



THE UNIVERSITY *of* EDINBURGH

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

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

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

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

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

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



THE UNIVERSITY
of EDINBURGH

Evidential Externalism

Giada Fratantonio

Doctor of Philosophy
University of Edinburgh
2018

Abstract

It is widely accepted, amongst epistemologists, that evidence plays an important role in our epistemic life. Crucially, there is no agreement on what evidence is. Following Silins, we can cash out the disagreement around the notion of evidence in terms of the opposition between Evidential Internalism and Evidential Externalism (Silins, 2005). Evidential internalists claim that evidence supervenes on one's non-factive mental states, such as, beliefs, impressions (BonJour, 1999, Audi, 2001). Evidential Externalists deny that. In this Thesis, first, I contrastively assess the plausibility of two prominent contemporary externalist theories: Duncan Pritchard's *Epistemological Disjunctivism*, the thesis on which one's evidence in perceptual cases is truth-entailing and reflectively accessible (Pritchard, 2012), and Timothy Williamson's $E=K$, the thesis on which one's evidence is all and only the propositions one knows (Williamson, 2000). Second, I develop a novel externalist account of evidence that I call *Ecumenical Evidentialism*. I show how Ecumenical Evidentialism is able to bring together some of the benefits of both Pritchard's Disjunctivism and Williamson's $E=K$. This Thesis is structured into three sections, each of which addresses the following three questions respectively:

- Does the Access Problem represent a real threat to Evidential Externalism?
- Is Evidential Externalism committed to a sceptical variety of Infallibilism?
- How does Evidential Externalism understand the relation between evidence and epistemic justification?

I argue that neither Epistemological Disjunctivism nor $E=K$ are fully satisfying Externalist accounts of evidence. On one hand, I argue that Disjunctivism captures the orthodox intuition on which justification is a matter of being evidence-responsive, but it does so on pain of facing the so-called Access Problem. On the other hand, by rejecting any strong accessibility thesis, Williamson's $E=K$ is better positioned to resist both the Access Problem as well as the Infallibility Problem, but it does not vindicate the orthodox intuition on which justification is a matter of being evidence-responsive. Finally, I show that, while retaining the main commitments of Williamson's theory of evidence, such as, $E=K$, my Ecumenical Evidentialism is able to capture the orthodox *responsiveness* intuition about epistemic justification.

Declaration

I declare that this thesis was composed by myself and that the work contained therein is my own, except where explicitly stated otherwise in the text. (*Giada Fratantonio*)

Table of Contents

Introduction	1
I Evidential Externalism and The Access Problem	15
1 The Access Problem for E=K	19
1.1 Introduction	19
1.2 The Access Problem	20
1.3 What is Armchair Knowledge?	23
1.3.1 On Silins' Notion of Armchair Knowledge	23
1.3.2 A Clarification	25
1.4 A First Response: A Restriction of Armchair Access	26
1.5 A Second Response: On the Nature of Armchair Knowledge	29
1.5.1 Experience, Counterfactuals, and Imagination	32
1.5.2 Second-order Knowledge Through Imagination	35
1.5.3 Why Knowing the Environment from the Armchair is not Absurd	37
1.6 Easy Second-order Knowledge	38
1.7 Conclusion	40
2 The Access Problem for Epistemological Disjunctivism	43
2.1 Introduction	43
2.2 Epistemological Disjunctivism and The Access Problem	45
2.3 Pritchard's Response to The Access Problem	47
2.4 On the Metaphysical and Folk Interpretation of Essential Dependency	49
2.5 On the Epistemic Interpretation of Essential Dependency	51
2.5.1 Why The Epistemic Interpretation is Problematic	53
2.5.2 Restricting Closure	56
2.6 An Alternative Diagnosis: Tim Kraft's Source/Content Distinction . . .	60

2.7	Conclusion	65
II Evidential Externalism and The Infallibility Problem		67
3	E=K and the Infallibility Problem	71
3.1	Introduction	71
3.2	The Argument from E=K to Scepticism	72
3.3	What is Infallibilism?	74
3.3.1	Varieties of Infallibilism (and Underdetermination)	77
3.3.2	Comments on Varieties of Infallibilism	81
3.3.3	E=K: what Infallibilism? What Underdetermination?	83
3.4	What is Underdetermination?	87
3.4.1	Varieties of Underdetermination (and Infallibilism)	90
3.4.2	Comments on Varieties of Underdetermination	92
3.4.3	E=K: What Underdetermination? What Infallibilism?	95
3.5	Back to Dodd and Why He is Mistaken	95
3.5.1	Why we know propositions about the future and propositions about the (unperceived) present although they are underdeter- mined!	102
3.6	Conclusion	105
III Evidential Externalism and Epistemic Justification		109
4	Evidence-first, Knowledge-first, and Ecumenical Evidentialism	113
4.1	Introduction	113
4.2	Evidence-first Evidentialism	114
4.2.1	Propositional and Doxastic Justification	115
4.2.2	The Relation between Roles of Evidence and Justification	120
4.3	Knowledge-first Evidentialism	123
4.3.1	Knowledge-first Evidentialism: on Propositional Justification	124
4.3.2	Knowledge-first Evidentialism: on Doxastic Justification	126
4.3.3	How to make sense of Knowledge-first Evidentialism	128
4.4	Evidence-first VS Knowledge-first Evidentialism	129
4.5	Ecumenical Evidentialism	137
4.5.1	A two-tiered account of knowledge	137

4.5.2	Meta-epistemological Pluralism	142
4.5.3	Non-traditional Foundationalism	146
4.6	Conclusion	149
	Conclusion	156
	References	157

Acknowledgements

In 2013 I moved to London. I was officially enrolled at University of Milan, but I was visiting KCL as an Erasmus student. I decided to take a course called “General Philosophy” which included, amongst other things, a series of lectures on epistemology with Clayton Littlejohn. It is in that occasion that I had my first encounter with contemporary epistemology. Although it took me a while to understand why everyone seemed to be so obsessed with red barns (red barns are not really a thing in Italy as far as I know), I soon realised that many of the questions epistemologists were asking were questions I cared about. The summer after that year in London, I wanted to read more about contemporary epistemology, so I checked what book philosophers took to be influential in the contemporary epistemological debate. That summer I bought and read copy of *Knowledge and Its Limits*. As you can imagine, it was a pretty intense summer! In a sense, this thesis wouldn’t have happened if I hadn’t gone to London, and if I hadn’t gone to Littlejohn’s lectures.

I have written most of this Thesis while in Edinburgh and, although it hasn’t always been a very smooth process, I was very lucky to meet great philosophers who had an incredibly positive impact on my PhD experience.

I’d like to thank everyone who at some point came by Office 2.16 at DSB. In no particular order, I would like to thank Kegan Shaw, Michel Croce, Matt Jope, Joe Dewurst, Lukas Schwengerer, Rie Iizuka, Hadeel Naeem, and Carlos Portales. Thanks also to Guy Fletcher, Michela Massimi, Debbie Roberts, Ellie Mason, Lani Watson, Joey Pollock, Alistair Isaac, Pauline Phemister, Alix Cohen, Natalie Ashton, and super Fiona! A special thank goes to Di Yang, for her constant encouragement and for the long, intense but wonderful philosophical discussions on epistemology, law, and vegan cheese. I am confident she is going to make the world a better place.

Thanks to my little “Edinburgh family”: Rosa Erica, Di Yang, Dan Calder, Elli Shackelton, Silvan Wittwer, Ana-Maria Cretu, Anna Ortin, Nick Rebol, Iliana Lytra and George Tentokallis. Each of them have contributed in their own way to make my

PhD life colourful and “homey”, and I am very grateful for that.

Over the last four years I presented most of the material in this thesis to various conferences across Europe and UK. This was made possible by the generous Research Grants that I have been awarded by PPLS of Edinburgh and that I would thus like to thank. I would also like to thank all the philosophers I have met in these occasions and that have provided me with very useful feedback on my work: Jie Gao, Davide Fassio, Rachel Fraser, Adam Carter, Giorgio Volpe, Luca Moretti, Tommaso Piazza, Simon Wimmer, Claire Field, Liz Jackson, Mona Simion, Clotilde Calabi, Giacomo Melis, Josh Habgood-Coote, Jaakko Hirvela, Lewis Ross, Julien Dutant, Cameron Boulton, Federica Malfatti. Thanks to Gloria Andrada DeGregorio for being an amazing pal of silence.

I would like to thank Lucy Campbell and Alex Greenberg for their advice and for always cheering for me. Thanks to Lucy for all the great philosophical discussion at the Head of the River in Oxford, and for making me discover gently bubbling Pimm’s. Thanks also to Olga and Nikola.

I would also like to thank The Phoenix Souls and Gospel Choir of Edinburgh. I was very lucky to sing with a bunch of wonderful music lovers. It was definitely beneficial for my mental and physical health and I couldn’t have imagined a better way of keeping my music passion alive while doing my PhD. In particular, thanks Adam for being such a cheerful musician, and Kat for being a great leader and for finding us those amazing gigs. Singing together at the Usher Hall of Edinburgh was definitely an amazing adventure!

Thanks also to Nik Felbab and Cathleen O’ Gradaigh who have been a real source of inspiration over the last years. Thanks to Mette, Jonas and Hella for all the walks in Leith and Portobello and the great time spent together.

Thanks to my wonderful Emma Syea, brilliant philosopher and great friend. I am sure Nietzsche would be proud of you.

I would also like to thank my Italian friends: Cappe, Cate, Marta, Livia, Nene, Lori. Thanks to my favourite astrophysicist of all, Francesco Coti Zelati, and for the long chats about science and evidence which would always make me realise how much there is to know and how much I do not understand. Thanks to Giovanna Marino for welcoming back every time I am back to Pescate. Thanks to Nicola Vassena, Blessio, Uil, and Roberto Meloni. Thanks to Luca Conca for allowing me to stay in Tuscany during Summer 2017. Thanks to the one and only Michele Dozio. The thought of our longstanding friendship always makes me very happy.

Thanks to Teresa, Massimo, Ale, Fausto, Nini, Donatella, Minuzzo, Nonna Gina, Lucia, Marco, and Simona Lugli.

Thanks to my highschool philosophy teacher, Maria Luisa Montagna. Her fascinating lessons surely played an important role in my decision to become a philosopher in the first place!

Thanks to Queen Bee and Nina Simone for being the most incredible soundtrack of many days of my PhD.

A special thanks goes to Timothy Williamson for supervising me during my time in Oxford as a visiting PhD student. I would like to thank him for all the useful advice and feedback on my work, and for the constant support and encouragement over the last year. Needless to say, this thesis wouldn't exist if I hadn't read his *Knowledge and Its Limits*, that Summer back from London. I would also like to thank the PPLS Edinburgh for the generous research grant towards my staying in Oxford.

I am very grateful to my supervisors Aidan McGlynn, Duncan Pritchard, and Martin Smith. Thanks to Aidan McGlynn for teaching me the importance of the details and the importance of being charitable to the philosophers we read. Thanks to Duncan Pritchard for always reminding me to look at the broad picture. Thanks to Martin Smith for his incredible support, for letting me audit his Philosophy of Law lectures, and for always encouraging the creative aspect of my philosophical writing. I have learnt a lot from each of you.

Of course, there is absolutely no way I could have written this thesis without the constant support of my family: my mum Minou, my dad Gioacchino, and my two brothers Yanick and Dominique (I know – what a bunch of names, right?). There are no words to explain how grateful I am for all your neverending love. In a sense, for me philosophy starts at home, with you all. It is at home, with you, that I have learnt the importance of critical and rational thinking, the importance of asking questions, and the value of discussions. It is at home, with you, that I have learnt that we should not be sad for our limits and our lack of knowledge, but we should accept them as an opportunity to learn. Thank you for teaching me that failure is part of life and it is just as important as success. Thanks for always reminding me what the very important things are.

Finally, I would like to thank Alberto. I couldn't have imagined that moving to Edinburgh would bring me such an amazing and sensitive soul. Your unconditional passion for knowledge and understanding makes you a philosopher – although you don't know that yet. Thank you for being so close to me, for being my best friend, and

for all the love you show me every day. I can't wait to have many more adventures together.

Introduction

Evidence plays a central role in our epistemic life. Crucially, however, there is no agreement, amongst epistemologists, on how to understand the notion of evidence. Following Silins (2005), we can cash out the disagreement around the notion of evidence in terms of Evidential Internalism and Evidential Externalism. Let Evidential Internalism be the thesis on which one's evidence supervenes on one's non-factive mental states (roughly put, those mental states that do not necessarily entail anything about the external environment, e.g., beliefs, impressions, sensations). To say that evidence *supervenes* on non-factive mental states amounts to saying that if there is a change in one's evidence, then there is a change in one's non-factive mental states as well. Let Evidential Externalism be the denial of Evidential Internalism.

Traditionally, epistemologists have been Evidential Internalists. For instance, great empiricists of the twentieth century like Ayer and Russell used to embrace an account of evidence that we can call, following Williamson's terminology, the "phenomenological conception of evidence" (Williamson, 2000: 173).¹ On this view, evidence is reducible to sense data that the subject can be immediately conscious of (cf. Ayer, 1936). But the view is also popular amongst contemporary epistemologists.² Audi, for instance, seems to embrace an account of evidence on which evidence is reducible to non-factive mental states, such as impressions, and seeming:

"[E]vidence is [...] ultimately internal. It centrally involves my sensory states, memory impressions, inferences, and the like." (Audi, 2001: 47)

Similarly, Bonjour seems to have in mind something like Evidential Internalism when defending his Foundationalist picture of justification (assuming justification is a function of one's evidence):

¹For a great overview of the historical reasons and motivations behind the popularity of this phenomenal conception of evidence see Kelly, 2008.

²One might be tempted to classify Conee and Feldman's 'mentalism' as an instance of Evidential Internalism. However, although they notoriously define their view as a version of epistemic internalism, they actually do not restrict the supervenience thesis between evidence and the mental to non-factive mental states (cf. Conee & Feldman, 2004: 57).

“[T]he obvious and [...] correct thing to say is that basic or foundational beliefs are justified by appeal to experience.” (BonJour, 1999: 230)

Recently, however, there has been a new trend in the literature, and Evidential Externalism has been gaining more and more consensus amongst epistemologists³. While Evidential Internalists have claimed that whatever evidence one has is fixed by the non-factive mental states one is in, Evidential Externalists deny that. Instead, they maintain that features of the external environment are going to be relevant in determining what evidence one has.

To see how Evidential Internalism differs from Evidential Externalism in some important respect, Silins invites us to consider Gary and Barry. Gary and Barry are “internal twins”, namely, they share all non-factive mental states. For example, they both have an experience as of two bananas in a bowl, and they both believe that there are two bananas in a bowl. However, Gary and Barry differ in some very important way:

“Although Gary is now and then mistaken about some matter of fact, Barry is a radically deceived brain in a vat, as deceived as can be given that he has the same non-factive mental states as Gary. Gary is in the good case; Barry is in the bad case.” (Silins, 2005: 375)

We can ask: do Gary and Barry have the same evidence? Evidential Internalists will answer this question affirmatively. Evidential Externalists will answer this question negatively.

