that there are even stronger forms of scepticism than those considered by Comesaña and Klein.

[12] Thus, on this way of thinking about scepticism, scepticism roughly corresponds to what is sometimes called 'anti-realism'. However, the identification of scepticism and anti-realism is complicated by the fact that the distinction between 'realism' and 'anti-realism' is itself contentious (Wright 1992: 1). Moreover, on some ways of understanding that distinction, anti-realists merely deny that the entities in question exist *objectively* but concede that they exist in some other sense.

epistemological content. To distinguish it from the epistemological dogmatic conception, we can call it '*negative dogmatic*', and the corresponding type of scepticisms 'ND-scepticism.'[13]

Finally, scepticism is sometimes understood as *a practice, a method, or an ability leading to judgement suspension.* For instance, Sextus Empiricus (PH 1 8)[14] famously calls scepticism 'an ability to set out oppositions among things which appear and are thought of in any way at all, an ability by which, because of the equipollence in the opposed objects and accounts, we come first to suspension of judgement and afterwards to tranquillity.' Since, according to it, scepticism is a practice, we can call this way of understanding scepticism '*practical*', and the corresponding type of scepticism '*P-scepticism'*.

1.2 Evidence and Justification

Epistemological dogmatic scepticism can be supported in several ways. One family of arguments for ED-scepticism, inclusive of arguments from underdetermination, appeals to sceptical scenarios along the lines of those made famous by Descartes (1996 [1641]). Sceptical scenarios are 'epistemic disaster scenarios', as Weatherson (ms) aptly calls them, namely situations in which our cognitive position is considerably worse than we normally think it is – at least those of us who are not sceptics. These scenarios – like the deceiving demon hypothesis, the brain-in-a-vat hypothesis etc. – have several features that the sceptic can put to work in a sceptical argument. To name a few: in a sceptical scenario ('the bad case'), things appear the same as in a corresponding non-sceptical scenario ('the good case'), but in the sceptical scenario the appearances are misleading. In a sceptical scenario one has the same beliefs as in the corresponding

[13] Notice that every instance of ED-scepticism counts as an instance of ND-scepticism. After all, by denying the existence/possibility of some positive epistemic condition C, ED-sceptics deny the existence of a certain entity, namely C. Thus, by denying the existence/possibility of knowledge, knowledge ED-sceptics deny the existence of a certain entity, namely knowledge. However, the converse is not true: not every instance of ND-scepticism is an instance of ED-scepticism, because some denials do not have an epistemological content. Alternatively, one can think of ED-scepticism as the denial of a certain class of truths, namely ascriptions of C. Every instance of ED-scepticism counts as an instance of ND-scepticism, but not the other way around, under this way of understanding it, too.

[14] References to *Outlines of Pyrrhonism* follow the conventional format *PH x y*, where *PH* stands for the work's original title *Purrhōneioi Hupotupōseis*, *x* for the number of the book (in the ancient sense of 'chapter'), and *y* for the number of the paragraph. All translations of *PH* are from Sextus Empiricus (1990 [2nd century CE]).

non-sceptical scenario, but in the sceptical scenario one's beliefs are false.[15] Further, it seems impossible to discriminate one's situation from the sceptical scenario. Each of these features can be used to construct a sceptical argument.

The underdetermination sceptic argues from the sameness of evidence in the good and the bad case to the underdetermination of our beliefs by the evidence, and from the latter (via an underdetermination principle) to the conclusion that we are not justified in believing the target proposition. Underdetermination sceptical arguments posit an important epistemic difference between evidence and the content of the target beliefs: the former is taken to be epistemically prior to the latter. Whereas we are supposed to have a direct or unproblematic cognitive access to the evidence, this is not true for the target propositions. Whatever cognitive access we have to the target propositions is supposed by the sceptic to be mediated by the evidence. The target propositions are thus removed from the immediate reach of our cognitive powers. For this reason, some authors have maintained that underdetermination scepticism about the external world rests on the foundationalist assumption of the epistemic priority of sense experience over ordinary beliefs about the external world (Stroud 1984; Davidson 1986; Williams 1991).

In this thesis, I shall make a number of assumptions about justification, evidence, confirmation, and their relationship, but I shall also remain neutral on a number of topics. Naturally, some of the assumptions will be controversial, but not only do we need to start somewhere, I will also highlight to what extent my assumptions affect the thesis as a whole (for more on this, see chapter 6, section 2).

1.2.1 An Evidentialist Conception of Justification

When I talk about justification, I am concerned with epistemic justification, i.e. justification by epistemic rather than practical reasons: in my view, one is justified in believing that a proposition P is true just in case one's total evidence E supports or confirms

[15] This is not true of all sceptical scenarios. Some authors distinguish between two sorts of sceptical scenarios: scenarios in which our beliefs are true (or, at any rate, scenarios in which the truth value of our beliefs is left undetermined and thus might be true) and scenarios in which our beliefs are false (Roush 2010; Freitag 2013).

P to a sufficient degree.[16] Accordingly, I shall take evidence to be that which can make a difference to the justification of beliefs – or to what is reasonable or rational to believe – by making their contents more or less probable or by indicating that they are true. In my usage, the concept of (a piece of) evidence and the concept of epistemic reason are thus interchangeable. So are the concept of justification and (epistemic) rationality.

It is standard to distinguish between propositional and doxastic justification (Firth 1978; Turri 2010; Silva and Oliveira forth.). Propositional justification is a property of propositions, while doxastic justification is a property of beliefs. Famously, propositional and doxastic justification can come apart. Roughly, one is propositionally justified to believe that P just in case one has enough reasons to believe that P. Thus, one can have propositional justification to believe that P even if one does not justifiably believe that P (either because one does not believe that P at all or because one's belief is not responsive to what gives one justification). In contrast, doxastic justification requires having a belief properly based on that which gives one propositional justification for that belief. The issues I am concerned with in this thesis are independent of this distinction. Accordingly, there is no need to pay much attention to it here.

Further, it is standard to distinguish between a graded and a categorical notion of justification (Brown 2018: 73; Hawthorne and Logins 2021). In the categorical sense, a belief (or proposition) is either justified or it is not, while, in the graded sense, a belief can be more or less justified. The graded notion of justification allows for comparative evaluations of justified beliefs in a way in which the categorical notion does not. In the categorical sense, two beliefs might be justified/unjustified although in the graded sense one is more justified than the other. I identify the graded notion of justification with the notion of evidential support or confirmation.

More generally, I shall think of justification in evidentialist terms, where

[16] One's justification to believe that P depends on one's *total* evidence rather than a subset thereof. Even if a piece of evidence e is sufficient to believe that P when considered in isolation, one need not be justified in believing that P as soon as one has e, for a piece of evidence can be undermined by further evidence. In other words, one might possess some additional evidence e^* such that believing that P given $e + e^*$ is unjustified. When this happens, we say that e^* defeats the justification that e confers on P, or that e^* defeats e. Examples abound: evidence that I am under the effect of a drug that impairs memory might defeat my memory that the French Revolution happened in 1789 (Cf. Kelly 2008a: 938).

evidentialism is the claim that what one is justified in believing is determined by one's evidence (Feldman and Conee 1985: 15). In fact, evidentialism applies to doxastic attitudes in general, not just to beliefs: a doxastic attitude is justified to the extent that it fits the evidence.

I shall understand evidentialism as a determination thesis. In my understanding, evidence determines justification in the sense that every difference in justification holds in virtue of some difference in evidence. In other words, you and I can differ in what is reasonable to believe for us just in case there is a difference in the evidence we have, and this difference in evidence is responsible for the difference in justification. Further, whether evidence justifies one in believing something does not depend only on what evidence one has, but also on what one's evidence supports. In the following, I shall assume that the evidential support relationship (i.e. the relationship of confirmation) holds necessarily for all subjects, times, and places.

It is common to understand 'evidentialism' as the thesis that facts about justification *supervene* on facts about evidence (McCain 2014: 1-2). In my understanding, evidentialism is stronger than that claim – the claim that there cannot be any difference in justification without a difference in evidence –, for, unlike the supervenience claim, the determination thesis reflects the asymmetric dependence of justification on evidence. Thus, according to my understanding of 'evidentialism', differences in justification hold in virtue of differences in evidence, but not the other way around (Smithies 2019: 200).

Finally, I take justification to be non-factive. It is a platitude that rational thinkers respect their evidence. One's beliefs are justified or epistemically rational insofar as they are responsive to the evidence. But respecting one's evidence is no guarantee that one's beliefs are true; sometimes, one's evidence is misleading, which is to say that evidence is a guide to truth but not a prefect guide. Accordingly, one's beliefs can be justified, if at all, despite being mistaken. Thus, on this point, I part ways with some authors who believe that justification is factive – that a belief is justified only if it is true (Sutton 2007; Williamson 2013).

1.2.2 A Mentalist Conception of Evidence

In this thesis, I shall think of perceptual evidence (and other evidence as well) as

restricted to non-factive mental states. In this, I part ways with some authors who endorse 'generous' conceptions of perceptual evidence – as far as I know, the apt name is due to Brown (2018: 1) – on which perceptual evidence can include true propositions about the external world (McDowell 1982; Williamson 2000, 2007; Schellenberg 2013; Pritchard 2015; Schönbaumsfeld 2016) and, more generally, with several authors who think about evidence in an externalist way.

Some readers of an externalist bent will think that this assumption grants too much to the sceptic – that once we accept a phenomenal or mentalist conception of evidence, scepticism becomes irresistible. To these readers, I say: bear with me. I cannot do full justice to this assumption here, but I shall say a few words about its rationale now and return to it in chapter 4 and in the conclusion.

In any case, although I do not think that I am able to convince a committed Externalist to abandon externalism, I tend to agree with these readers that, ultimately, a mentalist conception of evidence pushes us towards scepticism (although it is not trivial to show that it does). Thus, I invite dogmatic externalist readers to read this thesis as a *reductio ad absurdum* of mentalist conceptions of evidence, which ultimately may strengthen their own dialectical position, if they wish to (as for me, I do not think that the fact that mentalism might lead to scepticism is an argument in favour of externalism).

So, why think that mentalism is true? The reason I assume a mentalist conception of evidence is that I am thinking of evidence as the kind of thing to which our beliefs, in particular our perceptual beliefs, are responsive. And it is tempting to think that the kind of thing to which our (perceptual) beliefs are responsive are non-factive mental states. This point is made vivid by comparing seeing and hallucinating. Consider a case where you are looking at a tower in front of you at one moment. Upon seeing the tower, you believe that the tower is in front of you. A moment later, the tower ceases to exist, but you retain the same visual experience you had initially throughout the process. Even if the tower does not exist anymore, you still believe that the tower is in front of you. What this suggests is that your beliefs are not responsive to your seeing the tower or to the tower itself, but to whatever seeing and hallucinating the tower have in common, that is a non-factive experience (Bergmann 2021: 24-5). The idea is easily extended to non-perceptual beliefs and non-perceptual evidence.

In a similar vein, Schoenfield (2015) argues that mentalism is motivated by considerations about the causal role played by mental states in our cognitive lives, because mental states act as causal mediators between states of the world and beliefs. To paraphrase Schoenfield (2015: 257), mentalist evidence is important because this is the evidence that we should expect to respond to as a result of responding to whatever we take evidence to be.

To elaborate: mentalism is motivated by the idea that non-factive mental states are more causally proximate to our beliefs than whatever non-mental condition the externalist takes evidence to be. In virtue of this, 'mental' evidence acts as a constraint on 'non-mental' evidence in the sense that we should expect to respect the latter only if we should expect to respect the former.

As an illustration consider the following case:

> A BEACON IN THE NIGHT
>
> Paul Revere and his fellow revolutionaries are considering how to spread information about the direction from which the British will be coming. Revere proposes to light one lantern if the British come by land, and two lanterns if the British come by sea – call this 'the Revere Plan'. But someone questions the wisdom of the Revere Plan: they are worried about British decoys, that the British will make it appear as if they are coming by land when they are actually coming by sea, or the other way around. Paul Revere objects that there is no need to worry: the plan is to light lanterns according to what is the case, not according to what *appears* to be the case. (Adapted from Schoenfield 2015: 257)

There is something odd about Paul Revere's reply to the worry he might be misled by the British. Although the plan is to light lanterns according to what is the case, clearly, Revere will light lantern according to what appears to him to be the case. In other words, the result of making the Revere Plan is that Revere will conform to another plan: to light one lantern if it *appears* that the British are coming by land, and two lanterns if it *appears* that the British are coming by sea. Thus, if the chances that the British will use decoys are sufficiently high, the Revere Plan is not a good plan. The case illustrates how 'mental' evidence acts as a constraint on 'non-mental' evidence. If

Revere wants to form an accurate belief about the direction of movement of the British troops, he should expect to form an accurate belief on the basis of what appears to him, not just to what is the case (Schoenfield 2015: 262).

Nothing I said here implies that there must be something wrong with externalist conceptions of evidence. In fact, I am quite happy to concede that there is more than one acceptable way to use the term 'evidence' and that externalist conceptions of evidence might capture some senses of 'evidence' better than internalist conceptions do.[17] Still, I think these considerations show that there is an important sense in which mentalist conceptions of evidence are right. As I will explain in the conclusion of this thesis, this is a problem for externalists.

Finally, although I too tend to side with the venerable tradition of epistemologists that have objected to *irresponsible* beliefs and to a certain kind of epistemic hubris (Johnson King 2022: 10-1), I shall remain neutral on the controversy between accessibilists and non-accessibilists about evidence and justification. Accessibilists think that having evidence/justification requires some kind of awareness of that evidence/justification. Non-accessibilists deny this. Although accessibilism is more congenial to scepticism than non-accessibilism – after all, ceteris paribus, accessibilism requires more of one's beliefs to be justified than non-accessibilism does – in my view, underdetermination scepticism cuts across the distinction between accessibilism and non-accessibilism. In any case, the distinction between accessibilism and non-accessibilism is important for my thesis, because in chapter 4 I shall deal with the charge that the Underdetermination Argument for scepticism relies on a mistaken luminosity claim.

1.3 Underdetermination

The main focus of this thesis is a modally strong version of underdetermination scepticism about *perceptual* justification – justification whose source is sensory experience – or about the justification of *external world propositions* – propositions about the material reality which we take to inhabit, a reality made of bodies of different shapes, sizes,

[17] For the suggestion that that the concept of 'evidence' plays a number of distinct roles inside and outside philosophy and that some of these roles cannot be easily reconciled, see Kelly (2016).

colours, textures etc.[18] However, I shall say several things that apply to other kinds of underdetermination scepticism, too. It is still useful to focus on one familiar kind of underdetermination scepticism as this helps making the discussion more vivid.

The distinctive character of underdetermination scepticism is the idea that our beliefs are underdetermined by the evidence. The phenomenon of underdetermination has long been studied by philosophers of science. In fact, the idea of underdetermination is at the heart of an influential argument against scientific realism, the view that 'the central claims of our best scientific theories about how things stand in nature must be at least probably and/or approximately true' (Stanford 2006: 6). This argument is just a version of the sceptical argument from underdetermination described at the start of this chapter, where beliefs about unobservable entities are the target beliefs and observational data are the evidence.

As Okasha (2003: 618) remarks, although it often goes under the name of 'anti-realism', Van Fraassen's constructive empiricism is also a form of underdetermination scepticism. Van Fraassen's view – that 'science aims to give us theories which are empirically adequate; and acceptance of a theory involves as belief only that it is empirically adequate' (1980: 12) – is largely motivated by the thought that even our best scientific theories have *empirically equivalent rivals*, i.e. theories that have the same empirical consequences but make incompatible claims about unobservable entities. According to Van Fraassen, since empirical equivalent theories make the exact same predictions, they are underdetermined by any possible body of evidence.

Douven (2022: 56-7) characterises underdetermination as a relationship between distinct classes of propositions:

> 'We might for instance say that one class of propositions, C_1, <know, know>-underdetermines another class of propositions, C_2, if and only if knowing every member of C_1 is not enough to know any member of C_2. Similarly, C_1 <know, justifiedly believe>-underdetermines C_2 if and only

[18] It is common to equate external world scepticism and scepticism about perceptual justification, and I shall do the same for simplicity's sake. However, note that, strictly speaking, they are not equivalent. If there are non-perceptual ways to acquire justification about external world propositions, external world scepticism and perceptual justification scepticism can come apart. The same is true, if perception provides justification for other kinds of propositions (Bergmann 2021: 26-7).

> if knowing every member of C_1 is not enough even to be justified in believing any member of C_2.'

Unlike Douven, I shall think of underdetermination as a relationship of confirmation between a class of propositions (or beliefs) and a body of evidence. I shall say that a body of evidence underdetermines a class of propositions just in case that body of evidence supports equally well every member of that class. I shall also say that some proposition is underdetermined by the evidence just in case there is some incompatible proposition that is equally well supported by the evidence.

Thus, generally speaking, underdetermination claims are claims to the effect that the evidence is neutral about which beliefs we should adopt in response to it. On an alternative but related formulation, they are claims to the effect that the evidence is neutral about the truth value of the relevant theories/hypotheses. After all, if the evidence is neutral about the truth value of the relevant theories/hypotheses, it is also neutral about which theory/hypothesis we should accept in response to it.

Underdetermination claims can vary in character and strength. In line with Duhem (1914), Stanford (2017) suggests a distinction between two varieties of underdetermination: holist and contrastive underdetermination. Holist underdetermination is intimately connected to confirmational holism, the idea that scientific theories/hypotheses cannot be tested in isolation, but only in conjunction with a set of auxiliary hypotheses or background assumptions. The idea is that a scientific theory/hypothesis does not by itself carry any implication about the evidence we should expect in case it is true. Rather, it has testable implications only given certain background assumptions (e.g. claims about the workings of scientific instruments and claims about the relationship between the objects of study and their environment). Holist underdetermination occurs when, due to the holistic nature of confirmation, recalcitrant evidence (e.g. a failed prediction) is not sufficient to determine which claim we should revise in light of it: in the case of an experiment, should it be the hypothesis we wanted to test in the first place or one of the auxiliary assumptions we used to derive the testable consequence?

Contrastive underdetermination is so-called because it questions the ability of the evidence to confirm any given hypothesis/theory *against* its alternatives. Contrastive

underdetermination occurs when the evidence confirms more than one theory/hypothesis equally well. When this is the case, the evidence is not sufficient to determine what one should believe. To give an example (due to Stanford 2017), consider one of the motivations for the familiar methodological insight that correlation does not imply causation. Suppose you observe a strong correlation between F and G. If F causes G, you should expect a strong correlation between F and G. Should you then believe that F causes G? That would be hasty. If G causes F, you should also expect a strong correlation between F and G. Further, you should expect a strong correlation between F and G if F and G are caused by some third factor or if there is a bidirectional causal relationship. Other things being equal, the strong correlation confirms each causal hypothesis equally well. Thus, in this sort of situation, the evidence alone underdetermines what one should believe about the causal relationship between F and G.

Holist and contrastive underdetermination are intimately related, but contrastive underdetermination seems more fundamental in the following sense: all cases of holist underdetermination can be redescribed as cases of contrastive underdetermination, but not the other way around.

Here is why. Suppose a hypothesis H/theory T is h-underdetermined with respect to some piece of disconfirming evidence *e*. Since H/T is h-underdetermined with respect to *e*, *e* is not sufficient to determine whether we should revise H/T or one of the auxiliary hypotheses A_1, A_2, ..., A_n tested in conjunction with H. But that is the same as saying that *e* is not sufficient to determine whether we should prefer {~H/~T, A_1, A_2, ..., A_n} or {H/T, ~A_1, A_2, ..., A_n} or {H/T, A_1, ~A_2, ..., A_n} etc. That is: it is the same as saying that {~H/~T, A_1, A_2, ..., A_n} or {H/T, ~A_1, A_2, ..., A_n} or {H/T, A_1, ~A_2, ..., A_n} etc. are c-underdetermined with respect to *e*. Thus, holist underdetermination can be cashed out in terms of contrastive underdetermination.

However, contrastive underdetermination can occur even in the absence of holist underdetermination, because contrastive underdetermination does not require the truth of confirmation holism. This is illustrated by a well-known example in the literature. For any finite set of points on the Cartesian plane, there is an infinite number of functions describing different curves that fit all the points in the set. Adding new points to the initial set rules out an infinite number of functions that previously fitted all the points. However, notice that, no matter how many points we add to the initial set, as

long as the new set is finite there will always be an infinite number of functions that define curves fitting all the points.

If we interpret points as the evidence and functions defining curves as the competing hypotheses, we have a simple model of contrastive underdetermination without holist underdetermination. For, in this setup, each new point we add to the finite set of points, i.e. each new piece of evidence we acquire, is sufficient to determine which functions, i.e. which hypotheses, should be discarded. Thus, the evidence is sufficient to determine which hypotheses should be given up in light of it and no holist underdetermination occurs. However, each new point we add to the initial set also leaves an infinite number of functions uneliminated. Thus, the evidence is not sufficient to determine which hypothesis should be picked out for belief among the available ones and contrastive underdetermination occurs.

More generally, even if one knows which hypothesis to blame in light of new disconfirming evidence, there might still be some alternative hypotheses equally well confirmed by the evidence. Thus, although one is not subject to holist underdetermination, one might still be subject to contrastive underdetermination.

Underdetermination claims can vary in strength: our beliefs can be *transiently* underdetermined by the *actual* evidence*,* or they can be *permanently* underdetermined by *all possible* evidence (Stanford 2017). When two theories/hypotheses are transiently underdetermined by the present evidence, the existing evidence confirms them equally well, but some future evidence might support one better than the other. That is: a theory/hypothesis is transiently underdetermined by the evidence just in case there is some alternative hypothesis/theory that is confirmed equally well by the existing evidence despite their being empirically inequivalent (Sklar 1975).

In the philosophy of science, discussions of underdetermination have focused on the search for and assessment of empirical equivalent rivals as a means to establish permanent underdetermination. However, empirical equivalence and permanent underdetermination should not be conflated; empirical equivalent rivals provide only one path to permanent underdetermination. In fact, permanent underdetermination can be achieved by means of transient underdetermination as long as the latter is recurrent (Sklar 1981; Stanford 2001).

1.4 The Underdetermination Principle

Underdetermination scepticism turns on the idea that evidence does not confirm our beliefs *against* some doxastic alternatives. Thus, underdetermination scepticism turns on the supposed contrastive underdetermination of our beliefs. To reach the sceptical conclusion, the sceptic invokes an underdetermination principle. The principle says that if the evidence equally supports two incompatible beliefs/propositions, we are not justified in believing either. The sceptic might as well arrive at the conclusion that knowledge is impossible, if justification is necessary for knowledge, as I shall assume henceforth. At any rate, to the extent that it poses a genuine challenge, underdetermination scepticism about justification seems more fundamental as well as more central to our concerns than underdetermination scepticism about knowledge, because underdetermination is an obstacle to knowledge insofar as it is an obstacle to justification (Vogel 2004: 432) and because 'we can live with the concession that we do not, strictly, *know* some of the things we believed ourselves to know, provided we can retain the thought that we are fully justified in accepting them' (Wright 1991: 88; emphasis in the original).

It is easy to appreciate the appeal of the underdetermination principle. The principle captures the sane-sounding idea that having justification for belief in P requires having a rational basis that favours P over competing hypotheses. The idea can be expressed in multiple ways: that justification for belief in P requires being in a position to non-arbitrarily reject competitors to P, or that if you are justified in believing that P then P has more epistemic merit than its rivals (Vogel 2004: 427). In each case, the principle seems unimpeachable.[19]

[19] One might think that, as I have formulated it, the underdetermination principle is too strong, because it requires that one's beliefs be superior to any rival belief. A better version of the principle would restrict the requirement to competitors that are relevant or proper to consider, like this: if one is justified in believing that P, the evidence supports P over its *relevant* rivals. Add the assumption that sceptical hypotheses are not relevant competitors to our beliefs, and the amended principle would block the sceptical reasoning. This anti-sceptical strategy is similar to some famous attempts to counter closure scepticism (Dretske 1970; DeRose 1995; Lewis 1996). I shall put aside this way of responding to underdetermination scepticism. Although I cannot do full justice to these accounts here, the reason is this: they offer putative solutions to the sceptical problem that do not address its main root, namely the idea that our beliefs are underdetermined by the evidence. If our beliefs are underdetermined by the evidence, insisting that our beliefs are justified because sceptical hypotheses are irrelevant seems like a desperate

Further, the underdetermination principle seems to underpin our (sound) judgement in a vast range of cases (Vogel 2004: 427; Pritchard 2015: 31). Consider for instance the following situation: you witness a hit-and-run accident. A cab hits a pedestrian and leaves. It is dark and the accident happens fast. Your visual experience does not favour the cab's being green over the cab's being blue, but you know that the cab cannot be both. Assume this is all the information you have about the accident. Are you justified in believing that the cab is, say, green? I do not think you are. The underdetermination principle explains why: since you have no rational basis to favour the cab's being green over the cab's being blue, you are not justified in believing that the cab is green.

The underdetermination principle ties the justification of a belief to the degree of epistemic support or confirmation. However, it does not do so by determining the minimal degree of confirmation a belief must enjoy to count as justified. Rather, it does so by putting a constraint on the belief's comparative degree of support relative to its rival beliefs. As Yalçin (1992: 11) observes, this is a very weak constraint. It is compatible with the sufficiency of any absolute degree of support for justification. For all the underdetermination principle says, even a proposition that has very low epistemic support can be justifiably believed.[20] Further, it is compatible with the sufficiency of any degree of superiority over rival beliefs for justification. For all the underdetermination principle says, even a belief that edges its competitors in confirmation status by a very small margin can be justified.[21]

For all its plausibility, the underdetermination principle is not uncontroversial. Vahid (2005: 136-7) argues that the underdetermination principle only appears plausible because it is confused with a plausible but distinct principle: that if the evidence for believing that P does not favour P over a competitor belief, then we are not justified in believing that P *rather than* the competitor belief is true. The confusion is easily explained by the fact that the two principles are superficially similar. In other words,

---

move. To the extent that our beliefs are underdetermined by the evidence, they are rationally defective, and it is that very defect that should bothers us (Cf. Feldman 1999; Kornblith 2000; Pritchard 2015).

[20] Although, as we will see, one of the better arguments for the underdetermination principle exploits the idea that believing something requires believing that it is more likely than not.

[21] In fact, it is plausible that to be justified a belief must edge its competitors by a substantial margin.

according to Vahid, underdetermination is an obstacle to the justification of believing the comparative content 'P *rather than* the competitor belief is true' but is compatible with the justification of believing the non-comparative content 'P is true.'

Vahid's contention is grounded in the rejection of the following principle: that evidence cannot justify incompatible hypotheses (for more on this principle see chapter 2, sections 2.2.2 and 2.3.1). This amounts to a rejection of the Uniqueness Thesis, the claim that 'a body of evidence justifies at most one proposition out of a competing set of propositions (e.g. one theory out of a bunch of exclusive alternatives) and it justifies at most one attitude toward any particular proposition' (Feldman 2007: 205).

It is easy to see why rejecting the uniqueness of justification paves the way to rejecting the underdetermination principle. The underdetermination principle ties a lack of evidential favouring to a lack of justification. However, notice that, by itself, the fact that the evidence does not favour P over one of its competitors does not entail that the evidence does not justify a belief in P. After all, S's evidence might justify both P *and* its competitor. Thus, the underdetermination principle seems to presuppose the truth of the Uniqueness Thesis.

Vahid's case against the uniqueness of justification rests on a series of (putative) counterexamples: cases in which the evidence at hand seems to justify a pair of incompatible hypotheses. Vahid tries to boost the reader's intuition by considering confirmation/evidential support first. Clearly, the same body of evidence can confirm a pair of incompatible hypotheses. If you know that one of John, Jim, and Jill is a spy, the observation that Jill is not a spy raises the likelihood of the two incompatible hypotheses 'John is a spy' and 'Jim is a spy' (Vahid 2005: 131). This should help the reader see that the same is true when evidence raises the likelihood of the two incompatible hypotheses to a degree sufficient for justification.

I find Vahid's counterexamples to the uniqueness of justification unconvincing. Moreover, the idea that evidence can raise the likelihood of two *incompatible* hypotheses to a degree sufficient for justification is probabilistically incoherent, given the plausible constraint that believing something requires believing that it is more likely than not. For if the evidence raises the likelihood of one hypothesis over .5, the likelihood of the other must drop below .5.

I cannot settle the status of the Uniqueness Thesis here.[22] Instead, I shall assume that it is true. However, I want to suggest that rejecting the uniqueness of justification does not provide us with a satisfactory solution to the problem of underdetermination scepticism. For suppose that the evidence underdetermines our belief that P relative to some sceptical competitor. If evidence can justify incompatible hypotheses, then either (a) the evidence justifies neither P nor its sceptical competitor or (b) the evidence justifies both. If (a), scepticism is true. Thus, rejecting uniqueness while endorsing (a) does not solve the problem of underdetermination scepticism. If (b), scepticism is false. However, if the evidence justifies both P and its sceptical competitor, our beliefs are on a par with beliefs in sceptical hypotheses. This raises two problems. First, there seems to be a lingering sense in which our beliefs are arbitrary.[23] Thus, there seems to be a lingering sense in which underdetermination scepticism is true after all. Second, most anti-sceptics believe not only that our beliefs are justified, but that they are superior to beliefs in sceptical hypotheses. Thus, rejecting uniqueness while endorsing (b) does not provide a satisfactory solution to the problem of underdetermination scepticism.

1.5 The Underdetermination Argument

I can now bring together the points presented in sections 1-4 and introduce the underdetermination argument that I shall discuss in the rest of this thesis. Let S be a cognitive agent like us, and let P be a proposition about the external world, like the proposition that S is sitting under a tree, that it is a dry and lovely morning, etc. Let the good case be a situation in which it appears to S that P is true, and P is true. Let the bad case be a situation in which it appears to S that P is true, just like it does in the good case, but P is false. The good case and the bad case are phenomenologically indistinguishable: S is in exactly the same non-factive mental states in the two cases, and S has an experience in one case if and only if it has the same experience in the other. On the assumption that S's perceptual evidence is exhausted by (a subset of) her non-factive mental states, this entails that S has the same evidence in the two cases.

[22] See Kopec and Titelbaum (2016) for an overview of the debate.

[23] This problem is related to a much-discussed objection to Permissivism by White (2005).

This is the Sameness of Evidence Lemma:

(SEL) S has the same evidence in the good and the bad case.

The underdetermination sceptic infers from the Sameness of Evidence Lemma that S's evidence underdetermines the doxastic choice between P and the proposition that S is in the bad case (I discuss the exact nature of this inference in chapter 3). In other words, according to the sceptic, S's evidence is neutral between the two incompatible propositions, it supports them equally well. Thus:

(~FAV) S's evidence does not favour P over her being in the bad case.

Notice that, strictly speaking, ~FAV is not an underdetermination claim but a consequence thereof. After all, 'S's evidence does not favour P over her being in the bad case' is true even if S's evidence favours her being in the bad case over P, but 'S's evidence supports P and her being in the bad case equally well' is not. In this context, the difference between ~FAV and the corresponding underdetermination claim does not matter because ~FAV follows from the underdetermination claim. In fact, I shall often talk as if ~FAV were an underdetermination claim, although, strictly speaking, it is only a consequence thereof. The simplification is harmless and in line with current practice in the literature.

At this point, the underdetermination sceptic invokes a general principle that ties lack of justification and underdetermination (or, more precisely, lack of evidential favouring):

For any incompatible propositions φ, ψ, if S is justified in believing that φ, then S's evidence favours φ over ψ.

Since, by definition of the bad case, P is incompatible with S's being in it, the following instance of the principle is true:

(UP) If S is justified in believing that P, S's evidence favours P over her being in the bad case.

Putting together UP and ~FAV, the sceptic derives the sceptical conclusion:

(SCEPT) S is not justified in believing that P.

Since in this argument 'S', 'P', and 'evidence' are just placeholders for any combination

of a cognitive agent like us, a proposition about the external world, and a total body of evidence, ~FAV expresses a claim of permanent underdetermination, and SCEPT is the same as a modally strong version of justification scepticism about the external world, i.e. the claim that justification of external world propositions is impossible for creatures like us.

Notice that the issue with this kind of scepticism is not that our evidence favours our beliefs over the incompatible sceptical hypotheses albeit not enough for them to be justified. Rather, underdetermination scepticism is similar in spirit to Kornblith's 'Full-Blooded Scepticism' (2000: 25) – the view that we have no degree of justification at all for our beliefs – and to Rinard's 'Evidential Scepticism' (2022: 435)– the view that we have no good evidence for our beliefs – in that we lack justification because our evidence gives no support whatever to our beliefs relative to the competing sceptical hypotheses. Thus, our beliefs are not even marginally confirmed by our evidence against their alternatives.

1.6 Scepticism, Paradox, and Self-Refutation Charges

Underdetermination scepticism presents us with a compelling piece of reasoning that starts with a set of plausible premises and ends with an implausible conclusion. Thus, underdetermination scepticism has a paradoxical character (Stroud 1984; Wright 1985; Fumerton 1995; Pritchard 2015). This feature of the sceptical problem constrains what counts as an adequate solution of the problem of underdetermination scepticism. It is often noted that to solve the paradox it is not sufficient to blame one of the Underdetermination Argument's premises/inferential steps. Since the argument has an air of plausibility, any countermove requires some motivation. Further, to be rationally compelling, a countermove would need to explain why the contested claim/inference appears persuasive despite its being mistaken.

Here, I want to suggest that the paradoxical nature of underdetermination scepticism puts an additional constraint on what counts as an adequate solution to the problem of underdetermination scepticism. In particular, I want to argue that a family of approaches that aim to show that underdetermination scepticism is rationally self-undermining is less promising that it initially appears. To make my case against this family of approaches I will use an anti-sceptical argument due to Rinard (2018) as an example.

Say that you endorse scepticism on the basis of the Underdetermination Argument if either the Underdetermination Argument is your reason for believing that scepticism is true, or the Underdetermination Argument is your reason for suspending judgement on scepticism. Arguably, the Underdetermination Argument puts us under some pressure to endorse scepticism (Mooreans disagree; I discuss Mooreanism in chapter 5). But is it rational *all things considered* to endorse scepticism on the basis of the Underdetermination Argument? Rinard (2018) presents an argument for the claim that rationality requires us (sceptics included) to disbelieve certain forms of scepticism. In fact, if Rinard is right, it is not rational to endorse scepticism *on the basis of the Underdetermination Argument*.

To put it simply, Rinard's argument consists of three parts. In the first part, Rinard argues that accepting a certain argument for external world scepticism is irrational, because accepting that argument commits one to accept other forms of scepticism in a self-undermining way. In the second part, she argues on similar grounds that suspending judgement on the basis of that argument is irrational, too. In the third part, she argues that situations in which rationality prohibits every doxastic attitude are impossible. Thus, Rinard concludes, rationality requires disbelieving scepticism.

To elaborate: in the first part of her argument, Rinard makes the case that accepting a certain argument for external world scepticism is irrational, because accepting that argument commits one to accept other forms of scepticism in a rationally self-undermining way. More precisely, accepting that argument commits one by parity of reasons to accept a parallel argument for scepticism about the past. Further, scepticism about the past leads to scepticism about complex reasoning. However, Rinard notes, it would be rationally self-undermining to accept scepticism about complex reasoning on the basis of this chain of reasoning, because this chain of reasoning is itself complex. Thus, Rinard concludes, it is not rational to accept the initial argument for external world scepticism.

Notice that this is not sufficient to establish that rationality requires disbelieving external world scepticism. After all, rationality might require or allow for suspending judgement on scepticism rather than disbelieving it. With the second and third part of the argument, Rinard attempts to fill the gap between the irrationality of accepting scepticism and the rational requirement to disbelieve it.

In the second part of the argument, Rinard argues that suspending judgement on scepticism[24] is irrational, too, because it leads to the violation of certain rationality requirements. Rinard distinguishes between two ways of suspending judgement on scepticism: confident and unconfident suspending. While the confident suspender suspends judgement on scepticism (including scepticism about complex reasoning) and believes that it is rational to do so, the unconfident suspender suspends judgement on scepticism as well as on whether it is rational to do so. In each case, suspending judgement leads to problems similar to the ones noted for accepting scepticism. On the one hand, the confident suspender believes that it is rational to suspend judgement on scepticism *based on complex reasoning* while suspending judgement on scepticism about complex reasoning. On the other hand, the unconfident suspender suspends judgement on scepticism about the external world while suspending judgement on whether they ought to do so. In each case, the suspender is irrational because rationality requires agents to endorse their doxastic attitudes in the sense of believing them to be rational. Thus, it is not rational to suspend judgement on external world scepticism.

Since there are only three possible doxastic attitudes – belief, disbelief, and suspension of judgement – it might seem straightforward to derive that rationality requires disbelieving scepticism from the fact that rationality prohibits believing and suspending judgement on it. However, the derivation is unproblematic only on the assumption that the sceptical problem is not an instance of a doxastic dilemma, a situation in which rationality prohibits every doxastic attitude.

In the third part of her argument, Rinard concludes her case by arguing for the impossibility of doxastic dilemmas. Rinard's argument appeals to the idea that 'rationality constitutes an ideal to which one could coherently aspire, and by which one could be guided' (Rinard 2018: 259). If doxastic dilemmas were possible, rationality could not play its guidance-giving role, for in dilemmatic situations rationality could give us no practicable advice. Since its guidance-giving role is constitutive of rationality, doxastic dilemmas are impossible. And since doxastic dilemmas are impossible, and

[24] Henceforth, for reasons of readability, I leave the qualification 'on the basis of a certain argument' implicit.

rationality prohibits believing and suspending judgement on scepticism, rationality requires disbelieving scepticism.

Although Rinard's argument does not automatically generalise to every sceptical argument for external world scepticism (Rinard 2018: 241), it does apply to the Underdetermination Argument. For, by parity of reasons, if it is rational to accept the Underdetermination Argument for external world scepticism, it is rational to accept the Underdetermination Argument for scepticism about the past, too. More generally, Rinard's argument applies to any argument for scepticism that, by parity of reasons, can be extended to scepticism about the past.

Rinard's argument is interesting in several ways. On the one hand, her argument purports to appeal only to premises a sceptic could accept. Thus, if Rinard is right, it is possible to rationally persuade a sceptic that scepticism is false, contrary to a widespread opinion in contemporary epistemology (Pryor 2000: 517; Williamson 2000: 27; Byrne 2004: 301 ; Bergmann 2021: 145-6). Relatedly, the success of Rinard's argument would give new life to attempts to rationally persuade the sceptic. On the other hand, Rinard's argument does not attempt to identify the flaw in the sceptic's argument and explain its appeal. Rather, Rinard's argument brings out the irrationality of scepticism without identifying its source.

However, it should be noted that Rinard's argument has some limitations. For one thing, it is not obvious that Rinard's argument shows that underdetermination scepticism is rationally self-undermining. As Rinard grants, her argument relies on several controversial claims: claims about the role of memory in complex reasoning, about the rationality of certain types of conflict between first-order and higher-order attitudes, and about the possibility of doxastic dilemma (Rinard 2018: 241). Although these are claims that the sceptic *could* in principle endorse, they are not part of the Underdetermination Argument, and they are not claims that the sceptic *must* endorse. Thus, even granting that Rinard has shown that the Underdetermination Argument together with a number of controversial assumptions leads to a rationally self-undermining position, her argument does not show that underdetermination scepticism is self-undermining.