While epistemologists have dedicated plenty of attention to the traditional Internalism-Externalism debate over epistemic justification, the Evidential Internalism-Evidential Externalism debate is still widely overlooked. And yet, there are good reasons why, as epistemologists, we should care about this debate.

The first reason is meta-epistemological, and it concerns the fact that, as defined, Evidential Internalism and Evidential Externalism are relatively neutral theses. Let me explain. Defining a theory of evidence has never been an easy task, for there are at least three questions that a theory of evidence has to answer: i) what *is* evidence?, ii) what does it mean for one to *possess* evidence?, and iii) what does it mean for one’s evidence to *support* one’s beliefs or hypotheses? Although Evidential Externalism and, most of all, Evidential Internalism put a restriction on how we can answer these questions (particularly i) and ii)), they do not restrict the answer in any specific way. That is, they

³Examples of Evidential Externalists are McDowell, 1995, Williamson, 2000, Pritchard, 2012, Littlejohn, 2012, Littlejohn, 2018b, but also Alston, 1988 and Mitova, 2015.

do not force us to choose a specific account of evidence over the other. In this sense, Evidential Internalism and Evidential Externalism, so defined, are relatively silent on the above three important issues. Framing the debate around evidence in terms of Evidential Internalism and Evidential Externalism thus provides us with a framework that allows us to classify different theories of evidence into two broad categories: those that embrace the supervenience thesis between evidence and non-factive mental states, and those that reject it. This will allow for varieties of Evidential Externalist theses.

Second, it is worth emphasising that Evidential Internalism and Evidential Externalism as defined above do not entail anything about epistemic justification. As I will show in a moment, the traditional Internalism-Externalism discussion around epistemic justification, and the more recent one on Evidential Internalism and Evidential Externalism, are thus two different debates. However, as I will stress in Part III, I think it is plausible to take Evidential Internalism and Evidential Externalism as two ways in which one can embrace *Evidentialism*, the thesis on which justification supervenes on evidence (Conee & Feldman, 2004). Once we do so, then focusing on the Evidential Internalism-Evidential Externalism debate rather than on the more traditional Internalism-Externalism debate over epistemic justification might have some advantages. As Silins points out (Silins, 2005: 378), “an advantage of [Evidential Externalism] is that it does not entail reliabilist views” namely, the externalist view on which, one is justified in believing that p only if, roughly put, one’s belief that p is the result of some reliable belief-forming process⁴. The reason why Silins takes this to be an advantage of Evidential Externalism is that it avoids many of the problems notoriously associated with Process Reliabilism, such as, the so-called Generality Problem⁵. The Generality Problem is the problem of defining which type of process is the relevant type of reliable process that matters for assessing the justificatory status of a belief. By avoiding to talk about reliability of cognitive processes, Evidential Externalism does not face the same problem. At the same time, however, it looks like an Evidential Externalist can still account for the overall post-Gettier intuition that has motivated traditional externalist views of justification, namely, the idea that there has to be a connection between the truth of one’s belief and one’s justification for such a belief. As noted in Fratantonio and McGlynn (Fratantonio & McGlynn, 2018: 85), we can similarly expect Evidential Internalism to have some advantages over more tra-

⁴The locus classicus for Process Reliabilism is A. I. Goldman, 1979, and A. I. Goldman, 1986

⁵For a discussion on the Generality problem see A. I. Goldman, 1979, Feldman, 1985, and Conee & Feldman, 1998.

ditional versions of internalism, such as, Accessibilism. According to Accessibilism, if S is justified in believing that p , then S has access to whatever makes S's belief that p justified (call it the *justifier*)⁶. What does it mean for S to have *access* to the justifier of S's belief that p ? On one interpretation, it means to have knowledge by reflection alone of the relevant justifier (J). But if knowledge requires justification, then this further knowledge of whatever makes S's belief that p justified will also have to be the result of S having access to whatever makes S's belief that [J is the justifier of S's belief that p] justified, thereby leading to an infinite regress. Call this the Regress Problem. Evidential Internalism does not entail Accessibilism, for it does not pose any accessibility requirement. Therefore, Evidential Internalism seems to avoid some of the problems that are notoriously associated with Accessibilism, such as, the Regress Problem I have just described.^{7,8}

Third, and finally, Evidential Internalism seems to play an important role in some sceptical argument (Williamson, 2000: ch.8, Silins, 2005: 375-76). In a nutshell, the worry seems to be that if, as Evidential Internalism entails, Gary has just the same evidence as Barry's, who's radically deceived and therefore does not have knowledge, then how can Gary have knowledge on the basis of such evidence? One might thus want to explore whether there's some plausible Evidential Externalist view available, for that would mean having an account of evidence that is not vulnerable to the sceptical challenge.⁹

Epistemological Disjunctivism and E=K as varieties of Evidential Externalism

As the title of this Thesis suggests, my focus will be on Evidential Externalism. In particular, I will consider two prominent contemporary varieties of Evidential Externalism, namely, Duncan Pritchard's *Epistemological Disjunctivism* (Pritchard, 2012) and Timothy Williamson's *E=K* (Williamson, 2000), with a particular focus given to the latter. Before providing an overview of the chapters of this Thesis, however, let me briefly introduce Pritchard's Epistemological Disjunctivism and Williamson's E=K,

⁶For a traditional defence of Accessibilism see Ginet, 1975 and Chisholm, 1988.

⁷For a recent discussion on the Regress Problem for Accessibilism see Bergmann, 2006

⁸For an excellent survey of the Internalism-Externalism debate about epistemic justification see Littlejohn, 2012, Ch. 1.

⁹In his 'Deception and Evidence' Silins argues that, contrary to first appearances, Evidential Externalism does not have any strong advantage when it comes to the sceptical challenge (Silins, 2005: 395-96). I will not address this problem in this Thesis. However, see Fratantonio and McGlynn for a response to Silins' argument (Fratantonio & McGlynn, 2018: 97).

and let me clarify why we can classify these theories as instances of Evidential Externalism. By doing this, I will also be able to shed more light on what Evidential Externalism is, what it entails, and what, instead, it is silent on.

Let us start with Duncan Pritchard's view. In his book *Epistemological Disjunctivism* he writes:

“In paradigmatic cases of perceptual knowledge, an agent, S, has perceptual knowledge that p in virtue of being in possession of rational support, R [i.e., S's seeing that], for her belief that p which is both factive¹⁰ (i.e., R's obtaining entails p) and reflectively accessible.”

(Pritchard, 2012: 13).

Epistemological Disjunctivism focuses on cases of perceptual knowledge and it predicts that in epistemically favourable cases, one has *entailing* evidence for propositions concerning the external environment¹¹. It is thus easy to see why Epistemological Disjunctivism will classify as an instance of Evidential Externalism: it will say that Gary and Barry, despite sharing the same non-factive mental states, have different evidence. For only Gary, who's in the paradigmatic good case, has truth-entailing evidence that there are two bananas in the bowl (i.e., his *seeing that* there are two bananas entails that there are two bananas). On Epistemological Disjunctivism, factors concerning the external environment matter for determining what evidence one has.

In order to be an Evidential Externalist, however, one does not need to commit to the thesis that evidence is *entailing*, namely, that evidence entails the proposition it is evidence for (cf. Fratantonio & McGlynn, 2018: 85). Considering Williamson's E=K will help us seeing why this is so. On Williamson's E=K, one's evidence is constituted by all and only the propositions one knows¹². However, E=K does not entail that if one has evidence for a proposition p , then necessarily one's evidence entails p . All E=K predicts is that if one knows that e , then e is part of one's evidence and vice versa. That is, if a proposition e is part of one's evidence, then one knows that e and thus e

¹⁰Note that talking about 'factive' evidence can be confusing. For, in the literature, it often refers to the idea that evidence is constituted of true propositions. However, when Pritchard talks about 'factive' rational support he refers to the idea that rational support/evidence is truth-entailing. In this Thesis I will avoid talking about 'factive' evidence unless I specify what I mean by it.

¹¹As it appears from the reported quotation, it is worth noticing that Pritchard does not formulate his view in terms of evidence. However, I think we can have a version of Epistemological Disjunctivism in terms of evidence.

¹²Notoriously, John Hyman also defends E=K, although his view differs from Williamson's one in various respects. In particular, his defence of E=K is grounded on a notion of knowledge as "the ability to be guided" by evidence. I will not consider Hyman's version of E=K in this thesis. See Hyman, 1999 and Hyman, 2006.

is true. Of course, as I will stress in Chapter 3, one consequence of $E=K$ is that any proposition one knows is trivially entailed by one's total evidence. However, it is not a consequence of $E=K$ that, for every proposition p , if e is evidence *for* p then e entails p . It is compatible with $E=K$ that one's evidence e is evidence for a proposition p (e.g., by making p sufficiently likely) without entailing p . Nevertheless, Williamson's $E=K$ is a paradigmatic instance of Evidential Externalism: it will say that Gary and Barry, despite sharing the same non-factive mental states, have different evidence. For only Gary, who's in the paradigmatic good case, knows that there are two bananas in the bowl. On $E=K$, how things are in the external environment matters for determining what evidence one has.

However, Evidential Externalism is not even equivalent to the thesis on which evidence is constituted by only true propositions. One can have a view on which evidence is constituted by only true propositions and still be an Evidential Internalist. For instance, one can maintain that one's evidence is all and only true propositions about how things appear to one. This view will not classify as an instance of Evidential Externalism: it will say that Gary and Barry, who share the same non-factive mental states, have the same evidence. For things will appear exactly the same to both Gary and Barry and so the proposition that there seem to be two bananas in the bowl (and all propositions about how things look like) will constitute both Gary's and Barry's evidence. On this view, how things actually are in one's external environment will not matter for determining what evidence one has. But requiring one's evidence to be constituted by true propositions only is not even necessary for someone to be an Evidential Externalist. Think of someone who's sympathetic with Goldman's reliabilism (A. I. Goldman, 2009). As Fratantonio and McGlynn point out, one can defend a theory on which one's evidence is constituted by all and only the propositions one is justified in believing, while saying that justification depends on the reliability of one's belief-forming process (2018: 85). On this view, Gary might still have some false beliefs amongst his evidence for believing that there are two bananas in the bowl, but he would still end up having different evidence than Barry. For only Gary's belief would be the result of a reliable belief-forming process. Only Gary, given he's in a paradigmatic good case, would be justified in believing that there are two bananas in the bowl, and hence, Gary's "box of evidence" will at least include one instance of evidence that Barry lacks: the proposition that there are two bananas in the bowl.

Finally, note that Evidential Externalism does not entail any specific account of what evidence is made of. Once again, appealing to both Pritchard's Epistemological

Disjunctivism and Williamson's E=K will help see why. On one hand, on Williamson's E=K, one's evidence is constituted by all and only the *propositions* one knows. On the other hand, a natural evidentialist interpretation of Pritchard's Epistemological Disjunctivism takes evidence to be constituted by one's perceptual experience.

The moral of the story is the following. By merely denying that one's evidence supervenes on one's non-factive mental states, Evidential Externalism seems to be relatively neutral on important questions concerning the nature of evidence, the evidential-support relation, and the question of evidence possession, thereby allowing for varieties of Evidential Externalist theories. As we will see in this Thesis, Epistemological Disjunctivism and E=K provide us with very different answers to these questions, and, nevertheless, they can both be thought of as two instances of Evidential Externalism.

Outline of the Thesis

Now that we have clarified why and to what extent Pritchard's Disjunctivism and Williamson's E=K are instances of Evidential Externalism, let me briefly provide an overview of this Thesis. This Thesis is divided into three big parts, each of which addresses the following three questions respectively:

1. Does the so-called Access Problem represent a challenge for Evidential Externalism?
2. Is Evidential Externalism committed to some sceptical variety of Infallibilism?
3. How does Evidential Externalism understand the relation between evidence and epistemic justification?

Part I. *Evidential Externalism and The Access Problem*

In Part I I consider the so-called Access Problem for Evidential Externalism. In particular, I consider the Access Problem as it arises for Williamson's E=K and Pritchard's Epistemological Disjunctivism in Chapter 1 and Chapter 2 respectively. I argue that, while the Access Problem represents a real challenge for Pritchard's Epistemological Disjunctivism, it does not constitute a serious threat to Williamson's E=K.

Part II. *Evidential Externalism and The Infallibility Problem*

In Part II I evaluate whether Evidential Externalism entails some sceptical form of Infallibilism. I refer to this as The Infallibility Problem. More precisely, I evaluate whether Evidential Externalism is committed to some form of infallibilism that is nevertheless compatible with the plausible thesis that almost all our beliefs are underdetermined by our evidence. My focus in Part II will be mainly on Williamson's E=K. This is because, as I explain in the Introduction of Part II, even if Pritchard's Epistemological Disjunctivism is committed to some version of infallibilism, this will not have any important sceptical implications insofar as his view is restricted to paradigmatic cases of *perceptual* knowledge. Williamson's E=K, instead, is not restricted to cases of perceptual knowledge, so here we should take the Infallibility Problem seriously. Fortunately for the defender of E=K, I believe there is nothing to worry about. By developing varieties of Infallibilism and Underdetermination Theses, I argue, in Chapter 3, that The Infallibility Problem does not represent a threat for E=K.

Part III *Evidential Externalism and Epistemic Justification*

In Part III I evaluate how Evidential Externalism understands the relation between evidence and justification. As I mentioned above, Evidential Externalism is silent on what we should say about epistemic justification. However, it is plausible to take Evidential Externalism as one way of embracing *Evidentialism*, roughly put, the thesis that takes one's justification to be a function of one's evidence. If Epistemological Disjunctivism and E=K represent two different ways of being an *Evidential Externalist*, then, assuming Evidentialism, we should expect these views to provide us with two different ways in which one can be an Evidentialist in an externalist way. Pritchard's Epistemological Disjunctivism exemplifies what I call *Evidence-first Evidentialism*, while E=K underpins what I call *Knowledge-first Evidentialism*. In Chapter 4, first I shed light on what Evidence-first and Knowledge-first Evidentialism amount to, and I clarify to what extent these views are different. Second, I define a novel externalist variety of Evidentialism, what I call *Ecumenical Evidentialism*. I show that, while being an instance of Evidential Externalism, my Ecumenical Evidentialism brings together (what I take to be) the best features of both Evidence-first and Knowledge-first Evidentialism.

Conclusion

The main conclusion of the Thesis is the following: neither Williamson's account of evidence and justification, nor Epistemological Disjunctivism are fully satisfying. However, Ecumenical Evidentialism seems to provide us with a novel promising way of being an Evidential Externalist.

A Methodological Remark

Before diving into Chapter 1, I want to make a methodological remark concerning the target object of this Thesis. As I mentioned above, most of my attention will be devoted to Williamson's E=K. However, it is important to note that Williamson's E=K is to be conceived as part of his broader "Knowledge-first" project. While traditionally epistemologists have been concerned with an analytic program of defining knowledge in terms of more basic components, e.g., justification, belief, and truth, Williamson's Knowledge-first Epistemology notoriously rejects this project, while taking knowledge to be a more fundamental and unanalysable notion. Part of the motivation behind this Knowledge-first shift in the debate derives from the impasse that the epistemological discussion seems to have reached after Gettier's famous paper (Gettier, 1963), (allegedly) showing that justification, truth, and belief are notion jointly sufficient conditions for knowledge. In particular, Gettier's examples aim to show that one can have justified true belief by luck, and yet, luck seems to be incompatible with the attainment of knowledge. Since 1963, much of the post-Gettier literature has thus been concerned with fixing the traditional tripartite account of knowledge as justified true belief in a way that resists the problem raised by Gettier. Some epistemologists have added a fourth condition, others have opted for a different notion of justification than the one underlying the traditional tripartite account, others have just dropped the justification condition. Unfortunately, sooner or later, any new account that has been suggested had to face objections and counterexamples. In a sense, Knowledge-first Epistemology, despite being extremely controversial, seems to offer a way out of the impasse, by rejecting the analytic project altogether and its assumption that knowledge is to be factorised into more basic components. Instead, it reverses the order of explanation and it uses the notion of knowledge to define and understand other epistemic notions, e.g., belief, justification, assertion, action, as well as evidence.¹³ The Knowledge-first project thus has a wide application. Therefore, the following disclaimers are in order.

¹³For a great critical analysis of the Knowledge-first Project see McGlynn (2014).

- In Part I and Part II of this Thesis, my focus will be almost exclusively on Williamson's Knowledge-first account of *evidence*. I will make some brief remarks about the general Knowledge-first picture in Part III, when discussing how it is possible to be Knowledge-first and embrace Evidentialism about epistemic justification.
- The chapters that follow do not require any previous familiarity with Williamson's Knowledge-first Epistemology.
- This thesis does not aim to provide an overarching assessment of all the objections that have been raised against Evidential Externalism. Instead, it is meant to provide an assessment of two of the most prominent contemporary Evidential Externalist theories with respect to how they relate to a cluster of problems and topics that have traditionally been relevant in the epistemological discussion: the nature of reflective knowledge, the problem of infallible knowledge, and the nature of epistemic justification.

Part I

Evidential Externalism and The Access Problem

Introduction

As I have pointed out in the Introduction of this Thesis, there are different varieties of Evidential Externalism. However, there is one thing every version of Evidential Externalism agrees on. That is, by denying that evidence supervenes on non-factive mental states, every Evidential Externalist theory will claim that what evidence one has depends on features of the external environment. Here's one worry one might have. Assume that if Evidential Externalism is true, we can know by reflection that Evidential Externalism is true. Given that it is plausible to also assume that we can have reflective knowledge of the evidence we have, and given that if Evidential Externalism is true, then our evidence is sensitive to the external environment, then it seems we can come to know by reflection and deductive reasoning alone features about the external environment. Crucially, these kinds of "worldly" propositions do not seem to be something we can know by reflection alone! Call this the Access Problem. In Part I, I consider the Access Problem for Evidential Externalism.

Part I is constituted of two Chapters, and an overall Conclusion. I first consider the Access Problem for E=K in Chapter 1. I then address a very similar Access Problem that arises for Epistemological Disjunctivism in Chapter 2. Given that, as mentioned above, Epistemological Disjunctivism and E=K are very different varieties of Evidential Externalism, the Access Problems that seem to threaten these views will be slightly different. Therefore, in each Chapter, I will clarify the specific kind of Access Problem that each view seems to face. The Conclusion of Part I shows the overall lesson that we can draw in the light of what I argue in Chapter 1 and Chapter 2: while the Access Problem represents a challenge for Pritchard's Epistemological Disjunctivism, it does not constitute a serious threat to Williamson's E=K.

Chapter 1

The Access Problem for E=K

1.1 Introduction

In his *Deception and Evidence*, Silins claims that there is an argument showing that “Evidential Externalism is false because it has unacceptable consequences about what one knows from the armchair” (Silins, 2005: 380). More precisely, Silins’ target is a specific variety of Evidential Externalism, namely, Timothy Williamson’s E=K, which equates one’s evidence to one’s knowledge (Williamson, 2000). Following Silins, I will here focus on Williamson’s E=K. I will go back to assess the Access Problem for Epistemological Disjunctivism in the next Chapter.

Silins’ argument aims to show that, given a (allegedly) plausible thesis about what we can know from the armchair, E=K leads to the unacceptable conclusion that we can have armchair knowledge of specific empirical propositions. Call this The Access Problem¹. In this Chapter, I argue that Silins’ Access Problem does not represent a challenge for Williamson’s E=K, for it relies on an account of armchair knowledge that Williamson should reject. First, I reconstruct Silins’ argument (Section 1.2), and I shed light on the nature of its main assumption (Section 1.3). Second, I develop two lines of response, both of which put pressure on Silins’ Armchair Access thesis. While the first response focuses on the scope of the thesis (Section 1.4), the second response focuses on its nature (Section 1.5). The second line of response is the most interesting one. This is because it represents the framework within which I develop a novel account of second-order knowledge through imagination. After resisting a possible objection to my argument (Section 1.6), I point out that these two lines of response are mutually supportive and jointly compatible. I thus conclude that the Access Prob-

¹Silins actually calls this problem the ‘Armchair Access Problem’.

lem does not represent a threat for Williamson's variety of Evidential Externalism, i.e., E=K (Section 1.7).

1.2 The Access Problem

The variety of access argument that represents the focus of this Chapter is one which does not rest on the following Luminosity thesis:

[Luminosity]: For any proposition p , if one is suitably alert and conceptually sophisticated, then one is in a position to know whether or not one's evidence includes p . (Silins, 2005: 380)

The Access Problem I will address here is one that arises even if we accept Williamson's thesis that evidence is not a luminous condition (Williamson, 2000: ch. 4). That is, even if we grant that it is not the case that if E is part of one's evidence, then one is always in a position to know that E is part of one's evidence. Silins is well aware of Williamson's anti-luminosity argument². In fact, he explicitly takes his Access Problem to be a "novel access argument which uses a very different thesis, one which is not touched by Williamson's argument against [luminosity]" (Silins, 2005: 380). As Silins points out, while the luminosity thesis represents a very strong claim about what one can have access to, this novel argument relies, as we will see shortly, on a weaker and more plausible claim, one which is compatible with the limited nature of one's knowledge of one's evidence. Let us now see how the argument goes.

Let us imagine an agent, Gary, who is in a room staring at a dial. Gary sees that the dial reads 0.4. Let us assume Gary is in a perceptually good case³. By seeing that the dial reads 0.4, he thereby comes to know that the dial reads 0.4. Crucially, if E=K is true, then the proposition that the dial reads 0.4 belongs to Gary's evidence. Let us now assume that Gary considers what evidence he has. The Access Problem against E=K runs – schematically – as follows:

1. Gary has armchair knowledge that his evidence includes the proposition that the dial reads 0.4.

²The target of Williamson's anti-luminosity argument, as presented in *Knowledge and Its Limits*, is all non-trivial conditions, where a condition is trivial when either it always obtains, or it never does. However, I will here narrow my focus on the non-trivial conditions of knowledge and evidence.

³I take a perceptually good case to be a paradigmatic good case in which the subject S has a veridical experience, her perceptual capacities are reliable and fully functioning, no luck is involved, etc. That is, one in which S obtains perceptual knowledge.

2. Gary has armchair knowledge that, if his evidence includes the proposition that the dial reads 0.4, then the dial reads 0.4.
3. Therefore, Gary is in a position to have armchair knowledge that the dial reads 0.4

(Silins, 2005: 382).

The foregoing argument is grounded on the following thesis about what we can know from the armchair:

(Armchair Access): It is sometimes the case that: one's evidence includes some proposition E, and one knows from the armchair that one's evidence includes E. (Silins, 2005: 381-2)

In a nutshell, Silins' argument is that, given that it is sometimes possible to know, merely from the armchair, that a proposition p is part of our evidence, then, given factivity of evidence and a plausible closure principle for armchair knowledge, we can competently deduce, and thereby know, what this propositional evidence entails. Given $E=K$, if p is part of my evidence, then I know that p . Therefore, assuming factivity of knowledge, if p is part of my evidence, then p is true. It follows that if I can know from the armchair that a proposition p is part of my evidence, then I can competently deduce, thereby coming to know merely from the armchair, that p is true⁴. But, Silins says, although "Gary can know what the dial reads, he cannot know such a proposition through armchair reflection" (Silins, 2005: 382). It is worth noticing that Silins develops this argument as a parallel to the more notorious McKinsey Paradox, concerning the incompatibility between Privileged Access and Semantic Externalism (McKinsey, 1991). More precisely, in his [1991] McKinsey argues that the following three claims are jointly incompatible, assuming reflective knowledge is closed under entailment⁵:

(Content Externalism) The content of our thoughts is determined by factors external to the agent.

(Privileged Access) It is possible to have reflective knowledge⁶ of one's own mental content.

⁴Note that not only is the thesis that if p is part of your evidence then p is true can be easily derived from $E=K$ and factivity of knowledge, but a priori philosophical arguments have also been provided in favour of this thesis. (see Littlejohn, 2012: ch. 3)

⁵A similar argument is also offered by Boghossian (Boghossian, 1998).

⁶Terminology is not very helpful here, for there are various notions used in this debate: a priori, reflective, and armchair knowledge. It is worth pointing out that McKinsey does not talk about reflective knowledge: he talks about *a priori* knowledge of one's mental content. However, the notion of a priori

(Knowledge of External World) It is not possible to have reflective knowledge of the external world.

If Content Externalism and Privileged Access Thesis are both true, then, given a plausible closure principle for reflective knowledge, and assuming that it is possible to have reflective knowledge of Content Externalism, it follows that one could come to know, merely by reflection and deductive reasoning alone, that a certain empirical proposition is true. Crucially, this contradicts Knowledge of External World, which is taken to be indisputably true⁷. Let RK stand for the epistemic operator for reflective knowledge. From Content Externalism, Privileged Access, and a plausible closure principle for reflective knowledge we can thus get the following notorious McKinsey argument⁸:

(PR1) RK [I have the concept 'water']

(PR2) RK [If I have the concept 'water', then water exists]

(CONC) RK [water exists]

Again, (CONC) looks problematic, for the existence of water seems to be something we can know only empirically⁹.

In what follows, I argue that Silins' argument does not undermine Williamson's E=K. In order to do so, first, I evaluate and shed light on Silins' notion of armchair knowledge as underlying the Armchair Access thesis. Second, I argue that Gary's knowledge of his evidence does not constitute an instance of armchair knowledge as defined by Silins.

knowledge in McKinsey's papers (and, more in general, in the post-McKinsey literature) is thought of as including introspection (contrary to what happens in the a priori/a posteriori debate, where introspective knowledge is generally considered to be empirical knowledge). See Nuccetelli, 1999, Brown, 2004, and Farkas, 2008 for relevant discussion. Silins mainly talks about about 'armchair' knowledge. Here I will use 'armchair knowledge' and 'reflective knowledge' interchangeably. I will specify later how and to what extent the notions of a priori knowledge and armchair/reflective knowledge are intertwined.

⁷Although most philosophers have taken this consequence to be plainly unacceptable, Sarah Sawyer notoriously resists the McKinsey paradox by biting the bullet: Content Externalism and Privileged Access Thesis jointly lead to the conclusion that we can know by reflection alone specific empirical propositions. However, she argues, this is not an unacceptable result. See Sawyer, 1998.

⁸See Putnam, 1975, McKinsey, 1991, and Boghossian, 1998

⁹For a collection of articles on the McKinsey problem for Content and Semantic Externalism see Nuccetelli, 2003.

1.3 What is Armchair Knowledge?

1.3.1 On Silins' Notion of Armchair Knowledge

To fully understand Silins' Access Problem, we need to clarify the notion of armchair knowledge he employs. But to begin with, I would like to point out that Silins' introduction of the first premise is too abrupt. Silins introduces the argument by stating that "Gary has armchair knowledge that his evidence includes the proposition that the dial reads 0.4" (Silins, 2005: 382). Crucially, given $E=K$, to know what propositions belong to one's evidence is, first of all, to know what one's knowledge is. That is, given $E=K$, knowledge of one's evidence is nothing more than knowledge of one's knowledge. If we then ask what is the nature of Gary's knowledge that the dial reads 0.4, the answer will be straightforward. Gary's first-order knowledge is an instance of perceptual knowledge. Gary sees that the dial reads 0.4, thereby coming to know that the dial reads 0.4. But if we ask what is the nature of Gary's knowledge that he knows, the answer seems to be less straightforward. The following worry arises: is Gary's knowledge of his evidence, and thus Gary's knowledge of his knowledge, an instance of armchair knowledge, as stated in premise 1.? In order to answer this question, it is worthwhile to shed light on what it means, according to Silins' formulation, to have armchair knowledge of a proposition. Silins defines *armchair knowledge* as follows:

"One has armchair knowledge of a proposition when one knows it, and one's justification for believing the proposition does not constitutively depend on one's having had any particular experience or sense experience" (Silins, 2005: 380)

The foregoing definition is of little help unless we get a clear idea of what the notion of *constitutive dependency* amounts to. According to Silins:

"[...] one can have armchair knowledge of a proposition, even if a background condition for having that knowledge is that one has had a certain experience or type of experience. For example, one might have armchair knowledge that redness is a colour, even if one knows that proposition only if one has had experiences of redness" (Silins, 2005: 381)

In order to have a better understanding of what notion of armchair knowledge Silins has in mind, let us consider the distinction, often used in the a priori/a posteriori debate, between the *enabling* role of experience and the *evidential* role of experience. The enabling role of experience is usually taken to be what enables us to understand the concepts involved in the target proposition. The latter form of experience is what

is sometimes required, besides the enabling experience, in order for one to know a proposition. As Williamson says:

“[e]xperience is held to play an evidential role in our perceptual knowledge that it is sunny, but a merely enabling role in our knowledge that if it is sunny then it is sunny: we needed it only to acquire the concept sunny in the first place” (Williamson, 2013: 293).

Silins’ words in the above-mentioned quotation seem to suggest that, even if there is a sense in which one requires a certain experience (e.g., having perceived redness, having interacted with a linguistic community) in order to know a proposition (redness is a colour), this instance of propositional knowledge should be considered as being *armchair knowledge* as long as the required experience is merely *enabling* one to understand the concept involved (redness). If what Silins has in mind is that armchair knowledge is compatible with having had enabling experience, then to say that armchair knowledge does not constitutively depend on experience is to say that it does not depend on evidential-justificatory experience. That is, Silins seems to take armchair knowledge to be based on enabling experience only¹⁰. If it is possible to have such armchair knowledge, then, according to Silins, we should expect Gary’s knowledge of his evidence to be a perfect candidate for armchair knowledge: “we should expect that Gary knows from the armchair that his evidence includes the proposition that the dial reads 0.4” (Silins, 2005: 381).