For another thing, even if Rinard's argument showed that underdetermination scepticism is rationally self-undermining, it is not clear what one should infer from this. According to Rinard, one should disbelieve underdetermination scepticism – in other

words, the finding that scepticism is rationally self-undermining is comforting. But that rosy view is not the only option.

On the contrary: if underdetermination scepticism presents us with a paradox, the fact that scepticism is also rationally self-undermining should be discomforting to us. For if we cannot see where the Underdetermination Argument goes wrong or how we can do without the assumptions that contribute to the paradox, we are left with the additional problem that our assumptions lead to scepticism *and* to a rationally self-undermining position. So, far from making things better, simply pointing out that underdetermination scepticism is self-undermining seems to make things worse for us.

The point can be generalised to any attempt to resist the Underdetermination Argument via showing that underdetermination scepticism has some unpalatable features, like being rationally self-undermining. The problem with such attempts is that they ignore the fact that, since the Underdetermination Argument seems very compelling, the problematic feature can become an integral part of the sceptical paradox. After all, it was clear all along that the Underdetermination Argument has an unpalatable consequence, namely scepticism. So, why should the fact that it has one further unpalatable consequence, namely making one's position rationally self-undermining, make the problem of underdetermination scepticism less rather than more acute?[25]

1.7 The Plan

The remainder of this thesis comprises four chapters and a conclusion. The thesis can be read as an argument for underdetermination scepticism, with chapters 1 and 2 focusing on the Underdetermination Argument as a whole and on the underdetermination principle, chapters 3 and 4 focusing on the underdetermination claim and the Sameness of Evidence Lemma, and chapter 5 offering a methodological reflection on the persuasiveness of the argument. At the same time, the thesis can be read as a collection of essays exploring the connections between scepticism and some important topics in epistemology, and as a contribution to the existing literature on those topics.

Chapter 2 explores the relationship between the Underdetermination Argument

[25] For similar points on the anti-sceptical import of self-undermining charges to scepticism see Wright (1985: 440-41) and Fumerton (1995: 50-1). For additional criticism of Rinard's argument, see Bergmann (2021: 59-60, 69-73).

and the Closure Argument for scepticism from within the evidentialist framework for justification described in section 1.2.1. The chapter throws light on whether and to what extent the two arguments are distinct sources of scepticism. I argue that the Underdetermination Argument and the Closure Argument depend on each other in the following sense: the falsity of one argument's premises makes the other argument unsound. More precisely, I argue that the premises of the Underdetermination Argument entail the premises of the Closure Argument, and that the premises of the Closure Argument entail one of the premises of the Underdetermination Argument but not the other. In this context, the most important result is that, on a probabilistic construal of evidential favouring and justification, the underdetermination principle and the closure principle entail each other. Further, I argue that the Underdetermination Argument and the Closure Argument are independent from each other in another sense: their premises can be plausibly motivated without an appeal to the other argument. The upshot is a new account of the logical and dialectical relationship between the Underdetermination Argument and the Closure Argument, according to which they are distinct sources of scepticism. Even so, it turns out that, given the evidentialist conception of justification I assumed, the Underdetermination Argument and the Closure Argument are closely related arguments and that they have sceptical consequences in exactly the same circumstances.

Chapter 3 discusses what I call 'the Infallibility Objection' and vindicates a conception of underdetermination scepticism as an interesting form of scepticism: one that is distinct from forms of scepticism that rely on very demanding epistemic standards, such as certainty and infallibility requirements. According to the Infallibility Objection, the Underdetermination Argument is a bad argument because the sceptical inference from sameness of evidence in the good and the bad case to underdetermination presupposes infallibilism. According to proponents of the Infallibility Objection, the sceptic's inference is illegitimate because sameness of evidence does not entail evidential parity: S's evidence can favour one of two incompatible hypotheses although both are compatible with the evidence. After criticizing two attempts to sidestep the Infallibility Objection found in the literature, I argue that, pace critics, the contentious sceptical inference does not presuppose infallibilism. Proponents of the Infallibility Objection have failed to recognise the reasoning pattern that underpins the sceptical inference. I

provide a rational reconstruction of that inference along Bayesian lines and defend it against two objections. The idea is that S's evidence does not favour P over SH – the rival sceptical hypothesis – because S has the same evidence regardless of which hypothesis is true (and hence: Pr(E|P)=Pr(E|SH)) *and* P and SH have the same prior probability.

Chapter 4 assesses the anti-sceptical import of Williamson's anti-luminosity argument and of anti-luminosity considerations more generally. According to Williamson, the sceptic argues from the claim that S's evidence is luminous – that S is always in a position to know what her evidence is – to the claim that S has the same evidence in the good and the bad case. However, per the anti-luminosity argument, no non-trivial condition is luminous. I argue that Williamson's anti-sceptical strategy fails for two reasons. First, I offer a new argument to the effect that the anti-luminosity argument does not undermine the sceptic's reasoning. The key idea for my argument is that luminosity failures of the kind highlighted by the anti-luminosity argument are irrelevant to failures of knowledge in the bad case. Second, Williamson's anti-sceptical strategy hinges on the mistaken idea that the sceptic needs a luminosity claim to defend the Sameness of Evidence Lemma. I present two arguments against this idea. The first argument appeals to evidence mentalism, the view that S's evidence is determined by S's non-factive mental states. The second argument appeals to comparative judgements about S's rationality in the good and the bad case. The upshot is that the sceptic's case for the Sameness of Evidence Lemma is stronger as well as more resilient to challenges than Williamson envisioned. In particular, the sceptic's case for the Sameness of Evidence Lemma is independent of evidence luminosity.

Chapter 5 discusses Moorean responses to the Underdetermination Argument. Moorean anti-sceptics believe that one should not abandon one's common-sense beliefs – like the belief that one's beliefs are justified – in response to sceptical arguments, because one's common-sense beliefs are epistemically superior to the premises of any philosophical argument to the contrary. Crucially, they think this is true even if one cannot identify any flaws in such arguments. In doing so, Moorean anti-sceptics expose themselves to the charge of dogmatism. After all, if one can find no flaws in an argument, how can one's refusal to accept that argument be anything other than dogmatic? In this chapter, I evaluate Mooreanism as a way to respond to the Underdetermination

Argument for scepticism in light of the dogmatism charge. I tackle Mooreanism from three angles, which correspond to three ways of motivating Mooreanism. First, I discuss attempts to motivate Mooreanism by appeal to an alleged difference between philosophy and science. Second, I discuss attempts to motivate Mooreanism via general principles about belief revision. Finally, I discuss Kelly's attempt to vindicate Mooreanism via an argument for its methodological superiority over the methodology employed by the sceptic. I shall argue that none of these motivations for Mooreanism succeeds; the upshot is the Mooreanism provides no sound basis to resist the Underdetermination Argument.

The conclusion reviews, generalizes and puts in context the results reached in the previous chapters. The upshot is that a significant set of objections against underdetermination scepticism fails. At the end of the day, we might have to take the possibility of living with scepticism seriously – or at least more seriously than we thought.

2

Underdetermination Scepticism and Closure Scepticism

Consider the following arguments:

> *Underdetermination Argument (UA)*
>
> (~FAV) S's evidence does not favour P over SH.
>
> (UP) If S is justified in believing that P, S's evidence favours P over SH.
>
> (SCEPT) S is not justified in believing that P. [from ~FAV and UP]

> *Closure Argument (CA)*[1]
>
> (~J) S is not justified in believing that ~SH.
>
> (CP) If S is justified in believing that P, then S is justified in believing that ~SH.
>
> (SCEPT) S is not justified in believing that P. [from ~J and CP]

As usual, S is a cognitive agent like us, and P is a proposition about the external world, like the proposition that S has hands. The good case is a situation in which it appears to S that P is true, and P is true; while the bad case is a situation in which it appears to S that P is true (just like it does in the good case), but P is false. SH is the proposition that S is in the bad case.

The Underdetermination Argument and the Closure Argument resemble each other. For one thing, they share their conclusion, a modally strong version of justification scepticism about the external world. For another thing, they rely on the idea that S is in a weak epistemic position relative to sceptical scenarios, situations in which things are not what they appear to be, and in which S cannot tell that this is so. But, at first glance, the two arguments also differ in important ways, since UP and CP seem

---

[1] The Closure Argument is so-called because CP receives its support from a closure principle. Since, by definition of the bad case, P entails ~SH, CP can be derived from the following closure principle by substituting P for φ and ~SH for ψ:

> For all φ, ψ, if S is justified in believing that φ, and φ entails ψ, then S is justified in believing that ψ.

The principle captures the idea that justification transmits across entailment. How to formulate the closure principle is a controversial matter – over the years, several counterexamples have been put forward against specific formulations of the closure principle, leading to increasingly sophisticated formulations. See Luper (2020) for an overview of this debate. However, the exact formulation of the closure principle does not matter here.

to express different principles. While UP captures the idea that justification is incompatible with contrastive underdetermination, CP captures the idea that a justification for a proposition is also a justification for its consequences.

Appearances notwithstanding, the relationship between the Underdetermination Argument and the Closure Argument is a contentious matter. Some argue that the Underdetermination Argument and the Closure Argument are distinct sources of scepticism. Among them, some argue that one argument depends on the other (Cohen 1998), while others deny this (Pritchard 2015). Yet others can be read as arguing that the two arguments come down to the same argument (Brueckner 1994).

The issue is rich of consequences. If the Underdetermination Argument and the Closure Argument are different arguments, they should be dealt with separately. Yet, if one argument depends on the other, refuting one argument could be sufficient for refuting the other. And if the Underdetermination Argument and the Closure Argument are the same argument, the two arguments stand and fall together.

Further, the overall case for scepticism, or just for one of the two arguments, will be stronger or weaker depending on whether the Underdetermination Argument and the Closure Argument provide independent motivations for scepticism. Epistemologists should be sensitive to these issues. Undoubtedly, of the two arguments, the Closure Argument has so far received the lion's share of attention. But if it turned out that the Closure Argument and the Underdetermination Argument are independent arguments, or if it turned out that the Closure Argument depends on the Underdetermination Argument in some important way, epistemologists ought to start paying more attention to the latter.

This chapter explores the relationship between the Underdetermination Argument and the Closure Argument from within the evidentialist framework for justification described in chapter 1 – a framework in which justification is sufficient degree of evidential support that obeys probabilistic constraints. The goal is to throw light on whether and to what extent the Underdetermination Argument and the Closure Argument are distinct sources of scepticism. To answer these questions, I shall rephrase them in terms of two more manageable questions about soundness and cogency. I start by looking at the relationship between the soundness of the Underdetermination Argument and the soundness of the Closure Argument, as doing so helps answering

questions about their cogency. Then, I discuss the relationship between their cogency. The upshot is a new account of the logical and dialectical relationship between the two arguments.

Here is the plan. Section 1 sets things up by introducing the notions of soundness and cogency and the questions that the chapter aims to answer. Section 2 answers the question about the relationship between the soundness of the Underdetermination Argument and the soundness of the Closure Argument. I shall argue that although the premises of the Underdetermination Argument entail ~J (i.e. premise 1 of the Closure Argument), the Closure Argument's premises do not entail ~FAV (i.e. premise 1 of the Underdetermination Argument). Then, I look at the controversial relationship between the underdetermination principle and the closure principle and argue that, on the evidentialist picture of justification presupposed here, they entail each other. Section 3 answers the question about the relationship between the Underdetermination Argument's and the Closure Argument's cogency. I argue that, in this respect, the two arguments are independent of each other. Section 4 elaborates these findings, teases out their consequences, and puts them in context to provide a new account of the Underdetermination Argument and Closure Argument's relationship.

2.1 Preliminaries: Soundness and Cogency

I start with some preliminary work on the idea that a deductive argument A depends on a deductive argument B. Typically, when we evaluate deductive arguments, we are interested in two features: soundness and cogency. A deductive argument is *sound* just in case all its premises are true and it is valid. And a deductive argument is valid just in case the premises entail its conclusion. In contrast, a deductive argument is *cogent* just in case its premises are plausible.

Unlike soundness, cogency comes in degrees: an argument can be more or less cogent, but an argument cannot be more or less sound. Further, soundness and cogency may come apart: a cogent argument may be unsound, and a sound argument may not be cogent. In the best-case scenario, the two go hand in hand: sound arguments are cogent and vice versa.

Deductive arguments can relate to each other with respect to both soundness and cogency. Say that an argument A *depends for its soundness* (or s-depends) on an argument

B just in case the soundness of A depends on the truth of B's premises. I shall say that the soundness of A depends on the truth of B's premises just in case there is a premise of B whose falsity makes A unsound.[2] Accordingly, A is independent of B for its soundness just in case A can be sound although all of B's premises are false.[3]

Say that an argument A *depends for its cogency* (or c-depends) on an argument B just in case the plausibility of A's premises depends on the plausibility of B's premises. In turn, the plausibility of a premise of A depends on the plausibility of a premise of B just in case A's premise cannot be plausibly motivated without an appeal to B's premise. Accordingly, A is independent of B for its cogency just in case A's premises can be plausible although B's premises are not.

2.1.1 The Soundness Question

The question that concerns me in this chapter is whether the Underdetermination Argument and the Closure Argument are distinct sources of scepticism. We can use the notions of s-dependence and c-dependence to rephrase this question in terms of two simpler questions and make it more precise.

Consider the *Soundness Question*:

> What is the relationship of s-dependence between the Underdetermination Argument and the Closure Argument?

Recall that an argument A is s-independent of an argument B just in case A can be

---

[2] Notice that saying that there is a premise of B whose falsity makes A unsound is not the same as saying that A is sound only if B is sound. 'A is sound only if B is sound' is equivalent to 'If B is unsound, A is unsound'. However, that there is a premise of B whose falsity makes A unsound does not entail that A is unsound whenever B is. As an example, consider the following pair of argument schemes:

(A*) Fa/∴ Fa v Gb
(B*) ∃xFx/ ∃xFx → ∃xGx /∴ ∃xGx

Although there is a premise of B* whose falsity makes A* unsound, namely ∃xFx, it is false that A* is sound only if B* sound. After all, B* might be unsound because of the falsity of ∃xFx → ∃xGx, which has no bearing on the soundness of A*.

[3] There is another way for the soundness of A to depend on the soundness of B. The soundness of A may depend on the truth of B's premises in the sense that the *truth* of B's premises would make A unsound. If this were the case, then A and B could not be sound at the time. Although I am not concerned with this type of relationship between the Underdetermination Argument and the Closure Argument here, what I say about the Underdetermination Argument and the Closure Argument in section 2.2 vindicates the claim that the two arguments can be sound at the same time.

sound although all of B's premises are false. Thus, the Soundness Question asks whether the Underdetermination Argument and the Closure Argument can be sound irrespective of the falsity of the other argument's premises.

A natural way to answer the Soundness Question is to check whether the premises of the two arguments entail each other, because the soundness of A depends on the truth of B's premises just in case A's premises entail some of B's premises. Therefore, the Soundness Question admits of four mutually exclusive and jointly exhaustive answers:

> (SQ1) The Underdetermination Argument and the Closure Argument s-depend on each other (mutual dependence).
> (SQ2) The Underdetermination Argument s-depends on the Closure Argument, but not vice versa (non-mutual dependence).
> (SQ3) The Closure Argument s-depends on the Underdetermination Argument, but not vice versa (non-mutual dependence).
> (SQ4) The Closure Argument and the Underdetermination Argument do not s-depend on each other (independence).

Since the Underdetermination Argument and the Closure Argument have multiple premises each, the dependence relationship in SQ1, SQ2, and SQ3 can be realised in multiple ways. SQ4 makes exception, because there is only one possible configuration with no entailment between the premises of the two arguments. SQ1 includes the case in which all the premises of the Underdetermination Argument entail all the premises of the Closure Argument and vice versa as well as cases in which the entailment relationship is not total. SQ2 and SQ3 comprise cases of non-mutual dependence, where the dependence relationship is asymmetrical.

Notice that to have a picture of s-dependence between two arguments, it might not be sufficient to look at the entailment of each premise taken individually. Instead, one might have to look at the entailment of their conjunction, too. The reason is that the fact that no single premise of A entails a premise of B does not establish the s-independence of A from B. After all, A's premises might jointly entail some of B's premises.

2.1.2 The Cogency Question

A lot of this can be applied *mutatis mutandis* to the other question we can ask about the Underdetermination Argument and the Closure Argument's relationship, the *Cogency Question*:

> What is the relationship of c-dependence between the Underdetermination Argument and the Closure Argument?

Recall that A c-depends on B just in case the plausibility of some of A's premises depends on the plausibility of some of B's premises. Thus, the Cogency Question asks whether the Underdetermination Argument and the Closure Argument can be motivated independently of each other.

Like the Soundness Question, the Cogency Question admits of four mutually exclusive and jointly exhaustive answers:

> (CQ1) The Underdetermination Argument and the Closure Argument c-depend on each other (mutual dependence).
> (CQ2) The Underdetermination Argument c-depends on the Closure Argument, but not the other way around (non-mutual dependence).
> (CQ3) The Closure Argument c-depends on the Underdetermination Argument, but not the other way around (non-mutual dependence).
> (CQ4) The Underdetermination Argument and the Closure Argument do not c-depend on each other (independence).

To summarise, we can figure out the answer to the question 'Are the Underdetermination Argument and the Closure Argument distinct sources of scepticism?' by figuring out the answer to these two questions:

> (*Soundness Question*) What is the relationship of s-dependence between the Underdetermination Argument and the Closure Argument?
>
> (*Cogency Question*) What is the relationship of c-dependence between the Underdetermination Argument and the Closure Argument?

2.2 Answering the Soundness Question

In this section, I answer the Soundness Question: what is the relationship of s-

dependence between the Underdetermination Argument and the Closure Argument? I shall first argue that the premises of the Underdetermination Argument entail ~J (i.e. premise 1 of the Closure Argument), but that the premises of the Closure Argument do not entail ~FAV (i.e. premise 1 of the Underdetermination Argument). Thus, the soundness of the Underdetermination Argument depends on the soundness of the Closure Argument relative to ~J, while the soundness of the Closure Argument is independent from the soundness of the Underdetermination Argument relative to ~FAV.

Then, I shall discuss the controversial relationship between the underdetermination principle and the closure principle. In the literature, we find a divide between those who believe that the underdetermination principle and the closure principle are equivalent (Brueckner 1994; McCain 2013), and those who think that the closure principle is the logically stronger principle (Cohen 1998; Pritchard 2005). At the start of the debate, a consensus emerged around the idea that the closure principle entails the underdetermination principle, but that consensus has been challenged (Vahid 2005). However, the more controversial issue remains whether the underdetermination principle entails the closure principle. It has also been argued that the premises of the Underdetermination Argument jointly entail the closure principle (Cohen 1998).

With respect to these issues, I shall argue that, given the evidentialist picture of justification presented in chapter 1 – a picture according to which S is justified in believing that P just in case S's total evidence supports or confirms P to a sufficient degree, and degrees of confirmation obey probabilistic constraints – the closure principle entails the underdetermination principle. Thus, I shall argue that the soundness of the Closure Argument depends on the soundness of the Underdetermination Argument relative to the underdetermination principle. I shall also argue that the underdetermination principle entails the closure principle. Thus, given that evidentialist picture of justification, the soundness of the Underdetermination Argument depends on the soundness of the Closure Argument relative to the closure principle.

2.2.1 ~FAV and ~J

Recall ~FAV and ~J:

(~FAV) S's evidence does not favour P over SH.

(~J) S is not justified in believing that ~SH.

The Underdetermination Argument's Premises Entail ~J

It is easy to show that the premises of the Underdetermination Argument entail ~J – that S is not justified in believing that ~SH. It is sufficient to notice that the variable 'P' in the Underdetermination Argument ranges over propositions about the external world, including ~SH. Thus, we can construct an underdetermination argument for justification scepticism about ~SH:

($\sim FAV_{SH}$) S's evidence does not favour ~SH over SH.
(UP_{SH}) If S's evidence does not favour ~SH over SH, S is not justified in believing ~SH.
(~J) S is not justified in believing ~SH.[4]

Cohen (1998: 149) suggests a different derivation of ~J:

> 'Now suppose E does justify [~SH] (the denial of [~J]). Then E must favor P over [SH] (the denial of [~FAV]). Where E justifies [~SH], E fails to favor P over [SH] only if E justifies not-P. But we are assuming E does not justify not-P. Thus the denial of (~J) entails the denial of [~FAV] and so (relative to our assumption) (~FAV) entails (~J).' (Changes to account for terminological differences marked by square brackets.)

The assumption that E, i.e. S's evidence, does not justify not-P is motivated by the thought that '(UP) and [(CP)] will have skeptical significance only for those cases where S's evidence does not justify not-φ' (Cohen 1998: 149). One advantage of my derivation over Cohen's is that it does without this assumption. This is a substantial advantage: strictly speaking, Cohen's derivation does not prove that the soundness of the Underdetermination Argument depends on the soundness of the Closure Argument relative to ~J, only that it does so given the assumption that E does not justify ~P.

[4] Dodd (2012: 342-3) suggests a similar argument that starts with the observation that if S's evidence does not favour P over SH (~FAV), then S's evidence does not favour ~SH over SH ($\sim FAV_{SH}$). Thus, assuming UP_{SH}, if ~FAV is true, S is not justified in believing ~SH.

The Closure Argument's Premises Do Not Entail ~FAV

It takes slightly more work to show that the Closure Argument's premises do not entail ~FAV – that S's evidence does not favour P over SH. To see why, it is necessary to appreciate that ~J – that S is not justified in believing that ~SH – does not entail that S's evidence does not favour P over SH. To see this, it is sufficient to construct a case in which S is not justified in believing that ~SH although S's evidence favours P over SH. Cohen (1998: 149) describes one such case:

> 'Suppose E (my evidence) favors P over [SH]. Still, as we just noted in section II, E may not favor P enough to justify either P or [~SH]. If, for example, Pr(P/E)=.5 while Pr([SH]/E)=.4, then E favors P over [SH], while E fails to justify [~SH]. But then [~J] is true, while [~FAV] is false.' (Changes to account for terminological differences marked by square brackets.) [5]

Generally, the idea is that evidence may favour a proposition over one of its competitors even if it does not confirm P enough to justify a belief in P. A lack of justification is compatible with a lack of underdetermination.

That ~J does not entail ~FAV is not sufficient to prove that the soundness of the Closure Argument is independent from the soundness of the Underdetermination Argument relative to ~FAV, because the Closure Argument has a second premise, namely CP, that can be used to derive more claims together with ~J. However, it should be obvious that the closure principle cannot be used to derive ~FAV from ~J. After all, the truth of the closure principle does not change the fact that underdetermination and lack of justification can come apart in the way described by Cohen's model.

2.2.2 The Underdetermination Principle and the Closure Principle

Now, let us consider the other two premises of the Underdetermination Argument and the Closure Argument, namely UP and CP:

(UP) If S is justified in believing that P, S's evidence favours P over SH.

[5] The reasoning assumes that a belief is justified only if it its evidential probability is greater than .5.

(CP) If S is justified in believing that P, then S is justified in believing that ~SH.

As I explained in chapter 1, section 5 and in footnote 1 of the present chapter, UP and CP derive their support from two general principles about underdetermination and closure, given the uncontroversial assumption that P and SH are incompatible. In fact, UP and CP can be used as proxies for the underdetermination principle and the closure principle respectively. In what follows, I won't sharply distinguish between them anymore as this makes the presentation considerably easier to follow. Nothing substantial hinges on this simplification.

The Closure Principle Entails the Underdetermination Principle

Brueckner (1994: 832), Cohen (1998: 150-1) and Pritchard (2005: 42) believe that CP – if S is justified in believing that P, then S is justified in believing that ~SH – entails UP – if S is justified in believing that P, S's evidence favours P over SH – on the basis of the following argument:[6]

1 (1) If S is justified in believing that P, then S is justified in believing that ~SH. [assumption]
2 (2) S is justified in believing that P. [assumption]
1, 2 (3) S is justified in believing that ~SH. [from 1 and 2]
4 (4) For all S, φ, if S has justification for believing that φ, S lacks justification for believing that ~φ. [assumption]
1, 2, 4 (5) S lacks justification for believing that SH. [from 3 and 4]
1, 2, 4 (6) S has justification for believing that P and S lacks justification for believing that ~SH. [from 2 and 5]
7 (7) For all S, φ, ψ, if S has justification for believing that φ and S lacks justification for believing ψ, then S's evidence favours φ over ψ. [assumption]
1, 2, 4, 7 (8) S's evidence favours P over ~SH. [from 6 and 7]
1, 4, 7 (9) If S is justified in believing that P, S's evidence favours P over

[6] This formulation of the argument resembles most Pritchard's. Brueckner's and Cohen's versions are more elliptic: they employ the principle that if S's evidence justifies φ, and S's evidence justifies ~ψ, then S's evidence favours φ over ψ. That principle is entailed by assumptions 4 and 7 (see further).

SH. [from 2, 8]

The argument shows that the underdetermination principle can be derived from the closure principle given two auxiliary assumptions about the justification of incompatible hypotheses, namely assumption 4, and the relationship between justification and favouring, namely assumption 7. Accordingly, whether CP entails UP depends on the plausibility of these assumptions.

Assumptions 4 and 7 do look very plausible. On the one hand, assumption 4 follows from the appealing idea *that one cannot have justification for believing incompatible propositions* – call it 'INC'.[7] On the other hand, assumption 7 captures the appealing idea that a difference in the justification of two propositions requires a difference in their confirmation status. Thus, the argument makes a seemingly strong case for the claim that CP entails UP as well as for the claim that the soundness of the Closure Argument depends on the soundness of the Underdetermination Argument relative to the underdetermination principle.

Vahid (2005) objects to Cohen's derivation of the underdetermination principle from the closure principle that it employs the following principle:

> (*Cohen's Principle*) For all S, φ, ψ, if S's evidence justifies φ, and S's evidence justifies ~ψ, then S's evidence favours φ over ψ.

According to Vahid, the principle cannot be used to derive the underdetermination principle because it entails that:

> For all S, φ, ψ, if S's evidence justifies φ, and S's evidence justifies ~ψ, *and φ and ψ are incompatible*, then S's evidence favours φ over ψ.

But that consequence of Cohen's Principle is 'epistemically and semantically too close' to the underdetermination principle to be used in a derivation of UP. In fact, that consequence is just a consequence of the underdetermination principle itself (Vahid

[7] See Chapter 1, section 4 for a defence of INC. According to McCain (2013: 294), INC is false because it is subject to counterexamples with the following structure: S has excellent reasons to believe that φ and ψ but no reason to believe that φ and ψ are incompatible. However, unbeknownst to S, φ and ψ are incompatible. In such cases, says McCain, S is justified in believing both φ and ψ. Assuming McCain is right, the obvious reply, also considered by McCain, is to restrict assumption 4 to cases in which S knows that φ and ψ are incompatible. See McCain (2013: 298) for discussion.

2005: 133).

In defence of Cohen, we may object to Vahid that he is misrepresenting things here. The consequence of Cohen's Principle may be too close to the underdetermination principle to be used in a derivation of UP. But Cohen's derivation does not use that consequence. Instead, it uses Cohen's Principle.

Moreover, Cohen's principle can be derived from the two assumptions 4 and 7. First, assume the antecedent of Cohen's Principle:

> S's evidence justifies φ, and S's evidence justifies ~ψ.

Using assumption 4 – if S has justification for believing that φ, S lacks justification for believing that ~φ –, we can derive that:

> S's evidence justifies φ and S lacks justification for ψ.

Using assumption 7 – if S has justification for believing that φ and S lacks justification for believing ψ, then S's evidence favours φ over ψ –, we can derive that:

> S's evidence favours φ over ψ.

Thus, assumption 4 and 7 entail that if S's evidence justifies φ, and S's evidence justifies ~ψ, S's evidence favours φ over ψ. As I just showed, pace Vahid, Cohen's derivation, let alone Pritchard's, remains standing.

The Underdetermination Principle Entails the Closure Principle

The more controversial issue is whether the premises of the Underdetermination Argument entail the closure principle. Cohen (1998: 153) provides an argument to this effect:

> 1 (1) S's evidence does not favour P over SH. (~FAV, i.e. premise 1 of the Underdetermination Argument)
> 2 (2) If S's evidence does not favour P over SH, S is not justified in believing that P. (UP, i.e. premise 2 of the Underdetermination Argument)
> 3 (3) S is justified in believing that P, but S is not justified in believing that ~SH. [assumption for *reductio*]
> 1, 2 (4) S is not justified in believing that P. [from 1 and 2]
> 3 (5) S is justified in believing that P. [from 3]

Line 5, which rests on assumption 3, contradicts line 4, which rests on assumptions 1 and 2. Hence, we can deny assumption 3 on assumptions 1 and 2. Notice that assumption 3 is the denial of the closure principle. Thus, the premises of the Underdetermination Argument jointly entail the closure principle.[8] If Cohen is right, the soundness of the Underdetermination Argument depends on the soundness of the Closure Argument relative to the closure principle.

Pritchard (2005: 45) thinks that we should be suspicious of Cohen's argument, because it establishes its conclusion on the assumption that anyone who denies the closure-based premise – CP – of the sceptical argument must affirm the existence of a justified belief that P. Thus, continues Pritchard, Cohen's argument has the surprising consequence that one cannot rationally endorse the Underdetermination Argument while denying the closure principle, because this is tantamount to believing contradictory propositions: that scepticism is true and that one has justification for P. Clearly, this cannot be right.

So, where is the mistake in Cohen's argument? According to Pritchard (2005: 45-6), Cohen's argument goes wrong in equating the negation of CP, i.e. the negation of the claim that if S is justified in believing that P then S is justified in believing that ~SH – with the conjunction of a justification for P and a lack of a justification for ~SH. Instead, according to Pritchard, the denial of CP is equivalent to affirming the *possibility* of such a conjunction:

> 'After all, it could be that a justification for belief in [P] might not entail a lack of justification for belief in ~SH (and thus that [CP] might fail) even though there is *never* a case in which a justification for belief in [P] is possessed because of the success of the (UP)-based sceptical argument. In such a situation [CP] would fail, and yet the (UP)-based sceptical argument would go through without any inconsistency.' (Pritchard 2005: 46; changes to account for terminological differences marked by square brackets.)

[8] To simplify: notice that CP is equivalent to the disjunction 'S is not justified in believing that P or S is justified in believing that ~SH'. The conclusion of the Underdetermination Argument is that S is not justified in believing that P. By disjunction introduction, the conclusion of the Underdetermination Argument becomes 'S is not justified in believing that P or S is justified in believing that ~SH'. Thus, the premises of the Underdetermination Argument jointly entail CP.

As it is, Pritchard's reply fails, because it trades on an equivocation. When Pritchard interprets the falsity of the closure principle as the *possibility* that a justification for P co-exists with a lack of justification for ~SH, he reads the closure principle as a *necessary* principle. Plausibly, this is the correct reading: if the closure principle is true, then it is necessarily true.

However, by the same token, the underdetermination principle is a *necessary* principle: if UP is true, then it is necessarily true. Further, the same applies to ~FAV – the claim that S's evidence does not favour P over SH. For the sceptic's point is not that, as a matter of fact, S's evidence does not favour P over SH. Rather, the point is that S's evidence cannot favour P over SH, given the nature of S's evidence and of the sceptical hypothesis. Therefore, properly understood, the conclusion of the Underdetermination Argument is that S's evidence *cannot* justify one's belief that P (see chapter 1, section 5.).

Accordingly, when the Underdetermination Argument and the closure principle are understood in modally strong terms, the premises of the Underdetermination Argument do entail the closure principle. The impossibility of a justification for P is incompatible with the possibility that a justification for P co-exists with a lack of justification for ~SH. Thus, far from rebutting Cohen's argument, Pritchard's criticism rests on fallacious reasoning: it applies a double standard in the interpretation of the Underdetermination Argument and of the closure principle.

Nonetheless, there is something to say in favour of Pritchard's criticism in this sense. Granting the soundness of Cohen's argument, the closure principle seems to follow from the Underdetermination Argument only in an uninteresting sense of 'follow'. In fact, the closure principle turns out to be true, but only vacuously so, because its antecedent cannot be satisfied if the Underdetermination Argument is sound. As a consequence, the relationship between the Underdetermination Argument and the closure principle seems to be only a superficial one. There is no substantial link between the premises of the Underdetermination Argument and the idea that justification is closed under entailment.

A different case is made by Brueckner (1994: 832-4). He argues that the closure principle can be derived from the underdetermination principle alone:

> 'Assume [UP] and [CP]'s antecedent. Then it follows that S's evidence for believing that φ favors φ over the incompatible ψ. This by itself is not sufficient to show [CP]'s consequent (that S has justification for believing that ~ψ). This is because S's evidence might favor one of two incompatible hypotheses without being strong enough to justify a belief of either. However, we have assumed the antecedent of [CP] in the course of the present attempted derivation, and thus we have it that S's evidence is sufficient to justify a belief that φ. Now it seems that S is justified in believing that ~ψ.' (Changes to account for terminological differences marked by square brackets.)

One might fear that Brueckner is begging the question here. Bruckner thinks that the evidential favouring of φ over its competitor ψ does not entail S's justification for believing that ~ψ because 'S's evidence might favor one of two incompatible hypotheses without being strong enough to justify a belief of either.' But one might think that this is not the only explanation available: perhaps, S's evidence favours one of two incompatible hypotheses φ and ψ while being strong enough to justify a belief in φ but not strong enough to justify a belief in ~ψ. Crucially, Brueckner does not consider this possibility, but he ought to, because this is what his argument is supposed to show to be false. A different way to put this concern is this: Brueckner is suggesting that either evidence favours one of two incompatible hypotheses without being strong enough to justify a belief in either or it favours one of two incompatible hypotheses and justifies belief in one and disbelief in the other. But why should we believe that this exhausts all the possibilities unless we already accept the closure principle?

Cohen (1998: 152-3), Pritchard (2005: 42-3), and Dodd (2012: 341) are attracted by this line of criticism. They suggest that the most one can derive from the underdetermination principle is that *if S is justified in believing that P, then S is not justified in believing that SH*. That claim is weaker than the closure principle, because 'S is not justified in believing that SH' does not entail 'S is justified in believing that ~SH.' Thus, in their view, Brueckner's argument fails.

McCain (2013: 293) has a different take on the issue. He argues by applying some simple rules of probability that CP and UP are equivalent given a plausible account of

the relationship between evidence and propositional justification and INC – the assumption that for all incompatible φ, ψ, if S has justification for believing that φ, S lacks justification for believing ψ.

However, McCain's argument contains a mistake: it invalidly derives line 4, i.e. that Pr(P|E)>.5, from lines 2 and 3. Line 3 says that S's evidence propositionally justifies P *if* Pr(P|E)>.5, but it does not say that S's evidence propositionally justifies P *only* if Pr(P|E)>.5, and that is what is needed to derive line 4 from line 2, i.e. the claim that S's evidence justifies P.

Luckily, McCain's argument is easily fixed, because line 4 follows from lines 1, 2, and 3, as my improved version of McCain's argument shows:

McCain's argument (fixed version)

1 (1) If S has justification for believing that φ, S lacks justification for believing any incompatible ψ. [INC]

2 (2) S's evidence justifies P. [antecedent of CP]

3 (3) If Pr(φ|E)=.5, S is justified in believing neither that φ nor that ~φ; if Pr(φ|E)>.5, S is justified in believing that φ; if Pr(φ|E)<.5, S is justified in believing that ~φ. [assumption]

1, 2, 3 (4) Pr(P|E)>.5 [from 1, 2, 3]

1, 2, 3 (5) Pr(~P|E)<.5 [from 4 and the Complement Rule]

1, 2, 3 (6) S is not justified in believing that SH. [from 1, 2]

1, 2, 3 (7) Pr(SH|E)≤.5 [from 3, 6]

1, 2, 3 (8) Pr(SH|E) ≤.5 and Pr(SH|E) ≤ Pr(~P|E) [from 7 and Monotonicity]

1, 2, 3 (9) Pr(SH|E)<.5 [from 5, 8]

1, 2, 3 (10) Pr(~SH|E)>.5 [from 9]

1, 2, 3 (11) S is justified in believing that ~SH. [from 3, 10]

1, 3 (12) If S is justified in believing that P, S is justified in believing that ~SH. [from 2, 11]

According to McCain (2013: 298), the argument's upshot is that either the closure principle and the underdetermination principle are equivalent or that INC is false. McCain's argument derives the closure principle from INC, and INC was used in the

derivation of the underdetermination principle from the closure principle (see above). So, if the soundness of the Closure Argument depends on the soundness of the Underdetermination Argument relative to the underdetermination principle given INC, the soundness of the Underdetermination Argument depends on the Closure Argument relative to the closure principle given INC, too. Notice, however, that McCain's derivation does not vindicate Brueckner's point that the underdetermination principle entails the closure principle, for the underdetermination principle plays no role in McCain's derivation.

Moreover, it should be noted that McCain's equivalence result depends on an assumption that changes the original terms of the debate, namely assumption 3. McCain (2013: 290) assumes that 'in order for S to have propositional justification that p her reasons simply have to support p to a higher degree than they do ~p'. That is the same as assuming that in order for S to have propositional justification for P her reasons simply have to favour P over ~P. But that assumption is explicitly rejected in the original debate.

As we have seen above, in the original debate between Brueckner (1994), Cohen (1998), and Pritchard (2005) the issue is with complete or adequate justification. And there is a gap between this sort of justification and favouring, because S's evidence can favour P over ~P without favouring P enough to justify P. Since McCain's derivation conflates evidential favouring and justification, McCain's argument does not show that the underdetermination principle and the closure principle are equivalent. For all McCain has shown, Cohen and Pritchard may be right that the two principles are not equivalent, and, given their rebuttal of Brueckner's derivation of the closure principle from the underdetermination principle, the Underdetermination Argument may not depend on the Closure Argument relative to the closure principle after all – as McCain himself recognises in passing (2013: 294).

In the following, I shall argue that the underdetermination principle – or UP – entails the closure principle – or CP – given a conception of justification as sufficiently high degree of evidential support that obeys probabilistic constraints. Thus, conditionally on that conception of justification, the soundness of the Underdetermination Argument depends on the soundness of the Closure Argument relative to the closure principle. Interestingly, INC is not required for the derivation. The argument

vindicates and helps make sense of Brueckner's argument for the claim that the underdetermination principle entails the closure principle and supports the claim that the underdetermination principle and the closure principle are equivalent.

To prepare the terrain for my argument, I shall prove that, given McCain's conception of propositional justification, McCain's result can be derived by assuming the underdetermination principle – if S is justified in believing that P, S's evidence favours P over SH.

Assume UP and the antecedent of the closure principle:

1 (1) If S is justified in believing that P, S's evidence favours P over SH.

2 (2) S is justified in believing that P.

Now, assume McCain's conception of the relationship between evidence and propositional justification. In particular, assume that:

3 (3) If $Pr(\varphi|E)>.5$, then S is justified in believing that φ.