As mentioned above, Silins is well aware of Williamson’s anti-luminosity argument. However, even if we are not *always* in a position to know from the armchair what our evidence (and thus our knowledge) is, it is plausible to claim that we are sometimes able to know from the armchair what knowledge and what evidence we have. The Armchair Access thesis seems compatible with the limits of our knowledge suggested by the anti-luminosity argument. One would thus expect that everyone – including Williamson – would be on board with that. In what follows, I argue that Silins’ argument does not represent a challenge to Williamson’s E=K, insofar as it relies on a very questionable notion of armchair knowledge. More precisely, I will put forward two lines of response to Silins’ argument. On the first line of response, I argue that, given Silins’ formulation of Armchair Access, there is a plausible restriction the externalist would build into Armchair Access thesis, one that prevents the argument from getting going. On the second line of response, I argue that, even if we grant that there

¹⁰Although Silins does not explicitly use this terminology, the text I’ve quoted seems to suggest, as I’ve just pointed out, that he has this distinction in mind.

is a sense in which Williamson should acknowledge the possibility of one's having knowledge "from the armchair" of one's evidence, this knowledge is not a problematic instance of armchair knowledge as understood by Silins. In order to do so, I take Williamson's cognitive-based account of imagination as a framework for understanding what's happening in Gary's scenario, thereby suggesting an imagination-based account of second-order knowledge. As I will show in the conclusion of this Chapter, these two lines of response are compatible and mutually supportive.

1.3.2 A Clarification

Before developing my responses, a clarification is in order. One can point out that there is a straightforward way of resisting the Access Problem. That is, one could argue that the conclusion 3. does not follow from the premises: Gary does not know from the armchair that the dial reads 0.4¹¹. The reason is simply that Gary already knew that the dial reads 0.4 by looking at the dial (by assumption). The Access Problem would thus not be a genuine problem, as Gary has an *empirical* basis for believing that the dial reads 0.4¹². However, I believe this is not an effective response for at least two reasons. First, note that the validity of the Access Problem as presented by Silins, depends on the truth of a plausible closure principle for armchair knowledge. Anyone who resists the argument by claiming that the conclusion 3. fails to follow from the premises 1. and 2., thus also has to reject closure for armchair knowledge, while providing a satisfying explanation of why Gary's empirical basis in 1. prevents the entailment from 1. to 3.. Second, I believe this reply fails to capture the real worry underpinning the Access Problem. Everyone – including Silins – would grant that Gary originally had empirical knowledge of the target proposition. However, it still seems problematic to claim that he can reach knowledge of the same empirical proposition *from the armchair*. As Silins says:

“[I] know that I am having a certain sense experience only if I am having the sense experience, but that does not show that I lack armchair knowledge that I am having the experience. The source of my knowledge is arguably still reflection instead of experience” (Silins, 2005: 381).¹³

¹¹Thanks to Tommaso Piazza for raising this objection to me in conversation.

¹²Pritchard, 2012 provides a similar response as a solution to an analogous access problem that seems to threaten his Epistemological Disjunctivism. However, in the next Chapter, I will show that this kind of solution is problematic for Pritchard's view.

¹³Note that someone could further claim that what Silins is saying here is not even possible. That is, one could object that, in order for the source of one's knowledge to be reflection instead of experience, one's belief in the target proposition must be completely based on the non-empirical reasoning only,

Analogously to the original McKinsey challenge, I believe we can appreciate the (*prima facie*) problematic nature of the Access Problem once we state it as a paradox. We start from a very plausible claim about what we can know from the armchair (the Armchair Access thesis), and from a plausible distinction between enabling and evidential experience, and we reach a very counterintuitive conclusion, namely, that we can gain knowledge of empirical propositions that is grounded on enabling experience only. Even if one could (rightly) highlight Gary's original empirical knowledge, I still believe that we need to say something more in order to account for the above-mentioned paradox. What follows should thus be conceived in the light of the paradoxical structure under which the Access Problem can be presented. More importantly, the rest of the Chapter should be understood in the light of the following question: what is the role experience plays in so-called armchair knowledge?

1.4 A First Response: A Restriction of Armchair Access

As seen in the previous sections, Silins' argument relies on the following thesis regarding our access to our evidence:

(Armchair Access): It is sometimes the case that: one's evidence includes some proposition E, and one knows from the armchair that one's evidence includes E. (Silins, 2005: 381)

I have also already mentioned that Silins takes Armchair Access to be overwhelmingly plausible as well as compatible with Williamson's anti-luminosity argument. We can ask: is there a restriction that Williamson, and the externalist more in general, could plausibly build into Armchair Access as defined by Silins? Silins' answer is negative. Presumably, the externalist would restrict Armchair Access to non-environmentally sensitive propositions¹⁴ in order to resist Silins' argument. Crucially, according to Silins, any restriction of Armchair Access to non-environmentally sensitive propositions would be an ad hoc move. It is important to stress that Silins' claim is not that there is no plausible distinction between propositions that can be known from

rather than on sense experience. However, I will not develop this line of response here.

¹⁴I take environmentally sensitive propositions to be propositions whose truth-value depends on how the environment is. By contrast, I take non-environmentally sensitive propositions to be propositions whose truth-value does not depend on how the environment is.

the armchair (as defined by Silins) and propositions that cannot, where the latter category includes environmentally sensitive propositions. The relevant question, instead, is whether there is a plausible distinction between evidence one has and that can be known from the armchair (as defined by Silins) and evidence one has and that cannot be known from the armchair, where the latter category includes environmentally sensitive propositions. The question is thus whether it is plausible to say that we can have armchair knowledge of *some* instances of our evidence, but not of all of them. As I said, Silins's answer is negative, for he takes this to be an ad hoc move, something the Externalist is forced to do merely in order to resist the challenge. However, it is unclear why it would be ad hoc. Silins offers no significant reason for thinking that any relevant restriction would turn out to be implausible. In fact, the contrary seems true. First, remember that "armchair knowledge" as conceived by Silins is knowledge based on enabling experience only. But if this is the notion of armchair knowledge Silins has in mind, then claiming we can have enabling-based knowledge of *some* evidence is thus a very plausible move regardless of whether it resists the Access Problem or not. A move that, in fact, might also be made in the context of the a priori/a posteriori debate. Second, note that, by drawing a parallel with the McKinsey Paradox, Silins has formulated Armchair Access as a rather weak existential claim (i.e., Armchair Access is introduced by: "it is sometimes the case that"), and this is exactly what makes both McKinsey's and Silins' access theses look quite plausible¹⁵. Silins' concern about the restriction being ad hoc would thus be weighty only if he were committed to a stronger claim than he actually is. That is, only if he were committed to an access thesis that has the strength of a universal claim, one on which it is *always* the case that: one's evidence includes a proposition E, and one knows from the armchair that one's evidence includes E. Finally, note that Silins himself points out that, unlike the more notorious Luminosity Argument, his novel access argument rests on a weaker access thesis, thereby focusing on *how* one can know one's evidence, rather than putting pressure on *when* one can know one's evidence (Silins, 2005: 380). Given as things stand, it is overall unclear why restricting armchair knowledge to some (but not all) instances of our evidence is supposed to be an ad hoc move.

But, if we restrict Armchair Access to non-environmentally sensitive propositions, what are we left with? We need to find a set of propositions that can be known from the armchair, whilst not being environmentally sensitive. I take it a plausible candi-

¹⁵In fact, as mentioned above, McKinsey's Privileged Access is the thesis that *typically* one is in the position to know by reflection the content of one's mental states. A similar point is made in Fratantonio & McGlynn, 2018.

date to be represented by logical and mathematical truths. Traditionally, mathematical knowledge has been thought of as being a priori. It might be an open question whether visual thinking, as well as the use of external representations (e.g., symbols, arrows, diagrams, computer simulations, graphs) should seriously cast doubts on the a priori nature of mathematical knowledge (Cf. Burge, 1998). Nevertheless, the standard view has it that knowledge of mathematical truths depends on enabling experience only. For example, while Williamson claims that knowledge of some mathematical proofs seems to rely on more than merely enabling experience, at the same time, he is not willing to embrace an account of mathematical knowledge as a posteriori (Williamson, 2013). On one hand, Williamson maintains that the way in which we come to know some mathematical axioms in set-theory, for instance, is not very different from the way in which we come to know other propositions traditionally taken to be a posteriori. On the other hand, however, he also grants that this does not entail that *no* mathematical truths can be known a priori. My purpose here is not to clearly define which mathematical truths are knowable a priori and which ones are not. It is rather to point out that it is plausible to maintain that there are *some* propositions that are part of our evidence, that are knowable on the basis of enabling experience only, and that are not environmentally sensitive. More precisely, I claim that it is plausible to say that there is a subset of our mathematical knowledge that is based on enabling experience only. The point I want to make is that, although it might be very difficult, perhaps impossible, to exactly define which non-environmentally sensitive propositional evidence can be known “from the armchair” and which ones cannot be known “from the armchair”, it is still plausible to think that there indeed is such a distinction. Mathematical and logical truths (or at least some of them) seem to represent a very plausible candidate for a restriction of Armchair Access^{16,17}. Williamson, and the Externalist more generally, could thus restrict Armchair Access to non-environmentally sensitive propositions, without leaving Armchair Access an unmotivated and empty claim.

Once we realise that a plausible restriction of Armchair Access thesis to non-

¹⁶Note that it is not entirely appropriate to talk about a “restriction” of Armchair Access. As I mentioned above, Armchair Access is formulated as an existential claim. This is what makes it very plausible. What I am doing here is to pick some propositions as instances of that existential claim. However, given Silins himself calls this sort of move to be a “restriction” of Armchair Access, I here use his terminology.

¹⁷Although I believe it is very plausible to assume that mathematical and logical knowledge can be evidence for something, one might instead find this idea questionable. However, note that, by putting forward the Access Problem, Silins is offering a *reductio ad absurdum* of $E=K$. Once we assume $E=K$, any mathematical and logical truths become part of one’s evidence merely in virtue of one coming to know such truths. Thanks to an anonymous referee for pointing out this potential worry to me.

environmentally sensitive propositions is available, the argument does not go through: given Silins' notion of armchair knowledge, Williamson would be happy to deny that Gary has armchair knowledge of the fact that his evidence includes the proposition that the dial reads 0.4. If Silins wants to show that $E=K$, when joined with Armchair Access, brings about the undesirable conclusion, he has to consider a different example, one which Williamson would be happy to grant as a case of armchair knowledge as defined by Silins, namely, as a case of knowledge based on enabling experience only.

1.5 A Second Response: On the Nature of Armchair Knowledge

In the previous Section, I have argued that Silins is too dismissive in taking any restriction of Armchair Access to be ad hoc. The conclusion of my first line of response was that it is far from clear why Williamson should accept the first premise of Silins' argument:

1. Gary has armchair knowledge that his evidence includes the proposition that the dial reads 0.4.

However, someone might not find this first line of response as intellectually satisfying as it should be. For one could still object that it is plausible to say that there is a sense in which we can sometimes know "from the armchair" what our evidence is, even when our evidence includes a proposition about the environment. After all, even Williamson explicitly allows for extended but not unlimited knowledge of our knowledge, as well as of our evidence. For instance, Williamson himself grants the possibility of knowledge of our evidence when he writes:

"Our extensive but not unlimited ability to know that we know without further observation whether we know something is what enables us to use knowledge as evidence. It constitutes an extensive but not unlimited ability to know without further acquisition of evidence whether something is part of our present evidence." (Williamson, 2000: 15)

Similarly, Williamson acknowledges that, although no non-trivial mental state is luminous, "none of this is to deny that in favourable cases one can know without observation whether one is in a given mental state" (Williamson, 2000: 14). It might thus be tempting to argue that there is a sense in which even Williamson has to grant Gary

knowledge of what his evidence is. Nonetheless, as I will show in the second half of this Chapter, I believe that we should be careful in assimilating what Williamson takes to be “knowledge without further observation” with Silins’ notion of armchair knowledge.

In what follows, I put forward a second line of response against Silins’ argument. I argue that, although we do need to account for the fact that it is sometimes possible to know what our evidence is merely “from the armchair”, Silins’ argument relies on a notion of armchair knowledge that Williamson should reject. That is, I argue that Gary’s knowledge of his evidence (and, given $E=K$, Gary’s knowledge of his knowledge) is not an instance of armchair knowledge as understood by Silins, namely, it is not an instance of knowledge depending on *enabling* experience only. This will enable us to resist the Access Problem. For reasons of simplicity, instead of directly addressing Gary’s knowledge of his evidence, I will consider Gary’s second-order knowledge. This is because $E=K$ allows me to apply the conclusions I draw on the nature of one’s second-order knowledge to cases of one’s knowledge of one’s evidence.

The main challenge taken up in this Chapter is to provide a plausible explanation of why, in cases of higher-order knowledge where one’s first order knowledge is empirical, I consider second-order knowledge to not be armchair knowledge in any problematic sense¹⁸. I will not defend any specific theory of second-order knowledge, for an evaluation of which account is the correct one is beyond the scope of this paper. However, for the purpose of this Chapter, I will consider cases in which a subject has second-order knowledge in virtue of: i) having first-order knowledge (assuming factivity of knowledge), and ii) having conducted a further inquiry on her epistemic status (more precisely, on her first-order knowledge). This further inquiry must not be empirical, and it must, of course be properly conducted in order to constitute knowledge. It is important to stress that I do not want to suggest that second-order knowledge can *only* be achieved by means of a further inquiry on one’s epistemic status. In fact, it is plausible to think that, in some cases, one might be able to easily achieve second-order knowledge merely by “being exposed” to first order knowledge¹⁹. In Section 1.6, I will address more explicitly these cases of “easy second-order knowledge”, and I will argue that they do not constitute a threat to my overall strategy. However, my focus here will

¹⁸Note that, in this Chapter, I am only concerned with cases of second-order knowledge (or higher-order knowledge). That is, I am interested in a person’s knowledge of her knowledge, and not in her (psychological) self-knowledge, understood more broadly as the person’s knowledge of her own mental states.

¹⁹Thanks to Timothy Williamson and to an anonymous referee for raising this point to me.

be mainly on cases of second-order knowledge that involve reflection on one's epistemic status, for this represents Silins' target as well. In particular, his being concerned with these "reflective" cases of second-order knowledge is clear when he describes the Gary scenario. He says:

"[W]e can focus on a *thoughtful* subject in the good case. Suppose Gary sees that the dial reads 0.4, and *considers* what evidence he has and what his having certain evidence entails." (Silins, 2005: 381. Italics are mine).