The proof's strategy is to derive $Pr(\sim SH|E)>.5$ from assumptions 1 and 2 and to use assumption 3 to derive that S is justified in believing that ~SH. However, assumptions 1 and 2 entail that S's evidence favours P over SH, and that alone does not entail that $Pr(\sim SH|E)>.5$. So, how can we derive $Pr(\sim SH|E)>.5$ from assumptions 1 and 2? By applying some simple rules of probability. For *reductio*, assume the negation of $Pr(\sim SH|E)>.5$:

4 (4) $Pr(\sim SH|E)\leq.5$

By applying the complement rule for probabilities, we can derive that:

4 (5) $Pr(SH|E)\geq.5$

Since P entails ~SH we can also derive that:

4 (6) $Pr(P|E)\leq.5$

As a consequence of lines 5 and 6:

4 (7) Either $Pr(SH|E)>.5$ and $Pr(P|E)<.5$ or $Pr(SH|E)=Pr(P|E)=.5$

Since for S's evidence to favour P over SH is simply for S's evidence to support P to a higher degree than it does SH (McCain 2013: 291), line 7 entails that:

4 (8) S's evidence does not favour P over SH.

However, assumptions 1 and 2 entail the negation of line 8:

1, 2 (9) S's evidence favours P over SH.

So, a contradiction follows from the conjunction of assumptions 1, 2, and 4. Then, we can deny the latter assumption by *reductio ad absurdum*:

1, 2 (10) $Pr(\sim SH|E)>.5$

Since $Pr(\sim SH|E)>.5$, given McCain's conception of the relationship between evidence and propositional justification:

1, 2, 3 (11) S is justified in believing that ~SH.

By discharging the antecedent, we obtain the result that:

1, 3 (12) If S is justified in believing that P, S is justified in believing that ~SH.

Like McCain's original derivation, this argument conflates evidential favouring and justification. Thus, this argument does not show that the underdetermination principle and the closure principle are equivalent: after all, S's evidence can favour P over ~P without favouring P enough to justify P.

However, it is easy to show that the argument works even if we do not assume, like McCain does, that if $Pr(\varphi|E)>.5$, then S's evidence propositionally justifies φ. Instead, assume that S's evidence justifies φ if $Pr(\varphi|E)> t$, where t is some arbitrary value greater than .5. Then, we can run the same argument with 't' in place of '.5':

1 (1) If S is justified in believing that P, S's evidence favours P over SH. [UP]

2 (2) S is justified in believing that P. [assumption]

3 (3) If $Pr(\varphi|E)> t$, then S is justified in believing that φ. [assumption]

4 (4) $Pr(\sim SH|E)\leq t$ [assumption for *reductio*]

4 (5) $Pr(SH|E)\geq t$ [from 4]

4 (6) $Pr(P|E)\leq t$ [from 4]

4 (7) $Pr(SH|E)\geq t$ and $Pr(P|E)\leq t$ [from 5 and 6]

4 (8) S's evidence does not favour P over SH. [from 7]

1, 2 (9) S's evidence favours P over SH. [from 1, 2]

1, 2 (10) Pr(~SH|E)> t [from 8, 9]

1, 2, 3 (11) S is justified in believing that ~SH. [from 3, 10]

1, 3 (12) If S is justified in believing that P, S is justified in believing that ~SH. [from 2, 11]

The argument derives the closure principle from the underdetermination principle. Moreover, one cannot object that it conflates justification with evidential favouring, for t can be any value greater than .5.

The argument helps us make sense of Brueckner's derivation of the closure principle from the underdetermination principle in a way that his critics (Cohen, Pritchard) have been unable to. To rehearse, Cohen and Pritchard object to Brueckner's derivation that the most one can derive from UP – the claim that if S is justified in believing that P, S's evidence favours P over SH – is that if S is justified in believing that P, then S is not justified in believing that SH. After all, S's evidence can favour one two incompatible hypotheses φ and ψ while being strong enough to justify a belief in φ but not strong enough to justify a belief in ~ψ. What my derivation shows is that this claim is false under a conception of justification as sufficiently high degree of evidential support that obeys probabilistic constraints.

Assuming this conception of justification, Brueckner's derivation makes perfect sense: start with the underdetermination principle and the antecedent of the closure principle. Together, they entail that S's evidence favours P over SH. Notice, however, that S's evidence must favour P over SH to a degree sufficient for S to have justification in believing that P, because we are assuming the antecedent of CP. Further, if evidential support obeys the rules of the probability calculus, the degree of evidential support for ~SH available to S must be at least equal to the degree of evidential support for P, because P entails ~SH. Thus, if P is confirmed to a degree sufficient for justification, so is ~SH. In the end, the soundness of the Underdetermination Argument depends on the soundness of the Closure Argument relative to the closure principle, because the underdetermination principle entails the closure principle.

2.3 Answering the Cogency Question

I shall follow the same blueprint I used to answer the Soundness Question in order to

answer the Cogency Question. Recall that an argument A c-depends on an argument B just in case A's premises cannot be plausibly motivated without an appeal to B's premises. Thus, the Cogency Question is a way of asking whether the plausibility of the premises of the Underdetermination Argument is conditional on the plausibility of the premises of the Closure Argument, and vice versa.

There is an obvious sense in which the plausibility of the Underdetermination Argument's premises depends on the plausibility of the Closure Argument's premises if what I said in section 2 of this chapter is correct. If the soundness of the Underdetermination Argument depends on the truth of the Closure Argument's premises, then the plausibility of the Underdetermination Argument stands and falls together with the plausibility of the Closure Argument. Similarly, if the soundness of the Closure Argument depends on the truth of the underdetermination principle, the plausibility of the Closure Argument stands and falls together with the plausibility of the underdetermination principle. In fact, all this is just another way to say that the Closure Argument is less refutable than the Underdetermination Argument, that any refutation of the Closure Argument is a refutation of the Underdetermination Argument, but not vice versa.

This is not the sense of plausibility that concerns the Cogency Question. The Cogency Questions concerns the cogency, i.e. *the plausibility of the premises' motivation*, rather than *the refutability of the argument*. Just as it makes sense to distinguish cogency and soundness, so it makes sense to distinguish cogency and refutability. The Underdetermination Argument seems more vulnerable than the Closure Argument, but this does not mean that it is less cogent (Cohen 1998: 156-7).

In this section, I shall first argue that the premises of the Underdetermination Argument can be plausibly motivated independently of the Closure Argument. Then, I shall consider whether the same is true of the premises of the Closure Argument relative to the Underdetermination Argument. Some authors (Brueckner 1994; Boult 2013; Bergmann 2021) have argued that this is not so, but their arguments fail to establish this conclusion. In fact, I shall argue that the premises of the Closure Argument can

be plausibly motivated independently of the Underdetermination Argument.[9]

2.3.1 Motivating the Underdetermination Argument's Premises

~ FAV

As explained in chapter 1, the underdetermination sceptic argues that S's evidence does not favour P over SH, because S has the same evidence in the good and the bad case. I shall say more about this inference in the next chapter. But here I want to say this: insofar this is a plausible motivation for ~FAV, the Closure Argument seems to play no role in it. In fact, in section 2.1, I argued that the Closure Argument's premises do not entail ~FAV, because failures of justification do not entail failures of underdetermination. The Closure Argument can be used to derive that S lacks justification for believing that ~SH (in fact, that is just one of the argument's premises) and that S lacks justification for believing that P. But it is hard to see how one might derive that S's evidence does not favour P over SH from those claims.

The Underdetermination Principle

In section 2.2, we have seen that premise 2 of the Underdetermination Argument – or the underdetermination principle – can be derived from premise 2 of the Closure Argument – or the closure principle. This does not mean that the underdetermination principle's plausibility derives from the plausibility of the closure principle. After all, the plausibility of the underdetermination principle might derive from our grasp of the concepts of evidential favouring and justification. Alternatively, it might derive from its ability to explain our judgement over a range of cases (Vogel 2004: 427; Pritchard

[9] This is compatible with the claim that the Underdetermination Argument provides a natural motivation for ~J. Thus, it is compatible with the claim that the Underdetermination Argument plays a central role in the motivation of the Closure Argument. Indeed, the standard considerations in favour of ~J in terms of subjective indistinguishability can be traced back to an underdetermination argument. Pritchard (2015: 11) provides a representative example:

> 'The initial plank in the case for skepticism comes from the contention that one cannot know that one is not a BIV. Such a claim seems entirely compelling. After all, since the BIV scenario is ex hypothesi subjectively indistinguishable from normal perceptual conditions, it is hard to see how one might come to know such a thing. What kind of rational ground might one have for such a belief, given that there is no subjective basis on which one can discern that one is not in a radical skeptical scenario?'

2015: 31). At any rate, suppose that the underdetermination principle's plausibility derives from some other principle's plausibility. As Cohen (1998: 157) observes, it still does not follow that the closure principle plays any role in the motivation of the underdetermination principle, because the plausibility of the underdetermination principle might derive from other principles.

Cohen (1998: 157-8) suggests that the underdetermination principle can be derived from a principle about evidential parity and from INC – a principle about the justification of incompatible hypotheses. Suppose that S's evidence does not favour φ over ψ. By itself, this does not entail that S is not justified in believing that φ. After all, S's evidence might justify both φ and ψ. Rather, the assumption that S's evidence does not favour φ over ψ entails that if S's evidence justifies φ, then it justifies ψ, too.[10] This motivates the first principle for the derivation of the underdetermination principle:

> (PARITY) For all φ, ψ, if S's evidence does not favour φ over ψ, and S is justified in believing that φ, then S is justified in believing that ψ.

Further, it seems that evidence cannot justify incompatible hypotheses. For, as noted in chapter 1 section 4, if the evidence raises the likelihood of one hypothesis over .5, the likelihood of the other must drop below .5. On the plausible assumption that justification requires an evidential likelihood greater than .5, this motivates the second principle needed for the derivation:

> (INC) For any incompatible φ, ψ, if S is justified in believing that φ, then S is not justified in believing that ψ.

PARITY and INC jointly entail the underdetermination principle. This can be shown via an indirect proof, by substituting P and SH for φ and ψ:

> 1 (1) If S's evidence does not favour P over SH and S is justified in believing that P, then S is justified in believing that SH. (from PARITY)
> 2 (2) If S is justified in believing that P, S is not justified in believing that SH. (from INC)
> 3 (3) S is justified in believing that P and S's evidence does not favour P

[10] The assumption entails the stronger biconditional that S's evidence justifies φ if and only if it justifies ψ. However, for the derivation of the underdetermination principle, the weaker conditional suffices.

over SH [assumption for *reductio*]

1, 3 (4) S is justified in believing that SH. [from 1, 3]

2, 3 (5) S is not justified in believing that SH. [from 2, 3]

1, 2, 3 (6) S is justified in believing that SH and S is not justified in believing that SH. [from 4, 5]

Since assumptions 1, 2 and 3 entail a contradiction, 1 and 2 entail the negation of 3. But the negation of 3 is just the underdetermination principle. Thus, PARITY and INC entail the underdetermination principle.

Pritchard (2005: 48) argues that the underdetermination principle can be derived from INC alone, given the assumption that justification is evidential in character:

1 (1) If S is justified in believing that P, S is not justified in believing that SH. (from INC)

2 (2) S is justified in believing that P. [assumption]

1, 2 (3) S is not justified in believing that SH. [from 1 and 2]

1, 2 (4) S is justified in believing that P and S is not justified in believing that SH. [from 2, 3]

1, 2 (5) S's evidence favours P over SH. [from 4]

1 (6) If S is justified in believing that P, S's evidence favours P over SH. [from 2, 5]

Notice that Pritchard's argument presupposes the principle used in the derivation of UP from CP in section 2.2:

> For all S, φ, ψ, if S has justification for believing that φ and S lacks justification for believing ψ, then S's evidence favours φ over ψ.

However, crucially, neither Cohen's nor Pritchard's argument employs the closure principle itself. Thus, the underdetermination principle can be motivated without any appeal to the closure principle. Further, the two derivations appeal to at least *prima facie* plausible principles. Thus, the underdetermination principle can be motivated independently of the Closure Argument.

2.3.2 Motivating the Closure Argument's Premises

The Closure Principle

What I said about the motivation of the underdetermination principle applies *mutatis mutandis* to the closure principle. The closure principle's plausibility need not derive from the plausibility of the underdetermination principle, or from the plausibility of any other principle. But suppose that the closure principle is plausible only if it can be derived from some other principle. It does not follow that the closure principle's plausibility derives from the plausibility of the underdetermination principle just because the underdetermination principle entails the closure principle. In fact, in section 2.2, we saw that McCain has provided an argument that derives the closure principle independently of the underdetermination principle. Granted, that argument changes the terms of the debate by assuming that φ is justified by S's evidence as long as $Pr(\varphi|E)>t$, where t is .5. But, trivially, McCain's argument works even if one changes the value of t to any value greater than .5. Thus, McCain's argument shows that, within an evidentialist framework, the closure principle can be derived independently of the underdetermination principle.

~J

Perhaps, the plausibility of ~J – the claim that S is not justified in believing that ~SH – does not depend on our ability to put forward an argument for ~J; perhaps, ~J enjoys some sort of intuitive plausibility (Cohen 1998: 147). However, several authors think that an argument is required to support ~J, because ~J is itself a sceptical claim (Feldman 2003: 127-8; Conee and Feldman 2004: 290; Greco 2008: 111; Byrne 2004: 303-4). Brueckner (1994: 830; 2005: 388; 2011: 75-6) goes a step further to claim that the support for ~J should be construed as an underdetermination argument. If Bruckner is right, the plausibility of the Closure Argument depends on the Underdetermination Argument, because the Closure Argument cannot be plausibly motivated independently of the Underdetermination Argument. Further, since the Underdetermination Argument can reach the conclusion of the Closure Argument independently of the Closure Argument, this would mean that the Closure Argument is superfluous. In a similar vein, Bergmann (2021: 20) claims that 'the Closure Argument has a skeptical

premise that is itself in need of support by an underdetermination argument. This is a good reason to set aside the closure argument until after first examining the underdetermination argument'.

Although Brueckner and Bergmann are right that ~J – the claim that S is not justified in believing that ~SH – *can* be supported by the Underdetermination Argument, they provide no argument for the claim that ~J *must* be supported by the Underdetermination Argument. Thus, notwithstanding their claims, ~J might be plausibly motivated independently of the Underdetermination Argument. In this section, I shall argue that this is in fact the case. I shall examine three proposals to motivate ~J independently of the Underdetermination Argument and argue that two of them provide a *prima facie* motivation for ~J that is independent of the Underdetermination Argument.

First, Cohen (1998, 146) argues that ~J can be motivated by the thought that if SH were true, it would explain S's evidence. Thus, Cohen suggests that the sceptic can appeal to the following principle:

> For all φ, if the truth of φ would explain S's evidence, then S's evidence does not justify ~φ.

However, Cohen's suggestion faces some problems. The stated principle seems false. Suppose that two incompatible hypotheses φ and ψ explain S's evidence. Further, suppose that while φ does a very poor job at explaining S's evidence, ψ does it very well. Under these circumstances, it is plausible that S's evidence justifies ~φ.

The natural fix is to modify the principle to account for the quality of the explanation provided by φ compared to other hypotheses:

> For all φ, if the truth of φ would explain S's evidence as well as any hypothesis incompatible with φ, then S's evidence does not justify ~φ.

As Boult (2013: 1132) observes, it is not obvious that the new principle can be used to motivate ~J, because it is not obvious that ~SH explains S's evidence equally well as its non-sceptical competitor.[11] More importantly in this context, it is not clear that the

[11] Vogel (1990), McCain (2014) and (2019), Douven (2022) argue that ordinary non-sceptical hypotheses are better explanations of the evidence than SH. See Rinard (2017) and Bergmann (2021) for some criticism of abductive responses to scepticism.

new principle makes a distinct claim from the underdetermination principle rather than being just a different way of expressing the underdetermination principle. After all, it is tempting to interpret the claim that two hypotheses explain the evidence equally well as the claim that the evidence does not favour one over the other (Wright 1985: 431).

Second, might the claim that S is not justified in believing that ~SH be motivated by appeal to a sensitivity condition for justification instead?[12] Consider:

> (SEN) If S is justified in believing that φ, then: if φ were false, S would not believe that φ.

SEN seems to lend support to ~J, because if ~SH were false, S would still believe that ~SH. Further, it seems that someone who believes ~J because of sensitivity considerations could also believe that S's evidence favours ~SH over SH. But if this is correct, then the underdetermination principle plays no role in the justification of ~J via SEN (Cohen 1998: 147).

Brueckner (1994: 828) contends that the sceptic cannot use SEN to motivate ~J, because SEN makes the closure principle false.[13] However, the objection does not hold scrutiny. As Shatz (1987: 248-9) argues, sensitivity is compatible with closure. Although Shatz's point is that sensitivity is compatible with closure for anti-sceptical purposes, the same is true *mutatis mutandis* for sceptical purposes (Williamson 2000: 150-1). Finally, Briesen (2010: 226-7) points out that SEN is incompatible with the closure principle only on the assumption that sensitivity is *sufficient* (possibly together with other conditions met by S's belief that ~SH) for justification. But the sceptic need not share that assumption.

Perhaps, the best case against using SEN to motivate ~J is simply that SEN is not a genuine necessary condition for justification – something I cannot settle here. But even if, in the end, SEN is not a genuine necessary condition for justification, SEN is *prima facie* plausible: it captures the idea that if S is justified in believing that φ, then S

[12] This is modelled after Nozick's (1981) sensitivity condition for knowledge.

[13] At the general level, the idea is that a belief that P can be sensitive even if a belief in one of its logical consequences Q is not. This is possible because the closest possible world in which P is false – the world picked out by the antecedent of the counterfactual 'if P were false, one would not believe that P' – need not be the closest possible world in which Q is false – the world picked out by the antecedent of the counterfactual 'if Q were false, one would not believe that Q'.

can discriminate between the truth of φ and its falsity.

Finally, Briesen (2010: 227-8) and Weatherson (2007: 172) suggest that ~J can be motivated via what Weatherson calls the 'exhaustive argument'. According to this argument, S is justified in believing that ~SH only if S is justified in believing it either *a posteriori* or *a priori*. But S is justified in believing that ~SH neither *a posteriori* nor *a priori*. Hence, S is not justified in believing that ~SH.

According to Briesen, the argument motivates ~J without appealing to the Underdetermination Argument.[14] On the one hand, S is not justified in believing that ~SH *a priori,* because ~SH is not the kind of claim that can be supported without some empirical evidence. On the other hand, S is not justified in believing that ~SH by way of empirical evidence, because 'this kind of evidence cannot rationally be regarded as any stronger than one's independent reason for supposing that the procedure in question has been executed properly' (Briesen 2010: 227; the point is due to Wright 2004: 168).

Again, like in the case of SEN, it might be that, ultimately, the reasoning above is flawed. But this is not at issue in the present context. Rather, the point is that the reasoning above is independent of the Underdetermination Argument as well as *prima facie* plausible. That is all is needed to support the claim that ~J can be plausibly motivated independently of the Underdetermination argument.

2.4 Concluding Remarks

It is time to take stock. Consider again the following arguments:

Underdetermination Argument (UA)

(~FAV) S's evidence does not favour P over SH.

(UP) If S is justified in believing that P, S's evidence favours P over SH.

(SCEPT) S is not justified in believing that P. [from ~FAV and UP]

[14] Boult (2013: 1128) objects that 'a natural question to ask Briesen (and Weatherson) is: what motivates the claim that S is not justified by way of empirical evidence in believing that [~SH]? Is this to be understood as primitively plausible? Or is it plausible in virtue of implicit commitments to an underdetermination principle? I suggest that there must be a reason for thinking that S isn't justified empirically in her beliefs about [~SH], and a good candidate is the reason given by Brueckner below' (Changes to account for terminological differences marked by square brackets). It is a puzzling objection, because Briesen does provide a motivation for the claim that S is not justified by way of empirical evidence in believing that ~SH (see further).

Closure Argument (CA)

(~J) S is not justified in believing that ~SH.

(CP) If S is justified in believing that P, then S is justified in believing that ~SH.

(SCEPT) S is not justified in believing that P. [from ~J and CP]

In section 2, I have explored the relationship between the soundness of the Underdetermination Argument and the soundness of the Closure Argument. I have argued that the Underdetermination Argument and the Closure Argument depend on each other in the following sense: the falsity of one argument's premises makes the other argument unsound. More precisely, I have argued that the Underdetermination Argument's premises entail ~J – i.e. premise 1 of the Closure Argument – and that the underdetermination principle entails the closure principle. Thus, the Underdetermination Argument is unsound if one of the premises of the Closure Argument is false. In contrast, the closure principle entails the underdetermination principle, but the premises of the Closure Argument do not entail ~FAV – i.e. premise 1 of the Underdetermination Argument. Thus, the Closure Argument is unsound if the underdetermination principle is false. The upshot is that any refutation of the Closure Argument is a refutation of the Underdetermination Argument, but not the other way around. *In this sense, the Underdetermination Argument is more refutable than the Closure Argument.*

These results come with some *caveats* and should be put in context. First and foremost, they are achieved within an evidentialist framework for justification according to which evidence and evidential support are all that matters to justification, and evidential support or confirmation obeys the probability calculus. These assumptions are especially important to the relationship between the underdetermination principle and the closure principle. *Without them, the closure principle does not entail the underdetermination principle, and the underdetermination principle does not entail the closure principle.* If evidence and evidential support do not matter at all to justification, the closure principle might be true although the underdetermination principle is false. And if evidence and evidential support is not all that matters to justification, or if evidential support does not obey the probability calculus, the underdetermination principle might be true although the closure principle is false.

Second, the greater vulnerability of the Underdetermination Argument need not be significant. In particular, it is to be seen whether the Closure Argument remains available to the sceptic once ~FAV – the claim that S's evidence does not favour P over SH – is shown to be false. For instance, consider responses to the Underdetermination Argument according to which ~FAV is false because S is justified in believing that P and the underdetermination principle is true. Such responses make the Closure Argument unavailable to the sceptic, because if S is justified in believing that P and the underdetermination principle entails the closure principle, then ~J – the claim that S is not justified in believing ~SH – must be false.

In section 3, I have explored the relationship between the cogency of the Underdetermination Argument and the cogency of the Closure Argument. I have argued that the Underdetermination Argument and the Closure Argument are independent from each other in the following sense: their premises can be plausibly motivated without an appeal to the other argument's premises. More precisely, I have argued that premise 1 of the Underdetermination Argument can be derived from the claim that S has the same evidence in the good and the bad case, and that the Closure Argument plays no role in the derivation. Moreover, the underdetermination principle appears plausible because of its ability to explain our judgement over a range of cases, and it can be derived from some plausible principles without any appeal to the Closure Argument. Thus, the cogency of the Underdetermination Argument is independent of the cogency of the Closure Argument.

What I said about the underdetermination principle is true *mutatis mutandis* of the closure principle. In particular, the closure principle can be derived without the Underdetermination Argument from INC – the claim that evidence cannot justify incompatible hypotheses. As for ~J, it can be motivated without the Underdetermination Argument by appeal to a sensitivity condition for justification or to considerations about the kind of evidence available to S. Thus, the cogency of the Closure Argument is independent of the cogency of the Underdetermination Argument.

How should we evaluate these results in light of the question whether the Underdetermination Argument and the Closure Argument are distinct sources of scepticism? In my view, they support an affirmative answer: that underdetermination and closure scepticism present us with different sceptical problems. This is obvious in the case of

my answer to the Cogency Question, since it says that the Underdetermination Argument and the Closure Argument can be motivated independently of each other. It is perhaps less obvious in the case of my answer to the Soundness Question.

In fact, one might object that my considerations better support the opposite claim: that the Underdetermination Argument and the Closure Argument are the same argument. First, the fact that the Underdetermination Argument and the Closure Argument can be motivated independently of each other is compatible with their being the same argument. It would simply mean that the argument in question does not display a certain form of dialectical circularity. Second, one might argue that my answer to the Soundness Question strongly indicates that the Underdetermination Argument and the Closure Argument are the same argument. After all, I have argued not only that the underdetermination principle and the closure principle are equivalent, but also that they entail each other.

Although this line of reasoning is tempting, it should be resisted. This is not just for the trivial reason that, strictly speaking, the Underdetermination Argument and the Closure Argument cannot be the same argument because ~FAV – the claim that S's evidence does not favour P over SH – and ~J – the claim that S is not justified in believing that ~SH – really are different claims. Rather, this reasoning overlooks the fact that the underdetermination principle and the closure principle only entail each other given the evidentialist conception of justification described above. But, as already noted, when one forgoes that picture of justification their truth values can come apart. This is as it should be, since the two principles articulate different ideas – that justification requires evidential favouring over competing hypothesis and that a justification for a proposition is also a justification for its consequences. The inference from the equivalence or mutual entailment of the two principles to their identity is a fallacious one.

What that line of reasoning gets right, however, is that, given the evidentialist conception of justification assumed here, the Underdetermination Argument and the Closure Argument are closely related arguments. Since, given that picture of justification, the underdetermination principle and the closure principle entail each other, the underdetermination principle and the closure principle will have sceptical consequences in exactly the same circumstances.

If what I have argued in this chapter is correct, epistemologists interested in scepticism should change their ways. Although they should not stop working on closure scepticism, they should start working more on underdetermination scepticism. For one thing, the Underdetermination Argument is often (if only implicitly) appealed to as a motivation for one of the Closure Argument's premises. For another thing, the Underdetermination Argument represents a different route to scepticism than the Closure Argument: a route that, for all we have seen, is just as promising as the Closure Argument. Finally, given that the Underdetermination Argument and the Closure Argument are closely related on certain assumptions about justification, studying the Underdetermination Argument promises insights into the nature and tenability of the Closure Argument.

3

The Infallibility Objection

When it comes to sceptical arguments, it is standard to separate the wheat from the chaff. Where does the Underdetermination Argument for scepticism belong? The underdetermination sceptic infers from the sameness of evidence in the good and the bad case that our evidence does not favour our beliefs over rival sceptical hypotheses. It has been suggested (Brueckner 1994, 2005; Byrne 2004; Dodd 2012, 2014; Smith 2022) that the Underdetermination Argument is a bad argument because that inference presupposes infallibilism, the view that justification requires evidence that guarantees the truth of the justified belief. Call this 'the Infallibility Objection'.

The Infallibility Objection has consequences for epistemological theorising because some philosophical arguments trade on the fear of underdetermination scepticism. For instance, one argument for epistemological disjunctivism – the view that, under normal circumstances, the rational support for one's perceptual beliefs is both factive and reflectively accessible – is that it provides an answer to the Underdetermination Argument (Pritchard 2015). If the Infallibility Objection is correct, however, the most straightforward way to rebut the Underdetermination Argument is to reject infallibilism. Consequently, the argument for epistemological disjunctivism becomes considerably weaker. Similarly, Williamson (2000, 2007) argues against mentalism – the view that one's evidence consists of non-factive mental states – and accessibilism – the view that one is always in a position to know what one's evidence is – because they make the Underdetermination Argument irresistible. Again, if the Infallibility Objection is correct, this need not be the case. More generally, if the Infallibility Objection is correct, no controversial philosophical theory is needed to answer the underdetermination sceptic, at least as long as fallibilism is taken for granted. Thus, a great deal hinges on the correctness of the Infallibility Objection.

Following Briesen (2010), the sceptic has two strategies available to deal with the Infallibility Objection: confrontation and sidestepping. *Confrontation* attempts to show that the Infallibility Objection is misguided. It consists in arguing that infallibilism plays no role in the contentious inference or that if it does this is no reason to dismiss it. In contrast, *sidestepping* attempts to evade the Infallibility Objection. It consists in finding

arguments for a lack of evidential favouring that rely neither on the Sameness of Evidence Lemma – the claim that S has the same evidence in the good and the bad case – nor on infallibilism. In this chapter, I argue that the Infallibility Objection is misguided: the sceptic's inference from sameness of evidence to underdetermination does not presuppose infallibilism.

Here is the plan for the chapter. In section 1, I define fallibilism and infallibilism and sketch the standard story about their relationship to scepticism. In section 2, I present the Infallibility Objection and some considerations in support of it. In section 3, I describe two strategies available to the sceptic to respond to the Infallibility Objection: confrontation and sidestepping. Section 4 examines two proposals to sidestep the Infallibility Objection using an explanation and an entailment principle due to Briesen (2010) and finds them unsuccessful. Section 5 argues that the sceptic can successfully defend the Underdetermination Argument from the Infallibility Objection by pursuing a confrontation strategy. I offer a rational reconstruction of the sceptical inference from sameness of evidence to underdetermination that dispenses with infallibilism. Pace critics, the Infallibility Objection is incorrect because the contentious sceptical inference does not presuppose infallibilism.

3.1 Fallibilism, Infallibilism, and Scepticism

In philosophical jargon, 'fallibilism' and 'infallibilism' refer to views about knowledge and justification. In keeping with this thesis' focus on justification scepticism, I restrict my attention to fallibilism and infallibilism about justification. In keeping with its evidentialist assumptions, I characterize them in terms of the relationship between evidence and truth. To remain as neutral as possible on whether evidence is propositional, I conceive of this relationship as guarantee rather than entailment.

Unlike entailment, guarantee can obtain between propositional and non-propositional items alike. I opt for a definition in terms of guarantee over one in terms of entailment, because, in my discussion, I want to leave it open whether evidence includes non-propositional items like experiences. That said, guarantee is easily traced back to entailment. If S's experiences guarantee that P, there is a derivative sense of 'entail' in which they can be said to entail that P. For if S's experiences guarantee that

P, then the proposition that S has those experiences entails that P.[1]

Infallibilism is the view that:

> (INF) If S is justified in believing that P, then S's evidence guarantees the truth of P.

Equivalently, if S is justified in believing that P, then S could not have the same evidence in situations where P is false. Justification requires truth-guarantee by one's evidence.

Fallibilism is the negation of infallibilism:

> (FAL) S can be justified in believing that P although S's evidence does not guarantee the truth of P.

In other words, S can be justified in believing that P even if S could have the same evidence in situations where P is false. Justification does not require truth-guarantee by one's evidence.[2]

Claims as to whether justification requires truth-guarantee should be distinguished from claims as to whether one's evidence ever guarantees the truth of the relevant proposition. As illustrated by Kraft (2012: 54), these claims can be combined in four different ways:

A. Justification requires evidence that guarantees that P, and S's evidence sometimes guarantees that P.
B. Justification requires evidence that guarantees that P, but S's evidence never

[1] Other common definitions of fallibilism/infallibilism employ the concepts of ruling out or eliminating alternative possibilities (Lewis 1996). Whether they are equivalent to definitions in terms of entailment/guarantee depends on how the concept of ruling out every possibility in which ~P is true is specified. If to rule out every possibility in which ~P is true is for one's justifications to guarantee/entail that every alternative to P is false, then the formulations are equivalent. However, if to rule out every ~P-possibility is to know ~P not to obtain, then they are not. For then fallibilists will treat ruling out as fallible relation, so that one could rule out every ~P-possibility even if one's justifications did not guarantee/entail that P (Leite 2010: 370-1).

[2] Definitions of fallibilism and infallibilism are controversial. On this point see Brown (2018: 15-21). By adopting these definitions, I do not intend to deny the interest or legitimacy of other ones. Rather, my choice was guided by two considerations. First, the literature I use assumes either these definitions or closely related ones. Thus, these definitions connect my discussion with the relevant body of literature. Second, despite their well-known shortcomings (e.g. in dealing with necessary truths), which do not matter in the present context, they are elegant and easy to use because the notion of guarantee can be cashed out in terms of the regimented notion of entailment.

guarantees that P.

C. Justification does not require evidence that guarantees that P, but S's evidence sometimes guarantees that P.

D. Justification does not require evidence that guarantees that P, and S's evidence never guarantees that P.

What is the relationship between fallibilism, infallibilism, and scepticism? Clearly, B entails scepticism, since B combines infallibilism, a claim about a necessary condition for justification, with the denial that that condition is ever met. In contrast, scepticism does not follow from A, C, and D. On the one hand, A shares B's commitment to infallibilism, but, unlike B, it states that the infallibilist condition is sometimes met. On the other hand, unlike A and B, C and D deny that possessing truth-guaranteeing evidence is a necessary condition for justification.

Further, it should be noted that although A, C, and D do not entail scepticism, they do not entail that scepticism is false, either. In other words, A, C, and D are compatible with scepticism. After all, A, C, and D leave it open whether among the necessary conditions for justification is one that we never meet.[3] There are two lessons to learn from this. First, infallibilism alone does not entail scepticism. Rather, it entails scepticism when it is combined with the claim that the infallibilist necessary condition cannot be met. Second, fallibilism alone does not rule out scepticism.

In light of this, it might look as if fallibilism and infallibilism are on a par when it comes to their sceptical import: neither entails scepticism. However, this is not the standard tale about fallibilism, infallibilism, and scepticism. Many would argue that the parity is only apparent, for there seems to be a gap between our evidence and what we claim to be justified in believing. Examples abound: we form our beliefs about the world via sense experience, testimony, statistical reasoning, inference to the best explanation *etc.* but none of these processes seems to guarantee the truth of our beliefs. And since our evidence falls short of guaranteeing the truth of our beliefs, infallibilism

[3] Thus, Kraft (2012: 55) is wrong in claiming that A, the position he calls 'Cartesian Infallibilism', entails that we do know/are justified in believing something. Kraft is mixing necessary and sufficient conditions. However, on an evidentialist picture of justification according to which evidence is all that matter to justification, A and C have anti-sceptical import.

paves the way to scepticism:[4]

> *Infallibility argument*
>
> (INF) If S has justification for believing that P, then S's evidence guarantees the truth of P.
>
> (~G) S's evidence does not guarantee the truth of P.
>
> (SCEPT) S lacks justification for believing that P.

Thus, most would grant the truth of ~G but argue that this claim provides no basis for scepticism. Rather, ~G would be a reason to deny infallibilism, for infallibilism imposes excessively high demands on justification. In other words, the Infallibility Argument for scepticism invites the reply that infallibilism does not impose a genuine requirement on justification (Reed 2012: 585).[5]

3.2 The Infallibility Objection

Some have suggested that, in spite of initial appearances, the Underdetermination Argument is not more interesting than the Infallibility Argument, because the Underdetermination Argument presupposes infallibilism, too. The underdetermination sceptic argues from the sameness of evidence in the good and the bad case to the claim that S's evidence does not favour P over SH – a suitable sceptical hypothesis –, and from the latter claim via the underdetermination principle to the conclusion that S is not justified in believing that P (see chapter 1, section 5). Consider the inference:

> (SEL) S has the same evidence in the good and the bad case.
>
> ---
>
> (~FAV) S's evidence does not favour P over SH.

The Infallibility Objection is the idea that the Underdetermination Argument is a bad argument because that inference presupposes infallibilism. In this section, I present the motivations that proponents of the Infallibility Objection have adduced to support it.

---

[4] However, there is significant disagreement on this point due to the recent resurgence of anti-sceptical infallibilism. See chapter 1 in Brown (2018).

[5] Not only that: it is tempting to think that the disagreement between the infallibility sceptic and the anti-sceptical fallibilist is merely verbal. After all, both parties agree that our evidence does not guarantee the truth of P. Thus, one might suspect that the parties agree about the matters of fact and merely disagree about whether one should be said to have any justification (Stroud 1984: 40; Kraft 2012: 50).

3.2.1 Brueckner

Why think that the contentious sceptical inference presupposes infallibilism? Consider the Sameness of Evidence Lemma:

(SEL) S has the same evidence in the good and the bad case.

The Sameness of Evidence Lemma says that it is possible for S's evidence E in the good case to be present when SH, a suitable sceptical hypothesis, is true. In other words, the Sameness of Evidence Lemma points to the fact that E is consistent with SH, that E does not guarantee that P. But why should the fact that E does not favour P over SH *in this sense* entail that S lacks justification for believing that P? A fallibilist will reject that line of thinking. It seems that the sceptic's strategy of using the Sameness of Evidence Lemma to establish an underdetermination claim and using that underdetermination claim to derive the sceptical conclusion presupposes infallibilism, because S's lack of justification for P ultimately derives from the fact that S's evidence does not guarantee the truth of P (Brueckner 1994: 835, 2005: 390, 2011: 86-7).[6]

To put it another way: the underdetermination sceptic faces a dilemma. Say that S's evidence *strongly underdetermines* the doxastic choice between P and SH just in case S's evidence does not favour one over the other. Instead, say that S's evidence *weakly underdetermines* the doxastic choice between P and SH just in case S's evidence is logically consistent with both P and SH. On the one hand, it is uncontroversial that the Sameness of Evidence Lemma entails that S's evidence weakly underdetermines the choice between P and SH. However, weak underdetermination is too weak to derive the sceptical conclusion from the underdetermination principle. That derivation will appear plausible only if one accepts infallibilism. On the other hand, strong underdetermination is strong enough to derive the sceptical conclusion from the underdetermination principle. After all, S's evidence strongly underdetermines the doxastic choice between P and SH only if ~FAV is true. However, it is controversial whether the Sameness of

[6] I made some changes to Brueckner's presentation. First, Brueckner (2005, 2011) presents the Infallibility Objection in the context of using the Underdetermination Argument to motivate one of the Closure Argument's premises – that S is not justified in believing ~SH. Thus, in Brueckner's version of the Underdetermination Argument, P is just ~SH. However, Brueckner's reasoning applies to the Underdetermination Argument in general. Second, Brueckner reserves the term 'evidence' for propositional justifiers. Nothing hinges on these differences in presentation.

Evidence Lemma entails that S's evidence strongly underdetermines the doxastic choice between P and SH. That entailment will appear plausible only to those who accept infallibilism about evidential favouring:

(INF*) S's evidence favours P over SH only if S guarantees the truth of P.

Either way, some infallibilist assumption is needed for the contentious sceptical inference to work.

3.2.2 Dodd

Dodd (2012, 2014) sees the same problem with the sceptical inference from sameness of evidence to underdetermination as Brueckner does. Namely, for the sceptical inference to go through the sceptic must assume some form of infallibilism. But infallibilism is unacceptable, as the following examples make clear:

> LOTTERY
> Imagine a fair lottery where every ticket has four numbers followed by a letter. Suppose S owns lottery ticket 5382L. Then, S learns that the winning ticket has the number 5382 but ignores which of the twenty-six tickets starting with the number 5382 is the winning ticket. (Adapted from Dodd 2012: 347)

Now consider the two hypotheses 'S's ticket is the winning ticket' and 'One of the twenty-five other tickets with the number 5382 is the winning ticket.' In LOTTERY, S's evidence is the same regardless of whether S has the winning ticket or not. But S's evidence does favour one hypothesis over the other, namely the hypothesis that one of the twenty-five other tickets is the winning ticket. After all, given S's evidence, the probability that hypothesis is true is 25/26, while the probability that S has the winning ticket is only 1/26.

If infallibilism about evidential favouring were true, things would be different. Since S's evidence does not guarantee that one of the tickets S does not own is the winning ticket, S's evidence does not favour that hypothesis over the hypothesis that S has the winning ticket. Clearly, that is absurd.