In the Gary scenario, Gary knows that p by having a paradigmatic visual experience that p . Given what I have said above, if Gary knows that he knows that p , then Gary knows that p and Gary has conducted a further inquiry on his epistemic status, namely, on whether he indeed had a paradigmatic knowledge-yielding visual experience that p . What follows can thus be understood as indicating a way of cashing out what this "further inquiry" involves. In particular, I shall address two questions. First, what is the object of this inquiry? Second, what is the nature of this inquiry? Before providing an answer to these questions, I will consider Williamson's imagination-based epistemology of counterfactuals and how, on Williamson's view, the employment of these imaginative capacities involves experience playing a role which is "more than merely enabling and less than strictly evidential" (Williamson, 2016). I will then go back to address the questions regarding the object and the nature of Gary's further inquiry on his first-order knowledge. By developing a new application of Williamson's cognitive-based account of imagination, namely, an imagination-based account of second-order knowledge, I will argue that we should understand Gary's further inquiry on his epistemic status as involving those imaginative capacities that Williamson takes to be in place when we evaluate counterfactuals, or when we come to know some propositions, allegedly taken to be known a priori. I will conclude that, if my imagination-based account of second-order knowledge is correct, then there are good reasons to believe Williamson would resist the major assumption of Silins' argument: Gary's knowledge of his evidence is not "armchair" in any problematic sense, namely, it is not based on *enabling* experience only. A general point can be made: once we consider seriously Williamson's stand on the a priori/a posteriori distinction, Silins' objection can be undermined. Even if there is a sense in which our second-order knowledge is "armchair", namely, it is not merely grounded on *evidential* experience, it is not however an instance of *armchair knowledge* as underlying Silins' interpretation of the Armchair Access thesis. That is, our second-order knowledge is not grounded on *enabling* experience only.

1.5.1 Experience, Counterfactuals, and Imagination

Recently Williamson has offered an imagination-based account of epistemology of counterfactuals²⁰. That is, he argues that in evaluating a counterfactual conditional, namely, in considering both its antecedent and its consequent, we employ our imaginative cognitive capacities²¹. In order to understand what Williamson has in mind, let us consider his own example as introduced in *The Philosophy of Philosophy*:

“You are in the mountains. As the sun melts the ice, rocks embedded in it are loosened and crash down the slope. You notice one rock slide into a bush. You wonder where it would have ended if the bush had not been there. A natural way to answer this question is by *visualising* the rock sliding without the bush there, and then bouncing down the slope into the lake at the bottom. Under suitable background conditions, you thereby *come to know* this counterfactual: [...] if the bush had not been there, the rock would have ended in the lake.” (Williamson, 2007b: 142; Italics added)

What underpins Williamson’s evaluation of the foregoing scenario is the idea that imaginative exercises are knowledge-yielding cognitive processes that involve “a general human capacity to transpose ‘online’ cognitive skills originally developed in perception into corresponding ‘offline’ cognitive skills subsequently applied in imagination” (Williamson, 2013: 296)²², ²³. Without performing any kind of empirical experiment, one can come to know a counterfactual by ‘visually’ imagining its antecedent in one’s mind and by spontaneously developing its consequent. These two moments in the evaluation of a counterfactual mirror, on Williamson’s view, the two moments involved in our imaginative exercise: a first one in which imagination is exercised voluntarily, and a second one in which it is employed involuntarily. The former *modus*

²⁰See Williamson, 2007b. It is worth pointing out that Williamson’s imagination-based epistemology of counterfactuals should be understood within Williamson’s aim of showing epistemology of metaphysical modalities as being a subset of the epistemology of counterfactuals. An evaluation of Williamson’s epistemology of metaphysical modality is beyond the scope of this Chapter. For criticisms, see Jenkins, 2008.

²¹It is worth stressing that the notion of imagination here is not restricted to the Aristotelian notion of imagination as *phantasia* that might feature in dreaming and daydreaming (see Hicks & Aristotle, 1907, iii). Rather the notion of imagination at stake here is that of a broad capacity that shows itself in heterogeneity of knowledge-yielding cognitive processes, e.g., entertaining possibilities and alternative scenarios, evaluating counterfactuals, and making mental comparisons (cf. Williamson, 2016).

²²Similar ideas can also be found in Williamson, 2007b, Williamson, 2016.

²³Unfortunately, Williamson leaves the notion of “offline” rather unexplained in his texts. However, it is plausible to understand Williamson’s use of the notion of “offline” to be similar to the notion used by defenders of simulation theory (e.g., A. I. Goldman, 1989). That is, the imagination is the result of “offline” cognitive skills insofar as, roughly put, these cognitive skills work with surrogated and idealised situations and they can be employed in the absence of the actual object/scenario imagined. This also involves the idea that one can update one’s belief even in the absence of the new evidence.

operandi of imagination is exercised when we make ourselves imagine a possible scenario. That is, we voluntarily consider a specific possibility and we try to imagine what would happen if we were in that possible scenario. However, the imagination works involuntarily when it comes to fill the target possible scenario with details²⁴. In Williamson’s words:

“[H]aving forced the initial conditions, [one] lets the rest of imaginative exercise unfold without further interference” (Williamson, 2016: 116)

But what enables one to “involuntarily” develop the consequent of a counterfactual in such a way that it produces knowledge? That is, what makes imagination a knowledge-yielding cognitive process? On Williamson’s view, our imagination is always constrained by our background knowledge, where this includes our past experiences. More precisely, our past experiences determine the operation of imagination in at least two ways²⁵. First, when addressing the imagined scenario, our ‘offline’ cognitive skills minimise the changes with respect to our background knowledge of similar cases. That is, the imagined scenario remains tuned with how reality is. Second, our background knowledge has a causal role in our evaluation of the counterfactual. That is, when spontaneously developing the possible scenario in the light of our background knowledge, the way in which our imagination operates is reality-oriented. This is what makes imagination a reliable and predictive cognitive process. In fact, it is in virtue of its being constrained by previous experiences that the imagination is selective and truth-oriented, thereby “spontaneously” developing the possible scenarios in question^{26, 27}. In the epistemology of counterfactuals, and in our imaginative exercises in general, sense experience thus plays an important role. As Williamson writes:

“[I]n our imagination-based knowledge of counterfactuals, sense experience can play a role that is neither strictly evidential nor merely enabling.

²⁴See Williamson, 2007b, Williamson, 2016.

²⁵Thanks to Aidan McGlynn for stressing this point.

²⁶It is important to point out that although imagination is reality-oriented, it is far from being an infallible way of gaining knowledge. Nonetheless, the reliability of this cognitive capacity is enough, on Williamson’s view, to become of necessary use in our daily life, for example, as applied in our decision-making, or in the activity of mindreading, as well as in the way we gain knowledge about the future.

²⁷I believe that Goldman’s simulation theory can offer a framework for a better understanding of Williamson’s account of imagination. For instance, Goldman explicitly links imagination to simulation by defining the so-called enactment-imagination: “Enactment-Imagination is a matter of creating or trying to create in one’s own mind a selected mental state, or at least a rough facsimile of such a state, through the faculty of imagination” (A. L. Goldman, 2008). However, I take our imaginative capacities to be something broader than what Goldman takes our simulative capacity to be. See also Gordon, 1986 and Heal, 1986 who were the first ones to propose simulation as an account of human psychology.

For, even without serving as part of our total evidence, it can mold our habits of imagination and judgment in ways that go far beyond a merely enabling role.” (Williamson, 2007b: 165)

To sum up: given Williamson’s cognitive-based account of imagination, sense experience plays an interesting epistemic role in imagination. Although not being strictly evidential, experience seems to play a role that is more than merely enabling one in understanding the concepts involved in the counterfactual. Given the *less than strictly evidential yet more than merely enabling* role that experience plays in the employment of our imaginative capacities, and given the imagination-based epistemology of counterfactuals that Williamson defends, it follows that, on Williamson’s view, knowledge of counterfactuals should not be classified as being a priori, nor a posteriori. The point can thus be generalised to any instance of knowledge involving our imaginative capacities²⁸. As a consequence, given that much of the knowledge traditionally thought of as *a priori* involves the employment of our imaginative capacities, it actually fails to fit within that category, for experience plays a role that is more than merely enabling²⁹. An example is provided in ‘How Deep is The Distinction Between A Priori and A Posteriori Knowledge?’, where Williamson considers the following two truths:

- (a) All crimson things are red
- (b) All recent volumes of Who’s Who are red.

As Williamson points out, the standard view classifies one’s knowledge of (a) as a priori knowledge, and one’s knowledge of (b) as a posteriori knowledge. However, he argues that the cognitive process underlying one’s knowledge that (a) and the cognitive process underlying one’s knowledge that (b) are almost the same (Williamson, 2013). More precisely, he argues that what underlies one’s knowledge in both cases is an imaginative process. As a consequence of his cognitive-based epistemology of imagination just described, Williamson argues that the experience plays a similar role both in one’s knowledge of (a) as well as in one’s knowledge of (b). On the standard view, one’s knowledge of (a) relies on enabling experience only, thereby constituting a priori knowledge, while one’s knowledge of (b) relies on evidential experience, thereby constituting an instance of a posteriori knowledge. *Contra* the standard view,

²⁸For relevant discussion on the evidential-justificatory role of imagination see, for example, Jackson, 2016, Jackson, 2018.

²⁹Williamson further claims that this shows that the a priori/a posteriori distinction lacks any relevant explanatory power. I do not want to commit myself to such a claim here.

Williamson argues that the role that experience plays, both in one's knowledge of (a) and in one's knowledge of (b), is "more than pure enabling and less than strictly evidential". On one hand, in one's knowledge of (a), the experience of redness not only enables one to understand the proposition (a), it also provides one with the "skills in applying the terms 'crimson' and 'red' to the point where he could carry out the imaginative exercise successfully". On the other hand, in one's knowledge of (b), the only role the experience plays is that it grounds "his skill in recognising and imagining such volumes" (Williamson, 2013: 297).

1.5.2 Second-order Knowledge Through Imagination

In the previous Section, I have reconstructed Williamson's account of imagination and how it is widely involved in different ways in our daily epistemic practises, e.g., in our evaluation of counterfactuals. I have pointed out how, on Williamson's view, imagination, broadly understood as one of our cognitive capacities, is sensitive to our past experiences: in our imaginative exercises, sense-experience plays a role which, although less than strictly evidential, is more than merely enabling.

Where does this leave us? In this Section, I argue that Gary's knowledge of his knowledge (and hence, given $E=K$, his knowledge of his evidence) involves the evaluation of counterfactual conditionals, and hence the employment of imaginative exercises. That is, first, I argue that we should understand second-order knowledge as partly entertaining and mentally visualising error-possibilities; second, I argue that, given the more-than-enabling role experience plays in the imagination, and given my imagination-based account of second-order knowledge, Gary's knowledge of his knowledge does not represent an instance of armchair knowledge in any problematic sense. Given $E=K$, the same conclusion can be applied to Gary's knowledge of his evidence. Crucially, if this is so, the Access Problem does not represent a threat for $E=K$. Even if it is true that there is a sense in which Gary has "knowledge from the armchair" of what his evidence is, Gary's knowledge is not an instance of *armchair knowledge* as understood by Silins, and as underlying the Armchair Access thesis.

Remember that we are concerned here with a case of second-order knowledge in which Gary is a thoughtful subject. That is, Gary has second order knowledge that the dial reads 0.4 in virtue of the following two things: i) he has first-order knowledge that the dial reads 0.4; ii) he has conducted a further inquiry on his first-order knowledge³⁰.

³⁰I am here assuming failure of the KK-principle as originally defended by Hintikka, 1962. Daniel Greco has recently offered a novel defence of the KK-principle (D. Greco, 2014). However, the vast

In the previous section, I have addressed the following two questions: first, what is the object of this further inquiry? Second, what is the nature of this inquiry? Providing an answer to the first question will enable us to answer the second question, thereby shedding light on why Gary's first-order knowledge is empirical, yet his second-order knowledge is not an instance of armchair knowledge as intended by Silins.

In the scenario originally described, Gary knows that the dial reads 0.4 simply by looking at it. For Gary to conduct a further inquiry on his epistemic status means to check, without further empirical inquiry, whether he indeed had a knowledge-yielding visual experience. That is, Gary will consider whether the conditions for a paradigmatic good case of perceptual knowledge obtained. For example, he will consider whether his sight was good or not at the moment of the perceptual experience, whether the lights in the room were good, and so on. In a nutshell, for Gary to conduct a further inquiry on his epistemic status (in order to determine whether he had knowledge of a specific empirical proposition p), is to conduct a further inquiry on whether his visual experience that p was a reliable one, thereby constituting a knowledge-yielding process.

It is important to stress that, as I mentioned above, this further inquiry must not involve any empirical check in the environment, yet it must be put forward merely "by reflection" alone. If this is so, then what, in practice, does it mean for Gary to conduct this further inquiry on the reliability of his visual experience? It means that Gary would probably consider the possibility of having had a non-reliable and non-paradigmatic visual experience. In practice, by addressing error-possibilities, Gary would consider what would be the case, if he had not had a paradigmatic visual experience. In addressing this possibility, he would appeal to his background knowledge and he would try to remember previous cases in which he had a visual experience of a dial, perhaps in the same room, or in a different room. He would then make a comparison between past experiences and the visual experience he has just undergone. That is, he would "make up his mind" on his epistemic situation without conducting any further empirical check. Gary's further inquiry would thus consist of addressing and evaluating counterfactual conditionals such as the following:

*If I had not had a paradigmatic visual-experience of a dial reading 0.4,
then it would not look like there was a dial reading 0.4*³¹

majority of epistemologists would reject such a principle nowadays. In particular see Williamson, 2000, ch. 5.

³¹I am here ignoring any skeptical hypothesis.

But, as I have said above, evaluating such a counterfactual involves considering error-possibilities, mentally visualising past experiences and making comparisons with recent ones.

We can now go back to the second question I have addressed: what is the nature of this further inquiry? Given the above description of what Gary's second-order knowledge involves, I suggest we should understand Gary's inquiry on his epistemic status (required in order for him to have second-order knowledge) as involving the employment of those imaginative capacities that Williamson takes to be involved in our evaluating counterfactuals. In fact, as described in the above scenario, Gary seems to conduct this required inquiry by using much of the imaginative exercises that Williamson has taken to represent specific instances of the more general and heterogeneous cognitive capacity of transposing skills "offline". Crucially, I have pointed out that, according to Williamson, the employment of imaginative exercises is constrained by past experiences in such a way that makes imagination relying on more-than-enabling experience. It follows that Gary's second-order knowledge involves cognitive capacities relying on experience that does not merely have an enabling role. Although Gary's past experience fails to have a strictly evidential role, it still constrains Gary's imaginative process in a significant way. But, if Gary's second-order knowledge is not relying on merely enabling experience, then it is not an instance of armchair knowledge as understood by Silins. That is, it is not an instance of armchair knowledge in any problematic sense. Given $E=K$, the same can be said about Gary's knowledge that his evidence includes the proposition that the dial reads 0.4.

1.5.3 Why Knowing the Environment from the Armchair is not Absurd

In the previous Section, I have put forward a second line of response against Silins' Access Problem. I have questioned the main assumption underlying Silins' argument, namely, the Armchair Access thesis. I have argued that, although Gary's knowledge of his evidence is not empirical, it is not "armchair" in any problematic sense: Gary's inquiry on his epistemic situation is somehow constrained by experience, where this plays a *quasi-evidential* role. More precisely, along this second line of response, I have questioned the nature of Silins' notion of armchair knowledge as underlying the Armchair Access thesis. The Access Problem mistakenly relies on a picture according to which the enabling and the evidential roles experience can play are mutually exclusive.