Similarly, consider the following scenario:

> FLAT TYRE
> Upon approaching her car, S sees that it has a flat tyre. (Adapted from Dodd 2012: 348-9)

Now consider the two hypotheses 'S's car's tyre was punctured' and 'S's car's tyre was not punctured but was removed from S's car and replaced with a different tyre that was flat.' In FLAT TYRE, S's evidence – her visual as well as her background evidence – is the same regardless of which hypothesis is true. But it seems obvious that this does not stop S's evidence from justifying the hypothesis that the car's tyre was punctured.

Again, if infallibilism about justification were true, things would be different. Since S's evidence does not guarantee that S's car's tyre was punctured, S's evidence does not justify the hypothesis that S's car's tyre was punctured. Clearly, that is absurd.

3.2.3 Smith

Although Smith (2022) frames his discussion in terms of rational support and mental states rather than evidential support and evidence, his criticism of the Underdetermination Argument is in line with Brueckner's and Dodd's. Smith examines a series of principles the sceptic can put forward to justify the underdetermination claim that S lacks rational support that favours P over SH (this is Smith's equivalent of ~FAV, i.e. the claim that S's evidence does not favour P over SH) and finds them all faulty.

Smith (2022: 147) starts by considering the following principle about rational support:

> (RS) If S has rational support that favours a proposition φ over any proposition that is inconsistent with φ, then S's current subjective experiences entail that φ is true.

RS has a distinctively infallibilist flavour. Given that S's subjective experiences do not entail that P is true, it is a short step from RS to scepticism. However, RS is of no help to the sceptic because it is 'a premise that obviously stacks the deck in favour of scepticism, and that no non-sceptic should be willing to accept' (Smith 2022: 147).

According to Smith, the underdetermination sceptic argues that S does not have rational support that favours P over SH because S's experiences in the good and the bad case are subjectively indistinguishable (this is Smith's equivalent of the Sameness of Evidence Lemma, i.e. the claim that S has the same evidence in the good and the bad case). In other words, Smith explains, the sceptic argues that S does not have rational support that favours P over SH because S's experiences are consistent with both P and SH. Thus, it seems that to derive an underdetermination claim the sceptic is implicitly appealing to the following general principle:

> (U0) For any two inconsistent propositions φ and ψ, if ψ is consistent with S's current subjective experiences, then S does not have rational support that favours φ over ψ. (Smith 2022: 149)

According to Smith (2022: 150), U0 is logically stronger than RS. Thus, U0 is at least as problematic as the question-begging principle RS. However, Smith is mistaken about which principle is logically stronger.[7] On reflection, RS is the logically stronger principle. To see why, let us reformulate the two principles in order to make the comparison easier:

> (RS) For any two inconsistent propositions φ and ψ, if S has rational support that favours φ over ψ, then S's current subjective experiences entail that φ is true.

---

[7] For the benefit of readers without access to Smith's text, it might be useful to quote the relevant passage in full:

> 'If I have rational support that favours X over an inconsistent proposition Y then, according to U0, Y must be inconsistent with my subjective experiences. As a result, if I have rational support that favours X over *every* proposition that is inconsistent with X then, according to U0, every proposition that is inconsistent with X would also have to be inconsistent with my subjective experiences, which is just to say that my subjective experiences would have to entail X. Though it may not look it at first, U0 is actually *logically stronger* than [RS] – the question-begging premise from the simple sceptical argument.' (Smith 2022: 150; emphasis in the original; changes to account for terminological differences marked by square brackets.)

Notice that the claim that U0 is logically stronger than RS is not substantiated by what Smith says. If Smith is right, U0 entails RS. But that is not sufficient to establish that U0 is logically stronger than RS, because RS may also entail U0. In any case, I shall argue that Smith is wrong about U0's entailing RS (see further).

> (U0) For any two inconsistent propositions φ and ψ, if S has rational support that favours φ over ψ, S's current subjective experiences are inconsistent with ψ.

RS and U0 have the same antecedent, so any difference in logical strength, if any, will be determined by the logical strength of their consequents. But, clearly, the consequent of RS is logically stronger than the consequent of U0. After all, if S's current subjective experiences entail that φ is true, then they will be inconsistent with any ψ that is inconsistent with φ. That is: the consequent of RS entails the consequent of U0. Thus, U0 cannot be logically stronger than RS. Further, the consequent of U0 does not entail the consequent of RS: S's current subjective experiences can be inconsistent with ψ, i.e. entail that ψ is false, without also entailing that φ is true. After all, φ and ψ need not exhaust the space of possibilities; S's subjective experiences can rule out ψ without thereby entailing φ, because S's subjective experiences can leave uneliminated some further proposition χ that is inconsistent with φ.[8]

Although Smith's claim about the logical strength of U0 does not survive close inspection, he also notes that there are several counterexamples to U0:

> ROLL
> S rolls a fair 20-sided die, and it lands just out of view. (Adapted from Smith 2022: 150)

> RAIN
> Someone tells S that it is raining outside, and S has no reason to doubt their word. (Adapted from Smith 2022: 150)

Clearly, in ROLL, S has rational support that favours the proposition that the die landed on a number between 1 and 19 over the proposition that the die landed on the number 20. However, both propositions are consistent with S's subjective experiences. Similarly, in RAIN, it seems that S has rational support that favours the proposition that it is raining over the proposition that it is not, although these propositions are

[8] My argument assumes a reading of 'inconsistent' according to which inconsistent propositions need not exhaust the space of possibilities. I think this is the correct reading of Smith. In any case, even assuming that inconsistent propositions always exhaust the space of possibilities, Smith's claim that U0 is logically stronger than RS is mistaken. Rather, the two principles are equivalent.

both consistent with S's subjective experiences.

Smith (2022: 151) observes that SH, i.e. the relevant sceptical hypothesis, is not simply consistent with S's subjective experiences, it entails them. So, perhaps, the underdetermination sceptic can appeal to the following principle to motivate the underdetermination claim:

> (U0*) For any two inconsistent propositions φ and ψ, if ψ entails S's subjective experiences, then S does not have rational support that favours φ over ψ.

U0* quickly runs into problems because it still commits the sceptic to the question-begging principle RS. By contraposition, U0* says that if S has rational support that favours φ over a ψ inconsistent with φ, ψ does not entail S's subjective experiences. Further, notice that if ψ is inconsistent with φ but consistent with S's subjective experiences, then there is a further proposition ψ* that entails ψ, is inconsistent with φ, and entails S's subjective experiences. Then, if U0* is true, if S has rational support that favours φ over ψ, ψ cannot be consistent with S's subjective experiences. But that is just to say that if S has rational support that favours φ over any ψ inconsistent with it, S's subjective experiences entail φ. Unsurprisingly then, the counterexamples to U0 can serve as counterexamples to U0*, too (Smith 2022: 151).

What if the sceptic refuses to offer any motivation for the underdetermination claim? As Smith (2022: 151) remarks, anyone putting forward an argument, sceptics included, must employ some premises that are not derived from other things, sooner or later. So why is a general principle needed to derive the underdetermination claim?

Smith concedes that, in ordinary situations, we can judge that our evidence does not favour either one of two incompatible hypotheses without committing to any general principle about favouring. So, why can the sceptic not do the same without committing to a general principle like U0 or U0*? According to Smith, the reason is that SH is special, it is unlike rival hypotheses we encounter in day-to-day life. On the one hand, it is 'artificially designed'; on the other hand, it is artificially designed 'to be consistent with all of our subjective experiences.' And it is this latter, abstract feature that is doing all the sceptical work. Thus, one cannot accept the underdetermination claim without also accepting the general claim captured by U0 or U0* (Smith 2022: 152).

3.3 Confrontation and Sidestepping

The sceptic has two strategies available to deal with the Infallibility Objection: confrontation and sidestepping. *Confrontation* consists in arguing that infallibilism plays no role in the sceptic's motivation of the underdetermination claim via the Sameness of Evidence Lemma, or that if it does this is no reason to dismiss it. *Sidestepping* consists in finding a way to motivate the underdetermination claim that dispenses with the Sameness of Evidence Lemma altogether and does not rely on infallibilism. Thus, confrontation takes the Infallibility Objection head on and attempts to show that it is misguided, while sidestepping attempts to evade it by finding alternative ways to support the underdetermination claim.

Although they serve a common goal, confrontation and sidestepping accomplish different things and are valuable in their own way to the sceptic. If confrontation succeeds, the sceptic has salvaged a *prima facie* appealing rationale for underdetermination scepticism from a dangerous objection. If sidestepping succeeds, the sceptic has shown that there is more than one way to motivate underdetermination scepticism. Thus, it is in the sceptic's interest to pursue both: if the two strategies are successful, the case for underdetermination scepticism comes out stronger than before.

The rest of this chapter argues that the underdetermination sceptic can successfully confront the Infallibility Objection. First, I look at two unsuccessful attempts to sidestep the Infallibility Objection due to Briesen (2010). I argue that one problem affecting Briesen's proposals is that they implicitly rely on the Sameness of Evidence Lemma. Thus, they do not satisfy one of the success conditions for a working sidestepping strategy: dispensing with the Sameness of Evidence Lemma. Further, Briesen's proposals face some independent problems. This makes the task of defending the contentious sceptical inference more urgent, although nothing I say here entails that the sceptic has no other sidestepping strategy available. Then, I argue that, whether or not the sceptic can sidestep the Infallibility Objection, the sceptic can successfully confront it, because the inference from sameness of evidence to underdetermination does not presuppose infallibilism.

3.4 Two Unsuccessful Attempts to Sidestep the Infallibility Objection

Briesen (2010: 231-2) suggests two ways to motivate the underdetermination claim and sidestep the Infallibility Objection: *via* an explanation principle and *via* an entailment principle. If these principles offer ways to sidestep the Infallibility Objection, then they figure in arguments for the underdetermination claim that do not rely on infallibilism and the Sameness of Evidence Lemma. Here, I shall not attempt to reach a conclusive answer as to whether these arguments are sound – although I will suggest that one of the arguments is seriously flawed and probably unsound. More modestly, my goal is to decide whether they can be used to sidestep the Infallibility Objection. For this, it is sufficient that they provide a *prima facie* motivation for underdetermination without relying on infallibilism and the Sameness of Evidence Lemma. Thus, my goal is modest, but it is not too modest. It is commensurate to the task at hand, which is to decide whether the sceptic can sidestep the Infallibility Objection in the way Briesen suggests, not to decide whether the sceptic ultimately succeeds in establishing the underdetermination claim.

3.4.1 The Causal Explanation Principle

Briesen (2010: 231) suggests that the sceptic can argue for underdetermination using the following principle about competing causal explanations:

> (EXP) For all incompatible φ, ψ, if ψ causally explains S's evidence E at least as well as φ does, then E does not favour φ over ψ.

It seems to me that, if Briesen is right, the sceptic can argue for ~FAV using a simpler and more general principle about explanation:

> (EXP*) For all incompatible propositions φ, ψ, if ψ explains S's evidence E at least as well as φ does, then E does not favour φ over ψ.

Indeed, it would be puzzling if EXP were true and EXP* false because their only difference is that EXP is restricted to causal explanations. Since EXP and EXP* are epistemic principles that connect epistemic facts of the same kind in the same way – certain facts about the relative goodness of rival explanations with certain facts about evidential favouring – it would be puzzling if the non-causal nature of some

explanations were responsible for the falsity of EXP* when EXP is true.

Moreover, using EXP* might have the advantage of avoiding certain objections to EXP. For instance, one might object to EXP on the ground that two incompatible hypotheses might be on a par with respect to causal explanation but differ with respect to non-causal explanation. Thus, S's evidence might favour one hypothesis over the other although they causally explain S's evidence equally well. Although EXP* comes with these benefits for the sceptic, I will stick to EXP for the sake of conformity to Briesen's formulation. In any case, what I say about the EXP-based argument for ~FAV is true *mutatis mutandis* for the EXP*-based argument.

According to Briesen, the sceptic has a simple and valid argument from EXP to ~FAV – the claim that S's evidence does not favour P over SH. First, since P and SH are incompatible, EXP entails:

> (EXP_{SH}) If SH causally explains S's evidence at least as well as P does, E does not favour P over SH.

To arrive to ~FAV, the sceptic only needs to assume:

> (E2) SH causally explains S's evidence at least as well as P does.

Then, by *modus ponens*, the sceptic can infer:

> (~FAV) E does not favour P over SH.

Does the argument from EXP provide a promising motivation for ~FAV that dispenses with the Sameness of Evidence Lemma and avoids infallibilism? Briesen has something to say in favour of this. First, he observes that EXP is very plausible. Further, EXP is compatible with fallibilism. Finally, although E2 is controversial, there is something compelling about it, too, 'as many failed anti-sceptical attempts based on the inference to the best explanation illustrate' (Briesen 2010: 231).

Notice that EXP is similar to the principle Cohen (1998: 146) considers as a possible motivation for ~J, the claim that S is not justified in believing that ~SH:[9]

> For all φ, if the truth of φ would explain S's evidence as well as any

[9] More precisely, it is similar to the improved version of the principle Cohen considers. See chapter 2, section 2.3.2.

hypothesis incompatible with φ, then S's evidence does not justify ~φ.

As I observed in the last chapter, section 2.3.2, it is not clear that this principle can be used to motivate ~J independently of the Underdetermination Argument, because it is tempting to interpret the claim that two hypotheses explain the evidence equally well as the claim that the evidence does not favour one over the other. For the same reason, one might worry that the EXP-based argument fails to motivate ~FAV.

Further, reflecting on E2 reveals a problem for the prospects of using the EXP-based argument to sidestep the Infallibility Objection. For a natural question to ask is: what motivates E2 – the claim that SH causally explains S's evidence at least as well as P does? A natural answer is that the Sameness of Evidence Lemma is part of the motivation for accepting E2. After all, if S does not have the same evidence in the good and the bad case, then either in the good case S has some evidence that S lacks in the bad case or in the bad case S has some evidence that S lacks in the good case (or both). In either case, E2 loses its plausibility. For why would it still be true that SH causally explains S's evidence in the good case at least as well as P does? In the first case SH would explain too little while in the second too much (if both, then SH would explain too little as well as too much).

Could the sceptic insist that the plausibility of E2 is primitive and thus independent of the Sameness of Evidence Lemma? The move is not compelling. For the reasons just exposed, when one supposes that S does not have the same evidence in the good and the bad case, E2's immediate plausibility disappears. Thus, regardless of whether it is primitive or not, it seems to depend on the Sameness of Evidence Lemma. Of course, this does not entail that the plausibility of E2 must be bound to the Sameness of Evidence Lemma. Perhaps, one could find a motivation for E2 that makes it plausible regardless of it. However, until one such argument is presented, the EXP-based argument seems to rely on the Sameness of Evidence Lemma and hence is ill-suited to sidestep the Infallibility Objection.

3.4.2 The Entailment Principle

Let us now turn to the second argument suggested by Briesen (2010: 231-2) in order to sidestep the Infallibility Objection. According to Briesen, the sceptic can argue for ~FAV using the following entailment principle:

(ENT) For all incompatible φ, ψ, if ψ entails that S has evidence E and φ does not entail that S has E, then E does not favour φ over ψ.

With ENT, the sceptic has a simple and valid argument for ~FAV. First, given the uncontroversial claim that P and SH are incompatible, ENT entails:

(ENT_{SH}) If SH entails that S has evidence E and P does not entail that S has E, then E does not favour P over SH.

To arrive to ~FAV, the sceptic only needs to assume:

(E2*) SH entails that S has evidence E and P does not entail that S has E.

Then, the sceptic can infer by *modus ponens*:

(~FAV) E does not favour P over SH.

Does the ENT-based argument provide a promising motivation for ~FAV that dispenses with the Sameness of Evidence Lemma and avoids infallibilism? Again, Briesen has something to say in favour of this. First, according to Briesen, ENT is even more plausible than the already plausible EXP. Second, ENT is compatible with fallibilism. Finally, although E2* is controversial, it looks very plausible, at least as long as one conceives of S's evidence as consisting of non-factive mental states.

Pace Briesen, the ENT-based argument does not fare better than then EXP-based argument when it comes to sidestepping the Infallibility Objection. In fact, it arguably fares worse. One set of problems relates to ENT, which is far from being a plausible principle. First, ENT runs into problems when ψ is inconsistent. Suppose that ψ is inconsistent and that, while φ is consistent, φ does not entail that S has E. Since ψ is inconsistent, ψ entails both ~φ as well as that S has evidence E. Thus, according to ENT, E does not favour φ over ψ. Clearly, that is the wrong result.

As an illustration, consider the following toy model:

P: 'most swans are white.'

Q: 'most swans are white & most swans are not white & S only perceived white ones.'

E: S's perceptual experiences as of white swans.

In this model, E favours P over Q. However, P and Q are incompatible, and Q entails

that S has E while P does not. Thus, the antecedent of ENT is true, while its consequent false.

Second, ENT runs into problems when ψ is more complex than φ. To make the problem for ENT as perspicuous as possible, suppose that ψ is a maximally consistent conjunction that entails ~φ as well as that S has E (that is, ψ is a complete description of the world including the propositions ~φ and the proposition that S has E). Further, suppose that φ does not entail that E. Since ψ is incompatible with φ and E2* is true, according to ENT, E does not favour φ over ψ. Clearly, that cannot be right.[10]

But even ignoring these objections, ENT could not be used to sidestep the Infallibility Objection, at least if one motivates ENT in the way Briesen does. To see why, let us consider how Briesen illustrates the plausibility of ENT (Briesen 2010: 232). Let P be the proposition that most swans are white, Q the proposition that most swans are black, and Q* the proposition that most swans are black but S only perceived white ones. Briesen asks us to suppose that S's evidence E consists in perceptual experiences of white swans. Since neither P nor Q entails that S has E, Briesen observes, ENT allows that S's evidence favours P over the incompatible alternative Q. Thus, ENT is compatible with fallibilism. However, it also seems that E does not favour P over Q* (if you are not convinced, says Briesen, remember that E – S's perceptual experiences of white swans – exhausts S's evidence. That is: S has no additional evidence in the form of background assumptions that bears on P and Q*). And this is exactly what ENT predicts.

Let us take a closer look at Briesen's reasoning. Briesen bolsters the plausibility of ENT by observing that ENT makes a correct prediction: that E does not favour P over Q*. But why should we think that that prediction is correct? What motivates the claim that S's evidence does not favour P over Q*? That claim is not immediately plausible. On the contrary, it is natural to think that the claim is false because it is

[10] As an illustration, consider the following toy model:

> P: 'most swans are white.'
> Q: a maximally consistent conjunction that includes ~P as well as the proposition that S has E.
> E: S's perceptual experiences as of white swans.

In this model, E favours P over Q. However, P and Q are incompatible, and Q entails that S has E while P does not. Thus, the antecedent of ENT is true, while its consequent false.

natural to assume that S's background assumptions shift the evidential balance in favour of P. That is why Briesen needs to remember the reader that E exhausts S's evidence. But why should the fact that E exhausts S's evidence make the claim that E does not favour P over Q* more plausible? Of course, it is hard to say.

However, here is an explanation. Notice that Q* features some similarities to a sceptical hypothesis: in a Q*-world things appear like they do in a P*-world but are not that way. This suggests that what grounds the judgement that E does not favour P over Q* is just the thought that S has the same evidence in P-worlds and Q*-worlds. Does it sound familiar? This is just a variation of the sceptical reasoning from the Sameness of Evidence Lemma to ~FAV. However, if the ENT-based argument is to be used to sidestep the Infallibility Objection, the plausibility of ENT cannot depend on the Sameness of Evidence Lemma.

In addition to these difficulties, the ENT-based argument faces a further problem. The problem is that it is not clear whether the argument poses a challenge to the anti-sceptic at all. In fact, the argument works only as long as E2* is true, that is as long as SH is a proposition that entails that S has E and P a proposition that does not entail that S has E. However, just like the sceptic can easily find a proposition that combines the falsity of P with S's having E, namely SH, so the anti-sceptic can easily find a proposition P* that combines the truth of P with S's having E. Since P* entails both that S has evidence E and that P is true, the ENT-based argument does not work against P*. But then how can it work against P? In other words, the suspicion is that the ENT-based argument appears to work only because the sceptic is ignoring that the anti-sceptic can appeal to P*.[11] Thus, if the sceptic wants to use the ENT-based argument to sidestep the Infallibility Objection, she owes the anti-sceptic an explanation

[11] This point is related to a point Okasha (ms) makes against the Underdetermination Argument. According to Okasha, the inference from sameness of evidence to underdetermination fails because P and SH are not empirically equivalent rivals. After all, SH entails that S has whatever experiences she has, but P does not. As Okasha recognizes, the obvious fix is to replace P with P* in the Underdetermination Argument. But Okasha argues that the fix is not good enough, because the Underdetermination Argument conflicts with the dictates of evidential holism, the idea that propositions are responsible to evidence only in group, not in isolation. Okasha's treatment of the Underdetermination Argument is well-deserving of attention, but, due to its length and complexity, I cannot do it justice here.

for why the anti-sceptic cannot appeal to P*.[12]

3.5 Confrontation: A Bayesian Account of the Sceptical Inference

I now turn to the second strategy the sceptic can employ in response to the Infallibility Objection: confrontation. As noted above, confrontation can be carried out in two ways: by arguing that the contentious sceptical inference does not depend on infallibilism, and by arguing that even if it does this is no reason to dismiss it. Although both options are worth exploring, I shall focus on the first one. This is for two reasons. First, arguing that the contentious sceptical inference does not depend on infallibilism requires no defence of infallibilism. Second, if the contentious inference does not depend on infallibilism the Underdetermination Argument is deeply and not just superficially different from the Infallibility Argument for scepticism.

I shall argue that the sceptic can successfully confront the Infallibility Objection. Pace the objectors, the contentious sceptical inference does not presuppose infallibilism. The key idea for my argument is that the sceptical inference is similar to some acceptable inference from sameness of evidence to underdetermination that does not presuppose infallibilism. Hence, the sceptical inference from sameness of evidence to underdetermination is an acceptable inference that does not presuppose infallibilism. More precisely, my argument goes like this:

> 1. **Σ** is an acceptable inference scheme that does not presuppose infallibilism.
>
> 2. The sceptical inference is an instance of **Σ**.
>
> 3. The sceptical inference is an acceptable inference that does not presuppose infallibilism.

---

[12] Can the sceptic avoid the problem by modifying ENT to accommodate P*? Consider:

> (ENT)* For all incompatible φ, ψ, if ψ entails that S has evidence E, φ entails that S has evidence E, then E does not favour φ over ψ.

From ENT* there is a simple and valid argument to ~FAV that parallels the ENT-based argument. Further, unlike the ENT-based argument, the new argument is not vulnerable to the objection from P*. However, ENT* is also subject to counterexamples in the style of those levelled by Smith against the principles U0 and U0*.

3.5.1 Establishing Premise 1

As Dodd and Smith argue, it seems that S's evidence sometimes favours one proposition φ over one of its competitors ψ although S has the same evidence in φ-worlds and ψ-worlds. For instance, recall:

> ROLL
>
> S rolls a fair 20-sided die, and it lands just out of view.

Clearly, in ROLL, S's evidence favours the proposition that the die landed on a number between 1 and 19 over the proposition that the die landed on the number 20. Yet S's evidence is the same regardless of whether S rolls a 20.

Cases like ROLL are instructive: they teach us that sameness of evidence is not the same as underdetermination. As a consequence, we should be careful in inferring from the former to the latter. Importantly, that lesson is not a prohibition to ever make that inference. For one thing, it is compatible with the claim that in some circumstances that inference is legitimate. For another thing, even proponents of the Infallibility Objection should grant that in some cases where S has the same evidence regardless of whether φ or ψ is true, we correctly judge that S's evidence does not favour φ over ψ. Thus, although there are cases like ROLL, i.e. cases in which S's evidence favours φ over an incompatible ψ and S has the same evidence in φ-worlds and ψ-worlds, we should not forget that there are cases in which S has the same evidence in φ-worlds and ψ-worlds and S's evidence does not favour φ over ψ, too.

Although this does not establish premise 1 – that some acceptable inference scheme from sameness of evidence to underdetermination does not presuppose infallibilism – , it lays the ground for it. For if we can reconstruct our correct judgement about cases in which S has the same evidence in φ-worlds and ψ-worlds and S's evidence does not favour φ over ψ as the result of an acceptable inference from sameness of evidence to underdetermination, we may be able to vindicate the truth of premise 1. I shall follow this strategy to argue that premise 1 is true.

First of all, let us consider a case in which S has the same evidence in φ-worlds and ψ-worlds and S's evidence does not favour φ over an incompatible ψ:

CINEMA

S lives in a city with two cinemas, A and B, that belong to the same franchise. As it sometimes happens with franchises, A and B are almost identical inside: equally large, same number of seats, same chairs, same walls, same lights etc. As a birthday gift, her friends reserve cinema A to give S a surprise screening of her favourite film. On her birthday, S is taken blindfolded to cinema A. Finally, when she is in the cinema, S takes the blindfold off. What S sees is all the evidence she has about her whereabouts.

Arguably, in CINEMA, S's evidence does not favour her being in cinema A over her being in cinema B and S has the same evidence regardless of which cinema she is in. Now, if this judgement can be reconstructed as the result of an acceptable inference from sameness of evidence to underdetermination, then, minimally, the fact that S has the same evidence in A and in B has some bearing on the truth of the corresponding underdetermination claim.

In fact, I think some linguistic evidence lends support to this claim because a very natural answer to the question 'why is S's evidence neutral between her being in A and her being in B?' is that S has the same evidence regardless of which cinema S is in. But even if you do not think the linguistic evidence should bear much weight in this case, that sameness of evidence accounts for S's lack of evidential favouring is easily explained by observing that if S's experiences differed in A and B, then S's evidence in A would favour her being in A over her being in B. Thus, observing that S has the same evidence in A and B contributes to explain why S's evidence does not favour A over B by ruling out the case in which S's evidence in A differs from S's evidence in B.

Importantly, this cannot be the whole story: while sameness of evidence accounts for underdetermination in CINEMA, it cannot be the only thing that accounts for it. Otherwise, sameness of evidence would account for underdetermination in ROLL, too. But, as we have seen, there is no underdetermination in ROLL. So, what explains the difference between ROLL and CINEMA?

A simple Bayesian analysis turns out to be illuminating. In ROLL, S's evidence favours the proposition that the die landed on a number between 1 and 19 – call this $P_{1\text{-}}$

$_{19}$ – over the proposition that the die landed on 20 – call this P_{20}. In formal terms, $Pr(P_{1\text{-}19}|E)>Pr(P_{20}|E)$. Further, $P_{1\text{-}19}$ and P_{20} fit S's evidence equally well in the following sense: $Pr(E|P_{1\text{-}19})=Pr(E|P_{20})$. After all, S has the same evidence regardless of which hypothesis is true. Thus, by Bayes' theorem, if $Pr(P_{1\text{-}19}|E)>Pr(P_{20}|E)$, this can only be because the prior probability of $P_{1\text{-}19}$ is greater than the prior probability of P_{20}. And, in fact, the prior probability of $P_{1\text{-}19}$ – 19/20 – is greater than the prior probability of P_{20} – 1/20.

Now, consider CINEMA. In CINEMA, like in ROLL, the two rival hypotheses fit S's evidence equally well in the sense that $Pr(E|P_A)=Pr(E|P_B)$, where P_A is 'S is in A' and P_B is 'S is in B'. However, unlike in ROLL, S's evidence favours neither P_A over P_B nor P_B over P_A: $Pr(P_A|E)=Pr(P_B|E)$. Thus, by Bayes' theorem, if $Pr(P_A|E)=Pr(P_B|E)$, this can only be because the prior probability of P_A is equal to the prior probability of P_B. And, in fact, in CINEMA the prior probability of P_A is equal to the prior probability of P_B.

Another way to think about this is that CINEMA is disanalogous to ROLL but analogous to:

> TOSS
>
> S flips a fair coin, and it lands just out of view.

In TOSS, S's evidence is the same regardless of whether the coin landed head or tail. Further, the prior probability of the coin landing heads is equal to the prior probability of the coin landing tails. Thus, by Bayes' theorem, S's evidence does not favour the coin landing heads over the coin landing tail. Indeed, this is the correct result.

If this analysis is correct, what explains the difference between ROLL and CINEMA is that in the latter the prior probability of the rival hypotheses coincides. In turn, this provides the missing piece for reconstructing the correct judgement about CINEMA as the result of an acceptable inference from sameness of evidence to underdetermination:

> S has the same evidence regardless of which cinema S is in.
>
> The prior probability of being in A is the same as the prior probability of being in B.
>
> ∴ S's evidence does not favour her being in A over her being in B.

More generally, we can extrapolate an acceptable inference scheme that underpins the inference just reconstructed:

Σ

S has the same evidence regardless of whether φ or ψ is true.

φ and ψ have the same prior probability.

∴ S's evidence does not favour φ over ψ.

Importantly, this reasoning is free of infallibilist commitments. The point is not that S's evidence does not favour her being in A over her being in B because S has the same evidence regardless of which cinema she is in and S's evidence does not guarantee that S is not in cinema B. Rather, the point is that S's evidence does not favour her being in A over her being in B because S has the same evidence regardless of where she is *and* being in A and being in B have the same prior probability. The reasoning is just an application of Bayes' theorem.

This concludes my defence of premise 1: that some acceptable inference scheme from sameness of evidence to underdetermination, namely **Σ**, does not presuppose infallibilism.

3.5.2 Establishing Premise 2

Premise 2 of my argument says that the contentious sceptical inference, the inference from sameness of evidence to underdetermination, is an instance of the acceptable inference scheme **Σ**. Thus, what I am suggesting is that the sceptical inference should be reconstructed like this:

S has the same evidence in the good and the bad case.

Being in the good case and being in the bad case have the same prior probability.

∴ S's evidence does not favour P over SH.

The idea is that S's situation with respect to the good case and the bad case is the same as S's situation with respect to A and B in CINEMA. Since the good case and the bad case fit S's evidence equally well, and since the good case and the bad case have the same prior probability, by Bayes' theorem, S's evidence does not favour P over SH. The parallel between the rational reconstruction of our correct judgement

about CINEMA and the rational reconstruction of the contentious sceptical inference as an instance of **Σ** is at least *prima facie* plausible. Does it withstand scrutiny?

There are several ways to dispute the claim that the sceptical inference is an instance of **Σ**. Some will be ready to grant that S has the same evidence regardless of which cinema she is in while disputing that S has the same evidence in the good and the bad case (Pritchard 2015; Williamson 2000). However, notice that this sort of reply is not available to proponents of the Infallibility Objection in the present context, for we are assuming (controversially) that S's evidence consists of non-factive mental states. The assumption is in the spirit of the Infallibility Objection as its proponents typically assume a mentalist picture of evidence according to which S has the same evidence in the good and the bad case. In fact, one of the more attractive features of the Infallibility Objection is that it blocks the sceptical inference while remaining neutral between mentalist and non-mentalist conceptions of evidence. If the contentious sceptical inference can be reconstructed without any appeal to infallibilism, as I argue, the Infallibility Objection would remain misguided even if the Sameness of Evidence Lemma were false. Thus, I will ignore this sort of challenge to the claim that the sceptical inference is an instance of **Σ**.

Instead, I shall focus on two promising lines of attack that are available to proponents of the Infallibility Objection. The first one challenges the sceptic's claim that being in the good case and the bad case have the same prior probability on the ground that she must thereby appeal to a problematic principle. The second one challenges the idea that the sceptical inference is an instance of **Σ** on the ground that the sceptical inference involves three propositions but **Σ** only two.

The Principle of Indifference

A promising line of attack to my reconstruction of the sceptical inference from sameness of evidence to underdetermination focuses on the claim that being in the good case and being in the bad case have the same prior probability. Dodd (2012) suggests that what explains the appeal of that claim is the so-called 'Principle of Indifference', according to which equal possibilities are to be assigned equal probabilities (Van Fraassen 1989: 293), or if one has no more reason to believe φ than ψ, and no more reason to believe ψ than φ, then one's credence in φ should equal one's credence in ψ (Rinard

2014: 110), or in the absence of evidence to prefer one hypothesis over another one should assign them equal probability (Eva 2019:390).

The Principle of Indifference is a strikingly intuitive and plausible principle. As Eva (2019: 390) points out, the Principle boasts an impressive number of independent justifications ranging from considerations about epistemic utility (Pettigrew 2014), risk aversion (Williamson 2007), evidential support (White 2009), informativity (Jaynes 1957), and the Principal Principle (Hawthorne et al. 2015). However, it also faces a famous paradox, originally discussed by Bertrand (1889), and exemplarily illustrated by Van Fraassen (1989).

According to Dodd (2012: 345-6), one may be inclined to apply the Principle of Indifference to the good and the bad case and judge them equiprobable. But as Dodd observes following Bertrand and Van Fraassen, the Principle of Indifference is highly problematic, because there are different and incompatible ways of applying it.

A well-known illustration of the problem is due to Van Fraassen (1989: 303-4), who asks the reader to imagine a factory that produces cubes with edge length $L \leq 2$. Given the Principle of Indifference, the probability that $0 \leq L \leq 1$ is the same as the probability that $1 \leq L \leq 2$, namely 1/2. Now, consider the area A of a face of the cube. Since $L \leq 2$, $A \leq 4$. Given the Principle of Indifference, the probability that $0 \leq A \leq 1$ is the same as the probability that $1 \leq A \leq 2$, $2 \leq A \leq 3$, and $3 \leq A \leq 4$, namely 1/4. Where is the problem? Notice that $0 \leq L \leq 1$ and $0 \leq A \leq 1$ are logically equivalent, hence they have the same probability. However, the Principle of Indifference assigns them different probabilities, namely 1/2 and 1/4. Thus, the Principle of Indifference leads to a paradox: applying it can lead to assigning different probabilities to logically equivalent propositions.

The sceptic has several replies available here. First, the sceptic may distinguish between the idea that one ought not to prefer one hypothesis to another unless one has reason to do so, or that before having any evidence one ought to be equally confident in each hypothesis, and different ways of articulating that idea. That is, the sceptic may distinguish between the Principle of Indifference and formulations of it. Arguably, these formulations are attempts to capture the Principle of Indifference and spell it out in more rigorous terms. Some of these formulations run into the paradox just

mentioned. However, it is open to the sceptic to reject these problematic formulations of the Principle of Indifference while sticking to the plausible idea that motivates them.

For an analogy, consider knowledge closure, the idea that knowledge can be extended by deductive reasoning. Although that idea is very plausible, it is difficult to state it without incurring in counterexamples or falling prey to paradox, as witnessed by the extensive literature on knowledge closure principles (Luper 2020). Nonetheless, few epistemologists reject knowledge closure. On the contrary, while they recognise that it is difficult to give it a precise statement, most epistemologists continue accepting it. The sceptic can adopt the same attitude with respect to the Principle of Indifference.

Second, the sceptic may draw a distinction between problematic and non-problematic applications of the Principle of Indifference. The sceptic may agree with Dodd that the Principle of Indifference is a problematic principle because there are incompatible ways of applying it. But the sceptic may insist that while this fact invites caution when using the Principle of Indifference, some applications of the Principle of Indifference are perfectly legitimate. In particular, the sceptic may insist that this is her case. In fact, it is hard to see how to generate an analogue of Bertrand paradox for the good and the bad case.[13]

Third, the sceptic may insist on the parallel between being in cinema A and being in cinema B, on the one hand, and being in the good case and being in the bad case, on the other, to argue that Dodd's objection overgeneralises. More precisely, the sceptic may point out that, by parity of reasons, Dodd's objection is an objection to the sceptical inference only insofar it is an objection to the claim that being in cinema A and being in cinema B have the same prior probability in CINEMA. But there seems to be nothing wrong with that claim. Thus, the sceptic may insist that there is nothing wrong with her claim that being in the good case and being in the bad case have the same prior probability.

[13] The sceptic may go for a more ambitious reply to Dodd's objection than the ones I have presented here. She may endorse some specific formulation of the Principle of Indifference and argue that it is not subject to the paradox. The sceptic has a large number of proposals to choose from in the literature, although many of them avoid the paradox only at a high cost. Perhaps, an exception is the recent proposal by Eva (2019) to formulate the Principle of Indifference in terms of abstinence from comparative confidence judgements. Whether the sceptic can successfully defend an appeal to the Principle in this way is an interesting question I cannot settle here.

Finally, even if none of these replies succeeds, it is not clear that Dodd's objection is very useful to the anti-sceptic. After all, once we reject the Principle of Indifference, we have no other principle for applying priors to competing hypotheses. So, either any distribution of priors goes, as subjective Bayesians think, or no distribution of priors does. If the former is true, the probability distribution chosen by the anti-sceptic is on a par with the probability distribution chosen by the sceptic. Thus, there seems to be an important sense in which the anti-sceptic's beliefs are arbitrary and the sceptic is right. If the latter is true, the probability distribution chosen by the anti-sceptic is illegitimate, and the anti-sceptic's beliefs are irrational. Either way, this is not good news for the anti-sceptic.

A Fallacy of Three Terms?

Some may object to the claim that the sceptical inference is an instance of **Σ** on the ground that three hypotheses figure in the sceptical inference, but only two hypotheses figure in **Σ**. In my reconstruction, the sceptical inference's conclusion reads 'S's evidence does not favour P over SH', but if the sceptical inference were an instance of **Σ** it ought to read 'S's evidence does not favour being in the good case over SH.' Whereas SH is just a shortcut for 'S is in the bad case', one cannot equate P and being in the good case, for being in the good case requires more than the truth of P. Thus, one may object that my argument fails to establish that the sceptical inference is an acceptable inference that does not presuppose infallibilism because the sceptical inference subtly shifts from talk of favouring being in the good case over being in the bad case to talk of favouring P over being in the bad case.

The point is well taken, but it is less damaging than it may initially appear. The reason is that if S's evidence does not favour her being in the good case over her being in the bad case, then S's evidence does not favour P over SH. For recall that P is compatible with being in the good case but incompatible with being in the bad case. Further, being in the good case and being in the bad case ought to be assigned the same prior probability. Thus, if S's evidence favours P over SH, then S's evidence favours being in the good case over being in the bad case as well (see chapter 2, section 2).

3.6 Concluding Remarks

Let me recap. Infallibilism expresses the idea that the epistemic grounding sufficient for justification or evidential favouring precludes the possibility of error, i.e. is truth-guaranteeing. The underdetermination argument faces an objection I have called 'the Infallibility Objection'. The underdetermination sceptic infers from the sameness of evidence in the good and the bad case that S's evidence does not favour P over her being in the bad case. That sceptical inference has come under attack after some critics have argued that it presupposes infallibilism.

Why think that the sceptical inference presupposes infallibilism? On the one hand, the claim that S has the same evidence in the good and the bad case is tantamount to the claim that S's evidence is compatible with ~P. Thus, it seems that the sceptic's strategy of using the Sameness of Evidence Lemma to establish the underdetermination claim and using the underdetermination claim to derive the sceptical conclusion presupposes infallibilism, because S's lack of justification for P ultimately derives from the fact that S's evidence does not guarantee the truth of P. On the other hand, proponents of the Infallibility Objection have pointed to a series of counterexamples to the sceptical inference: situations in which S's evidence favours one of two incompatible hypotheses although S has the same evidence regardless of which is true.