There is, instead, a third role experience can play, one that is neither strictly evidential, nor merely enabling. I have argued that, once we take seriously Williamson’s cognitive-based epistemology of imagination, and once we see how pervasive the employment of this cognitive capacity is, then we can think of a plausible imagination-based account of second-order knowledge (and, thus, of knowledge of one’s evidence). What follows is that, even if there is a sense in which Gary has “armchair knowledge” of what his evidence is, Gary’s knowledge of his evidence is based on this third *quasi-evidential* role of experience, thereby failing to bring about any disastrous conclusion. As Williamson says:

“[I]t should be no surprise if we turn out to have armchair knowledge of truths about the external environment.” (Williamson, 2007b: 269)

Before considering whether cases of “easy” second-order knowledge (i.e., cases that do not involve any inquiry) might affect my overall argument, let me make a final remark on a potentially interesting application of my imagination-based account of second-order knowledge. As mentioned earlier in this Chapter, the Access Problem is developed by Silins as a parallel to the McKinsey Paradox which establishes a tension between Content Externalism and Privilege Access. Now that we have a way of resisting the Access Problem, one might thus wonder whether a similar response is also available to the Content Externalist in order to escape the McKinsey Paradox. That is, perhaps the Content Externalist could accept that one can know by reflection contingent facts about the external environment, whilst arguing that this is not a problematic result, insofar as this ‘reflective knowledge’ is based on *quasi-evidential* experience³². This potential line of response is one that deserves careful examination and that I have no time to address here. Nevertheless, I believe it is an option worth exploring in the future.

1.6 Easy Second-order Knowledge

It could be objected that my account of second-order knowledge is not very plausible after all. One could argue, for instance, that no epistemic work is required for one to have second-order knowledge. According to this line of thought, my explanation of Gary’s scenario would thus fail to show that Gary’s knowledge of his evidence (and

³²For an interesting response to the McKinsey Paradox that accepts the conclusion of the paradox while arguing that it is not absurd see Sawyer, 1998. For an excellent survey of the various responses to the McKinsey Paradox see Kallestrup, 2011a, Kallestrup, 2011b, ch. 5.

thus of his knowledge) is based on more than merely enabling experience. For, if Gary is able to gain second-order knowledge without engaging in counterfactual thinking, and without dismissing error-possibilities, then it seems Gary is not employing his imaginative capacities after all. But remember that my conclusion that Gary's second-order knowledge is grounded on quasi-evidential experience is a consequence of the fact that, as the story goes, it is plausible to take Gary as employing his imaginative capacities. The worry would thus be that, if second-order knowledge does not require a further epistemic inquiry, yet it comes "for free" together with first-order knowledge, then Silins seems right after all: Gary's knowledge is supposedly based on enabling experience only. Previously in this Chapter, I have already noted that my account is compatible with the fact that there might be cases in which one can gain second-order knowledge easily by merely being exposed to one's first-order knowledge. Furthermore, I have pointed out that the cases Silins himself is interested in are cases in which the target subject is thoughtful, and in which he actively considers what evidence he has and what his evidence entails (Silins, 2005: 381). In this Section, I argue that, even if one insisted that we should read Gary's second-order knowledge as not involving any inquiry, this interpretation of Gary's case would still fail to constitute a real threat to my response.

First of all, if, as one might argue, one's second-order knowledge comes almost "for free" together with one's first-order knowledge, then the following question arises: what is this second-order knowledge grounded on? One plausible answer could be that our second order-knowledge is rooted in the same epistemic basis our first-order knowledge is grounded on. However, if this turned out to be the most plausible account of second-order knowledge available, I would happily bite the bullet here, for it would basically concur with my conclusion that, contrary to what Silins claims, Gary's second-order knowledge is based on more than merely enabling experience. In fact, if this account of second-order knowledge is correct, we would have even a more direct response to Silins, given that Gary's second-order knowledge would turn out to be just as empirically based as his first-order knowledge is³³. A different way of answering my question could be to say that Gary's second-order knowledge does not require any further inquiry, and, at the same time, is not epistemically based on the same empirical

³³One might also worry that the story I have provided, one that understands second-order knowledge as involving evaluation of counterfactual conditionals, is incompatible with Williamson's "knowledge-first" project. However, note that in *The Philosophy of Philosophy* Williamson extensively argues for his epistemology of counterfactuals within a knowledge-first framework. Here I am just applying that framework to cases of second-order knowledge.

ground on which Gary's first-order knowledge is based. Crucially, defenders of this account owe us an explanation of where this second-order belief is gaining the justificatory support needed for this belief to constitute knowledge. But, even if a plausible explanation is provided, I believe it would not affect my overall argument. On one hand, all I wanted to argue for is that there is a plausible explanation of how we sometimes acquire knowledge of what our evidence is (and hence of our knowledge), one that sheds light on the role experience plays in these cases of second-order knowledge. It follows that, even if there can be cases in which one gains second-order knowledge without conducting any active inquiry on one's epistemic status, what matters is that this second-order knowledge is not based on *merely* enabling experience only. On the other hand, my account does not completely rule out cases of second-order knowledge that are genuine instances of "armchair knowledge" as intended by Silins. In fact, I take Section 1.4 of this Chapter to show exactly this.

1.7 Conclusion

In this Chapter, I have considered Silins' Access Problem. This is supposed to represent a novel access argument against $E=K$, one which does not rely on a luminosity claim. The argument aims to show that $E=K$ is jointly incompatible with Armchair Access, i.e., the thesis that it is sometimes possible to know from the armchair what our evidence consists of. Williamson's $E=K$, together with Armchair Access, seems to lead to the disastrous conclusion that it is possible to have armchair knowledge of a specific empirical proposition. Given that Armchair Access is a very plausible claim, the argument (allegedly) represents a *reductio ad absurdum* of $E=K$.

This Chapter has rejected the Access Problem as a genuine problem for $E=K$ ³⁴. More precisely, I have put forward two lines of response. According to the first line of response, Silins' formulation of Armchair Access as an existential claim allows for a plausible restriction of Armchair Access to non-environmentally sensitive propositions. Williamson should thus reject the first premise of Silins' argument, thereby

³⁴Silins' paper has been very influential within the epistemological debate on Externalism and Internalism about evidence. However, most of the discussion that followed his *Deception and Evidence* mainly focused on his so-called 'Supervenience Argument' and his novel Sceptical Argument. For relevant discussion on these arguments see Kennedy, 2010, Dunn, 2012, McGlynn, 2014, Fratantonio & McGlynn, 2018. Less attention has been devoted to the Access Problem. Nevertheless, Littlejohn, 2011 constitutes an exception. Note that my response differs from the one provided by Littlejohn for various reasons. In particular, while Littlejohn rejects the principle of armchair access by arguing that $E=K$ (and Externalism more in general) can account for the intuitions allegedly motivating such a principle, this Chapter aims to reconcile the principle of armchair access with $E=K$.

claiming that Gary does not have armchair knowledge that his evidence includes the proposition that the dial reads 0.4. According to the second line of response, everyone – including Williamson – should grant that there is a sense in which Gary can know “from the armchair” that his evidence includes the proposition that the dial reads 0.4. However, Gary’s second-order knowledge is not an instance of *armchair knowledge* as underpinning Silins’ Access Problem. In fact, the Armchair Access thesis as understood by Silins, relies on the assumption that experience can play either an enabling or an evidential role, where these roles are mutually exclusive. That is, Silins’ formulation of Armchair Access overlooks the possibility of experience having a third role: one that is more than merely enabling and less than strictly evidential (what I have called *quasi-evidential*). Furthermore, I have offered an explanation of why I believe Gary’s knowledge of what his evidence is (and thus, given E=K, Gary’s knowledge of his knowledge) represents an instance of knowledge based on *quasi-evidential* experience. More precisely, I have argued for a novel imagination-based account of second-order knowledge, one according to which we gain second-order knowledge by means of evaluating counterfactual conditionals, thereby employing our imaginative exercises. Crucially, if we take seriously Williamson’s cognitive-based account of imagination, according to which imagination relies on more than merely enabling experience, then we are forced to re-think the role experience plays in second-order knowledge. Second-order knowledge seems to rely on experience that, although less than strictly evidential, plays more than a mere enabling role.

Finally, let me highlight why these two lines of response are compatible and mutually supportive. If we understand “armchair knowledge” as knowledge based on *enabling* experience only (as Silins understands it) then Williamson should be happy to accept the Armchair Access thesis, while saying that one can only have armchair knowledge of non-environmentally sensitive propositions. It follows that Gary does not have armchair knowledge of the fact that his evidence includes the proposition that the dial reads 0.4. This is what the first line of response has shown. At the same time, Williamson could grant that there is a sense in which Gary has knowledge “from the armchair” of what his evidence is, while rejecting the idea that Gary has armchair knowledge (as defined by Silins) of the target proposition. The conclusion 3. that Gary knows merely “from the armchair” the specific empirical proposition that the dial reads 0.4 thus follows from the premises 1. and 2.. However, this does not represent a disastrous consequence, for Gary’s knowledge is based on more than mere enabling experience. This is what the second line of response has shown.

The moral of the story is that the Access Problem rests on an ambiguity over the notion of armchair knowledge underpinning its main assumption (i.e., Armchair Access thesis). Once we get clear on which notion of armchair knowledge we are considering, we realise that the Access Problem does not represent a challenge for E=K.

Chapter 2

The Access Problem for Epistemological Disjunctivism

2.1 Introduction

In the previous Chapter, I have considered the so-called Access Problem for E=K, and I have argued that, contrary to first appearances, it does not represent a threat for Williamson's Externalist view. Are all varieties of Evidential Externalism able to resist the Access Problem? This Chapter shows that the answer to this question is 'no'. In this Chapter, I consider the Access Problem for Epistemological Disjunctivism and I argue that it represents a challenge for this view.

Before getting into the details of the problem, let me first introduce what *Epistemological Disjunctivism* is, and why one might find this view appealing. According to Pritchard's *Epistemological Disjunctivism*:

“In paradigmatic cases of perceptual knowledge, an agent, S, has perceptual knowledge that p in virtue of being in possession of rational support, R [i.e., S's seeing that], for her belief that p which is both factive (i.e., R's obtaining entails p) and reflectively accessible.” (Pritchard, 2012: 13).

In a context in which epistemology has been centred on the opposition between internalism and externalism about epistemic justification, *Epistemological Disjunctivism* (henceforth, (ED)) seems to offer (if true) a middle way which would not force us to choose one side over the other¹. By requiring one's rational support to be both reflectively

¹Note that I am not going to consider whether epistemological disjunctivism is also committed to metaphysical disjunctivism. For a defense of epistemological and metaphysical disjunctivism See McDowell, 1995

tively accessible and factive², ED combines the insights underpinning both epistemic internalism and epistemic externalism respectively. On one hand, the accessibility requirement enables Pritchard to account for the (internalist)³ intuition that the notion of justification should be considered in relation to the notion of epistemic responsibility (Pritchard, 2012: 2). On the other hand, the factivity requirement accommodates the (externalist) intuition that there must be a connection between the truth of the proposition and the reason why one believes such a proposition. It is in virtue of this factivity requirement that, as we've seen in the Introduction of this Thesis, Pritchard's view can be classified as an Evidential Externalist theory: on Epistemological Disjunctivism, two subjects can differ in the evidence they have even if they have the same non-factive mental states. Remember our Gary and Barry. Remember that Gary and Barry have the same non-factive mental states but Barry is radically deceived while Gary is in a paradigmatic good case. Epistemological Disjunctivism will say that Gary and Barry differ in the evidence they have, for only Gary has *factive* evidence. This brings me to a further reason one might find Pritchard's ED very appealing. That is, given that on ED it is only in the good case that one has reflective access to one's factive rational support, this view combines internalism and externalism in such a way that it (allegedly) offers a neo-Moorean-style solution to the underdetermination-based sceptical paradox, while leaving our common pre-theoretical intuitions untouched (Pritchard, 2012, Pritchard, 2016a). If true, ED would thus represent, according to Pritchard, the *holy grail* of epistemology. However, in his book, Pritchard anticipates what he takes to be three *prima facie* problems for his view, namely, the Access Problem, the Basis Problem, and The Distinguishability Problem. While lots of the critics have focused on the last two problems, as well as on whether Pritchard's disjunctivism is really able to solve a version of the sceptical paradox, not much attention has been devoted to the so-called Access Problem⁴. In this Chapter, I will focus on this more neglected argument, thus

²As I mentioned previously in this Thesis, it is worth noting that Pritchard's use of 'factive' here differs from what epistemologists often have in mind when talking about 'factive' evidence/justification. While epistemologists traditionally refers to the idea that evidence/justification is constituted by true propositions, Pritchard refers to the truth-entailing nature of one's rational support. Whenever I talk about 'factivity' in Chapter I have in mind the way in which Pritchard talks about factivity.

³Conee & Feldman, 2004 notoriously defend a variety of internalism about justification, namely, mentalism, whilst rejecting the accessibility requirement. However, I will not be concerned with mentalism here. Rather, following Pritchard, I will consider internalism as traditionally understood, namely, as accessibilism.

⁴For a discussion on the Basis Problem see, for example, Pritchard, 2011a, Ranalli, 2014 and Ghislen, 2015. For a discussion on the Distinguishability Problem see Dennis, 2014 and again Ranalli, 2014. For a recent discussion on whether Pritchard's Disjunctivism is able to resist the Sceptical Paradox see Ashton, 2015, Zalabardo, 2015, and Smith, 2016b

hoping to foster and expand the debate around this issue. After clarifying the nature of the problem (Section 2.2), and considering Pritchard's response to this challenge, I point out that in order to assess whether Pritchard's response is a satisfying one, we first need to have a better idea of what it takes for someone to have 'Reflective Access' of one's rational support R (Section 2.3). In particular, Pritchard has to provide an account of reflective knowledge that predicts (i) that one can have reflective knowledge of one's empirical rational support (*seeing that*), and (ii) that one does not have reflective knowledge of what one's rational support entails (*p*). This turns out to be a very difficult task. After considering three possible ways of cashing out the notion of Reflective Access, I show that none of them enables us to satisfyingly resist the Access Problem (Section 2.4 and Section 2.5). Finally, I consider Tim Kraft's paper (Kraft, 2015), in which he offers a different diagnosis of why Pritchard's response to the Access Problem is unsatisfying. I argue that the source/content distinction he appeals to is unmotivated and leads to undesirable results (Section 2.6). This shows why my diagnosis should be preferred. I conclude that, as it stands, the Access Problem represents a real challenge for Pritchard's *Epistemological Disjunctivism*. More interestingly, I conclude that the arguments I have provided shed light on a general troublesome feature of Pritchard's *Epistemological Disjunctivism*, namely, the lack of an unambiguous and unproblematic account of reflective knowledge (Section 2.7).