If correct, the Infallibility Objection ought to change our conception of underdetermination scepticism as a distinctive and interesting form of scepticism. According to the standard conception, while some arguments for scepticism trade on very high epistemic standards (such as certainty and infallibility), the Underdetermination Argument does not. But if the proponents of the Infallibility Objection are right, the Underdetermination Argument is neither special nor interesting. It is not special because, like many other arguments for scepticism, it imposes high demands on justification, and it is not interesting because such demands are implausibly high.

Relatedly, if the Infallibility Objection is correct, there is very little to learn, philosophically speaking, from engaging with the Underdetermination Argument. For if the Infallibility Objection is correct, no controversial philosophical theory is needed to answer the underdetermination sceptic, at least as long as fallibilism is taken for granted. Thus, several 'indispensability' arguments for interesting but controversial

claims – arguments that attempt to establish the truth of some claim by way of showing that said claim is needed to avoid underdetermination scepticism – are considerably weakened.

In this chapter, I have argued that the sceptic can successfully confront the Infallibility Objection. Pace objectors, the contentious sceptical inference does not presuppose infallibilism. In fact, proponents of the Infallibility Objection have displayed a lack of philosophical imagination and failed to recognise the reasoning pattern that underlies the sceptical inference. I provided a rational reconstruction of this reasoning pattern and argued that infallibilism plays no role in it.

The key idea for my argument is that the sceptical inference is similar to some acceptable inferences from sameness of evidence to underdetermination that do not presuppose infallibilism. Hence, the sceptical inference is an acceptable inference that does not presuppose infallibilism. More precisely, according to my reconstruction, the sceptical inference instantiates the following argument scheme:

> S has the same evidence regardless of whether φ or ψ is true.
> φ and ψ have the same prior probability.
> ∴ S's evidence does not favour φ over ψ.

Importantly, the argument scheme does not rely on infallibilism. The point is not that S's evidence does not favour φ over ψ because S has the same evidence in φ-worlds and ψ-worlds and S's evidence does not guarantee the truth of φ. Rather, the point is that S's evidence does not favour φ over ψ because S has the same evidence regardless of which hypothesis is true (and hence: $Pr(E|\varphi)=Pr(E|\psi)$) *and* φ and ψ have the same prior probability. The sceptical inference is just an instance of this valid Bayesian reasoning:

> S has the same evidence in the good and the bad case.
> Being in the good case and being in the bad case have the same prior probability.
> ∴ S's evidence does not favour P over SH.

Here is the take-home message. We should think of the sceptical inference from sameness of evidence to underdetermination as an example of Bayesian reasoning. If

this is correct, the Infallibility Objection fails because it misrepresents that inference. Three consequences follow. First, fallibilism has no greater anti-sceptical merit than infallibilism when it comes to underdetermination scepticism. Second, underdetermination scepticism is interestingly different from other kinds of scepticism that trade on high epistemic standards. Third, for all it has been shown in this chapter, philosophical arguments that trade on the fear of underdetermination scepticism retain their original appeal.

4

Sameness of Evidence and Anti-Luminosity

The previous chapter discussed one way to challenge the Underdetermination Argument: challenging the inference from the Sameness of Evidence Lemma, i.e. the claim that S has the same evidence in the good and the bad case, to the underdetermination claim that S's evidence does not favour P over SH. This chapter discusses another: challenging the Sameness of Evidence Lemma by challenging the claim that S's evidence is luminous, i.e. that S is always in a position to know what her evidence is.

Williamson (2000) has put forward an influential argument against the luminosity of any non-trivial condition. According to Williamson, the argument provides a way to resist the sceptic's motivation for the Sameness of Evidence Lemma. In a nutshell, the idea is that the sceptic's argument for the Sameness of Evidence Lemma depends on the assumption that S's evidence is luminous. But no non-trivial condition is luminous, not even the condition of having evidence. Thus, the sceptic's argument for the Sameness of Evidence Lemma is unsound.

This chapter evaluates the anti-sceptical import of Williamson's anti-luminosity argument and assesses the merits of Williamson's anti-sceptical strategy. I shall argue that Williamson's anti-sceptical strategy fails. On the one hand, the anti-luminosity argument does not undermine the sceptic's argument for the Sameness of Evidence Lemma. On the other hand, even if it did, this would not be sufficient to give a satisfactory reply to the sceptic. The reason is that the sceptic does not need the luminosity claim to motivate the Sameness of Evidence Lemma.

Here is the plan for the chapter. Section 1 presents Williamson's anti-luminosity argument. Section 2 describes Williamson's account of the sceptic's argument for the Sameness of Evidence Lemma and explains how the anti-luminosity argument is meant to bear on it. Section 3 summarises and offers a preview of my overall argument against Williamson's treatment of scepticism. Section 4 offers a new argument for the claim that the anti-luminosity argument does not undermine the sceptic's argument for the Sameness of Evidence Lemma described by Williamson. Section 5 offers two additional arguments for the Sameness of Evidence Lemma that dispense with the luminosity assumption altogether. The concluding remarks tie things together and discuss

the relationship between my results and Williamson's E=K-thesis.

4.1 The Anti-Luminosity Argument

Williamson's anti-luminosity argument purports to show that there are no non-trivial luminous conditions. In other words, the argument purports to show that, for any condition that obtains in some but not all cases, there is a possible case in which it obtains but one is not in a position to know that it obtains. Candidate luminous conditions of interest to philosophers include phenomenal states, appearances, meanings, rule following, normative guidance, knowledge, and justification among others. In this section, I present Williamson's anti-luminosity argument.

4.1.1 Luminosity and Being in a Position to Know

According to Williamson (2000: 95), a condition C is luminous just in case:

> (LUM) For every case α, if in α C obtains, then in α S is in a position to know that C obtains.[1]

Following Smithies (2019: 346), it is helpful to enrich the taxonomy by distinguishing positive, negative, and strong luminosity. A condition C is positively luminous just in case:

> (LUM+) For every case α, if in α C obtains, then in α S is in a position to know that C obtains.

A condition C is negatively luminous just in case:

> (LUM–) For every case α, if in α C does not obtain, then in α S is in a position to know that C does not obtain.

A condition C is strongly luminous just in case:

> (LUM±) For every case α, in α S is in a position to know whether C obtains.

What is it to be in a 'position to know'? According to Williamson (2000: 95), if S is in a position to know that C obtains, and S has done what she is a position to do to

[1] Cases are centred possible worlds – triples comprising a world, a subject, and a time (Williamson 2000: 94).

determine whether C obtains, S knows that C obtains. As Rosenkranz (2021: 37) puts it, being in a position to know lies somewhere in between knowing and having the capacity to know. To be in a position to know is not the same as knowing, because, like having the capacity to know, S might be in a position to know something without actually knowing it. But, like knowing and unlike having the capacity to know, being in a position to know is factive: if S is in a position to know that C obtains, then C obtains. In other words, all that is lacking for S to know that P when she is in a position to know that P, is that S forms the target belief by basing it in the right way on the available evidence.[2]

4.1.2 The Argument

There are several acceptable presentations of Williamson's anti-luminosity argument. In fact, Williamson presents the argument in different ways depending on the context. On one reconstruction due to Rosenkranz (2021: 77), the argument proceeds from three independent assumptions. The first assumption states that:

> (CHANGE) For any non-trivial condition C, C can gradually change in an arbitrarily small way across a sequence of cases α_0, α_1, α_2, …, α_n, ranging from a case α_0 in which C clearly obtains to a case α_n in which C clearly does not obtain.

The second assumption states that:

> (INDISCRIMINABILITY) If S believes that C obtains in a case α_i, then S could easily believe that C obtains in the immediately succeeding case α_{i+1}.

According to Williamson (2000: 127), the assumption is motivated by the idea that S's belief that C obtains is not perfectly discriminating. As a consequence, S could easily have the same belief in very similar situations.

The third assumption states that:

> (SAFE) If S is in a position to know that C obtains in a case α_i, then S

[2] See chapters 3 and 4 of Rosenkranz (2021) for an extensive discussion of the notion of being in a position to know and its logic.

> could not easily believe falsely that C obtains in the immediately succeeding case α_{i+1}.

The assumption follows from a safety-from-error requirement on knowledge and is motivated by considerations of reliability. The idea is that S knows that C obtains only if S's belief is sufficiently sensitive to the non-occurrence of C (Williamson 2000: 127).

The three assumptions can be used to prove that C is not luminous. By classical logic, CHANGE guarantees that the sequence α_0, α_1, α_2, …, α_n contains two adjacent cases such that one is the last case in which C obtains while the other is the first case in which C does not obtain. Consider the last case in which C obtains and suppose that S has done what she is in a position to do to determine whether C obtains. As a consequence, if S were in a position to know that C obtains, S would know that C obtains. Thus, if S does not believe that C obtains, then she is not in a position to know that C obtains.

What if S believes that C obtains instead? By SAFE, if S is in a position to know that C obtains in the last case in which C obtains, then S could not easily believe falsely that C obtains in the immediately succeeding case. But, by INDISCRIMINABILITY, if S believes that C obtains in the last case in which C obtains, then S could easily believe that C obtains in the first case in which C does not obtain. In other words, S could easily believe falsely that C obtains. Thus, if S believes that C obtains in the last case in which C obtains, S is not in a position to know that C obtains. Either way, S is not in a position to know that C obtains. Hence, C is not luminous.

On another reconstruction due to Srinivasan (2015: 296-7), the anti-luminosity argument describes a counterexample to the putative luminosity of the condition of feeling cold. Since analogous counterexamples can be constructed for any non-trivial condition, the argument shows that no such condition is luminous.

Consider the following case:

> COLD MORNING
> S wakes up at dawn feeling freezing cold. As the sun rises, S slowly warms up and feels very hot at noon. The change is very gradual, and S is unaware of any change in her experience over one millisecond. Throughout the process, S is doing what she is in a position to do to determine whether

she feels cold. (Williamson 2000: 97)

COLD MORNING is a counterexample to the luminosity of being cold, because if being cold is a luminous condition, then one can prove that S is cold at noon. But that is false by description of COLD MORNING, so being cold is not a luminous condition.

Let $t_0, t_1, t_2, \ldots, t_n$ be a series of times at one millisecond intervals. Let α_i be the case at time t_i. Now assume for *reductio* that C='being cold' is a luminous condition, that is:

(LUM) If in α_i C obtains, then in α_i S is in a position to know that C obtains.

Further, assume the following margin of error principle for S's knowledge that she feels cold:

(MAR) If in α_i S knows that C obtains, in α_{i+1} C obtains.

Like SAFE, the principle is motivated by considerations of reliability and the fact that S's powers of discrimination are limited (Williamson 2000: 12, 97, 103-4). Finally, by description of COLD MORNING, the following statements are true:

(BEG) In α_0 C obtains.

(END) In α_n C does not obtain.

LUM, MAR, BEG, and END are incompatible. Since S is doing what she is in a position to do to determine whether C obtains, by LUM, we can derive from BEG that in α_0 S knows that C obtains. Thus, by MAR, in α_1 C obtains. By repeating the reasoning, we can derive that in α_n C obtains. But that contradicts END, which is true by description of COLD MORNING. Hence, LUM is false: S is not always in a position to know that she feels cold whenever she is. By extension, no non-trivial condition is luminous.

4.2 From Luminosity to Sameness of Evidence

According to Williamson (2000, chapter 8), the sceptic argues from the claim that one's evidence is luminous to the claim that one has the same evidence in the good and bad case, and from the latter claim to scepticism. According to Williamson, the idea is that if S has the same evidence in the good and the bad case, then S's evidence in the good case is insufficient for the truth of P in the sense that S could falsely believe

that P with the very same total evidence. But then S seems to know that P only in a stretched and weakened sense of 'know' to be contrasted with the sense in which S knows the evidence itself (Williamson 2000: 174).

I have my reservations about Williamson's reconstruction of the sceptic's argument because that argument seems to presuppose infallibilism about knowledge. If that is correct, it is susceptible to the knowledge-version of the Infallibility Objection. As I have argued in chapter 3, the Infallibility Objection rests on a misunderstanding of the Underdetermination Argument, and so does Williamson's reconstruction.[3]

Further, fallibilists ought to resist Williamson's contention that S knows that P at best in a stretched and weakened sense of 'know'. It is a platitude, but nonetheless true, that we know some things better than others. Surely, S can know that P in the full sense of 'know' even if she is in a weaker epistemic position with respect to P than with respect to her evidence for P (Brown 2018: 19). In fact, it is to be expected that evidence is known better than what it supports because the evidence is known directly but what it supports is known through the evidence, if at all (Roush 2005: 158; Kelly 2008a: 942).

However, the details about how the sceptic arrives at the sceptical conclusion are unimportant in the present context – Williamson is less concerned with how the sceptic can ground scepticism in the Sameness of Evidence Lemma than with how she can motivate the Sameness of Evidence Lemma itself.[4] In the present context, it also does not matter that Williamson is thinking about knowledge scepticism rather than justification scepticism. If his criticism of the Sameness of Evidence Lemma is correct, the Lemma cannot be used to support either kind of scepticism. In this section, I present Williamson's reconstruction of the sceptic's rationale for the Sameness of Evidence Lemma and explain how the anti-luminosity argument is supposed to undermine that rationale.

[3] Still, it should be noted that Williamson is an infallibilist, so, from his perspective, there is nothing wrong with the sceptic's appeal to infallibilism.

[4] However, these details are important for the larger aims of Williamson's chapter: to show the untenability of certain conceptions of evidence (mentalism and accessibilism) on the ground that they make scepticism unavoidable.

4.2.1 The Sceptic's Argument

How does the claim that one's evidence is luminous figure in the sceptic's argument for the Sameness of Evidence Lemma? In a nutshell, the sceptic's argument goes like this. Assume for *reductio* that S does not have the same evidence in the good and the bad case. Then in the bad case S is in a position to know that she is not in the good case, because she is in a position to compare her evidence in the bad case with her evidence in the good case. But that is absurd: in the bad case S is not in a position to know that she is not in the good case. As Williamson nicely puts it, 'part of the badness of the bad case is that one cannot know just how bad one's case is' (Williamson 2000: 165). Therefore, S has the same evidence in the good and the bad case.

For the argument to work, Williamson observes, the sceptic must assume that in the bad case S is in a position to know what her evidence is. Otherwise, in the bad case S would not be in a position to compare her evidence with her evidence in the good case. For the same reason, the sceptic must assume that in the bad case S is in a position to know what her evidence is in the good case.

To bring out more clearly the role that S's access to the evidence plays in the argument for the Sameness of Evidence Lemma, let us rewrite the argument in a more rigorous way. Let e be a piece of evidence and let C='having a piece of evidence e'. We can rewrite the sceptic's argument as follows.

First, the sceptic assumes that S's evidence is strongly luminous:

> (LUM±) For every case α, in α S is in a position to know whether C obtains.

What is the motivation for (LUM±)? According to Williamson, LUM± is motivated by the thought that 'rationality requires one to respect one's evidence and cannot require one to respect something unless one is in a position to know what it is' (Williamson 2000: 173). Therefore, if rationality can impose genuine requirements on one's beliefs, one's evidence must be luminous.

Second, the sceptic assumes that S's knowledge in the bad case extends to her evidence in the good case:

> (CONCEPT) In the bad case S is in a position to know what her

evidence is in the good case.

At first glance, the assumption may look implausibly strong. How can S be in a position to know what her evidence is in one case from within another? But there need not be anything especially demanding about S's knowledge, because S may know what her evidence is in the good case by unpacking her descriptive concept of the good case. The sceptic is happy to grant that S is sufficiently conceptually sophisticated to possess a concept of the good case that specifies what evidence she has in it. The idea is that it does not matter how much S knows about what the good case is like because S can never recognise her actual case from within the bad case (Williamson 2000: 171). Moreover, if it turned out that S has the same evidence in the good and the bad case when S is *that* conceptually sophisticated, nothing would be gained by assuming that S is not that sophisticated. If anything, S's epistemic position would seem to get worse.

Third, the sceptic assumes the uncontroversial claim that:

(BAD) In the bad case S is not in a position to know that she is not in the good case.

Finally, she assumes for *reductio* that S's evidence in the good and the bad case differs:

(DIFF) In the bad case S lacks *e* and in the good case S has *e*.

LUM±, CONCEPT, BAD, and DIFF are incompatible. Assume for *reductio* that DIFF is true. Since in the good case S has *e*, by CONCEPT, in the bad case S knows that in the good case S has *e*. Further, since in the bad case S lacks *e*, by LUM±, in the bad case S knows that she lacks *e*. But if in the bad case S knows that in the good case she has *e* while knowing that in her actual case she lacks *e*, then in the bad case S is in a position to know that she is not in the good case. Thus, LUM±, CONCEPT, and DIFF entail that BAD is false. But BAD is an uncontroversial truth. Thus, DIFF must be false: either in the bad case S has *e* or in the good case S lacks *e* (Williamson 2000: 172).

As it is, the argument does not establish that S has the same evidence in the good and the bad case. For DIFF is false when S has *e* in the bad case *and* S lacks *e* in the good case, in which case the Sameness of Evidence Lemma is false, too. However, *mutatis mutandis*, the argument would go through even if DIFF read 'in the bad case S

has *e* and in the good case S lacks *e*.' In which case the conclusion would read 'either in the bad case S lacks *e* or in the good case S has *e*.' The conjunction of the two conclusions is equivalent to 'S has *e* in the bad case if and only if S has *e* in the good case.' Thus, LUM±, CONCEPT, and BAD jointly entail that S has the same evidence in the good and the bad case (Williamson 2000: 172).

Before we proceed, one important remark. Notice that the sceptic's argument goes through even if the LUM± is weakened to its consequence '*in the bad case* S knows what her evidence is.' So, one might wonder why Williamson assumes LUM± instead of its weaker consequence on behalf of the sceptic. Williamson does not say, but presumably his motivation is that if the luminosity claim is not the reason why the sceptic thinks that in the bad case S know what her evidence is, it is not clear what it is (Williamson 2000: 172-3). As we will see in section 4.4.2, this assumption is problematic.

4.2.2 Williamson's Reply

We are now in a position to appreciate Williamson's rebuttal of the sceptic's motivation of the Sameness of Evidence Lemma. According to Williamson, the sceptic's argument fails because LUM± is false: S is not always in a position to know what her evidence is. The anti-luminosity argument can be easily adapted to the condition of possessing evidence.

Consider the following situation:

> SUNRISE
>
> Before dawn, it is pitch dark and S's evidence includes the proposition that it is dark. As S watches the sun rise, her visual experience changes very gradually. At noon, there is a bright daylight and S's evidence does not include the proposition that it is dark anymore. S is unaware of any change in her experience over one millisecond. Throughout the process, S is doing what she is in a position to do to determine what her evidence is. (Williamson 2000: 175)

Let $t_0, t_1, t_2, \ldots, t_n$ be a series of times at one millisecond intervals. Let α_i be the case at time t_i. By description of SUNRISE, in α_0 S's evidence includes the proposition that it is dark. Assume for *reductio* that S's evidence is luminous: whenever S's evidence

includes the proposition that it is dark, S is in a position to know that her evidence does. Since S is doing what she is in a position to do to determine what her evidence is, and her evidence is luminous, in α_0 S knows that her evidence includes the proposition that it is dark. Further, since S's ability to detect changes in her visual experience is limited, S's knowledge of her evidence is not perfectly discriminating, it has a margin for error. In particular, if S knows that her evidence includes *e* in α_i, then S's evidence includes *e* in α_{i+1}. Then, since in α_0 S knows that her evidence includes the proposition that it is dark, in α_1 S's evidence includes the proposition that it is dark. By repeating the reasoning, we can derive that at noon S's evidence includes the proposition that it is dark. But that contradicts the description of COLD MORNING. Hence, LUM is false: S is not always in a position to know that she has *e*, whenever she has it. *A fortiori* LUM± is false: S is not always in a position to know what her evidence is.

The argument does not depend on any specific theory of evidence. What matters is that having evidence is a non-trivial condition that can vary gradually in an arbitrarily small way across a spectrum of cases (Williamson 2000: 178). Further, the argument generalizes to sceptical arguments that assume that S is using the same rule to arrive at her beliefs in the good and the bad case rather than the Sameness of Evidence Lemma, because that assumption is motivated by the claim that S is always in a position to know what rule she is using (Williamson 2000: 181-3).

If Williamson is right, the anti-luminosity argument undermines the sceptic's rationale for the Sameness of Evidence Lemma. Importantly, the argument does not prove that the Sameness of Evidence Lemma is false. For all the argument shows, S may have the same evidence in the good and the bad case. Incidentally, Williamson believes that the Sameness of Evidence Lemma is false. This is a consequence of his view that E=K, i.e. that one's evidence and one's knowledge are the same thing, together with his anti-scepticism. Since, according to Williamson, S knows that P in the good case and S does not know that P in the bad case, S does not have the same evidence in the good and the bad case. But the view that E=K does not follow from the anti-luminosity argument (Williamson 2000: 180-1). I return on the anti-sceptical import of E=K in the concluding remarks.

4.3 A Summary and a Preview

To recap: Williamson aims to undermine the Underdetermination Argument by undermining the sceptic's argument for the Sameness of Evidence Lemma. Supposedly, the anti-luminosity argument does the job because the sceptic derives the Sameness of Evidence Lemma from a luminosity claim – the claim that S is always in a position to know what her evidence is – and the anti-luminosity argument proves that the luminosity claim is false. The rest of this chapter evaluates the impact of the anti-luminosity argument on underdetermination scepticism.

I shall argue that Williamson's anti-sceptical strategy fails for two reasons. First, pace Williamson, the anti-luminosity argument does not undermine the sceptic's argument for the Sameness of Evidence Lemma. To support this claim, I offer a new argument that presents Williamson with a dilemma. The key idea for the argument is that failures of the luminosity of S's evidence of the kind highlighted by the anti-luminosity argument are irrelevant for failures of knowledge of S's evidence in the good and the bad case.

Second, Williamson's anti-sceptical strategy does not undermine the Underdetermination Argument because Williamson's anti-sceptical strategy depends on the idea that the sceptic needs a luminosity claim to defend the Sameness of Evidence Lemma. That idea is flawed. I present two arguments for the Sameness of Evidence Lemma that do not use the luminosity claim on behalf of the sceptic. The first one appeals to mentalism about evidence, the view that S's evidence is determined by S's non-factive mental states. The second appeals to comparative judgements about the rationality of S in the good and the bad case and the link between rationality and the evidence available to S.

4.4 What the Anti-Luminosity Argument Does Not Show

Since its appearance, Williamson's anti-luminosity argument has drawn numerous objections. Some authors have argued that the argument employs soritical assumptions (Wong 2008). Others have argued that the argument fails when one assumes that there is a constitutive connection between the obtaining of C and believing that C obtains (Weatherson 2004; Berker 2008; Ramachandran 2009) – something many proponents of the luminosity of the mental are inclined to believe. Some have argued that

Williamson's anti-luminosity argument does not apply to certain non-trivial conditions, like conditions specified with direct phenomenal concepts (Duncan 2018), a certain kind of higher-order ignorance (Rosenkranz 2021), or mental *events* and *processes* – rather than states – like judgement and deliberation (Jenkins 2021). Some have argued that the argument appeals to an implausible safety requirement (Leitgeb 2002; Vogel 2010) while others have rejected safety altogether (Brueckner and Fiocco 2002; Neta and Rohrbaugh 2004; Comesaña 2005; Conee 2005). Yet others have suggested that Williamson's argument is problematically circular (Blackson 2007; Cohen 2010). The list goes on.

I tend to agree with some of these objections. In particular, as Srinivasan (2015: 306) concedes, Williamson's talk of limited powers of discrimination invites charges of question-begging. It is part of Williamson's description of COLD MORNING that 'one is not aware of any change in [her feelings of cold or hot] over one millisecond' (Williamson 2000: 97). Similarly, Williamson claims that 'the main idea behind the argument against luminosity is that our powers of discrimination are limited' (Williamson 2000: 12). However, these claims are too close to a statement of the argument's conclusion to be reasons to accept the anti-luminosity thesis.[5]

Further, at least on its simplest formulation, Williamson's argument does not work when there is a constitutive connection between C and believing that C obtains such that

(CON) S believes that she is in C if and only if S is in C.

Remember that, according to Williamson, no non-trivial condition C is luminous because INDISCRIMINABILITY is true: if S believes that C obtains in a case α_i, then S could easily believe that C obtains in the immediately succeeding case α_{i+1}. In particular, INDISCRIMINABILITY entails that if S believes that C obtains in the last case in which C obtains, S could easily believe that C obtains in the immediately succeeding case. If CON is true, that consequence of INDISCRIMINABILITY is false. For, given CON, in the last case in which C obtains, S believes that C obtains, but in the first case

[5] See Srinivasan (2015: 304-7) for a discussion of this issue and an attempt to evade the charge of begging the question by appealing to a claim about what creatures like us are disposed to believe in extremely similar situations.

in which C does not obtain S does not believe that C obtains. And neither does she in any case sufficiently similar to it. Thus, if CON is true, INDISCRIMINABILITY is false.[6]

These objections to the anti-luminosity argument provide as many potential ways for the sceptic to resist Williamson's argument. If Williamson's anti-luminosity argument is unsound, it does not undermine the sceptic's argument for the Sameness of Evidence Lemma. Here, I won't pursue this way of challenging Williamson's anti-sceptical strategy. For one thing, it is not obvious that the sceptic will emerge on top. Perhaps, Williamson's anti-luminosity argument really does work. For another thing, even if Williamson's anti-luminosity argument does not work, there might be other decisive reasons to accept the anti-luminosity thesis (Schwitzgebel 2008, 2011; Greenough 2012; Goldstein and Waxman 2021). Thus, it is interesting to ask whether the sceptic can respond to the anti-luminosity argument even if that argument is sound, or if the anti-luminosity thesis is true.

In this section, I shall argue that Williamson's anti-luminosity argument does not undermine the sceptic's reasoning. In a nutshell, the reason is that failures of the luminosity of S's evidence *in adjacent cases* are irrelevant for failures of knowledge of S's evidence *in the good and the bad case*. The result lays the groundwork for showing that the sceptic can motivate the Sameness of Evidence Lemma without the luminosity claim.

Why does the anti-luminosity argument not undermine the sceptic's reasoning? Consider the following argument:

1. Either the differences in S's evidence between adjacent cases are like the differences in S's evidence between the good and the bad case or they are not.

2. If the differences in S's evidence between adjacent cases are like the differences in S's evidence between the good and the bad case, the anti-luminosity argument does not undermine the sceptical argument.

3. If the differences in S's evidence between adjacent cases are not like

[6] Srinivasan (2015: 309), drawing on Williamson (2000: 99), defends the anti-luminosity argument from this objection by appealing to a safety requirement in terms of degree of confidence. See Vanrie (2020) for a critical discussion.

the differences in S's evidence between the good and the bad case, the anti-luminosity argument does not undermine the sceptical argument.

4. Either way, Williamson's anti-luminosity argument does not undermine the sceptical argument.

Clearly, the argument is valid. Is it also sound? Obviously, the issue is whether premises 2 and 3 are true. I turn to these in the next subsections.

4.4.1 Adjacent Cases and the Good and the Bad Case

Let me start with premise 2. I shall argue that if the differences in S's evidence between adjacent cases – any pair of cases α_i and α_{i+1} in a sequence of cases α_0, α_1, α_2, …, α_n where the evidence changes gradually in an arbitrarily small way – and the differences in S's evidence between the good and bad cases are exactly similar, the anti-luminosity argument leads to scepticism. But if the anti-luminosity argument leads to scepticism, it cannot be used for anti-sceptical purposes to undermine the sceptical argument. Hence, if the differences in S's evidence between adjacent cases and the differences in S's evidence between the good and the bad case are exactly similar, the anti-luminosity argument does not undermine the sceptical argument.

Assume that the differences in S's evidence between adjacent cases and between the good and the bad case are exactly similar. What follows? Consider SUNRISE (section 4.2.2). Williamson (2000: 175-6) thinks that, given S's limited powers of discrimination, it is consistent with what S knows in α_i that S is in α_{i+1}.

> 'In such cases, (3_i) [that it is consistent with what S knows in α_i that S is in α_{i+1}] is obvious in roughly the way in which it is obvious that it is consistent with what I know by sight when I am in fact looking at a distant tree *i* millimetres high that I am looking at a tree only *i* –1 millimetres high. From premises which I know on the basis of sight to the conclusion that I am not looking at a tree only *i* –1 millimetres high, there is no hope of constructing a valid deduction, not even one which I am somehow not in a position to carry out.'

Notice that Williamson's reasoning cuts both ways. Just as, given S's limited powers of discrimination, it is consistent with what S knows in α_i that S is in α_{i+1}, so it is

consistent with what S knows in α_{i+1} that S is in α_i. In other words, there is a symmetry of epistemic accessibility between adjacent cases.

Now, if the differences in S's evidence between adjacent cases and between the good and the bad case are exactly similar, there is a symmetry of epistemic accessibility between the good case and the bad case, too. I.e. just as it is consistent with what S knows in the bad case that S is in the good case, so it is consistent with what S knows in the good case that S is in the bad case.

However, a symmetry of epistemic accessibility between the good and the bad case entails scepticism. Since it is consistent with everything S knows in the bad case that she is in the good case, and in both cases S knows that the truth of P is incompatible with S's being in the bad case, if in the good case S knew that P, the bad case would not be epistemically accessible from the good case. Thus, if the good and the bad case are epistemically accessible from each other, S does not know that P in the good case.

Therefore, if the differences in evidence between adjacent cases and between the good and the bad case are exactly similar, the anti-luminosity argument leads to scepticism. Hence, it cannot be used for anti-sceptical purposes to undermine the sceptic's reasoning.

One might object to my argument for premise 2 that if the differences in S's evidence between adjacent cases and between the good and the bad case are exactly similar, the anti-luminosity argument undermines the Underdetermination Argument, after all. For some adjacent cases are *marginal* cases of being in C: pair of cases in which the non-trivial condition C changes from obtaining to not obtaining. If the differences in S's evidence between marginal cases and between the good and the bad case are exactly similar, then S does not have the same evidence in the good and the bad case: just as in marginal cases S does not have the same evidence, so S does not have the same evidence in the good and the bad case.

The point is well taken but of no help to the anti-sceptic who would like to use the anti-luminosity argument to undermine the Underdetermination Argument. Since marginal cases are adjacent cases, it follows from my argument that if the differences in S's evidence between marginal cases and between the good and bad case are exactly similar, scepticism follows.

The reason is that the anti-luminosity argument involves a trade-off between S's evidence and S's ability to tell close cases apart. The anti-luminosity argument guarantees that in marginal cases S has/lacks *e* although S is not in a position to know that S has/lacks *e*. But S's having/lacking *e* without being in a position to know that S has/lacks *e* comes at the cost of S's ability to discriminate between her case and adjacent cases. For *e* does not discriminate between adjacent cases and although S's having *e* distinguishes S's case from an adjacent case, that S has *e* is not evidence for S. Thus, whatever evidence S has in the good case which S lacks in the bad case, that evidence does not put S in a position to know that she is not in the bad case. In fact, it looks like that evidence has no epistemic value for S: *e* is epistemically idle. If anything, the point reinforces the sceptic: scepticism does not require that S has the same evidence in the good and bad case, provided that S's evidence in the good case is not discriminating enough to tell the two cases apart.

4.4.2 Small Differences in S's Evidence Do Not Matter

What about premise 3? Williamson (2000: 164) claims that the anti-luminosity argument is parallel to the sceptic's argument for the Sameness of Evidence Lemma. As my argument above shows, if the two arguments were perfectly parallel, scepticism would follow. Why does this matter?

Because if the differences in S's evidence between adjacent cases are of a different type than the differences in S's evidence between the good and the bad case, that luminosity fails for differences in evidence like those in adjacent cases does not imply that luminosity fails for differences in evidence like those between the good and the bad case. In other words, if the differences in evidence between adjacent cases are not relevantly similar to the differences in evidence between the good and the bad case, that the former can make a difference to S's access to her evidence does not show that the latter can.

In fact, I will argue that even if differences like those between adjacent cases can be responsible for luminosity failures, such failures pose no problem for the sceptic's reasoning. The idea is that since luminosity failures à la Williamson do not indicate luminosity failures in the good and the bad case, the anti-luminosity argument does not undermine the sceptic's argument. In other words, even if in some adjacent cases S

cannot know that she has *e*, in other cases S can still know that. After all, that S is *not always* in a position to know what her evidence is does not entail that S is *never* in a position to know what her evidence is. Thus, that S is not always in a position to know what her evidence is in some adjacent cases does not show that S is not in a position to know what her evidence is in the good and the bad case, unless the good and the bad case are like those adjacent cases.

One might object that the reason why the anti-luminosity argument undermines the sceptic's argument is that it shows that in a sequence of cases α_0, α_1, α_2, …, α_n C can obtain in one case while C fails to obtain in a different case in which S mistakenly believes that C obtains. For instance, in SUNRISE, at the start of the process S has *e* but later in the day, on the first occasion in which S lacks *e*, S mistakenly believes that she has *e*. Similarly, in the bad case S might mistakenly believe to have *e* although she has *e* in the good case. As Williamson (2000: 177) says, 'one's evidence in the bad case can appear exactly similar to one's evidence in the good case, not because it is almost exactly similar, but because it is so radically impoverished that one lacks evidence of its impoverishment.'

The objection does not hold scrutiny. The anti-luminosity argument does not show that C can obtain in one case while C fails to obtain in a different case in which S mistakenly believes that C obtains. It only shows that sometimes C obtains but S is not in a position to know that it does. But that is compatible with S's never believing that C obtains when it does not. After all, S might refrain from believing that C obtains when S is close to a case in which C does not obtain. Relatedly, it is not part of the description of SUNRISE that at the start of the process S has *e* but later in the day, on the first occasion in which S lacks *e*, S mistakenly believes that she has *e*.

Is there any reason to think that the differences in S's evidence between the good and the bad case and the differences in S's evidence between adjacent cases are not similar? Consider again SUNRISE. How do differences in S's evidence between adjacent cases in SUNRISE differ from differences in S's evidence between the good and the bad case? An obvious answer is this: the differences in S's evidence between adjacent cases are gradual, whereas the differences in S's evidence between the good and bad cases, if any, are not. Whereas S's visual experience changes gradually across SUNRISE cases, S's visual experience stays the same in the good and the bad case.

This is an important disanalogy between adjacent cases and the good and the bad case, because the fact that S's evidence changes gradually across SUNRISE cases is the reason why luminosity fails. Since S's evidence does not change gradually from the good case to the bad case, failures of luminosity in SUNRISE cases do not impinge on S's ability to access her evidence in the bad case.

In any case, for the purposes of my argument, I do not need to show that the differences in evidence between the good and the bad case and adjacent cases do differ. That they differ in a way that is relevant to the anti-luminosity argument is an assumption motivated by the fact that if they did not, scepticism would follow, as I showed in the previous subsection.

One might agree that failures of luminosity of the kind highlighted by the anti-luminosity argument do not automatically spread to S's knowledge of her evidence in the bad case. In other words, the anti-luminosity argument does not show that in the bad case S does not know what her evidence is. But one might insist that the anti-luminosity argument need not accomplish anything like that to undermine the sceptic's argument for the Sameness of Evidence Lemma.

Rather, so the objection goes, the anti-luminosity argument undermines the sceptic's argument simply because the sceptic's argument relies on a luminosity claim and the anti-luminosity argument shows that the luminosity claim is false. As I observed at the end of section 4.2.1, the sceptic's argument for the Sameness of Evidence Lemma works even if the luminosity claim is weakened to its consequence 'in the bad case S knows what her evidence is'. However, presumably, the reason to accept the latter claim is that S's evidence is luminous. Once the luminosity claim is off the table, the sceptic cannot simply assume that S knows what her evidence is in the bad case. Thus, my argument does not show that the anti-luminosity argument does not undermine the sceptic's reasoning.

I agree that what I have said so far only shows that if the differences in S's evidence between adjacent cases are not like the differences in S's evidence between the good and the bad case, Williamson's anti-luminosity argument is compatible with S's knowing what her evidence is in the bad case. That is not the same as showing that if the differences in S's evidence between adjacent cases are not like the differences in S's evidence between the good and the bad case, Williamson's anti-luminosity argument

does not undermine the sceptical argument.

In the rest of this section, I show how my argument can be completed to bear the conclusion that if the differences in S's evidence between adjacent cases are not like the differences in S's evidence between the good and the bad case, Williamson's anti-luminosity argument does not undermine the sceptical argument.

Recall that, according to Williamson (2000: 173), the rationale for the luminosity claim is that 'rationality requires one to respect one's evidence and cannot require one to respect something unless one is in a position to know what it is.' Notice that there is a tension between Williamson's account of the sceptic's rationale for the luminosity claim in terms of rational requirements and Williamson's attack to the luminosity claim in terms of undetectably small differences in S's evidence.

According to Williamson, luminosity fails because in some cases S has some piece of evidence *e* but in adjacent, i.e. very similar, cases S lacks *e*. As I noted above, in such cases there is a trade-off between S's evidence and S's ability to tell close cases apart. In fact, in such cases, it looks like *e* has no epistemic value for S – although S has *e*, it is as if S did not have *e*. Now, ask yourself whether rationality requires S to respect *e* in such cases.

The answer hinges on what it means for S to respect *e*. In keeping with the example hitherto used, namely SUNRISE, suppose that *e* is the proposition that it is dark. Further, suppose that if rationality required S to respect *e*, minimally, S should believe that it is dark. In this sense of respecting the evidence, it seems that rationality does not require S to respect *e*, because, since S is in a marginal case, *e* has no epistemic value for S. In other words, *e* is not the kind of evidence rationality can require S to respect; S should not believe that it is dark.

But now suppose that respecting the evidence simply consists in adopting the fitting attitude to the evidence. Perhaps, in this sense, rationality can require S to respect *e*. But, since *e* has no epistemic value for S, rationality must require S's beliefs to be insensitive to *e*. In other words, S should form her beliefs as if *e* were not part of her evidence. Thus, S should not believe that it is dark.

Either way, this is a problem for Williamson. For, in both cases, *e* seems to be evidence only in a stretched sense of the term 'evidence'. After all, in both cases, *e* has no epistemic value for S and thus has no rational impact on her beliefs. But, surely, when

the sceptic claims that rationality requires S to respect her evidence, she does not mean items like *e* that have no epistemic value for S.

Further, recall that *e* has no epistemic value for S precisely because the differences between S's evidence in adjacent cases are undetectably small. By assumption and lest one accepts scepticism, the differences in S's evidence between the good and the bad case are not like the differences in S's evidence between adjacent cases that determine failures of luminosity. Thus, the sceptic can continue to appeal to considerations about rationality requirements to motivate the claim that S knows what her evidence is in the bad case. If this is correct, the anti-luminosity argument does not undermine the sceptical argument.

In conclusion, Williamson faces a dilemma. On the one hand, if the differences in S's evidence between adjacent cases are like the differences in S's evidence between the good and the bad case, scepticism follows. But if scepticism follows from the analogy of the good and the bad case with adjacent cases, the anti-luminosity argument cannot be used to undermine the sceptical argument for anti-sceptical purposes. On the other hand, if the differences in S's evidence between the good and the bad case and adjacent cases are not similar, failures of luminosity of the kind highlighted by the anti-luminosity argument do not spread to the good and the bad case. In particular, the sceptic can still appeal to considerations about rationality requirements to argue that in the bad case S knows what her evidence is.