2.2 Epistemological Disjunctivism and The Access Problem

In his *Epistemological Disjunctivism* (2012), Pritchard addresses the so-called Access Problem as representing a *prima facie* challenge for his view. Before describing the Access Problem in detail, it is worth mentioning that, similarly to what Silins does when formulating the Access Problem for $E=K$, Pritchard takes this argument to mirror the more notorious McKinsey's paradox. Let us thus briefly recall what the McKinsey Paradox is about. As mentioned in the previous Chapter, the McKinsey paradox allegedly shows the incompatibility between Content Externalism and Privileged Access. I've taken Privileged Access to be, roughly put, the thesis that one can typically know by reflection the content of one's mental states. I've taken Content Externalism to be, roughly put, the thesis that the content of one's mental states is determined by factors in the external environment. Let us consider again the example involving "wa-

ter thoughts". Imagine that one is thinking that the water is wet. Privileged Access entails that one can know by reflection that one is thinking that water is wet. Content Externalism entails that if one has a thought about water, then one must have had interacted with H_2O . Assuming one can know by reflection that Content Externalism is true, and assuming a closure principle for reflective knowledge, the McKinsey Paradox aims to show that one can know by reflection a contingent fact about the external environment (e.g., one can know by reflection that one has interacted with H_2O), yet this is (supposedly) absurd.

By requiring the rational support to be both factive and reflectively accessible, and assuming the plausible thesis that it is possible to know by reflection alone that *seeing that p* entails *p*, Epistemological Disjunctivism (ED) seems to immediately face a problem analogous to the one faced by Content Externalism. Pritchard's formulation of this so-called Access Problem goes as follows:

*The Access Problem*⁵

(AP1) S can know by reflection alone that her reason for believing the specific empirical proposition *p* is the factive reason R.

(AP2) S can know by reflection alone that R entails *p*.

So,

(APC) S can know by reflection alone the specific empirical proposition *p*.

(Pritchard, 2012: 46)

The first premise AP1 is a direct consequence of Pritchard's ED. In particular, what matters for the foregoing argument to go through is the thesis, entailed by ED, that a subject S can know by reflection alone that R, where S's rational support R is S's *seeing that*. The second premise AP2 is a thesis about S's knowledge of the truth-entailing nature of R. This second premise is usually taken to be plausible, for it merely assumes that S can know by reflection that *seeing that p* entails *p*. However, the conclusion, which is the result of S's making a competent deduction, looks (allegedly) unacceptable. Before evaluating Pritchard's response to the Access Problem, let me highlight the crucial role that a Closure Principle for Reflective Knowledge plays in the foregoing argument. I take the Closure Principle for Reflective knowledge to be the following CRK:

⁵In evaluating the McKinsey-style challenge for (ED), Pritchard actually considers four different formulations of the Access Problem. In this Chapter, I will only refer to the main formulation, while assuming the agent being in a *good+* case, where a *good+* case is, following Pritchard's taxonomy of good-bad cases, a paradigmatic case of perceptual knowledge. See Pritchard, 2012, p. 29-30.

(CRK): If one knows by reflection that p , and one knows by reflection that p entails q , and one competently deduces q from p while retaining her knowledge that p , then one comes to know by reflection that q .^{6, 7}

The foregoing formulation of closure is what underpins the argument driving the Access Problem. That is, the argument's validity crucially depends on the truth of CRK. We can thus briefly summarise the Access Problem as follows: given that it is possible to know by reflection alone that R (AP1), and that it is possible to know by reflection alone that R entails p (AP2) then, if CRK is true, it follows that it is possible to know by reflection alone that p . If Pritchard wants to save ED he thus needs to block the argument from AP1 to APC.

2.3 Pritchard's Response to The Access Problem

As described above, what distinguishes Epistemological Disjunctivism is the conjunction of the following two constraints:

- (a) The possibility of having reflective access to one's rational support R .
- (b) The truth-entailing nature of one's rational support R .

These two constraints underpin AP1 and AP2 respectively. It is thus the conjunction of (a) and (b) that, together with CRK, gives rise to the Access Problem. Fortunately, Pritchard seems to have a solution to this problem. That is, he argues that:

“the access problem does not represent a challenge to [epistemological disjunctivism], because the conclusion of the above argument [...] fails to follow from the premises, contrary to first appearances”

(Pritchard, 2012: 47).

Insofar as Pritchard is not willing to reject either (a) or (b), his line of response can be defined as *compatibilist*⁸. That is, he maintains that the accessibility thesis (a) and factivity of one's rational support (b) are jointly compatible. Instead, given the

⁶This formulation of closure principle is built upon the one offered by Williamson, 2000 and Hawthorne, 2005.

⁷Insofar as it is plausible to assume that deductive reasoning is a way of coming to know something by reflection, the consequent of the conditional CRK might look redundant. However, for the sake of clarity I will use closure principle for reflective knowledge as stated in CRK.

⁸The terminology is here borrowed from Jessica Brown who, in her discussion of the McKinsey paradox for Content Externalism, takes a compatibilist response to the McKinsey Paradox to be one claiming that Content Externalism and Privileged Access are jointly compatible (Brown, 2004).

argument from AP1 to APC, he rejects the conclusion APC: given the possibility of reflective knowledge of one's rational support, and given factivity of one's rational support (as entailed by ED), it does not follow that one can have reflective knowledge of the empirical proposition p , thereby leaving ED untouched.

But what is Pritchard's motivation for rejecting the entailment from AP1 and AP2 to APC? The reason why the conclusion APC does not follow from the premises is, according to Pritchard, that one's rational and factive support for believing that p is that one sees that p . That is, although one can have reflective knowledge of one's rational support R, once we specify that one's rational support is the empirical reason that one sees that p , then it becomes clear that

“S's route to her acquisition of this (putatively) exclusively reflective knowledge of the target proposition *essentially depends* on the fact that she has empirical reason to believe this proposition.”

(Pritchard, 2012: 47, italics added).

This, in a nutshell, explains why S's knowledge in APC fails to constitute a genuine instance of reflective knowledge. More precisely, what Pritchard is claiming in the foregoing quotation is that, in order for one's instance of propositional knowledge to be in the market for reflective knowledge, it cannot be *essentially dependent* on one's empirical reasons. It is far from clear what it means for one's knowledge (and for one's propositional attitude more in general) to *essentially depend* on one's empirical reasons, and, unfortunately, Pritchard does not explicitly spell out which notion of dependency is in play here. Crucially, in order to assess whether Pritchard's response to the Access Problem is as effective as he hopes it to be, we need to have a better grasp of what it means to have reflective access. That is, we need an account of what it takes for an instance of knowledge to be *essentially dependent* on any empirical reasons one might have. More precisely, for Pritchard's response to be successful, we need a notion of dependency that meets the following two desiderata:

1. it predicts that one has reflective knowledge of one's empirical rational support R (i.e., one's *seeing that*).
2. it predicts that one does not have reflective knowledge of what the rational support R entails (i.e., p).

In what follows, I will consider three ways of understanding this 'essential dependency' as it underpins the notion of reflective access: a metaphysical understanding, a

folk understanding (Section 2.4), and an epistemic understanding (Section 2.5). While I will consider and develop all three possible readings, I will devote most of my attention to the epistemic understanding because, as I will point out later, I take this to be the most charitable interpretation of Pritchard's words. However, I will argue that none of these interpretations is able to meet both the above-mentioned desiderata.

2.4 On the Metaphysical and Folk Interpretation of Essential Dependency

As mentioned in the previous section, Pritchard's way of dealing with the Access Problem Argument is to reject its conclusion. Although we can have reflective knowledge of our rational support R (i.e., our *seeing that p*), we do not have reflective knowledge of what R entails (i.e., p). This is because, Pritchard says, our knowledge of p *essentially depends* on our having empirical reason R in the first place, and reflective knowledge is knowledge that does not *essentially depend* on empirical reasons. But what does it mean for a doxastic attitude to *essentially depend* on empirical reasons? The first interpretation I consider here is what I call the *metaphysical* notion of dependency:

(MD): S 's believing that p metaphysically depends on x for a subject S iff, S wouldn't believe that p if x had not been the case.

If MD is the correct way of understanding the notion of 'essential dependency', then, reflective knowledge amounts to the following:

(MDRF): S 's knowledge that p is reflective knowledge iff, S 's knowledge that p does not metaphysically depend on S 's seeing that p . That is, S 's knowledge that p is reflective knowledge iff, it is not the case that: if S hadn't seen that p , S wouldn't believe that p .

Can we rely on this metaphysical understanding of the notion of dependency in order to satisfyingly resist the Access Problem? If S comes to know that p by *seeing that p* , then it is plausible to say that S 's knowledge that p metaphysically depends (in the sense expressed by MD) on S 's *seeing that p* . For remember that Pritchard is concerned with paradigmatic cases of perceptual knowledge⁹. Therefore, if S knows

⁹The reason why it is important to stress that Pritchard is concerned with paradigmatic cases of perceptual knowledge is that one might be tempted to say that, if MD is the right interpretation of the notion of essential dependency in play, then S knows that p by reflection. After all, one might say, even if S hadn't seen that p , one might come to know that p in a different way, such as, by testimony. Appealing to paradigmatic cases of perceptual knowledge stops this move.

that p by seeing that p , then S wouldn't have believed that p if S hadn't seen that p . Thus, it looks like MD rightly predicts that one does not have reflective knowledge of the empirical proposition p . However, note that, if MD is the notion of essential dependency that we should consider, then, we can't allow for one to have reflective knowledge of one's rational support either. In fact, one's knowledge of one's *seeing that p* clearly metaphysically depends (in the sense expressed in MD) on S 's *seeing that p* . This first metaphysical interpretation of essential dependency is thus not a promising one, for it fails to meet one of the desiderata that a satisfying interpretation of essential dependency should have: although it successfully denies reflective knowledge of the empirical proposition one's rational support entails, it fails to allow for reflective knowledge of one's rational support.

Let us now move onto the second interpretation of essential dependency, namely, what I call the *folk* interpretation. Remember, once again, that, on Pritchard's view, one's knowledge can be classified as *reflective* knowledge iff it does not *essentially* depend on one's empirical reasons. However, a defender of Epistemological Disjunctivism might claim that we should weaken the foregoing notions of 'essential dependency' and 'reflective access'. That is, one might argue that we should understand 'reflective access' as referring to a very cheap and undemanding notion, one that is compatible with the layman's everyday usage of "knowing something by reflection alone and without further empirical enquiry"¹⁰. The idea here is that, even if one had to see that p before coming to know that one sees that p , one can acquire this knowledge merely by 'reflecting' on one's situation and without having to carry any further empirical check. The benefit of appealing to this folk understanding of the notion of reflective access is that, contrary to the metaphysical interpretation, it allows for reflective knowledge of one's rational support. The drawback of this interpretation is, however, that it entails that one can have reflective knowledge of the target empirical proposition entailed by one's rational support. After all, one can just "reflect" on what seeing that p entails without having to conduct an empirical check.

Perhaps Pritchard could thus bite the bullet here, thereby granting the possibility of "reflective knowledge" of a specific empirical proposition p . At the same time, he could say that, once we are not taking "reflective knowledge" to be any technical notion, this wouldn't be a disastrous conclusion after all. Although this might be a tempting move, this strategy is also doomed to failure, for it is incompatible with

¹⁰Littlejohn, 2012 and J. Greco, 2014 seem to have something like this in mind when talking about reflective knowledge.

Pritchard's overall project of offering Epistemological Disjunctivism as being the *holy grail* of epistemology, as well as being a non-revisionist theory. First, remember that Pritchard takes the strength of his Epistemological Disjunctivism to be the fact that it combines the core internalist intuition together with the core externalist intuition. More specifically, the reflective access requirement is supposed to account for the internalist intuition that (doxastic) justification should be considered in relation to the notion of epistemic responsibility. Crucially, a worry here is that a weakening of the notion of "reflective access" would fail to capture the internalist notion of access that has been usually traditionally understood in relation to the technical notion of doxastic responsibility¹¹. Second, and I think more importantly, Pritchard takes another advantage of his view to be that Epistemological Disjunctivism is not a revisionary theory, thereby leaving our pre-theoretical intuitions untouched. A biting-the-bullet strategy would thus fail to account for the general intuition that it is in fact absurd to have "reflective knowledge" of specific empirical propositions.

2.5 On the Epistemic Interpretation of Essential Dependency

In the previous Section, I have briefly considered two interpretations of the notion of essential dependency, namely, a metaphysical interpretation and a folk interpretation. I have shown that neither of them is satisfying. The former is too restrictive and it predicts that one can have reflective knowledge of neither one's rational support R nor what R entails. The latter is too liberal, for it predicts that one can have reflective knowledge of both one's rational support R as well as the empirical proposition *p* one's rational support R entails. In the remainder of this Chapter, I will focus on what I call the 'epistemic interpretation' of the notion of *essential dependency*. I take this to be the most charitable interpretation of Pritchard's words. Crucially, I will argue that, despite being more promising than the metaphysical and folk interpretations, the epistemic interpretation is nonetheless highly problematic. I will thus conclude that, as things stand, there is no plausible account of essential dependency (and thus of Reflective Access) that would enable us to resist the Access Problem, while remaining in line with the spirit of Pritchard's Epistemological Disjunctivism.

I believe that Pritchard's notion of dependency is to be understood in the light of

¹¹A similar concern is expressed by J. Greco, 2014

what constitutes his main target of interest, namely, the nature of perceptual knowledge and, most of all, the nature of a (rational) epistemic basis, together with his interest in doxastic justification¹². Bearing this in mind, I will here understand the notion of epistemic dependency as the following conditional (EPD):

(EPD) S's believing that *p* epistemically depends on *x* for a subject S iff, S's believing that *p* is based on *x*.^{13, 14}

If we take EDP to be the notion of dependency Pritchard has in mind, then the fact that S's knowledge that *p* *essentially* depends on S's having the empirical reason that she *sees that p* is to be understood as the fact that the (rational) reason why S is believing that *p* (and, granted that S is in a paradigmatic good case, S knows that *p*) is that she *sees that p*¹⁵. In other words, S's knowledge that *p* *epistemically* depends on S's *seeing that p*, insofar as S is *basing* her belief that *p* on the fact that she *sees that p*. We can recapitulate things as follows: Pritchard has to resist the argument from AP1 to APC in order to save his ED. Given he is committed to both premises AP1 and AP2, he rejects the conclusion APC. What is Pritchard's motivation for rejecting APC? The answer seems to be that, while S's knowledge that R is reflective, S's knowledge that *p* is not reflective after all, rather it is an instance of empirical knowledge, given that S's reason for believing that *p* is primarily S's seeing that *p*. That is S's knowledge that *p* is *based on S's seeing that p*.

¹²His being concerned with what constitutes the epistemic basis for perceptual knowledge becomes clear in his text, especially in the following two passages:

"[W]hat is key is that one is no longer basing one's belief on the empirical epistemic support once the competent proof has been conducted. Indeed, it is only if this is so that the resultant knowledge is properly classed as exclusively reflective." (Pritchard, 2012: 48)

"[...] the challenge we are raising for the access problem explicitly concerns a case where the agent continues to base her belief on the prior empirical epistemic support that she has [i.e., the fact that she sees that *p*]." (Pritchard, 2012: 48)

¹³I take the epistemic basis in (EPD) to be the notion of epistemic basis traditionally involved in the notion of *doxastic* justification. This is not an uncontroversial topic. However, an analysis of the notion of epistemic basis would be beyond the scope of this Chapter. For the purpose of this Chapter, it is sufficient to cash out the notion of "epistemic basis" in terms of S's reasons on the basis of which S believes a target proposition. I will go back to discussing the notion of basing and doxastic justification in Chapter 4.