4.5 Sameness of Evidence Without Luminosity

Williamson's anti-sceptical strategy is predicated on the idea that the sceptic's motivation for the Sameness of Evidence Lemma reduces to the argument reconstructed in section 4.2.1. That argument uses the claim that S is always in a position to know what her evidence is.

If the sceptic can motivate the Sameness of Evidence Lemma without any appeal to the claim that S's evidence is luminous, then Williamson's anti-sceptical strategy loses efficacy because the anti-luminosity argument cannot be used to undermine the sceptic's argument. Does the sceptic need the luminosity claim to establish that S has the same evidence in the good and the bad case?

Williamson (2000: 172-3) thinks so:

> 'If something like this argument [read: the argument for the Sameness of Evidence Lemma reconstructed in section 4.2.1] is not the reason for which sceptics and others think that one has the same evidence in the two cases, it is not at all clear what is.'

Against Williamson, I shall argue that the sceptic does not need the claim that S's evidence is luminous to motivate the Sameness of Evidence Lemma. In fact, there are at least two ways to motivate the Sameness of Evidence Lemma available to the sceptic that do not use the luminosity claim. Further, one might think that these two motivations better explain the appeal of the Lemma than the argument that Williamson attributes to the sceptic. The first motivation appeals to mentalism about evidence, the view that S's evidence is determined by S's non-factive mental states, while the second one appeals to comparative judgements about the rationality of S in the good and the bad case and the link between rationality and the evidence available to S.

4.5.1 Sameness of Evidence and Mentalism

Here is a simple argument for the claim that S has the same evidence in the good and the bad case that proves my point. Assume that S's evidence consists of non-factive mental states. A consequence of that claim is that 'mental duplicates' have the same evidence, i.e. that if any two individuals are exactly alike mentally, then they have the same evidence. Further, assume that good-case S and bad-case S are mental duplicates, i.e. S is in a non-factive mental state *m* in the good case if and only if *S* is in *m* in the bad case. The two assumptions entail that S has the same evidence in the good and the bad case.

Mental Duplicates and Semantic Externalism

It is hard to contest the claim that S is in the same total non-factive mental state in the good and the bad case. The claim seems a platitude that follows from the description of the good and the bad case. Famously, Putnam (1981: 7-14) has argued that a brain-in-a-vat lacks the necessary causal connections to grasp the propositions that a normal subject can grasp. However, that idea lacks generality. A recently envatted brain retains enough causal connections with her past environment to entertain all the propositions

that can be grasped in the good case (Smith 1984: 117).

Recently, Thorpe (2018) has challenged this way to respond to anti-sceptical considerations from semantic externalism. According to Thorpe, if the sceptic makes use of what he calls a 'non-radical sceptical scenario', like the recent envatment scenario, then there is a large number of empirical beliefs that is left unscathed by the sceptical argument. According to Thorpe, those empirical beliefs are part of S's evidence and indicate that being in the good case is more likely than being in the bad case. Thus, if S is in the same total non-factive mental state in the good case and in the corresponding non-radical sceptical scenario, the sceptic's argument won't work. The upshot is a dilemma for the sceptic:

> 'either the sceptic makes use of what I will call a non-radical sceptical scenario, and so enables me to appeal to my empirical beliefs to show that I am not in that scenario; or she makes use of what I will call a radical sceptical scenario, in which case I can use the vat argument [read: Putnam's considerations from semantic externalism] to show that I am not in that scenario.' (Thorpe 2018: 668)

Either way, S's evidence favours her being in the good case over her being in the bad case.

Thorpe's argument does not work. It rests on the false assumption that S can satisfy the demands imposed by semantic externalism on concept possession only in a non-radical sceptical scenario that preserves some empirical beliefs that indicate that being in the good case is more probable than being the bad case.

Whether a sceptical scenario is 'radical' in Thorpe's sense depends on the size of the gap between the truth of S's beliefs and the world. In a radical sceptical scenario, the gap is very large, because a very large number of S's beliefs are false. However, whether S satisfies the externalist conditions for concept possession is an entirely different matter: it depends on whether S stands in the right sort of relationship to the environment. In recent envatment scenarios S possesses the relevant concepts because in those scenarios S stood in the right sort of relationship to her environment prior to envatment.

Notice that whether S stood in the right sort of relationship to her environment prior to envatment does not determine what *beliefs* S has after envatment. It only determines what *concepts* she has. The upshot is that the sceptic can easily specify a sceptical scenario in which S satisfies all the externalist conditions for concept possession *and* nearly all of S's empirical beliefs are false. Thorpe's dilemma is a false dilemma.

Mentalism and Luminosity

What about the premise that S's evidence consists of non-factive mental states? On one way to carve out the distinction between internalism and externalism in epistemology the distinction is about the accessibility to or awareness of the basis for justification. According to internalism, having a justifier requires some kind of awareness of that justifier. Externalism denies this requirement. On another way to carve out the distinction it is about the nature of justifiers. According to internalism, a justifier is always a non-factive mental state. Externalism denies this (Madison 2010; Pappas 2017).

To avoid confusion, we can recast these distinctions in terms of accessibilism/non-accessibilism and mentalism/non-mentalism respectively. What is the relationship between accessibilism/non-accessibilism and mentalism/non-mentalism? At least in principle, mentalism can be paired with either accessibilism or non-accessibilism. So does non-mentalism. In fact, each of these four combinations has found some defenders in the literature. So, Chisholm (1977), Bonjour (1985), Huemer (2001), and Smithies (2019) subscribe to mentalism and accessibilism, while Pollock and Cruz (1999), Conee and Feldman (2001), and Wedgwood (2002) subscribe to the former but not to the latter. McDowell (1982), Alston (1988), and Pritchard (2012, 2015) reject mentalism and accept accessibilism, while Williamson (2000), Littlejohn (2012), and Srinivasan (2020) reject both.

The argument for the Sameness of Evidence Lemma I have just presented uses the claim that S's evidence consists of non-factive mental states. On an evidentialist picture of justification, that claim is equivalent to the mentalist claim that all justifiers are non-factive mental states. This matters for my argument because, as I have just explained, mentalism need not be paired with accessibilism. *A fortiori*, it need not be paired with the luminosity claim that S is always in a position to know what her evidence is, which

is a very strong form of accessibilism. Thus, the claim that S's evidence consists of non-factive mental states is compatible with the conclusion of Williamson's anti-luminosity argument.

Although evidence mentalism is compatible with the anti-luminosity of evidence, one might worry that the anti-luminosity of evidence undermines the case for evidence mentalism. This would be the case if the luminosity of evidence were the only reason to believe that evidence consists of non-factive mental states, as Williamson (2000: 173, 183) sometimes seems to suggest. The idea is that evidence luminosity drives the identification of one's evidence with non-factive mental states because being in a non-factive mental state is the best candidate for a luminous condition. But if no trivial condition is luminous, one is free to identify evidence in a different way.

Although the luminosity thesis might be a reason to endorse some form of evidence mentalism, the appeal of evidence mentalism is independent of luminosity considerations. In chapter 1, I explained that a mentalist conception of evidence is plausible given the idea that evidence is a cognitive input, that to which our beliefs are responsive. I made the case with respect to perceptual beliefs by comparing cases of hallucination and perception because I was focusing on perceptual evidence. However, the idea is easily extended to other kinds of belief and evidence (Conee and Feldman 2001; Bergmann 2021).

In a similar vein, Schoenfield (2015) argues that mentalism is motivated by considerations about the causal role played by mental states in our cognitive lives, because mental states act as causal mediators between mind and world. To paraphrase Schoenfield (2015: 257), mentalist evidence is important because this is the evidence that we should expect to respond to as a result of responding to whatever we take evidence to be (see chapter 1, section 1.2.2).

4.5.2 Sameness of Evidence, Rationality, and the New Evil Demon

The sceptic can argue for the Sameness of Evidence Lemma by resorting to comparative judgements about S's rationality in the good and the bad case. The idea is that S is equally rational in the good and the bad case. On the evidentialist assumption that epistemic rationality is a matter of what evidence S has, that idea implies that S has the same evidence in the good and the bad case.

Unlike Williamson's reconstruction of the sceptic's argument for the Sameness of Evidence Lemma, this argument does without the claim that S's evidence is luminous. It only requires the claim that S is equally rational in the good and the bad case in conjunction with the evidentialist assumption that epistemic rationality is determined by one's evidence, neither of which entails or is motivated by the luminosity claim.

I have talked about the evidentialist assumption elsewhere (chapter 1, section 1.2.1). Here I shall focus on the claim that S is equally rational in the good and the bad case. That claim is at the heart of the so-called 'New Evil Demon Problem'.

Originally, the New Evil Demon Problem was put forward as a problem for reliabilist theories of epistemic justification (Cohen and Lehrer 1983: 192-3; Cohen 1984: 281; Dutant forth.). Consider the good and the bad case. While in the good case S's beliefs are produced by a reliable process, in the bad case they are not. According to reliabilist theories of justification, S's beliefs are justified just in case they are the output of a (sufficiently) reliable process. Thus, according to reliabilist theories of justification, S is justified in believing that P in the good case, but S is not justified in believing that P in the bad case. The New Evil Demon Problem is this: the reliabilist verdict about S's justification in the good and the bad case is highly counterintuitive. Although the intuition is not universal, it seems to many that in the good case and the bad case S's beliefs are equally justified.

Can the sceptic appeal to the intuition that in the good and the bad case S's beliefs are equally justified? It might seem not, for, originally, the New Evil Demon Problem contained an anti-sceptical assumption. The idea was that in the good case and the bad case S's beliefs are equally justified because S's beliefs in the bad case are justified although they are not reliably formed. So, how can the sceptic appeal to the claim that S is equally rational in the good and the bad case?

Although in its original formulation the New Evil Demon Problem made an anti-sceptical assumption, it is not clear that that assumption is needed to generate it. In fact, it is not clear that an anti-sceptical intuition is driving the Problem at all. Some formulations of the New Evil Demon Problem sound compelling while dispensing with the claim that S's beliefs in the bad case are justified. Rather, the intuition behind

the Problem seems to be simply that S is equally rational in the good and the bad case.[7] But being equally rational in the good and the bad case is compatible with S's beliefs being *equally unjustified*.

The intuition driving the New Evil Demon Problem sits well with mentalism. Conversely, it conflicts with non-mentalism. For what the New Evil Demon intuition suggests is that one's total non-factive mental state is all that matters to justification (Wedgwood 2002: 349). As such, the argument for sameness of evidence from comparative judgements about S's rationality in the good and the bad case goes well with the argument for sameness of evidence from mentalism. The two arguments reinforce each other.

4.6 Concluding Remarks

Williamson (2000) aims to undermine the Underdetermination Argument using his well-known anti-luminosity considerations. According to Williamson, the sceptic argues from the claim that S's evidence is luminous – that S is always in a position to know what her evidence is – to the claim that S has the same evidence in the good and the bad case, and from the latter claim to scepticism.

Williamson's anti-sceptical strategy is to produce a counterexample to the luminosity claim, a possible case in which S is not in a position to know what her evidence is. The idea is that there are marginal cases of S's having some piece of evidence *e*. When S is in a marginal case – call it α –, there is a very similar case that is *not* a case of S's having *e* – call it β. If in α S believes that she has *e*, she fails to know that she has *e*, because S's true belief is made unreliable by an untrue belief that S has *e* in a very similar case, namely β. Thus, when S is in a marginal case of S's having *e*, S cannot

[7] Consider for instance Wedgwood's (2002: 349) statement of the relevant intuition:

> 'Consider two possible worlds, w_1 and w_2. In both worlds, you have exactly the same experiences, apparent memories, and intuitions, and in both worlds you go through exactly the same processes of reasoning, and form exactly the same beliefs. In this case, it seems, exactly the same beliefs are rational in both worlds, and exactly the same beliefs are irrational in both worlds. Now suppose that in w_1 you are bedeviled by an evil demon who ensures that many of your experiences are misleading, with the result that many of the beliefs that you hold in w_1 are false. In w_2, on the other hand, almost all your experiences are veridical, with the result that almost all the beliefs that you hold in w_2 are true. Intuitively, this makes no difference at all. Exactly the same beliefs are rational and irrational in both worlds.'

know that S has *e*: luminosity fails in marginal cases.

In this chapter, I have taken issue with Williamson's anti-sceptical strategy on two grounds. First, I have argued on the basis of new argument that the anti-luminosity argument does not undermine the sceptic's case for the Sameness of Evidence Lemma. More precisely, my argument presents Williamson with a dilemma.

On the one hand, if the differences in S's evidence between adjacent cases are like the differences in S's evidence between the good and the bad case, scepticism follows. After all, there is a symmetry of epistemic accessibility between adjacent cases, and if there is a symmetry of epistemic accessibility between the good and the bad case scepticism follows. But if scepticism follows, the anti-luminosity argument cannot be used for anti-sceptical purposes. Thus, if the differences in S's evidence between adjacent cases are like the differences in S's evidence between the good and the bad case, the anti-luminosity argument does not undermine the case for scepticism.

On the other hand, if the differences in S's evidence between the good and the bad case are not like the differences in S's evidence between adjacent cases, failures of luminosity of the kind highlighted by the anti-luminosity argument do not matter to S's knowledge of her evidence in the bad case. After all, according to Williamson, luminosity fails because in some cases S has some piece of evidence *e* but in adjacent, i.e. very similar, cases S lacks *e*. In such cases there is a trade-off between S's evidence and S's ability to tell close cases apart. In fact, in such cases, it looks like *e* has no epistemic value for S – although S has *e*, it is as if S did not have *e*, and *e* seems to be evidence only in a stretched sense of the term 'evidence'.

However, when the sceptic claims that rationality requires S to respect her evidence, she does not mean items like *e* that have no epistemic value for S. Thus, the sceptic can continue to appeal to considerations about rationality requirements to motivate the claim that S knows what her evidence is in the bad case. The upshot is that if the differences in S's evidence between the good and the bad case are not like the differences in S's evidence between adjacent cases, the anti-luminosity argument does not undermine the sceptic's argument for the Sameness of Evidence Lemma.

Second, I have argued that Williamson's anti-sceptical strategy fails because it is predicated on the idea that the sceptic cannot motivate the Sameness of Evidence Lemma without resorting to the luminosity claim. Pace Williamson, the sceptic can

motivate the Sameness of Evidence Lemma without the luminosity claim.

First, the sceptic can appeal to a mentalist conception of evidence according to which S's evidence consists of non-factive mental states. A consequence of that view is that 'mental duplicates' have the same evidence. Since good-case S and bad-case S are mental duplicates – S is in a non-factive mental state *m* in the good case if and only if *S* is in *m* in the bad case – mentalism about evidence entails that S has the same evidence in the good and the bad case.

Second, the sceptic can argue for the Sameness of Evidence Lemma by resorting to comparative judgements about S's rationality in the good and the bad case. The intuitive idea at the basis of the argument is that S is equally rational in the good and the bad case – an idea that informs the New Evil Demon Problem. On the evidentialist assumption that epistemic rationality is a matter of what evidence S has, that idea implies that S has the same evidence in the good and the bad case.

A few remarks. If I am right, Williamson's anti-luminosity argument does not undermine the Sameness of Evidence Lemma. In fact, if I am right, the sceptic has at least three options to motivate the Sameness of Evidence Lemma: the argument indicated by Williamson, which I have defended against the anti-luminosity argument, and the arguments from evidence mentalism and from the New Evil Demon intuition. Importantly, the three arguments work well together, and the sceptic need not choose one at the expense of the others. This makes the sceptic's case for the Sameness of Evidence Lemma stronger as well as more resilient to challenges than Williamson envisioned. In particular, it makes the sceptic's case for the Sameness of Evidence Lemma independent of the luminosity claim.

Another important upshot of my discussion is that it helps clarify the import of the anti-luminosity argument for the debate about mentalist and non-mentalist conceptions of evidence. One way to interpret Williamson's treatment of scepticism is to interpret it as an argument against mentalist conceptions of evidence. Williamson himself encourages this reading when he says that the anti-luminosity argument lays the groundwork for the equation of knowledge and evidence (Williamson 2000: 180-1). Roughly, the idea is that mentalist conceptions of evidence are motivated by the luminosity of evidence. The luminosity of evidence supports the Sameness of Evidence, which in turn drives evidence towards the mental (Williamson 2000: 173).

If I am right, the anti-luminosity argument provides little to no support to non-mentalist conceptions of evidence. This is for two reasons. First, the anti-luminosity argument does not undermine the sceptic's argument for the Sameness of Evidence Lemma. Thus, the anti-luminosity argument provides no basis for rejecting mentalism about evidence. Second, mentalist conceptions of evidence need not be motivated by the claim that S's evidence is luminous or that S has the same evidence in the good and the bad case. This is especially clear in the case of the argument for the Sameness of Evidence Lemma from evidence mentalism, where evidence mentalism is used to support the Sameness of Evidence Lemma without any appeal to the luminosity of evidence.

How do these results interact with Williamson's thesis that knowledge and evidence are the same thing (E=K)? The view that E=K conflicts with the Sameness of Evidence Lemma on the anti-sceptical assumption that S knows that P in the good case. If S knows that P in the good case, then P is part of S's evidence in the good case, given the equation of knowledge and evidence. By the same token, since in the bad case S does not know that P, P is not part of S's evidence in the bad case. As a consequence, if E=K and S knows that P in the good case, S does not have the same evidence in the good and the bad case. Thus, E=K might seem to provide a way to resist the Sameness of Evidence Lemma (Williamson 2000: 180).

In a sense, that does not matter, because my point still stands. First, the anti-luminosity argument does not undermine the sceptic's argument for the Sameness of Evidence Lemma. Second, Williamson has failed to identify two important ways in which the sceptic can motivate the Sameness of Evidence Lemma, namely via evidence mentalism and via comparative judgements about S's rationality, two ways that are not blocked by the anti-luminosity argument.

In any case, the prospects of using E=K for anti-sceptical purposes are poor. For one thing, this anti-sceptical move, like any move that tries to avoid scepticism by changing the sceptic's conception of evidence, makes itself vulnerable to charges of missing the point (I argue for this claim in chapter 6, section 1). For another thing, according to Williamson, E=K establishes an evidential asymmetry between the good and the bad case because P itself, namely the proposition targeted by the sceptic, is part of S's evidence in the good case. Thus, S's evidence favours P over the competing

sceptical hypothesis SH because S's evidence includes P. P is evidence for itself (Williamson 2000: 187).

Perhaps, in some cases P is evidence for itself – like in cases of perception or introspection. Still, in many other cases, it seems that P is not evidence for itself: the evidential basis for P does not include P. In these cases, evidence is used to justify propositions for which we have no direct way of ascertaining their truth. In other words, whatever epistemic access we have to the truth of these propositions is mediated by the evidence.

As some authors (Dodd 2007: 644-49; Brown 2018: 51-2; Alspector-Kelly 2019: 42-4) have pointed out, these cases constitute a problem for Williamson. The problem generalizes to any theory that, like E=K, attempts to avoid scepticism by making the target proposition P part of S's evidence for P (e.g. disjunctivism). Consider an underdetermination argument for scepticism, where P is a proposition whose epistemic access is mediated by evidence that does not include P itself. According to Williamson, if E=K is true and S knows that P in the good case, then P is part of S's evidence in the good case. By the same token, since in the bad case S does not know that P, P is not part of S's evidence in the bad case. Thus, S does not have the same evidence in the good and the bad case.

This is clearly the wrong result. By stipulation, P is a proposition whose epistemic access is mediated by evidence that does not include P. Thus, P is not part of S's evidence in the good case. *A fortiori*, it is false that S's evidence in the good case favours P over the competing sceptical hypothesis SH because S's evidence includes P.

Granted, Williamson could insist that in the good case S really does know that P and, therefore, that P is part of S's evidence given E=K. If the sceptic assumes otherwise, she is begging the question. But that reply misses the point, for the sceptic's point is that the evidential basis for P must be the same in the good and the bad case. To the extent that E=K conflicts with S's having the same evidence in the good and the bad case, that is a strike against E=K, not against the Sameness of Evidence Lemma.

Thus, if P is part of S's evidence, as Williamson thinks on the basis of E=K, then it is part of S's evidence in one sense, but not in the sense that concerns the sceptic. For the sense of 'evidence' that concerns the sceptic is the sense of 'evidence' in which P is not evidence for itself in cases like the ones I have described. And that is a perfectly

acceptable sense of 'evidence' the sceptic can leverage to draw a sceptical conclusion.

Finally, Williamson's appeal to E=K to reject the Sameness of Evidence Lemma will work only on the assumption that S knows that P in the good case. That is, it will work only on the assumption that scepticism is false. Williamson, like many other anti-sceptics, seems to think that this kind of anti-sceptical move is not problematic. The sceptic argues from the Sameness of Evidence Lemma to scepticism; Williamson argues from anti-scepticism to the negation of the Sameness of Evidence Lemma. The upshot is a dialectical stalemate, but that is no problem for the anti-sceptic, because the burden of proof is on the sceptic to prove that scepticism is true. This kind of move can be traced back to a common philosophical attitude, Mooreanism. It is the topic of the next chapter.

5

Underdetermination Scepticism and Common Sense

We live in an age of 'post-Moorean modesty' (Fine 2001: 2). Mooreans believe in the epistemic superiority of some common-sense beliefs over philosophical arguments. In particular, they believe that one should not abandon one's common-sense beliefs in response to any philosophical argument to the contrary. Crucially, they think this is true even if one cannot identify any flaws in such arguments. Mooreanism is well-known for its allegiance with anti-scepticism. The Moorean anti-sceptic believes that one should not abandon one's common-sense beliefs in response to sceptical arguments.

Mooreans run the risk of being dogmatic. After all, if one can find no flaws in an argument, how can one's refusal to accept that argument be anything other than dogmatic? Insofar dogmatism is incompatible with rationality, Moorean anti-sceptics face the charge of irrationality. Thus, the Moorean cannot simply dismiss a counterargument to her beliefs, she must explain why her dismissal is justified.

My main goal in this chapter is to evaluate Mooreanism as a way to respond to the Underdetermination Argument for scepticism in light of the dogmatism charge. The chapter builds on work by Thomas Kelly and Susanna Rinard on Mooreanism and philosophical methodology. I tackle Mooreanism from three angles, which correspond to three ways of motivating Mooreanism. First, I discuss attempts to motivate Mooreanism by appeal to an alleged difference between philosophy and science. Second, I discuss attempts to motivate Mooreanism via general principles about belief revision. Finally, I discuss Kelly's attempt to vindicate Mooreanism via an argument for its methodological superiority over the methodology employed by the sceptic. I shall argue that none of these motivations for Mooreanism succeeds; the upshot is that Mooreanism provides no sound basis to resist the Underdetermination Argument. Although my discussion centres around the anti-sceptical import of Mooreanism, I hope that my discussion will prove useful to those who wish to endorse or oppose

Mooreanism in other contexts too.[1]

Here is the plan for the chapter. Section 1 introduces Mooreanism, its supposed anti-sceptical import, and argues that Mooreans face a charge of dogmatism. Section 2 discusses the attempt to justify Mooreanism via general considerations about the epistemic differences between philosophy and science. It defends a positive argument due to Rinard (2013) for the claim that philosophy can overturn common sense from considerations about the role of experience in philosophy and science. Section 3 builds on work by Kelly (2005) and looks at Moorean appeals to plausibility, certainty, and knowledge to justify Mooreanism, and argues that they all fail. Section 4 introduces the idea of motivating Mooreanism using methodological considerations. I argue contra Kelly (2005) and in agreement with Rinard (2013) that methodological considerations do not favour Mooreanism over scepticism. The conclusion ties things up and gives an overview of the main results of the chapter.

5.1 Mooreanism, Scepticism, and Dogmatism

In this section, I present the Moorean view of common sense and philosophy, and how it offers a way to resist scepticism. I explain what this view does and does not involve. Then, I present Mooreans with a challenge. The challenge is to explain why the Moorean appeal to common sense does not qualify as dogmatic.

5.1.1 Mooreanism: What It Is and What It Is Not

Consider the following excerpt:

> 'What I want, however, finally to emphasize is this: Russell's view that I do not know for certain that this is a pencil or that you are conscious rests, if I am right, on no less than four distinct assumptions […] And what I can't help asking myself is this: Is it, in fact, as certain that all these four assumptions are true, as that I do know that this is a pencil and that you are conscious? I cannot help answering: It seems to me *more* certain that I

[1] An obvious example are metaphysical debates about the existence of composite objects like chairs, tables, trees, cars, people etc. Against common sense, some theorists (Van Inwagen 1990; Merricks 2000, 2001) hold that there are no inanimate composite objects. Others go a step further and deny that there are composite objects at all (Unger 1979; Cameron 2010; Sider 2013). For some appeals to common sense against these claims see Markosian (1998: 221) and Rea (1998: 348).

> *do* know that this is a pencil and that you are conscious, than that any single one of these four assumptions is true, let alone all four. That is to say, though, as I have said, I agree with Russell that (1), (2) and (3) *are* true; yet of no one even of these three do I feel *as* certain as that I do know for certain that this is a pencil. Nay more: I do not think it is *rational* to be as certain of anyone of these four propositions, as of the proposition that I do know that this is a pencil.' (Moore 1959: 226; emphasis in the original)

Similar statements can be found in Reid (2012 [1818]: 23), Chisholm (1982: 69-70), Armstrong (1983: 49), Lewis (1996: 549), Pollock and Cruz (1999: 7), Pryor (2000: 518), Fine (2001: 2), Lycan (2001: 42-3), Kelly (2005: 181) and (2008b: 56), Gupta (2006: 178) and Bergmann (2021: 115-6). What do these philosophers have in common?

They are anti-sceptics, and they are Mooreans: in a nutshell, they share the idea that some of our common-sense beliefs are resistant to rational revision by means of philosophical argument, and they ground their anti-scepticism in this idea. More precisely, they hold that some of our common-sense beliefs – beliefs that are deeply and widely held (Lemos 2004: 4) – are not rationally undermined by any philosophical argument to the contrary. In particular, they hold that some of our common-sense beliefs – beliefs like the belief that we know a lot or the belief that we are justified in believing many things – provide us with a rational basis to resist any sceptical counterargument. Because of their allegiance to common sense, their position is also known as Commonsensism (Chisholm 1982: 113).

Although there are differences, sometimes important ones, among the views of these philosophers, according to Lemos (2004: 5) the Moorean tradition has four main tenets:

1. Generally, we know what we think we know;
2. Some common-sense beliefs are common knowledge (in the sense that almost everyone knows them);
3. Some common-sense beliefs can be used as fixed points to evaluate philosophical theories;
4. It is more reasonable to accept some common-sense beliefs than any

philosophical argument to the contrary.

A few remarks about what Mooreanism so understood involves and what it does not will prove useful. First, despite their sympathy for *some* of our common-sense beliefs, Mooreans are not committed to the claim that *all* our common-sense beliefs are common knowledge, or that *all* can be used as fixed points to assess philosophical theories, or that *all* are immune from rational revision in response to any philosophical counterargument. Neither are Mooreans committed to the view that common-sense beliefs are knowledge or epistemically justified in virtue of their being common sense, i.e. in virtue of their being deeply and widely believed (Lemos 2004: 6).

Second, to say that some common-sense beliefs are resistant to rational revision by means of any philosophical argument is not to say that they are immune from rational revision altogether. Mooreanism allows for the rational revision of common-sense beliefs because it allows for common-sense beliefs to be corrected by empirical and scientific evidence (Lycan 2001: 40). Thus, although Mooreanism affirms the superiority of some common-sense beliefs over any philosophical argument, it does not promote the absolute invulnerability of any common-sense belief (Kelly 2005: 180).

Third, to endorse Mooreanism, i.e. to say that it is more reasonable to accept some common-sense beliefs than any philosophical argument to the contrary, is not to say that no philosophical argument can persuade us to abandon such beliefs. Suppose that, for some of our common-sense beliefs, no philosophical argument could persuade us to abandon them. That would not be a vindication of Mooreanism: after all, our resistance to abandon those beliefs could be a consequence of our stubbornness and dogmatism. Similarly, suppose that philosophical arguments could persuade us to abandon all our common-sense beliefs. That would not be a refutation of Mooreanism: after all, our persuasion might be a consequence of our gullibility.

The point is that persuasion is a matter of psychological force, and Mooreans are not making a point about the psychological force of philosophical arguments. Rather, they are making a point about whether philosophical arguments *ought to* persuade us, about whether it would be *reasonable* to follow them to their conclusion when they conflict with some of our common-sense beliefs. In other words, Mooreans are making a normative point (Kelly 2005: 181; 2008b: 54).

Fourth, Mooreans are not committed to the claim that philosophical arguments that run against epistemically superior common-sense beliefs are philosophically uninteresting (Kelly 2008b: 60). Thus, Mooreans who are anti-sceptics are not committed to the claim that sceptical arguments are philosophically uninteresting. What Mooreans are committed to, however, is that, when it comes to the acceptance of sceptical arguments, common sense provides us with a sufficient and decisive reason to reject them. As a consequence, the Moorean might engage in the activity of diagnosing what goes wrong in the sceptic's argument, but, according to the Moorean, individuating the mistake in the sceptic's argument is not necessary to reject it.

Finally, to endorse Mooreanism as I understand it here is not to take a stance on Moore's well-known 'proof of an external world' – gesturing that here is one hand, and here is another to conclude that there are external things (Moore 1959: 145-6). The status of Moore's Proof is a matter of debate. Famously, Wright (2002, 2007) holds that Moore's argument is flawed because it exhibits a certain form of vicious circularity. More precisely, Moore's Proof fails because Moore's warrant for the premises does not transmit to the conclusion – in fact, Moore's warrant for the premises presupposes an antecedent warrant for the conclusion. Others have argued that Moore's argument is a genuine proof but that it is dialectically ineffective (Pryor 2004).[2]

For what is worth, I am sympathetic to Wright's views on this issue. However, my point here is that questions about the status of Moore's Proof and questions about the alleged superiority of some common-sense beliefs over philosophical arguments and, more specifically, over the Underdetermination Argument are distinct. In particular, as I understand it here, a Moorean can but need not accept Moore's Proof of an external world (Kelly 2008b: 61). Thus, in the following, I will set aside Moore's Proof.[3]

[2] See Carter (2012) for an overview of the debate on the status of Moore's Proof.

[3] This is not to say that there are no points of contact between the issue in which I am interested and Moore's Proof. However, these points of contact are not my focus. In particular, I want to resist the suggestion that Moore's Proof is directly relevant to underdetermination scepticism because Moore's Proof is meant as a refutation of scepticism about the external world. Although this view is commonly aired in the literature, it is mistaken. Moore's Proof targets the existence of the external word, not one's knowledge of or justification for such existence. Hence, it is not meant as a reply to the sceptic (Moore 1942: 668).

5.1.2 Mooreanism and the Dogmatism Challenge

As noted above, Mooreans think that some common-sense beliefs cannot be rationally undermined by philosophical arguments. Further, they think that this is so even if one cannot identify any flaws in such arguments. In doing so, Mooreans expose themselves to the charge of dogmatism. After all, if one cannot find any flaws in an argument that conflicts with one's common-sense beliefs, how can one resist revising one's common-sense beliefs without being dogmatic? The question casts doubt on the anti-sceptical import of Mooreanism. It presents the Moorean with a challenge: the challenge to explain why one is not being dogmatic in refusing to revise one's common-sense beliefs in response to a philosophical argument.

As Kelly (2005: 182-3) explains, the Moorean might be able to fend off the dogmatism charge by pointing out that when one cannot find any flaws in an argument, there are two competing explanations available. Either the argument is flawless, or the argument is flawed but one is not knowledgeable or skilful enough to identify its flaws. Then, the Moorean could hold that in the relevant cases (e.g. the Underdetermination Argument) one's cognitive limitations are a better explanation for the apparent flawlessness of the argument than its actual flawlessness. As Kelly (2005: 183) observes, if this is dogmatism, then it is of a fairly modest variety, for one's attachment to one's own beliefs derives from the recognition of one's own cognitive limitations, not from the fact that such beliefs are deeply entrenched.

For instance, the Moorean might attempt to answer the challenge by pointing out that the following claim is false:

> if S cannot find any flaws in an argument that conflicts with her beliefs, and S does not revise her beliefs in response to the argument, S is dogmatic.

Surely, the claim is false. In some cases, it might be perfectly reasonable for S not to revise her beliefs although she cannot spot any flaws in an argument that conflicts with them. After all, sometimes, seemingly flawless arguments are merely that: seemingly flawless. And, in some cases, one can know that things are not what they seem. This is true for seemingly flawless arguments that conflicts with S's beliefs, too. But if S knows that a (to her) seemingly flawless argument that conflicts with her beliefs is

flawed, S should not revise her beliefs in response to it. Notice that there need not be anything mysterious in S's knowledge that a (to her) seemingly flawless argument is flawed. S might know that the argument is flawed because S knows that someone else has identified its flaws.

More generally, as Fantl (2018: xii) argues, there are cases in which knowledge can survive encounters with an apparently flawless counterargument. In fact, there are cases in which one should retain one's original confidence in the denial of the argument's conclusion. Think about trick calculations to the effect that 1=0 – surely, when one encounters such calculations, one should not decrease one's confidence that 1=0 is false, even if one cannot spot any flaw in them. *A fortiori*, there are cases in which beliefs can survive encounters with seemingly flawless counterarguments.

Unfortunately, the Moorean cannot dodge the charge of dogmatism so easily when it comes to the Underdetermination Argument, because there is an important disanalogy between the cases just mentioned and the latter's case. Although knowledge that a seemingly flawless argument is flawed is possible, it is doubtful that such knowledge is available in the case of the Underdetermination Argument. For while in other cases one may rely on one's knowledge that others have found or could easily find the flaws in the relevant counterargument, this does not seem to be the case with the Underdetermination Argument. When it comes to the Underdetermination Argument, the philosophical consensus is simply not there – one only needs to take note of the vast literature on the topic to realise this. Thus, the Underdetermination Argument seems to be impervious to decisive criticism from both philosophical laymen and philosophers, at least to criticism that can draw widespread agreement.

The Moorean might retort that things are not as grim as I purport them to be. Although there is no widespread agreement on *what* is wrong with the Underdetermination argument, there is widespread agreement among epistemologists that *something* is wrong. After all, almost every epistemologist agrees that scepticism is false.

Again, I think this kind of reply won't work. It is true that almost every epistemologist agrees that scepticism is false. But these epistemologists are anti-sceptics for different and often incompatible reasons: insofar as the reasons for their agreement conflict, the Moorean can hardly take comfort in this agreement. On the contrary, permanent disagreement on why scepticism is false might be viewed as evidence that

scepticism is true, after all. As an analogy, consider the history of attempts to provide a set of non-circular necessary and sufficient conditions for knowledge. Since Gettier (1963) refuted the analysis of knowledge as justified true belief, several attempts have been made to find an alternative analysis. Needless to say that, to this day, disagreement rules. This disagreement might be viewed as evidence that the project of providing a non-circular analysis of knowledge is misguided; perhaps no such analysis exists (Williamson 2000: 4).

Additionally, as Kelly (2008b: 61) recognizes, the Moorean faces the problem that, while most epistemologists agree with the Moorean that scepticism is false, many, perhaps most, would hold that there is something deeply objectionable about the Moorean response to the sceptic. This puts additional pressure on the Moorean to explain why her appeal to common sense in response to the Underdetermination Argument is not a case of dogmatism.

5.2 Philosophy, Science, and Common Sense

In this section, I discuss the attempt to justify Mooreanism and meet the dogmatism challenge via general considerations about the epistemic differences between philosophy and science. I argue that the prospects of success for the Moorean are bleak because an argument due to Rinard (2013) shows that science can overturn common sense only if philosophy can. Then, I defend Rinard's argument against the attempt to reinstate a significant difference between philosophy and science via considerations about the role played by experience in the confirmation of scientific and philosophical theories and argue that Rinard's argument is especially worrisome to the Moorean anti-sceptic.

5.2.1 Rinard's Argument

Typically, the Moorean opposes philosophy to science: while science can overturn common sense, philosophy cannot. So, some common-sense beliefs are immune to rational revision by philosophical argument, although they are not immune to rational revision by scientific argument (Lycan 2001: 41; Gupta 2006: 178; Kelly 2008b: 55-64). On these accounts, what explains the rationality of the Moorean attitude to common sense and philosophy is the difference between philosophy and science.

Rinard (2013) has challenged the Moorean claim that philosophy, unlike science, is powerless against common sense by offering a positive argument for the claim that philosophy can overturn common sense. The argument turns on the idea that science requires philosophy.

According to Rinard (2013: 187), science can overturn common sense, and if science can overturn common sense, philosophy can, too. So, philosophy can overturn common sense. The key idea in Rinard's argument is that science requires philosophy in the following sense: in several cases in which science has managed to overturn common sense (e.g. the theory of special relativity has managed to overturn the common-sense belief that simultaneity judgements are absolute), scientific arguments have relied on philosophical claims (e.g. simple hypotheses should be preferred to complex ones). But an argument is only as strong as its weakest premise. So, philosophical claims that figure in scientific arguments that overturn common sense must be sufficiently powerful to overturn common sense. But if a philosophical claim is sufficiently powerful to figure in an argument that overturns common sense, there could be an argument consisting entirely of philosophical claims that can overturn common sense. Thus, philosophy can overturn common sense.

5.2.2. A Significant Disanalogy Between Philosophy and Science?

The Moorean might object to Rinard's argument that the fact that some scientific arguments that involve philosophical assumptions can undermine common sense is at best weak evidence that philosophy *alone* can overturn common sense. For the Moorean can point to principled reasons to think that science is in a better position to undermine common sense than philosophy is. The idea is marshalled in an especially clear way by Kelly (2008b).

According to Kelly (2008b), the philosopher's arsenal is limited to arguments and theories, and, crucially, the assessment of arguments and theories is not independent of common-sense beliefs. In particular, the 'the more credible the proposition targeted by a given argument, the less credible it is that the argument is sound' (Kelly 2008b: 56). This makes it particularly hard for a philosophical argument/theory to undermine common sense, because the fact that it conflicts with common sense seems to count against its credibility.

In contrast, science has a better claim against common sense because the credibility of scientific theories and arguments is less dependent on our present beliefs than the credibility of philosophical theories and arguments. The reason is that scientific theories and arguments, unlike philosophical ones, are sensitive to experience. In particular, 'a scientific theory that seems surely false by our present lights can nevertheless come to be rationally accepted: namely, by making surprising predictions that are independently verified' (Kelly 2008b: 56).

The upshot is that there is an important disanalogy between philosophy and science. In philosophy, there is no mechanism analogous to prediction in science that can make a theory confirmed despite its being at odds with common sense. And that disanalogy may be used by the Moorean to insist that philosophy *alone* cannot overturn common sense.

I have two responses to this sort of considerations against Rinard's argument. The objection against Rinard's argument comes down to this: the fact that science in conjunction with philosophy can overturn common sense is at best weak evidence that philosophy alone can overturn common sense, because science in conjunction with philosophy can overturn common sense due to its sensitivity to experience. But philosophy alone is not sensitive to experience.

That objection is ineffective against Rinard's argument, because, in the examples Rinard gives, science overturns common sense in conjunction with philosophy but the philosophical claims at issue receive no support from experience. They are purely philosophical claims, and they are strong enough to figure in arguments that can overturn common sense. Thus, Rinard's examples are good evidence that philosophy alone can overturn common sense.

Second and in light of what I have just said, it is not true that in philosophy there is no mechanism analogous to prediction in science that can make a theory confirmed despite its being at odds with common sense. A philosophical theory or argument can overturn common sense provided that its premises are sufficiently powerful. In fact, those premises may be themselves part of common sense or be supported by arguments that draw on common sense.

If Rinard's argument is correct and what I have said is right, the Moorean cannot resist a philosophical argument that runs contrary to common sense on the grounds

that philosophy alone cannot overturn common sense. Rather, such philosophical arguments deserve scrutiny. Indeed, it would be dogmatic for the Moorean to dismiss a philosophical argument just because it conflicts with some deeply and widely held belief, even more so if the argument was seemingly flawless.

Even so, the Moorean anti-sceptic might take comfort in the idea that, although she cannot automatically dismiss the Underdetermination Argument, looking at the Underdetermination Argument will reveal that its premises are too weak to overturn our common-sense beliefs.

As Rinard (2013: 192) observes, the Moorean should not take comfort in this idea. The premises of the Underdetermination Argument – that S's evidence does not favour P over the incompatible sceptical hypothesis SH, and that justification requires evidential favouring – are indeed very plausible. Now compare these claims with the philosophical claim that simple hypotheses should be preferred to complex ones. As Rinard's argument shows, that claim is powerful enough to overturn common sense. But, at least at first look, that claim does not seem to be more powerful than the Underdetermination Argument's premises. Thus, there is some *prima facie* evidence that the Underdetermination Argument can overturn common sense.

Additionally, the problem for the Moorean anti-sceptic is especially acute when it comes to scepticism. Notice that the phrase 'to overturn common sense' is ambiguous. An argument for a conclusion C may be said to overturn common sense if C conflicts with some common-sense belief and one should accept C in response to the argument. Rinard's example fits this scheme. The theory of special relativity commands acceptance of the claim that simultaneity is not absolute, which conflicts with common sense.

However, there is another way in which an argument for a conclusion C may be said to overturn common sense, namely when C conflicts with some common-sense belief and one should suspend judgement on it in response to the argument.

The distinction is relevant in the case of the Underdetermination Argument. It is a common assumption in the literature on scepticism that the sceptic's purpose is to convince one to accept scepticism, i.e. to accept the claim that justification is impossible. But the assumption is problematic, because the sceptic's purpose might be to lead one to suspend judgement rather than convincing one to accept scepticism. And

although the sceptic's premises might be too uncertain to generate conviction in the sceptical conclusion, they might be sufficiently strong for suspending judgement on whether justification is possible. After all, the bar is lower for suspending judgement on a common-sense belief than for disbelieving it.

The distinction between accepting scepticism and suspending judgement echoes the distinction between Academic and Pyrrhonian Scepticism (Sextus Empiricus PH 1 3). Notice that suspending judgement on whether justification is possible is just as damaging to the anti-sceptic as believing that scepticism is true. For one thing, both suspending judgement on scepticism and accepting scepticism involve giving up the belief that scepticism is false. For another thing, both lead to extensive belief revisions (unless one can appeal to non-epistemic reasons).

If you believe that justification about whether P is impossible, you should suspend judgement on whether P. For if you believe that justification about whether P is impossible, then you ought to believe that believing in P/~P is unjustified. But it is irrational to believe that P/~P while also believing that believing in P/~P is unjustified. Similarly, if you suspend judgement on whether justification is impossible, you should suspend judgement on whether P. For if you suspend judgement on whether justification is impossible, you ought to suspend judgement on whether believing that P/~P is justified. But it is irrational to believe that P/~P while suspending judgement on whether believing that P is justified (Rinard 2018: 257; Smithies 2019: 286).

5.3 Mooreanism and the Norms of Belief Change

Even if the Moorean cannot appeal to a supposed difference between philosophy and science, there may be other ways to justify Mooreanism and avoid the dogmatism charge. One common strategy to motivate Mooreanism is by appeal to some general norm of belief revision (Lemos 2004: 1; Kelly 2005: 188). The idea is this: some of our common-sense beliefs have feature F, where, typically, F is an epistemic feature, like being known or being justified. Further, if a belief has F, one should not abandon that belief in response to any philosophical argument. Therefore, one should not abandon some of our common-sense beliefs in response to any philosophical argument.

In this section, I explore the idea that some norm of belief revision can be used to motivate Mooreanism while avoiding charges of dogmatism. In doing so, I look at

different proposals for different norms meant to accomplish this task and find them all unsuccessful.

5.3.1 Two Constraints

Any successful justification of Mooreanism via a norm of belief revision must satisfy (at least) two constraints. First, the Moorean must choose a feature F that is both a feature of our common-sense beliefs and a feature that makes such beliefs impervious to change in response to philosophical arguments. Although this constraint is often overlooked, it is, for reason that will become clear during my discussion, not easy to satisfy.

Second, F should provide a genuine motivation for Mooreanism. Consider the following belief revision principle:

> (MORE REASONABLE) One should not abandon a belief in response to an argument when it is more reasonable to hold that belief given one's total evidence.

Suppose it is more reasonable, given one's total evidence, to hold some common-sense beliefs than to abandon them in response to any philosophical argument. Then, some common-sense beliefs are not rationally undermined by any philosophical argument to the contrary.

The Moorean cannot appeal to this sort of reasoning to motivate Mooreanism. Surely, the Moorean is convinced that it is more reasonable, given one's total evidence, to retain some common-sense beliefs in response to any philosophical argument than to abandon them. But that is just a restatement of Mooreanism; it is what the Moorean needed to motivate in the first place.

5.3.2 Plausibility

Consider the following norm for belief change:

> (MORE PLAUSIBLE) One should not abandon a belief in response to a counterargument when the belief is more plausible than the argument's premises.

MORE PLAUSIBLE seems to be the norm that Fine (2001: 2) and Lycan (2001: 42)

have in mind when they endorse Mooreanism. It is an obvious candidate to motivate Mooreanism, because, at least to many, some common-sense beliefs seem more plausible than the premises of any philosophical argument to the contrary. Moreover, the idea that one should favour the more plausible option in cases of doxastic conflict is itself plausible.

However, MORE PLAUSIBLE faces two decisive objections, at least if plausibility is understood literally. First, MORE PLAUSIBLE seems false because it ignores that only marginal superiority in plausibility does not make a belief resistant to change. Suppose S believes that P, and P is very plausible. Now, suppose that S encounters a counterargument to P whose premises are only slightly less plausible than P. Since P is more plausible than the argument's premises, MORE PLAUSIBLE says that S should not abandon her belief that P. But surely this is the wrong result: the argument for ~P seems a very good argument. In the absence of additional considerations, S should revise her belief that P by lowering her confidence in P. In fact, it is conceivable that S should lower her confidence in P so much that she should suspend judgement on P and thus abandon her belief in P.[4]

For the Moorean, the natural response is to strengthen the antecedent of MORE PLAUSIBLE:

> (MUCH MORE PLAUSIBLE) One should not abandon a belief in response to a counterargument when the belief is *much* more plausible than the argument's premises.

MUCH MORE PLAUSIBLE raises the bar for a belief's resistance to revision by imposing a more exacting condition on the relevant belief than MORE PLAUSIBLE does: it is not sufficient that the belief be merely more plausible than the counterargument's premises. Instead, the belief must be *much* more plausible in comparison.

Unlike MORE PLAUSIBLE, MUCH MORE PLAUSIBLE does not deliver the wrong verdict when S meets a counterargument to P whose premises are only slightly less plausible than P, because MUCH MORE PLAUSIBLE remains silent on such cases. Thus, MUCH MORE PLAUSIBLE leaves open the possibility that S should

[4] The case is analogous to a case discussed by Hájek (2008: 94) in his defence of Hume's balancing principle – a principle linking belief and probability – in the context of Hume's no-miracles argument.

revise her belief when she encounters a counterargument whose premises are only slightly less plausible than the conflicting belief.

Notice that this comes at a cost for the Moorean. By raising the standards for resisting revision, MUCH MORE PLAUSIBLE has also made it harder for a belief to meet those standards. The problem is pressing because a feature of some of the arguments the Moorean would like to resist (e.g. the Underdetermination Argument) is that their premises enjoy a high degree of plausibility, or so I have suggested at various points in this thesis. Hence, even if some common-sense beliefs are more plausible than the arguments' premises, it is doubtful that they are *much* more plausible.

In any case, MORE PLAUSIBLE and MUCH MORE PLAUSIBLE face another decisive objection. To see why, it is sufficient to appreciate that the plausibility of P does not track whether it is rational to believe that P, but merely whether it *seems* rational to believe that P, or whether it is reasonable to believe it *according to first impressions or indications*. In other words, P is plausible to the extent that it seems true when one initially considers it. As Conee (2001: 57) and Kelly (2005: 189) stress, plausibility in this sense is a bad proxy for belief worthiness. Hence, comparative plausibility is a poor guide to belief revision.

That plausibility is a bad proxy for belief worthiness is nicely illustrated by an example provided by Kelly (2005: 189): the Unrestricted Comprehension Principle – the principle that for any property there is a set of all objects with that property – does not stop being plausible when one learns that it leads to Russell's Paradox. Nonetheless, when one learns that the Unrestricted Comprehension Principle leads to Russell's Paradox, it is not reasonable to continue believing it. Ironically, MORE PLAUSIBLE may itself be a case in point: MORE PLAUSIBLE might initially strike one as plausible, but reflection reveals that MORE PLAUSIBLE is false. More generally, plausibility is consistent with known falsity. Hence, plausibility is a bad proxy for belief worthiness.

From this, Kelly concludes that comparative plausibility is a poor guide to belief revision and that MORE PLAUSIBLE is false. But how can he derive that MORE PLAUSIBLE is false? After all, his example does not involve an *argument whose premises are less plausible than the conflicting belief*. Thus, some additional argument is needed to connect the claim that plausibility is a bad proxy for belief worthiness with the claim that comparative plausibility is a poor guide to belief revision.

Such an argument is readily available. Notice that plausibility is a bad proxy for belief worthiness in two senses: a plausible proposition can be known to be false and hence be belief-unworthy – think about the Unrestricted Comprehension Principle –, and an implausible proposition can be known to be true and hence be belief-worthy – think about cases where one is aware that one is subject to an illusion, like the Müller-Lyer effect.[5] Then, it is possible that an argument from implausible but known premises leads to the negation of a plausible and false proposition P. Surely, in such cases, one should abandon the belief that P although the premises of the argument are less plausible than P. Hence, MORE PLAUSIBLE is false.

For the same reasons for which comparative plausibility is a bad proxy for belief worthiness, *great* comparative plausibility is. A very plausible proposition can be known to be false and hence be belief-unworthy, and a very implausible proposition can be known to be true and hence be belief-worthy. Then, it is possible that an argument from very implausible but known premises leads to the negation of a very plausible and false proposition P. Surely, in such cases one should abandon the belief that P although P is much more plausible than the premises of the counterargument – the argument from the theory of special relativity to the negation of the absolute simultaneity judgements is a case in point.

What is the upshot for the Moorean response to the Underdetermination Argument? The Moorean anti-sceptic cannot appeal to a norm about comparative plausibility to justify her attitude towards the Underdetermination Argument. Even if we grant that the premises of the Underdetermination Argument are much less plausible than one's conflicting belief that one is justified in believing that P, one is not thereby entitled to dismiss the Underdetermination Argument, for, upon scrutiny, the Underdetermination Argument may be more worthy of belief than one's conflicting common-sense belief. The problem is especially pressing since, as I have shown in chapters 1 to 4, the sceptic typically puts forward arguments in support of the Underdetermination Argument's premises. So, it is not as if the anti-sceptic has nothing to work on except for the brute plausibility of the sceptic's argument premises and the brute

[5] The Müller-Lyer effect is an optical illusion in which two lines of equal length appear to be of different length due to their having arrowheads or arrow tails.

plausibility of her own common-sense beliefs. In light of this, appeals to plausibility look like the manifestation of dogmatic closed-mindedness.

5.3.3 Certainty

The Moorean might attempt to justify her anti-sceptical attitude by appealing to the certainty of her common-sense beliefs. In fact, this seems to be the rationale invoked by Moore (1959: 226), Armstrong (1983: 49), and Pollock and Cruz (1999: 7). Consider the following norm for belief change:

> (MORE CERTAIN) One should not abandon a belief in response to a counterargument when the belief is more certain than the argument's premises.

How should we understand MORE CERTAIN? As Kelly (2005: 190) points out, one option is to understand 'certainty' as 'psychological certainty' or 'confidence'. Another option is to understand 'certainty' in epistemic terms, i.e. as 'evidential certainty'. In the latter sense, certainty does not concern one's actual confidence in the truth of a proposition, but the confidence one *ought* to have in it given one's total evidence. In this sense, certainty is about what is rational or reasonable to believe in one's situation.

Let us look at both ways to articulate MORE CERTAIN in turn. Under the psychological interpretation of 'certainty', MORE CERTAIN is equivalent to:

> (MORE CONFIDENT) One should not abandon a belief in response to a counterargument when one has more confidence in the belief than in the argument's premises.

Like MORE PLAUSIBLE, MORE CONFIDENT is an obvious candidate to motivate Mooreanism, because many feel more confident in some of their common-sense beliefs than in the premises of any philosophical argument to the contrary. Moreover, the idea that one should favour beliefs of which one is more confident in cases of doxastic conflict has an initial air of plausibility.

It is easy to see that MORE CONFIDENT shares the same problems that afflict Moorean appeals to plausibility. For one thing, MORE CONFIDENT neglects the fact that marginal differences in confidence between P and the premises of an argument for ~P are not sufficient to maintain a high confidence in P. Further, any attempt

to remedy to this problem by adding a requirement that the confidence in P be much higher than the confidence in the argument's premises makes it more difficult for common-sense beliefs to meet the bar.

For another thing, Moorean appeals to confidence face the charge that comparative confidence is not a good proxy for what one ought to believe. As noted in section 5.1.1, sheer confidence has no normative import; it can be justified or be the result of dogmatism or gullibility. But, surely, a correct norm for belief revision cannot tell us to not abandon our dogmatic or gullible beliefs just because we are very confident in them.[6]

The obvious fix is to replace talk of sheer confidence with talk of rational confidence and talk of psychological certainty with talk of evidential certainty. Thus, the Moorean might attempt to replace MORE CONFIDENT with:

> (MUCH MORE RATIONALLY CONFIDENT) One should not abandon a belief in response to a counterargument when one is much more rationally confident in the belief than in some of the argument's premises.

MUCH MORE RATIONALLY CONFIDENT avoids the problems of appeals to plausibility and psychological certainty. Unlike them, rational confidence is a property that matters to what one *ought* to believe. Further, although marginal superiority in rational confidence might be insufficient to preserve a belief against a counterargument, great superiority in rational confidence is not. Indeed, MUCH MORE RATIONALLY CONFIDENT seems trivially true.

That MUCH MORE RATIONALLY CONFIDENT seems trivially true should raise suspicions. In fact, it faces the same problem as MORE REASONABLE does: it does not offer a genuine motivation for Mooreanism. Surely, the Moorean is convinced that one is more rationally confident in some of one's common-sense beliefs than in the premises of any philosophical argument to the contrary. But that is just a restatement of Mooreanism; it is what the Moorean needed to motivate in the first place.

[6] Thus, appeals to the alleged irresistibility of common-sense beliefs (Strawson 1985: 414) won't help the Moorean.

Again, what is the upshot for the Moorean anti-sceptic? I have examined several attempts to motivate Mooreanism via a norm for belief that appeals to certainty. Every attempt fails either because the suggested norm is false – it does not specify a property of one's belief in light of which it is rational not to revise one's beliefs in response to the Underdetermination Argument – or because it does not motivate Mooreanism in a genuine way – it begs the question it was meant to answer.

5.3.4 Knowledge

Finally, the Moorean might attempt to justify her resistance to the Underdetermination Argument by appealing to the fact that she *knows* that her common-sense beliefs are true. Consider the following norm:

> (KNOWN) One should not abandon a belief in response to a counterargument when one knows that the belief is true.

At first glance, KNOWN might look promising to the Moorean. On the one hand, unlike plausibility and certainty, being known does not come in degrees. Thus, KNOWN does not face the problem that a belief with F can be undermined by a counterargument whose premises have F to a slightly less degree. On the other hand, unlike psychological certainty, being known is a positive epistemic feature. Thus, being known might be a good candidate for a property that makes a proposition belief-worthy.[7]

I shall argue that KNOWN faces problems analogous to the ones Kelly (2005: 192-94) identifies for MORE CONFIDENT. Notice that KNOWN has two readings:

> [KNOWN (1)] One should not abandon a belief in response to a counterargument when *antecedent to becoming aware of the counterargument* one knows that the belief is true.
>
> [KNOWN (2)] One should not abandon a belief in response to a counterargument when *posterior to becoming aware of the counterargument* one knows

[7] In fact, some authors have suggested that knowledge is the norm of belief (Williamson 2000; Sutton 2007; Bach 2008).

that the belief is true.

Consider KNOWN (1). KNOWN (1) seems to condone neglecting one's total evidence, for one's knowledge might depend on the assumption that there is no powerful counterargument to one's belief. Hence, KNOWN (1) is false. To put if differently, it is well-known that even if one knows that P, one is not rationally entitled to dismiss every future piece of misleading evidence. To do so would be unreasonable and dogmatic (Kripke 2011; Harman 1973).

Kelly (2005: 72-3) believes that the falsity of KNOWN (1) is no obstacle to appealing to knowledge in response to the sceptic, because in some cases one is rationally entitled to dismiss misleading counterevidence on the basis of one's knowledge. Not all misleading counterevidence destroys one's antecedent knowledge. In so doing, Kelly seems to suggest that the Moorean could appeal to KNOWN (2) to justify Mooreanism.

However, KNOWN (2) is problematic because it does not give applicable advice. The point of a norm like KNOWN (2) is to settle what one should believe in cases of conflict between one's belief and a counterargument. But KNOWN (2) presupposes that one already knows what one knows when the conflict at issue is settled. So, the Moorean cannot appeal to KNOWN (2), either.

This concludes my review and assessment of Moorean appeals to norms of belief revision. The picture is bleak for the anti-sceptic: none of the norms I have examined can be used to resist the Underdetermination Argument in a non-dogmatic way. I return to these results in the concluding remarks.

5.4 Mooreanism and Philosophical Method

Kelly (2005) offers a defence of Mooreanism that draws on methodological considerations. According to Kelly, the sceptic goes wrong by imposing 'artificially demanding standards for what it takes to undermine a general philosophical principle' (2005: 202), because the sceptic is committed to the claim that no significant weight should be accorded to judgements about particular cases in the assessment of general philosophical principles. That commitment is not shared by Mooreans. Thus, Mooreanism is methodologically superior to scepticism.

In this section, I evaluate Kelly's argument for the methodological superiority of

Mooreanism over scepticism. I start with a review of the key methodological concepts in Kelly's argument: particularism, generalism or methodism, and reflective equilibrium. Then, I reconstruct Kelly's argument and argue that it fails.

5.4.1 Particularism, Generalism, and Reflective Equilibrium

Drawing on Chisholm (1982: 66), Kelly (2005: 197-8) describes two competing methodologies in epistemology: particularism and methodism or generalism. They are competing methodologies in that they prescribe different starting points and paths for building and assessing epistemological theories. How do they differ?[8]

Particularists believe that epistemology should begin with our judgements about particular cases and that such judgements should act as a constraint on epistemological theorising. On the contrary, generalists hold that epistemology should begin with our judgements about general epistemic principles and that such judgements should guide our evaluation of particular cases.

In other words, particularists give more weight to their judgements about particular cases, whereas generalists give more weight to their judgements about general principles. In fact, particularists use their judgements about particular cases to regulate judgements about general principles, whereas generalists use their judgements about the latter to regulate the former.

As Kelly (2005: 198) observes, *tertium datur*: particularism and generalism are not the only methodologies available. Instead of giving priority to particular cases over general principles or to general principles over particular cases, one could give priority to neither. That is, one could assign equal weight to judgements about particular cases and to judgements about general principles in pursue of a balance by way of mutual

[8] Chisholm (1982: 66) explains the difference between particularism and methodism in terms of how they address two basic questions:

> A: *What* do we know, or what is the *extent* of our knowledge?
> B: How are we to decide *whether* we know, or what are the *criteria* of knowledge?

Particularists start with an answer to A, and given an answer to A they proceed to answer B. Thus, particularists give methodological priority to their insights into particular cases of knowledge and ignorance over insights into general epistemic principles. Methodists have it the other way around: they start with an answer to B, and given an answer to B they proceed to answer A. Thus, methodists give methodological priority to their insights into general epistemic principles over insights into particular cases of knowledge and ignorance.

adjustments. This is the method of reflective equilibrium (Goodman 1953; Rawls 1972; Harman 2003).

So far, particularism, generalism, and reflective equilibrium have been described in their pure form. Particularists give no weight at all to judgements about general principles. For them, what determines the goodness of a general principle is just how well it fits their judgements about particular cases. Similarly, but in an opposite fashion, generalists give no weight at all to judgements about particular cases. For them, what determines the goodness of a judgement about a particular case is just how well it fits their judgements about general principles. Finally, the reflective equilibrium theorist gives exactly equal weight to judgements about particular cases and to judgements about general principles.

Of course, this is a simplification: in reality, particularism, generalism, and reflective equilibrium come in many shades. A better picture is this: we can think of the methods available to the epistemologist as lying on a spectrum. At the very ends of the spectrum lie particularism and generalism in their pure form. Borrowing Kelly's terminology, we can call them 'hyper-particularism' and 'hyper-generalism' or 'hyper-methodism'. In the exact centre of the spectrum lies the pure form of reflective equilibrium. All in between is a continuum of methods, that is moderate versions of particularism, generalism, and reflective equilibrium (Kelly 2005: 199).

5.4.2 Kelly's Argument

With this picture of methods in mind, let us turn to Kelly's argument for the methodological superiority of Mooreanism over scepticism. In its simplest form, Kelly's argument amounts to this: the sceptic must endorse hyper-methodism, the claim that no significant weight should be accorded to judgements about particular cases when assessing general principles. Yet hyper-methodism is a bad methodology. On the contrary, the Moorean is committed to the sane-sounding claim that it is reasonable to reject a principle when it conflicts with a great number of particular judgements. Hence, Mooreanism is methodologically superior to scepticism.

Why think that the sceptic must endorse hyper-methodism? Because the sceptic is committed to the claim that certain principles force a radical revision of our beliefs, and that it is no telling point against these principles that they are inconsistent with a

large number of judgements about particular cases. But how can that commitment be rational unless one accepts hyper-methodism? In other words, the sceptic must endorse hyper-methodism because it is the only methodology available to her (Kelly 2005: 199-201).

Established that the sceptic is committed to hyper-methodism, Kelly goes on to argue that hyper-methodism is a bad methodology because it promotes ignoring counterexamples (Kelly 2005: 201-2). Indeed, Kelly stresses, the methodology employed by the sceptic differs radically from the methodology employed in philosophy, epistemology included, when scepticism is not at issue. Thus, scepticism rests on bad methodology.

If Kelly's argument is successful, the Moorean response to the Underdetermination Argument is vindicated: one should not abandon one's common-sense beliefs in response to the Underdetermination Argument, because scepticism, but not Mooreanism, rests on bad methodological commitments. Indeed, if Kelly's argument is sound, the Moorean has a satisfactory answer to the dogmatism charge. Clearly, one need not be able to identify a flaw in the sceptic's argument to appreciate that the sceptic is following a radically mistaken methodology. And if the sceptic is following such a methodology, it is sensible to resist her arguments. In any case, even if hyper-methodism were not as objectionable as Kelly thinks, scepticism would be better off without such a substantial methodological commitment.

Further, if Kelly's argument is correct, the popular conception of the Moorean as someone committed to a strong meta-philosophical view, namely particularism, is mistaken. For the Moorean is only committed to the modest claim that it is reasonable to reject a principle when it conflicts with a great number of particular judgements. That idea is not exclusive to particularism, but it is shared by reflective equilibrium views and moderate form of generalism. As a consequence, rejecting particularism is not sufficient to reject Mooreanism. Ironically, it is the sceptic who is committed to a substantial and implausibly strong meta-philosophical view, not the Moorean (Kelly 2005: 202-3).

5.4.3 Scepticism Without Hyper-Methodism

Here is some good news for the sceptic: contrary to what Kelly thinks, the sceptic need

not endorse hyper-methodism. Hence, the sceptic is not committed to a bad methodology, the Moorean has no methodological advantage over the sceptic, and Kelly's argument fails to vindicate Mooreanism, or so I shall argue. My criticism of Kelly's argument parallels and draws on a similar criticism put forward by Rinard (2013).

My objection to Kelly's argument can be put succinctly. Let us assume that the sceptic gives no significant weight to judgements about particular cases that conflict with scepticism. Does this make the sceptic a hyper-methodist? I do not think it does. At least in principle, moderate particularists, reflective equilibrium theorists, and moderate methodists can adopt the sort of stance the sceptic adopts in the case of scepticism without thereby endorsing hyper-methodism. In other words, they can grant that in some cases judgements about particular cases that conflict with a general principle have no weight while rejecting the claim that this is always the case. The upshot is that the sceptic is not committed to hyper-methodism.

To better appreciate the point, consider these two cases due to Rinard (2013):

> GAMBLER'S FALLACY
>
> S tends to commit the gambler's fallacy. If S sees a fair coin land heads many times in a row, S judges that the coin is more likely to land tails than heads on the next toss. In fact, S has judged so many times. One day, S takes a probability theory class, and learns the principle 'independent tosses of a fair coin are equally likely to come up heads.' After class, S realizes that if she accepts this principle, she should no longer judge that heads is more likely than tails after a long series of tails. (Adapted from Rinard 2013: 204)
>
> CERTAINTY
>
> S has never carefully reflected on the epistemic status of her beliefs, and she is convinced that she is justified in being certain of a great deal of things. In fact, S is convinced that she has ruled out all the possibilities in which her beliefs are false. One day, S takes an epistemology class and learns the principle 'if one cannot eliminate a sceptical scenario in which ~P is true, one is not justified in being certain that P.' After class, S realizes that if she accepts this principle, she should stop believing that she is

> justified in being certain of a great deal of things. (Adapted from Rinard 2013: 204-5)

Notice that both situations are cases of conflict between a large number of individual judgements about individual cases and one's judgement about a general principle. How should S respond to this conflict?

To some extent, this depends on features of the two situations that have not been described, of course. Still, the crucial point is this. As Rinard (2013: 205) argues, for both cases, the rational response for S may well be to abandon her individual judgements about probability and certainty in favour of the general principles learned in class. For in both cases, S may learn about the general principles and thereby come to understand that her antecedent beliefs were systematically mistaken and confused.

GAMBLER'S FALLACY and CERTAINTY provide counterexamples to the claim that in cases of conflict between judgements about general principles and judgements about individual cases one should favour judgements about individual cases. In fact, in GAMBLER'S FALLACY and CERTAINTY, one ought to give priority to judgements about general principles. Further, this is something moderate particularists, reflective equilibrium theorists, and moderate methodists can and should recognize. Thus, although it is true that if the sceptic endorses hyper-methodism, then she gives no significant weight to any judgement about particular cases, it is false that if the sceptic gives no significant weight to any judgement about particular cases in the case of scepticism, then the sceptic endorses hyper-methodism. Kelly's argument fails.

In defence of Kelly, one might insist that the sceptic's case is different from GAMBLER'S FALLACY and CERTAINTY. Thus, it may be rational to favour the general principle although it conflicts with a large number of individual judgements in GAMBLER'S FALLACY and CERTAINTY, but it is not rational in the sceptic's case. Surely, if hyper-methodism is unacceptable, in most cases in which a principle conflicts with a large number of individual judgements one should favour the individual judgements to the principle. Cases in which it is rational to prefer a general principle to a large number of particular judgements are an exception. The case of scepticism is at least *prima facie* suspicious.

Notice that CERTAINTY puts considerable pressure on this line of reasoning. As

Rinard (2013: 205-6) observes, CERTAINTY is structurally analogous to the underdetermination scepticism case. In both cases, S becomes acquainted with some epistemic principle. And in both cases, one initially judges that one has a certain kind of epistemic access to the truth. Finally, just like the rational response in the certainty case may be to revise one's belief that one is justified in being certain in a great deal of things, the underdetermination sceptic claims that the rational response in the underdetermination case is to revise one's belief that one's beliefs are justified.

One might suggest that in GAMBLER'S FALLACY and CERTAINTY S comes to understand that her individual judgements were systematically mistaken and confused. Although it is easy to see how S's individual judgements might be systematically mistaken and confused in these cases, the same is not true in the case of underdetermination scepticism. Thus, GAMBLER'S FALLACY and CERTAINTY are not structurally analogous to the sceptic's case, after all.

Again, CERTAINTY puts considerable pressure on this line of thinking. S's judgements in CERTAINTY are systematically mistaken because S has never carefully reflected on the epistemic status of her beliefs. The underdetermination sceptic can point out that the same explanation is readily available in the case of underdetermination scepticism and justification. After all, most people, even most philosophers, rarely spend time, if at all, thinking carefully about scepticism and justification.

To summarise: Kelly claims that the underdetermination sceptic, unlike the Moorean, is committed to hyper-methodism, a mistaken philosophical methodology. Hence, the Moorean can resist the Underdetermination Argument based on methodological considerations that vindicate the Moorean approach. Drawing on Rinard (2013), I have argued that Kelly's vindication of Mooreanism fails. Just as one need not be a particularist to appreciate counterexamples, one need not be a hyper-methodist to appreciate the possibility of systematic belief revisions. For all Kelly has shown, the Moorean can claim no methodological superiority over the underdetermination sceptic.

5.5 Concluding Remarks

Mooreans believe in the epistemic superiority of some common-sense beliefs over philosophical arguments. According to Mooreans, some common-sense beliefs are not

rationally undermined by any philosophical argument to the contrary. As a consequence, one should not abandon some common-sense beliefs in response to any philosophical argument, regardless of one's ability to identify its flaws. Famously, Mooreanism has anti-sceptical applications: some of our common-sense beliefs – beliefs like the belief that we know a lot or the belief that we are justified in believing many things – provide us with a rational basis to resist any sceptical counterargument.

Mooreans face the charge of dogmatism. After all, if one can find no flaws in an argument, how can one's refusal to accept that argument be anything other than dogmatic? Thus, Mooreans face the challenge of explaining why common-sense beliefs enjoy such immunity from philosophical argument. The problem is especially pressing in the case of scepticism, since many sceptical arguments, like the Underdetermination Argument, are notoriously appealing as well as hard to refute.

In this chapter, I have looked at three attempts to motivate Mooreanism and avoid the dogmatism charge: by appeal to an alleged difference between philosophy and science, by appeal to general norms for belief revision, and by appeal to the methodological superiority of Mooreanism over scepticism.

First, the Moorean might attempt to justify Mooreanism by opposing philosophy to science: while science can overturn common sense, philosophy cannot. So, some common-sense beliefs are immune to rational revision by philosophical argument, although they are not immune to rational revision by scientific argument.

Following Rinard (2013), I have argued that the prospects for justifying Mooreanism in this way are poor. For Rinard has shown that there is a positive argument for the claim that philosophy can overturn common sense that turns on the idea that science can overturn common sense and that science requires philosophy. As a consequence, the Moorean faces an unhappy choice: either science cannot overturn common sense or philosophy can.

I have defended Rinard's argument against the objection that the argument fails because science can overturn common sense due to its sensitivity to experience, but philosophy *alone* is not sensitive to experience. When science overturns common sense in conjunction with philosophy, the philosophical claims at issue receive no support from experience. Thus, a purely philosophical theory or argument can overturn common sense provided that its premises are sufficiently powerful. In fact, those premises

may be themselves part of common sense or be supported by arguments that draw on common sense.

Again, the problem is especially acute in the case of the Underdetermination Argument. On the one hand, the premises of the Underdetermination Argument seem at least as plausible as philosophical claims that figure in scientific arguments that can overturn common sense. On the other hand, although the sceptic's premises might be too uncertain to generate conviction in the sceptical conclusion, they might be sufficiently strong for suspending judgement on whether justification is possible. Judgement suspension on the possibility of justification is all is needed to overturn common sense.

Second, the Moorean might attempt to justify Mooreanism by appeal to some general norm of belief revision. The idea is this: some common-sense beliefs have feature F, where, typically, F is an epistemic feature, like being known or being justified. Further, if a belief has F, one should not abandon that belief in response to any philosophical counterargument. Therefore, one should not abandon some common-sense beliefs in response to any philosophical argument.

I have looked at three candidates for F that may make a common-sense belief immune to belief revision by philosophical argument: being more plausible, being more certain, and being known. I have argued that each of these proposals faces problems that make it unsuitable for the Moorean. Comparative plausibility and psychological certainty are not good candidates because they do not have the epistemic import needed to vindicate Mooreanism – they do not make a belief that has them belief-worthy. Comparative evidential certainty and knowledge are not good candidates because they do not motivate Mooreanism in a genuine way or do not provide appropriate guidance.

More generally, a problem emerges for the Moorean project of identifying a feature F and a norm that vindicate Mooreanism, given that the Moorean cannot simply insist that her favoured beliefs are more belief-worthy than the premises of any philosophical counterargument. On the one hand, if F is not an epistemic property, a norm that favours beliefs with F over beliefs that lack F is not the right sort of norm to justify Mooreanism, because F and the corresponding norm lack epistemic import. On the other hand, if F is an epistemic property, a norm that favours beliefs with F over beliefs

without F does not provide guidance, since a belief's having F depends on the assumption that no powerful counterargument to one's belief is available. Either way, Mooreans are vulnerable to the dogmatism charge.

Finally, the Moorean might argue for Mooreanism drawing on methodological considerations. According to Kelly (2005), the sceptic is committed to hyper-methodism, the claim that no significant weight should be accorded to judgements about particular cases when assessing general principles. Yet hyper-methodism is a bad methodology. On the contrary, the Moorean is committed to the sane-sounding claim that it is reasonable to reject a principle when it conflicts with a great number of particular judgements. Hence, anti-sceptical Mooreanism is methodologically superior to scepticism.

Drawing on Rinard (2013), I have argued that Kelly's argument fails: contrary to what Kelly believes, the sceptic is not committed to hyper-methodism. Although it is true that if the sceptic endorses hyper-methodism, then she gives no significant weight to any judgement about particular cases, it is false that if the sceptic gives no significant weight to any judgement about particular cases *in the case of scepticism*, then the sceptic endorses hyper-methodism. At least in principle, moderate particularists, reflective equilibrium theorists, and moderate methodists can adopt the sort of stance the sceptic adopts in the case of scepticism without thereby endorsing hyper-methodism. One need not be a hyper-methodist to appreciate the possibility of systematic belief revisions.

6

Conclusion: Last Thoughts of a Provisional Sceptic

In this thesis, I have argued for underdetermination scepticism about the external world and defended the Underdetermination Argument against some popular objections. The two subgoals of this thesis – to elucidate and to defend underdetermination scepticism – go hand in hand and serve its overarching goal: to show that, despite its present unpopularity, underdetermination scepticism is a live option in epistemology.

This conclusive chapter is painted with a broad brush. It recaps the thesis' core results, generalizes them, and shows how they serve its main aim. It also discusses the thesis' limitations and describes some open questions to be pursued in future research. It gives the reader a synoptic view of the thesis by organising its content around some important ideas. The reader interested in detail may refer back to the individual chapters.

6.1 Summary

Scepticism about justification is the view that justification is impossible. Underdetermination scepticism is scepticism that turns on the idea that our beliefs are underdetermined by the evidence relative to certain sceptical hypotheses. More precisely, the underdetermination sceptic argues that S's evidence does not favour P over SH – a suitable sceptical hypothesis – because S has the same evidence in the good and the bad case and infers via an underdetermination principle that S is not justified in believing that P.

Chapter 1 set the stage for the thesis by laying out the Underdetermination Argument for scepticism. It did a number of things like clarifying the core concepts of the thesis (evidence, justification, contrastive underdetermination) and motivating the underdetermination principle and the principle that evidence cannot justify incompatible hypotheses. Most importantly, it narrowed down the thesis' argumentative scope by assuming (quite controversially) an evidentialist conception of justification, a mentalist conception of evidence, and by setting aside certain kinds of response to the Underdetermination Argument. However, these moves were not wholly arbitrary.

I defended mentalism about evidence – the view that evidence consists of non-factive mental states – on the grounds that such view sits well with a conception of evidence as the kind of thing to which our beliefs are responsive. More precisely, mentalism appears compelling in light of the causal role played by non-factive mental states in our cognitive lives: they are the causal mediators between mind and world.

I also argued that certain responses to underdetermination scepticism are doomed to failure because they do not address the root of the sceptical problem (contextualism and relevant alternatives theories), they only shift the problem (permissivism), or they rest on a mistaken conception of the sceptical challenge (self-undermining charges). Of course, my considerations are far from conclusive, and my few remarks can hardly do justice to the richness and complexity of those topics. Still, I had to start somewhere. So, rather than putting those responses to scepticism aside without explanation, I opted for explaining to the reader why I did so.

With chapter 2, I continued the work of elucidation and motivation of underdetermination scepticism started in chapter 1. The chapter is a contribution to the debate about the structure of sceptical arguments inaugurated by Brueckner almost three decades ago. Therein, I explored the logical and dialectical relationship between the Underdetermination Argument and the Closure Argument. I argued that the Underdetermination Argument and the Closure Argument depend on each other in the following sense: the falsity of one argument's premises makes the other argument unsound. More precisely, I argued that the premises of the Underdetermination Argument entail the premises of the Closure Argument, and that the Closure Argument's premises entail one of the premises of the Underdetermination Argument but not the other. In this context, the most important result was that, on a probabilistic construal of evidential favouring and justification, the underdetermination principle and the closure principle entail each other. Further, I argued that the Underdetermination Argument and the Closure Argument are independent from each other in another sense: their premises can be plausibly motivated without an appeal to the other argument.

The upshot was a new account of the logical and dialectical relationship between the Underdetermination Argument and the Closure Argument, according to which they are distinct sources of scepticism. Even so, it turned out that, if justification is construed as sufficient degree of evidential support, the Underdetermination

Argument and the Closure Argument are closely related arguments and that they have sceptical consequences in exactly the same circumstances.

One important upshot of this chapter for those interested in scepticism is that that they should start paying more attention to the Underdetermination Argument for scepticism, a lesson that is reinforced by this thesis as a whole. Undoubtedly, of the two arguments, the Closure Argument has so far received the lion's share of epistemologists' attention. But, in light of my discussion of their relationship, this should change. Since the Underdetermination Argument and the Closure Argument are closely related, careful study of the Underdetermination Argument will repay those interested in the Closure Argument. Further, the Underdetermination Argument is interesting on its own right.

Chapter 1 and 2 focused on showing that there is something to be said for the Underdetermination Argument, both in itself as well as in comparison to the Closure Argument. But, as far as a motivation of underdetermination scepticism goes, chapter 1 and 2 focused mainly on the underdetermination principle, showing that the principle sits well with our sound judgement about a vast range of cases and that it can be derived from minimal assumptions about confirmation and justification. In contrast, chapter 3 and 4 focused on the underdetermination claim, the claim that S's evidence does not favour P over SH.

Chapter 3 discussed what I called 'the Infallibility Objection', some unsuccessful attempts to meet the objection, and put forward a new solution on behalf of the sceptic. The Infallibility Objection is the idea that the Underdetermination Argument is a bad argument because the sceptical inference from sameness of evidence to underdetermination presupposes infallibilism. More precisely, according to proponents of the Infallibility Objection, the sceptic's inference is illegitimate because sameness of evidence does not entail parity of evidential support: S's evidence can favour one of two incompatible hypotheses although both are compatible with the evidence. In that, underdetermination scepticism resembles uninteresting forms of scepticism that trade on excessively high standards for justification.

Against these critics, I argued that the contentious sceptical inference does not presuppose infallibilism. Proponents of the Infallibility Objection have failed to recognise the sound reasoning pattern that underpins the sceptical inference. Roughly, the idea

is this: the sceptic argues from sameness of evidence to underdetermination because being in the good case and being in the bad case have the same prior probability. Since S has the same evidence in the good and the bad case, the posterior probability of being in the good case and being in the bad case remains the same no matter what evidence S has. The upshot is a reconstruction of the sceptic's reasoning and a motivation of the underdetermination claim along Bayesian lines.

Chapter 4 discussed the anti-sceptical import of Williamson's anti-luminosity argument. According to Williamson, the sceptic's argument for sameness of evidence relies on the claim that S's evidence is luminous – that S is always in a position to know what her evidence is. However, per Williamson's anti-luminosity argument, no non-trivial condition is luminous, not even the condition of having evidence.

I argued that Williamson's anti-sceptical strategy fails for two reasons. First, I offered a new argument to the effect that the anti-luminosity argument does not undermine the sceptic's reasoning. The key idea for my argument was that luminosity failures of the kind highlighted by the anti-luminosity argument are irrelevant to failures of knowledge in the bad case. Second, Williamson's anti-sceptical strategy hinges on the mistaken idea that the sceptic needs a luminosity claim to defend the Sameness of Evidence Lemma. I presented two arguments against this idea. The first argument appealed to evidence mentalism, the view that evidence consists of non-factive mental states. The second argument appealed to comparative judgements about one's rationality in the good and the bad case. The upshot was that the sceptic's case for the Sameness of Evidence Lemma is stronger and more resilient to challenges than Williamson envisioned. In particular, the sceptic's case for the Sameness of Evidence Lemma is independent of Williamson's case against evidence luminosity.

The results of this chapter generalize to the charge that underdetermination scepticism fails because it asks too much of our ability to recognize the evidence. Just as the sceptic does not need to appeal to the claim that evidence is luminous to argue that S has the same evidence in the good and the bad case, she does not need to appeal to *any* claim about the cognitive accessibility of evidence. What I showed then is that endorsing non-accessibilism about evidence is not a way to resist the underdetermination argument. Of course, this is not to say that accessibilism about evidence is false or that it does not lend itself to sceptical uses.

Chapter 1 and chapter 4 also bear on the anti-sceptical import of non-mentalism about evidence. Chapter 1 argued that 'mental' evidence matters because it is the kind of thing we should expect to respond to whenever we form a belief. That this is so, is guaranteed by the fact that non-factive mental states are the causal mediators between mind and world. Chapter 4 argued that mentalism about evidence is sufficient to ground the sameness of evidence claim.

If this is right, a powerful argument against the claim that endorsing a non-mental conception of evidence is key to avoid underdetermination scepticism is readily available. The sceptic endorses a mentalist conception of evidence because she thinks of evidence as the cognitive input of our beliefs and the cognitive input of our beliefs are non-factive mental states. The externalist attempts to replace the mentalist conception of evidence with a non-mentalist conception. However, in doing so, the externalist is either denying that evidence is the cognitive input of our beliefs or that the cognitive input of our beliefs are non-factive mental states. Neither denial goes very far.

Suppose that the externalist denies that evidence is the cognitive input of our beliefs. Of course, the externalist is entitled to use the term 'evidence' however she likes. But why think that such a denial constitutes a response to the sceptic? Rather, the externalist seems to be changing topic. To see why, let us grant that evidence is not the cognitive input of our beliefs. Now, let us introduce a new term: 'shmevidence'. Shmevidence is simply that which the sceptic took 'evidence' to be: the kind of thing to which our beliefs are responsive. Now the sceptic can run a new underdetermination argument with 'shmevidence' in place of 'evidence'. The new argument is just as threatening as the old argument was, because shmevidence matters to our epistemic rationality in a way that seems to make sameness of shmevidence a threat to the justification of our beliefs. Thus, insisting on a non-mentalist conception of evidence affords no anti-sceptical leverage.

Now suppose that the externalist accepts that evidence is the cognitive input of our beliefs but denies that such input consists of non-factive mental states. The move looks like desperation: there seems to be an undeniable sense in which our beliefs are sensitive to non-factive mental states rather than to condition 'external' to us. It is the sense in which our beliefs remain the same regardless of whether we are seeing or hallucinating something. And, to reiterate, this seems a perfectly legitimate sense in which

our beliefs are responsive to non-factive mental states.

Chapter 1, 2, 3, and 4 discussed objections to the Underdetermination Argument that take issue with some specific premise, inferential step, or rationale for the argument, be it the underdetermination principle, the inference from sameness of evidence to underdetermination, or the supposed motivation for the Sameness of Evidence Lemma. Chapter 5 discussed a different kind of response to underdetermination scepticism, one that, like self-undermining charges, does not try to diagnose the mistake in the sceptic's argument: Moorean responses to the Underdetermination Argument.

According to Moorean anti-sceptics, the Underdetermination Argument cannot rationally undermine one's confidence that scepticism is false because anti-sceptical beliefs are epistemically superior to the premises of any philosophical argument to the contrary. Crucially, this is true even if one cannot identify any flaws in the Underdetermination Argument. I argued that Moorean anti-sceptics thereby make themselves vulnerable to the charge of dogmatism. After all, if one can find no flaws in the Underdetermination Argument, how can one's refusal to accept that argument be anything other than dogmatic?

I asked whether the Moorean can avoid the dogmatism charge and, building on work by Thomas Kelly and Susanna Rinard, I arrived at a negative answer. I tackled Mooreanism from three angles. First, I discussed attempts to motivate Mooreanism by appeal to an alleged difference between philosophy and science and argued, following Rinard, that philosophy, and more specifically the Underdetermination Argument, can overturn common-sense.

Second, I criticized attempts to motivate Mooreanism via general principles about belief revision and identified a general problem for any attempt to justify Mooreanism along these lines. Let F be a candidate for a feature that may make a common-sense belief immune from belief revision by the Underdetermination Argument. On the one hand, if F is not an epistemic property, a norm that favours beliefs with F over beliefs that lack F is not the right sort of norm to justify Mooreanism, because F and the corresponding norm lack epistemic import. On the other hand, if F is an epistemic property, a norm that favours beliefs with F over beliefs that lack F does not provide guidance, since a belief's having F depends on the assumption that no powerful counterargument to one's belief is available. Either way, the Moorean's refusal to engage

with the sceptic's argument is unjustified.

Finally, I discussed Kelly's attempt to vindicate Mooreanism via an argument for its methodological superiority over the methodology employed by the sceptic. Again, I found no reason to think that the Moorean and the sceptic employ different methodologies. The upshot is the Mooreanism provides no sound basis to resist the Underdetermination Argument for scepticism.

Taken together, the five chapters push towards a conception of the sceptical challenge that is at odds with some common ways of looking at the sceptical problem. It has become standard to distinguish between different kinds of anti-sceptical project. Pryor (2000) distinguishes between an ambitious and a modest anti-sceptical project. The ambitious anti-sceptical project is 'to refute the skeptic on his own terms, that is, to establish that we can justifiably believe and know such things as that there is a hand, using only premises that the skeptic allows us to use' (Pryor 2000: 517). The modest anti-sceptical project is 'to establish to *our* satisfaction that we can justifiably believe and know such things as that there is a hand' (Pryor 2000: 517; emphasis in the original). It does so by attempting to diagnose the mistake in sceptical arguments and to explain away their appeal. Crucially, 'since this modest anti-skeptical project just aims to set our own minds at ease, it's not a condition for succeeding at it that we restrict ourselves to only making assumptions that the skeptic would accept' (Pryor 2000: 518).

In a similar vein, Byrne (2004) distinguishes between two ways to play the sceptical game: expose-the-sceptic and convince-the-sceptic. When we play expose-the-sceptic, all we try to do is 'to expose one of the sceptic's premises as false or highly dubious' (Byrne 2004: 300). Instead, in convince-the-sceptic, 'our objective is to try to convince the sceptic that we have knowledge of the external world, without 'begging the question' against her' (Byrne 2004: 301).

Again, Bergmann (2021) distinguishes between a proselytizing and an autodidactic approach to address scepticism. The goal of the proselytizing approach is 'outreach to skeptics that tries to provide them with evidence or arguments that will bring them into the non-skeptical fold' (Bergmann 2021: 345-6). Instead, an autodidactic approach focuses on 'getting one's own anti-skeptical house in order, without worrying too much about forcing others, on pain of irrationality, to accept the conclusions of one's

own reflections' (Bergmann 2021: 346-7).

In other words, anti-sceptics have framed the dialectical situation as one in which it is one thing to play by the sceptic's rules and another thing to play by our own rules. If I am right, this way of describing the dialectic between sceptics and anti-sceptics is profoundly mistaken, for underdetermination scepticism can be framed as a paradox that concerns *us*, quite independently of whether *anyone else* endorses scepticism. Thus, the project of convincing the sceptic cannot be easily separated from the project of giving an answer to the sceptic that satisfies us. The sceptic lives in *us*; *we* are the sceptic.

Another important upshot is that attempts to resist the Underdetermination Argument without diagnosing the mistake in the argument – e.g. by arguing that scepticism is false or that underdetermination scepticism is rationally self-undermining – do not look very promising. The problem with such attempts is that the Underdetermination Argument exerts considerable pressure on them. Thus, the risk with such attempts is that, even if they can undermine the Underdetermination Argument to some extent, one remains undecided between underdetermination scepticism and anti-scepticism – an outcome that goes in the sceptic's favour.

The point is important because it uncovers another widespread misgiving about scepticism. I began this thesis by asking the question 'What is scepticism?' and by distinguishing between three technical uses of the term 'scepticism' in philosophy: scepticism as a family of views about our cognitive limitations, scepticism as a family of negative views about the existence of certain entities or the truth of certain classes of propositions, and scepticism as a practice aimed at the suspension of judgement. Although underdetermination scepticism fits the first definition (and, *a fortiori*, the second definition), the Underdetermination Argument for scepticism can be used in the context of a practice aimed at the suspension of judgement. In other words, it would be wrong to assume, as critics of scepticism often do, that the underdetermination sceptic endorses scepticism in the sense of believing that scepticism is true on the basis of the Underdetermination Argument.

The point is important because that mistaken view of scepticism makes the sceptical challenge easier to meet than it actually is. Although the standing of the Underdetermination Argument's premises might be too weak to generate conviction in the argument's conclusion, it might be sufficiently strong to undermine confidence that

justification is possible and to lead to judgement suspension. So, understanding the threat of scepticism as the threat of being convinced that scepticism is true is misleading because it generates a false sense of security: that the sceptical challenge is met as soon as the sceptic cannot convince one that scepticism is true.

6.2 Open Questions

I have invited the reader to read my thesis as an argument for underdetermination scepticism. If I am right, underdetermination scepticism follows from a cogent argument, and some important objections against this argument fail. Further, the threat of scepticism is magnified by our meta-epistemological situation, a situation in which no rational agreement has been reached so far on the correct solution of the sceptical paradox. This is the kind of environment in which the Underdetermination Argument for scepticism can thrive. Still, the truth of underdetermination scepticism and the success of the Underdetermination Argument are, in light of my thesis, just a conjecture.

For this conjecture to become something more than just a conjecture, more must be achieved than I could achieve here. So, proponents of non-evidentialist accounts of justification will no doubt feel that my case for underdetermination scepticism can exert no pressure on them. They will be inclined to read my thesis as a *reductio* of evidentialist accounts of justification. Similarly, evidentialists who appeal to explanatory virtues (Vogel 1990; McCain 2014, 2019; Douven 2022) will feel that my case is incomplete since I had very little to say about abductive responses to underdetermination scepticism – although a lot has already been said by others about the shortcomings of this approach (Rinard 2017; Bergmann 2021). The same can be said about phenomenal conservatists (Pryor 2000; Huemer 2001) and entitlement theorists (Wright 2004). At this stage, the best I can do is to recognise these limitations in my work, while the best I can hope for is to address these issues in future work.

If my conjecture is true, we have no knowledge or justification – our cognitive lives are suffused with arbitrariness and ignorance. This leaves an outstanding question: how can we live such cognitively impoverished lives? It has been a commonplace of anti-scepticism since ancient times – and possibly one of the main reasons for opposition to scepticism since – that the sceptic is doomed to inaction. Roughly, the idea is that sceptics cannot act because they suspend judgement on matters on which one is

required to form an opinion in order to act. Thus, when sceptics act, they are susceptible to the charge of incoherence (Burnyeat 1983). A related worry is that scepticism robs us of the resources to evaluate actions. For whether an action is praiseworthy or not often depends on features of the relevant situation on which the sceptic should suspend judgement. Accordingly, an important component of the project of defending underdetermination scepticism is to explain how we can live with its truth, how we can be sceptics while avoiding practical and intellectual paralysis.

To conclude, I would like to point out what seems to me a promising suggestion to address this problem due to Rinard (2022): pragmatic scepticism. According to Rinard (2022), the pragmatic sceptic agrees with the underdetermination sceptic that justification is impossible because she accepts the Underdetermination Argument for scepticism. However, the pragmatic sceptic does not believe that underdetermination scepticism leads to suspending judgement on P. This is because she believes that there are *pragmatic* reasons for belief. In particular, she believes that we should retain our ordinary belief that P in response to the Underdetermination Argument, because we are better off doing so. Moreover, the pragmatic sceptic sees no reason to change the way we talk about evidence, justification, knowledge or other terms of epistemic appraisal in our day-to-day dealings: although we do not have good evidence for our beliefs and they neither are justified nor constitute knowledge, we have good *practical* reasons to talk as if things stood otherwise. So, endorsing pragmatic scepticism is a way to endorse underdetermination scepticism while avoiding practical and intellectual paralysis.

It is not without irony that, at the end of a thesis on scepticism, I should feel *confident* that the way to scepticism is a long and arduous way. And although I cannot see where it ends, at least for now it does not look like a dead-end. At any rate, it seems worth travelling on to see where it leads.

Bibliography

Alspector-Kelly, M. 2019. *Against Knowledge Closure*. Cambridge: Cambridge University Press.

Alston, W. P. 1988. 'An Internalist Externalism.' *Synthese* 74 (3): 265-83.

Annas, J. and Barnes, J. 1985. *The Modes of Scepticism*. Cambridge: Cambridge University Press.

Armstrong, D. 1983. *What is a Law of Nature?* Cambridge: Cambridge University Press.

Austin, J. L. 1946. 'Other Minds.' *Proceedings of the Aristotelian Society, Supplementary Volumes* 20 (1): 148-87.

Bach, Kent. 2008. 'Applying Pragmatics to Epistemology.' *Philosophical Issues* 18 (1): 68-88.

Bergmann, M. 2021. *Radical Skepticism and Epistemic Intuition*. New York: Oxford University Press.

Berker, S. 2008. 'Luminosity Regained.' *Philosophers' Imprint* 8 (2): 1-22.

Bertrand, J. 1889. *Calcul des probabilités*. Paris: Gauthier-Villars et Fils.

Blackson, T. 2007. 'On Williamson's Argument for (I_i) in His Anti-Luminosity Argument.' *Philosophy and Phenomenological Research* 74 (2): 397-405.

Boult, C. 2013. 'Epistemic Principles and Sceptical Arguments: Closure and Underdetermination.' *Philosophia* 41 (4): 1125-33.

Briesen, J. 2010. 'Reconsidering Closure, Underdetermination, and Fallibilism.' *Grazer Philosophische Studien* 80 (1): 221-34.

Brown, J. 2018. *Fallibilism: Evidence and Knowledge*. Oxford: Oxford University Press.

Brueckner, A. 1994. 'The Structure of the Skeptical Argument.' *Philosophy and Phenomenological Research* 54 (4): 827-35.

_____ 2005. 'Fallibilism, Underdetermination, and Skepticism.' *Philosophy and Phenomenological Research* 71 (2): 384-91.

_____ 2011. '~K~SK.' *Philosophical Issues* 21 (1): 74-89.

Brueckner, A. and Fiocco, M. O. 2002. 'Williamson's Anti-Luminosity Argument.' *Philosophical Studies* 110 (3): 285-93.

Burnyeat, M. F. 1983. 'Can the Skeptic Live His Skepticism?' In Burnyeat, M. F. (ed.) 1983. *The Skeptical Tradition*. Berkeley and Los Angeles: University of California Press.

Byrne, A. 2004. 'How Hard Are the Sceptical Paradoxes?' *Noûs* 38 (2): 299-325.

Cameron, R. 2010. 'How to Have a Radically Minimal Ontology.' *Philosophical Studies* 151 (2): 249-64.

Carter, J. A. 2012. 'Recent Work on Moore's Proof.' *International Journal for the Study of Skepticism* 2 (2): 115-44.

Cavell, S. 1979. *The Claim of Reason: Wittgeinstein, Skepticism, Morality, and Tragedy*. New York: Oxford University Press.

Chisholm, R. 1977. *Theory of Knowledge*. Englewood Cliffs, NJ: Prentice Hall.

_____ 1982. *The Foundations of Knowing*. Minneapolis: University of Minnesota Press.

Christensen, D. 2009. 'Disagreement and Evidence: The Epistemology of Controversy.' *Philosophy Compass* 4 (5): 756-67.

Cohen, S. 1984. 'Justification and Truth.' *Philosophical Studies* 46 (3): 279-95.

_____ 1998. 'Two Kinds of Skeptical Argument.' *Philosophy and Phenomenological Research* 58 (1): 143-59.

_____ 2010. 'Luminosity, Reliability, and the Sorites.' *Philosophy and Phenomenological Research* 81 (3): 718-30.

Cohen, S. and Lehrer, K. 1983. 'Justification, Truth, and Coherence.' *Synthese* 55 (2): 191-207.

Comesaña, J. 2005. 'Unsafe Knowledge.' *Synthese* 146 (3): 395-404.

Comesaña, J. and Klein, P. 2019. 'Skepticism.' *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition). Available at: https://plato.stanford.edu/archives/win2019/entries/skepticism/.

Conee, E. 2001. 'Comments on Bill Lycan's 'Moore against the New Skeptics.' *Philosophical Studies* 103 (1): 55-9.

_____ 2005. 'The Comforts of Home.' *Philosophy and Phenomenological Research* 70 (2): 444-51.

Conee, E. and Feldman, R. 2001. 'Internalism Defended.' *American Philosophical Quarterly* 38 (1): 1-18.

_____ 2004. *Evidentialism: Essays in Epistemology*. New York: Oxford University Press.

Davidson, D. 1986. 'A Coherence Theory of Truth and Knowledge.' In LePore, E. (ed.) 1986. *Truth and Interpretation. Perspectives on the Philosophy of Donald Davidson*. Cambridge: Blackwell.

DeRose, K. 1995. 'Solving the Skeptical Problem.' *The Philosophical Review* 104 (1): 1-52.

Descartes, R. 1996 [1641]. *Meditations on First Philosophy: with Selections from the Objections and the Replies* [transl. Cottingham]. Cambridge: Cambridge University Press.

Dodd, D. 2007. 'Why Williamson Should Be a Sceptic.' *The Philosophical Quarterly* 57 (229): 635-49.

_____ 2012. 'Evidentialism and Skeptical Arguments.' *Synthese*, 189 (2): 337-52.

_____ 2014. 'How to Motivate Skepticism.' In Dodd, D. and Zardini, E. (eds.) 2014. *Scepticism and Perceptual Justification*. Oxford: Oxford University Press.

Douven, I. 2022. *The Art of Abduction*. Cambridge, MA: MIT Press.

Dretske, F. 1970. 'Epistemic Operators.' *Journal of Philosophy* 67 (24): 1007-23.

Duhem, P. 1914. *La Théorie Physique. Son Objet et sa Structure* (2nd edition). Paris: Librairie Marcel Rivière.

Duncan, M. 2018. 'What It's Like to Have a Cognitive Home.' *European Journal of Philosophy* 26 (1): 66-81.

Dutant, J. (ed.) forth. *The New Evil Demon: New Essays on Knowledge, Justification, and Rationality*. Oxford University Press.

Eva, B. 2019. 'Principles of Indifference.' *Journal of Philosophy* 116 (7): 390-411.

Fantl, J. 2018. *The Limitations of the Open Mind*. Oxford, New York: Oxford University Press.

Feldman, R. 1999. 'Contextualism and Skepticism.' *Philosophical Perspectives* 13: 91-114.

_____ 2003. *Epistemology*. Upper Saddle River, NJ: Prentice Hall.

_____ 2006. 'Epistemological Puzzles About Disagreement.' In Hetherington, S. (ed.) 2006. *Epistemology Futures*. New York: Oxford University Press.

_____ 2007. 'Reasonable Religious Disagreements.' In Antony, L. (ed.) 2007. *Philosophers Without Gods*. Oxford: Oxford University Press.

Feldman, R. and Conee, E. 1985. 'Evidentialism.' *Philosophical Studies* 48 (1): 15-34.

Fine, K. 2001. 'The Question of Realism.' *Philosophers' Imprint* 1 (1): 1-30.

Firth, R. 1978. 'Are Epistemic Concepts Reducible to Ethical Concepts?' In Goldman. A. and Kim, J. (eds.) 1978. *Values and Morals: Essays in Honor of William Frankena, Charles Stevenson, and Richard Brandt*. Dordrecht: Kluwer Academic Publishers.

Frances, B. 2018. 'Scepticism and Disagreement.' In Machuca, D. E. and Reed, B. (eds.) 2019. *Skepticism: From Antiquity to the Present*. London, New York: Bloomsbury Academic.

Freitag, W. 2013. *I Know: Modal Epistemology and Skepticism*. Münster: mentis Verlag.

Fumerton, R. A. 1995. *Metaepistemology and Skepticism*. Lanham, Md.: Rowman & Littlefield.

Gettier, E. L. 1963. 'Is Justified True Belief Knowledge?' *Analysis* 23 (6): 121-23.

Goodman, N. 1953. *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.

Greco, J. 2008. 'Skepticism about the External World.' In Greco (ed.) 2008. *The Oxford Handbook of Skepticism*. New York: Oxford University Press.

Grundmann, T. 2019. 'How to Respond Rationally to Peer Disagreement: The Preemption View.' *Philosophical Issues* 29 (1): 129-42.

Gupta, A. 2006. *Empiricism and Experience*. Oxford: Oxford University Press.

Hájek, A. 2008. 'Are Miracles Chimerical?' In Kvanvig, J. (ed.) 2008. *Oxford Studies in Philosophy of Religion, Volume 1*. Oxford: Oxford University Press.

Harman, G. 1973. *Thought*. Princeton, NJ: Princeton University Press.

_____ 2003. 'Skepticism and Foundations.' In Luper-Foy (ed.) 2003. *The Skeptics: Contemporary Essays*. Aldershot: Ashgate Press.

Hawthorne, J. and Landes, J. and Wallmann, C. and Williamson, J. 2015. 'The Principal Principle Implies the Principle of Indifference.' *The British Journal for the Philosophy of Science* 68 (1): 123-31.

Hawthorne, J. and Logins, A. 2021. 'Graded Epistemic Justification.' *Philosophical Studies* 178 (6): 1845-58.

Huemer, M. 2001. *Skepticism and the Veil of Perception*. Lanham, MD: Rowman & Littlefield.

Jaynes, E.T. 1957. 'Information theory and statistical mechanics.' *The Physical Review* 106 (4): 620-30.

Jenkins, D. 2021. 'Luminosity in the Stream of Consciousness.' *Synthese* 198 (1): 1549-62.

Johnson King, Z. 2022. 'Radical Internalism.' *Philosophical Issues* (Online Version of Record): 1-19.

Kelly, T. 2005. 'Moorean Facts and Belief Revision, or Can the Skeptic Win?' *Philosophical Perspectives* 19 (1): 179-209.

_____ 2008a. 'Evidence: Fundamental Concepts and the Phenomenal Conception.' *Philosophy Compass* 3 (5): 933-55.

_____ 2008b. 'Common Sense as Evidence: Against Revisionary Ontology and Skepticism.' *Midwest Studies in Philosophy* 32 (1): 53-78.

_____ 2010. 'Peer Disagreement and Higher-Order Evidence.' In Goldman, A. I. and Whitcomb, D. (eds.) 2010. S*ocial Epistemology: Essential Readings*. Oxford: Oxford University Press.

_____ 2016. 'Evidence.' *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition). Available at: https://plato.stanford.edu/archives/win2016/entries/evidence/.

Kopec, M. and Titelbaum, M. 2016. 'The Uniqueness Thesis.' *Philosophy Compass* 11 (4): 189-200.

Kornblith, H. 2000. 'The Contextualist Evasion of Epistemology.' *Philosophical Issues* 10: 24-32.

Kraft, T. 2012. 'Scepticism, Infallibilism, Fallibilism.' *Discipline filosofiche* 22 (2): 49-70.

Kripke, S. 2011. 'On Two Paradoxes of Knowledge.' In Kripke, S. 2011. *Philosophical Troubles: Collected Papers Vol 1*. New York: Oxford University Press.

Leite, A. 2010. 'Fallibilism.' In Dancy J., Sosa, E. and Steup, M. (eds.) 2010. *A Companion to Epistemology* (2nd edition). Malden, Oxford: Wiley-Blackwell.

Leitgeb, H. 2002. 'Review of Timothy Williamson, *Knowledge and Its Limits*.' *Grazer Philosophische Studien* 65 (1): 207-17.

Lemos, N. 2004. *Common Sense: A Contemporary Defense*. Cambridge: Cambridge University Press.

Lewis, D. 1996. 'Elusive Knowledge.' *Australasian Journal of Philosophy* 74 (4): 549-67.

Levy, N. 2015. 'Moral Responsibility Skepticism: Meeting McKenna's Challenge.' In Buckareff, A., Moya, C. and Rosell, S. (eds.) 2015. *Agency, Freedom, and Moral Responsibility*. Basingstoke (Houndmills), New York: Palgrave Macmillan.

Littlejohn, C. M. 2012. 'Evidence and Armchair Access.' *Synthese* 179 (3): 479-500.

Luper, S. 2020. 'Epistemic Closure.' T*he Stanford Encyclopedia of Philosophy* (Summer 2020 Edition). Available at: https://plato.stanford.edu/archives/sum2020/entries/closure-epistemic/.

Lycan, W. G. 2001. 'Moore against the New Skeptics.' *Philosophical Studies* 103 (1): 35-53.

Madison, B. J. C. 2010. 'Epistemic Internalism.' *Philosophy Compass* 5 (10): 840-53.

Markosian, N. 1998. 'Brutal Composition.' *Philosophical Studies* 92 (3): 211-249.

McCain, K. 2013. 'Two skeptical arguments or only one?.' *Philosophical Studies* 164 (2): 289-300

_____ 2014. *Evidentialism and Epistemic Justification*. New York: Routledge.

_____ 2019. 'In Defense of the Explanationist Response to Skepticism.' *International Journal for the Study of Skepticism* 9 (1):38-50.

McDowell, J. 1982. 'Criteria, defeasibility, and knowledge.' *Proceedings of the British Academy* 68: 455-79.

Merricks, T. 2000. 'No Statues.' *Australasian Journal of Philosophy* 78 (1): 47-52.

_____ 2001. *Objects and Persons*. Oxford: Oxford University Press.

Moore, G. E. 1942. 'A Reply to My Critics.' In Schilpp, P. A. (ed.) 1942. *The Philosophy of G.E. Moore.* Chicago: Northwestern University Press.

_____ 1959. *Philosophical Papers*. New York: Routledge.

Neta, R. and Rohrbaugh, G. 2004. 'Luminosity and the Safety of Knowledge.' *Pacific Philosophical Quarterly* 85 (4): 396-406.

Nozick, R 1981. *Philosophical Explanations*. Cambridge, Mass.: Harvard University Press.

Okasha, S. 2003. 'Scepticism and Its Sources.' *Philosophy and Phenomenological Research* 68 (3): 610-32.

_____ (ms). 'Scepticism, Evidential Holism and the Logic of Demonic Deception.' Available at: http://philsci-archive.pitt.edu/20869/1/Logic%20of%20Demonic%20Deception.pdf.

Pappas, G. 2017. 'Internalist vs. Externalist Conceptions of Epistemic Justification.' *The Stanford Encyclopedia of Philosophy* (Fall 2017 Edition). Available at: https://plato.stanford.edu/archives/fall2017/entries/justep-intext/.

Pettigrew, R. 2014. 'Accuracy, Risk and the Principle of Indifference.' *Philosophy and Phenomenological Research* 92 (1): 35-59.

Pollock, J. and Cruz, J. 1999. *Contemporary Theories of Knowledge*. Oxford: Rowman and Littlefield.

Pritchard, D. 2002. 'Recent Work on Radical Skepticism'. *American Philosophical Quarterly* 39 (3): 215-57.

_____ 2005. 'The Structure of Sceptical Arguments.' *The Philosophical Quarterly* 55 (218): 37-52.

_____ 2012. *Epistemological Disjunctivism*. Oxford: Oxford University Press.

_____ 2015. *Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing*. Princeton, NJ: Princeton University Press.

Pryor, J. 2000. 'The Skeptic and the Dogmatist.' *Noûs* 34 (4): 517-49.

_____ 2001. 'Highlights of Recent Epistemology.' *The British Journal for the Philosophy of Science* 52 (1): 95-124.

——— 2004. 'What's Wrong with Moore's Argument?'. *Philosophical Issues* 14 (1): 349-78.

Putnam, H. 1981. *Reason, Truth, and History*. Cambridge: Cambridge University Press.

Ramachandran, M. 2009. 'Anti-Luminosity: Four Unsuccessful Strategies.' *Australasian Journal of Philosophy* 87(4): 659-73.

Rawls, John 1972. *A Theory of Justice*. Cambridge, MA: Belnkap.

Rea, Michael. 1998. 'In Defence of Mereological Universalism.' *Philosophy and Phenomenological Research* 58 (2): 347-60.

Reed, B. 2012. 'Fallibilism.' *Philosophy Compass* 7 (9): 585-96.

Reid, T. 2012 [1818]. *An Inquiry into the Human Mind on the Principles of Common Sense*. New York: Cambridge University Press.

Rinard, S. 2013. 'Why Philosophy Can Overturn Common Sense.' *Oxford Studies in Epistemology* 4: 185-213.

_____ 2014. 'The Principle of Indifference and Imprecise Probability.' *Thought* 3 (2): 110-14.

_____ 2017. 'External World Scepticism and Inference to the Best Explanation.' In McCain, K. and Poston, T. (eds.) 2017. *Best Explanations: New Essays on Inference to the Best Explanation*. New *York*: Oxford University Press.

_____ 2018. 'Reasoning One's Way Out of Skepticism.' In McCain and Poston (eds.) 2018. *The Mystery of Skepticism: New Explorations*. Leiden: Brill.

_____ 2022. 'Pragmatic Skepticism.' *Philosophy and Phenomenological Research* 104 (2): 434-53.

Rosenkranz, S. 2007. 'Agnosticism as a Third Stance.' *Mind* 116 (461): 55-104.

_____ 2021. *Justification as Ignorance*. New York: Oxford University Press.

Roush, S. 2005. *Tracking Truth. Knowledge, Evidence, and Science*. New York: Oxford University Press.

_____ 2010. 'Closure on Skepticism.' *Journal of Philosophy* 107 (5): 243-56.

Russell, B. 1912. *The Problems of Philosophy*. New York: Barnes & Noble.

Schellenberg, S. 2013. 'Experience and Evidence.' *Mind* 122 (487): 699-747.

Schönbaumsfeld, G. 2016. *The Illusion of Doubt*. Oxford: Oxford University Press.

Schoenfield, M. 2015. 'Internalism Without Luminosity.' *Philosophical Issues* 25 (1): 252-72.

Sextus Empiricus 1990 [2nd century CE]. *Outlines of Pyrrhonism* [transl. Bury]. Lanham, Md.: Prometheus Books.

Shatz, D. 1987. 'Nozick's Conception of Skepticism.' In Luper-Foy, S. (ed.) 1987. *The Possibility of Knowledge: Nozick and His Critics*. Lanham, Md.: Rowman and Littlefield.

Sider, T. 2013. 'Against Parthood.' In Bennett, K. and Zimmerman, D. W. (eds.) 2013. *Oxford Studies in Metaphysics*, *Volume 8*. Oxford: Oxford University Press.

Silva Jr., P. and Oliveira, L. R. G. forth. 'Propositional Justification and Doxastic Justification.' In Lasonen-Aarnio, M. and Littlejohn, C. (eds.) forth. *The Routledge Handbook of the Philosophy of Evidence*. Routledge.

Sinnott-Armstrong, W. 2006. *Moral Skepticism*. New York: Oxford University Press.

Sklar, L. 1975. 'Methodological Conservatism.' *Philosophical Review* 84 (3): 384-400.

_____ 1981. 'Do Unborn Hypotheses Have Rights?' *Pacific Philosophical Quarterly* 62 (1): 17-29.

Smith, M. 2022. 'Underdetermination and Closure: Thoughts on Two Sceptical Arguments.' In Jope, M. and Pritchard, D. (eds.) 2022. *New Perspectives on Epistemic Closure*. New York: Routledge.

Smith, P. 1984. 'Could We Be Brains in a Vat?' *Canadian Journal of Philosophy* 14 (1): 115-23.

Smithies, D. 2019. *The Epistemic Role of Consciousness*. New York: Oxford University Press.

Srinivasan, A. 2015. 'Are We Luminous?' *Philosophy and Phenomenological Research* 90 (2): 294-319.

_____ 2020. 'Radical Externalism.' *The Philosophical Review* 129 (3): 395-431.

Stanford, P. K. 2001. 'Refusing the Devil's Bargain: What Kind of Underdetermination Should We Take Seriously?' *Philosophy of Science* 68 (3): S1-S12.

_____ 2006. *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. New York: Oxford University Press.

_____ 2017. 'Underdetermination of Scientific Theory.' *The Stanford Encyclopedia of Philosophy* (Winter 2017 Edition). Available at: https://plato.stanford.edu/archives/win2017/entries/scientific-underdetermination/.

Strawson, P. F. 1985. *Skepticism and Naturalism: Some Varieties*. New York: Columbia University Press.

Stroud, B. 1984. *The Significance of Philosophical Scepticism*. New York: Oxford University Press.

Sutton, J. 2007. *Without Justification*. Cambridge, MA: MIT Press.

Thorpe, J. R. 2018. 'Closure Scepticism and the Vat Argument.' *Mind* 127 (507): 667-90.

Turri, J. 2010. 'On the Relationship Between Propositional and Doxastic Justification.' *Philosophy and Phenomenological Research* 80 (2): 312-26.

Unger, P. 1975. *Ignorance: A Case for Scepticism*. New York: Oxford University Press.

_____ 1979. 'There Are No Ordinary Things.' *Synthese* 41 (2): 117-54.

Vahid, H. 2005. *Epistemic Justification and the Skeptical Challenge*. New York: Palgrave Macmillan.

Van Fraassen, B.C. 1980. *The Scientific Image*. New York: Oxford University Press.

_____ 1989. *Laws and Symmetry*. New York: Oxford University Press.

Van Inwagen, P. 1990. *Material Beings*. Ithaca, NY: Cornell University Press.

Vanrie, W. 2020. 'Some Problems with the Anti-Luminosity Argument.' *Pacific Philosophical Quarterly* 101 (3): 538-559.

Vogel, J. 1990. 'Cartesian Skepticism and Inference to the Best Explanation.' *The Journal of Philosophy* 87 (11): 658-66.

_____ 2004. 'Skeptical Arguments.' *Philosophical Issues* 14 (1): 426-55.

_____ 2010. 'Luminosity and Indiscriminability.' *Philosophical Perspectives* 24 (1): 547-72.

Watson, G. 1977. 'Skepticism about Weakness of Will.' *The Philosophical Review* 86 (3): 316-39.

Weatherson, B. 2004. 'Luminous Margins.' *Australasian Journal of Philosophy* 82 (3): 373-83.

_____ 2007. 'The Bayesian and the Dogmatist.' *Proceedings of the Aristotelian Society* 107 (1): 169-85.

_____ (ms). 'Scepticism.' Available at: http://brian.weatherson.org/SkepticismNotes.pdf.

Wedgwood, R. 2002. 'Internalism Explained.' *Philosophy and Phenomenological Research* 65 (2): 349-69.

White, R. 2005. 'Epistemic Permissivism.' *Philosophical Perspectives* 19 (1): 445-59.

_____ 2009. 'Evidential Symmetry and Mushy Credence.' *Oxford Studies in Epistemology* 3: 161-86.

Williams, M. 1991. *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*. Oxford: Blackwell.

Williamson, J. 2007. 'Motivating Objective Bayesianism: From Empirical Constraints to Objective Probabilities.' In Harper, W. and Wheeler, G. (eds.) 2007. *Probability and Inference: Essays in Honor of Henry E. Kyburg Jr.* London: College Publications.

Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.

——— 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.

_____ 2013. 'Gettier Cases in Epistemic Logic.' *Inquiry* 56 (1): 1-14.

Wittgenstein, L. 1969. *On Certainty*. New York/London: Harper Torchbooks.

Wong, W. 2008. 'What Williamson's Anti-Luminosity Argument Really Is.' *Pacific Philosophical Quarterly* 89 (4): 536-43.

Wright, C. 1985. 'Facts and Certainty.' *Proceedings of the British Academy* 71: 429-72.

_____ 1991. 'Scepticism and Dreaming: Imploding the Demon'. *Mind* 100 (1): 87-116.

_____ 1992. *Truth and Objectivity*. Cambridge, MA: Harvard University Press.

_____ 2002. '(Anti-)Sceptics Simple and Subtle: G.E. Moore and John McDowell.' *Philosophy and Phenomenological Research* 65 (2): 330-48.

_____ 2004. 'Warrant for Nothing (and Foundations for Free)?' *Aristotelian Society Supplementary Volume* 78 (1): 167-212.

_____ 2007. 'The Perils of Dogmatism.' In Nuccetelli, S. and Seay, G. (eds.) 2007. *Themes from G. E. Moore: New Essays in Epistemology*. Oxford: Oxford University Press.

Yalçin, Ü. 1992. 'Skeptical Arguments from Underdetermination.' *Philosophical Studies* 68 (1): 1-34.