¹⁴Note that this is not related to what Pritchard has in mind when talking about "epistemic dependence" in Pritchard, 2015.

¹⁵I am here assuming that S's belief that *p* is based on *r* iff *r* is S's reason for believing that *p*. Note that this biconditional is not uncontroversial. In particular one could question the right to left side of the biconditional. However, this won't matter for the purpose of the Chapter.

2.5.1 Why The Epistemic Interpretation is Problematic

In the previous Section, I reconstructed Pritchard's response to the Access Problem in the light of an epistemic interpretation of the notion of dependency involved. *Prima facie*, it looks like this epistemic interpretation is able to meet the above-mentioned two desiderata. Crucially, at closer inspection, we can see that, even in the light of this more promising notion of 'essential dependency', Pritchard's response is problematic. In particular, I now argue that relying on this epistemic interpretation enables Pritchard to successfully resist the Access Problem only at the high cost of rejecting a plausible Closure Principle for Reflective Knowledge. Let's see why.

As I have mentioned previously in this Chapter, Pritchard's response to the Access Problem can be thought of as a compatibilist one, insofar as he argues that factivity of rational support and privileged access are jointly compatible (as follows from ED). Let us go back to the Access Problem as originally stated by Pritchard:

The Access Problem

(AP1) S can know by reflection alone that her reason for believing the specific empirical proposition *p* is the factive reason R.

(AP2) S can know by reflection alone that R entails *p*.

So,

(APC) S can know by reflection alone the specific empirical proposition *p*.

(Pritchard, 2012: 46)

If Pritchard wants to save ED, he needs to block the argument from AP1 and AP2 to APC. He cannot reject AP1 without giving up his view, for, as mentioned above, AP1 directly follows from ED. He cannot reject AP2 without embracing very counter-intuitive ideas, for, as mentioned above, AP2 is overwhelmingly plausible. Rejecting the conclusion APC, thereby questioning the argument's validity, seems thus to be the only way to go. As we have seen, the way Pritchard argues for the rejection of APC is by pointing out that S's knowledge that *p* essentially depends on S's having factive reason R (i.e., on S's *seeing that p*). On Pritchard's view, even if we grant the possibility of reflective knowledge of R and reflective knowledge of the truth-entailing nature of R, the allegedly disastrous conclusion APC does not follow, for S already knows the specific empirical proposition *p* on the basis of her factive rational support R. The most S can come to know by reflection alone is something about her rational support, and not the specific empirical proposition *p* (Pritchard, 2012: 51). As Pritchard maintains

when reformulating the Access Problem, what follows from AP1 and AP2 is, at best, the following modified conclusion (MAPC):

(MAPC) “[...] S can know by reflection alone that her reason for believing the specific empirical proposition p is the factive empirical reason R that entails p ” (Pritchard, 2012: 51).

But, is MAPC really the strongest claim you can get from AP1, AP2 and CRK, as presented in the Access Problem? I here argue that it is not.

In order to see clearly what is in play in Pritchard’s response, and why MAPC is not the strongest claim we can get from AP1 and AP2, let us reformulate Pritchard’s modified argument, namely, one which has AP1 and AP2 as its premises, and MAPC as its conclusion. Let RK_s be “S knows by reflection alone that”; let r_s be “S has factive rational support R” and let p be an empirical proposition:

(AP1*) $RK_s[r_s]$

(AP2*) $RK_s[r_s \supset p]$

(MAPC) $RK_s[r_s \& (r_s \supset p)]$

The above formulation of the argument shows that the conclusion MAPC results from the conjunction of one’s reflective knowledge of one’s having rational support R (AP1) with one’s knowledge of the truth-entailing nature of one’s rational support R (AP2). What Pritchard needs to assume in order to derive the conclusion MAPC from the premises is nothing as strong as a closure principle CRK, but rather a mere Principle of Agglomeration¹⁶. Crucially, if this is correct, Pritchard’s reformulated version of the argument fails to capture the real worry addressed by the original Access Problem¹⁷. The original Access Problem, as Pritchard himself presents it, is one which involves the possibility of knowing by reflection alone the conceptual implications that follow from one’s previous reflective knowledge, assuming that reflective knowledge is closed under known entailment. If Pritchard wants to save his ED from the Access Problem,

¹⁶I take a general Principle of Agglomeration for reflective knowledge to be the following conditional: $(KR_s[p] \& KR_s[q]) \supset KR_s[p \& q]$.

¹⁷In fact, it would also fail to capture McKinsey’s original worry. When describing the paradox, what McKinsey has in mind is a conceptual notion of dependency, and the problem is one involving a closure principle for knowledge. He writes:

“[I]f you could know a priori that you are in a given mental state, and your being in that state conceptually or logically implies the existence of external objects, then you could know a priori that the external world exists” (McKinsey, 1991: 16).

he should thus consider the original formulation of the Access Problem in the first place, one which rests on the above-mentioned closure principle CRK (and not one involving the Principle of Agglomeration). Crucially, once the Closure Principle for Reflective Knowledge is back on the table, it is easy to see that MAPC is no longer the strongest claim we can get from the premises AP1 and AP2. In fact, if we reformulate the argument involving a Closure Principle for Reflective Knowledge (CRK), we can derive the stronger conclusion APC, i.e., that it is possible to have reflective knowledge of a specific empirical proposition p :

(AP1*) $KR_s[r_s]$

(AP2*) $KR_s[r_s \supset p]$

(APC*) $KR_s[p]$ ¹⁸

Remember that what Pritchard needs is an account of ‘essential dependency’ (and thus of ‘reflective knowledge’) that meets the following two desiderata:

1. it predicts that one has reflective knowledge of one’s rational support R (i.e., one’s *seeing that*).
2. it predicts that one does not have reflective knowledge of what the rational support R entails (i.e., p).

What I have argued for in this section is that, although the epistemic interpretation of the notion of essential dependency seems more promising than the metaphysical and the folk interpretation, it is still unsatisfying. Merely appealing to this epistemic notion of essential dependency doesn’t enable us to satisfyingly meet the above-mentioned desiderata. Although understanding reflective knowledge in terms of epistemic dependency allows us to account for reflective knowledge of one’s rational support, it is not able explain why one cannot gain reflective knowledge of what the rational support entails. If the notion of reflective access underpinning Epistemological Disjunctivism is to be cashed out in terms of *epistemic* dependency then, Pritchard cannot resist the conclusion of the Access Problem, unless he is willing to reject a plausible Closure Principle for Reflective Knowledge. As I have shown, by questioning the argument’s validity, Pritchard has to account for the fact that the argument is valid iff the Closure Principle for Reflective Knowledge is true. To sum up: with an epistemic notion of essential dependency in place, the only way to resist the Access Problem is thus to reject the following plausible Closure Principle for Reflective Knowledge:

¹⁸It is worth noting that a single premise closure is actually enough to take us from (MAPC) to the undesirable consequence (APC*).

(CRK): If one knows by reflection alone that p , and one knows by reflection alone that p entails q , and one competently deduces q from p while retaining her knowledge that p , then one comes to know by reflection alone that q .

2.5.2 Restricting Closure

Defenders of Pritchard's version of ED could reply by biting the bullet once again: they could be willing to reject both APC as well as CRK. Crucially, this does not seem to be what Pritchard has in mind. If Pritchard were really committed to rejection of closure principle, then it is not clear why he never explicitly blames closure as being what generates this prima facie McKinsey-style problem¹⁹. As things stand, it seems Pritchard's disjunctivism is in trouble. Perhaps, however, we can think of a possible solution. One possible strategy could be to argue for the following restricted version of closure (CRK'):

(CRK') If one knows by reflection alone that p , and one competently deduces q from p while retaining one's knowledge by reflection that p , thereby coming to believe that q *for the first time* (on a reflective basis, or on any other basis), then one comes to know by reflection alone that q ²⁰.

Pritchard could argue that the only closure principle we should consider seriously is the restricted CRK'. The advantage of arguing for CRK' is straightforward: by appealing to CRK' Pritchard could provide an explanation of why APC fails to follow from the premises, while anyway retaining a version of closure. By posing an extra condition (i.e., one deductively comes to know that q *for the first time*), the weaker CRK' is entailed by CRK, thereby allowing Pritchard to reject the latter without rejecting the former version of closure. Crucially, while CRK is needed in order for the Access Problem argument to go through, CRK' does not give rise to the same problem. Given that one knows by reflection alone that one has R, and that R entails p , it is easy for one to bring about a competent deduction, thereby coming to know by reflection that p . However, this piece of deductive knowledge is merely a second way in which

¹⁹In his [2012: 51] Pritchard writes: "[T]he reasoning at issue in the access problem is revealed to be fallacious, in that for the premises to be true it simply cannot be the case that S's knowledge of the target proposition is exclusively reflective". One might read this passage as showing Pritchard's commitment to rejecting closure, yet he never explicitly commits himself to such a rejection. On the contrary, although Pritchard never mentions a Closure Principle for Reflective Knowledge in his discussion of the access problem, he does explicitly consider a very similar variety of closure principle for knowledge in other sections of his book. (see Pritchard, 2012: 68).

²⁰Remember that I am still assuming a notion of reflective knowledge as understood in Section 2.5, namely, as knowledge that does not *epistemically* depend on empirical reasons.

one knows that p . That is, one does not know that p on a reflective basis for the first time, for one first comes to know that p empirically by means of *seeing that p* . It follows that from AP1, AP2 and a weaker closure CRK', the conclusion APC does not follow. To recapitulate: Pritchard resists the Access Problem argument by rejecting its conclusion, namely, by questioning its validity. Since, as I have pointed out above, the validity of the argument driving the Access Problem depends on the acceptance of CRK, Pritchard is also forced to reject CRK. At the same time, there is another closure principle for reflective knowledge, namely CRK', that Pritchard can easily embrace without being forced to accept the Access Problem as a real challenge for ED.

Crucially, there are some problems with this restricted closure principle. First, it is far from obvious why we should embrace this restricted version of closure principle and reject the unrestricted one. In order to see why I take CRK' to be left unmotivated, let us imagine a paradigmatic case of perceptual knowledge. Let us suppose that S sees that q , thereby coming to know that q at time t_1 . Let us now suppose that, at a later time t_2 , S also gets testimony knowledge that p , and that p entails q . It seems intuitive to say that, given a plausible closure principle for testimonial knowledge, S also knows that q by testimony. Let us now consider the following restricted version of closure principle for testimonial knowledge:

(CTK') If one knows by testimony that p , and one knows by testimony that p entails q , and one competently deduces q from p while retaining her knowledge that p , thereby coming to believe that q *for the first time* (on a testimony basis, or on any other basis), then one comes to know by testimony that q .

If we were to embrace CTK' (instead of an unrestricted closure principle for testimonial knowledge), nothing will guarantee that the conclusion that [S has testimonial knowledge that q at time t_2] is true, insofar as S does not come to have testimonial knowledge that q *for the first time* (in fact, she already knew that q at t_1). This, however, seems counterintuitive, for it is plausible to say that S can know that q both perceptually and, at a later time, by testimony. We probably wouldn't be willing to reject an unrestricted closure principle for testimonial knowledge in favour of the restricted (CTK'): (CTK') does not look appealing. But then, why should we embrace a restricted closure for reflective knowledge, while rejecting restricted closure for testimonial knowledge? The point I want to make is that, we should expect any defender of a restricted closure principle, such as CRK', to have a general strong motivation for favouring restricted closure principles over unrestricted ones. Crucially, it is not

clear what reasons could be behind restricted closure principle. Besides avoiding the foregoing problem, what motivations underpin CRK' is left unclear, thereby making this restriction looking like an *ad hoc* move.

One could rightly object that as stated (CTK') is questionable²¹. However, I think we can raise a more general concern, one that is enough to show the implausibility of restricted closure principle even to those who might want to reject my first point by stressing that it is not clear whether testimony knowledge is closed under competent deduction. In fact, the previous consideration shows that, more generally, a restricted closure principle seems to strike against the intuition that it is indeed possible to know a specific proposition in more than one way. Consider the following scenario. John is a first-year student in math. One day John goes to the lecture and his professor, Jane, mentions that there are infinitely many prime numbers. John thus comes to know this proposition by testimony. Suppose that some later time, John is actually able to prove that there are infinitely many prime numbers. He reasons and he deductively proves that proposition. Thus, John now knows the same proposition by reflection and deductive reasoning. Crucially, although John's knowledge is the result of deductive reasoning, he does not come to know that proposition *for the first time*, as required by CRK'. It follows that, if CRK' (and not CRK) is the relevant variety of closure principle, then nothing will guarantee that John comes to know by reflection that there are infinitely many prime numbers. But this is highly counterintuitive²². Rejecting CRK, yet holding CRK' might thus have the effect of preventing someone from knowing something in more than one way. Once again, this restricted closure principle does not look appealing.

One might rightly point out that there is a crucial asymmetry between the scenario involving John and Jane, and the original scenario involved in the formulation of the Access Problem. While in the John and Jane scenario we have a combination of testimony and reflective knowledge, in the original scenario we have a combination of perceptual and reflective knowledge. It can then be argued that the way in which reflective knowledge depends on experience in the first scenario is not the same way in which reflective knowledge depends on testimony knowledge in the John and Jane scenario. In fact, while in the Access Problem S could have not had reflective knowledge that p without having had experience that p in the first place, the same is not

²¹In particular, one might say that one's basis for q isn't just testimony, but inference from testimonially-supplied premises. Thanks to Aidan McGlynn for stressing this point.

²²Note that we can get the same result when considering any conjunction of a contingent truth with a logical truth, e.g., the proposition that [the sun is shining $\&(p \vee \neg p)$].

true in the second scenario: John could have gained reflective knowledge that there are infinitely many prime numbers even without having had testimony knowledge of that proposition. But this line of objection seems to appeal to the notion of metaphysical dependency we have seen in the first half of this Chapter, one that, as I have shown, we should not use to cash out the notion of ‘reflective access’ underpinning Epistemological Disjunctivism. Furthermore, and more importantly, what my argument aims to show is merely that even if CRK’, as it stands, turns out to be true, it is nevertheless limited in its application and we should thus also embrace CRK.

It is worth noting, however, that Pritchard is well aware of this problem and he explicitly denies the idea that one can never have reflective knowledge of a proposition if one already had empirical knowledge of the proposition. He says:

“Note that the worry being raised here is not that one cannot come to have exclusively reflective knowledge of a proposition in cases where one already has an empirical basis for believing that proposition, since this is clearly false.” (Pritchard, 2012: 48)

In fact, he claims that the challenge posed by the Access Problem is more specific than that. After considering a case similar to the one I presented here, a case in which a student learns something first empirically, and then by means of a proof, he says:

“Accordingly, it had better not be the case that the difficulty we are raising for the access problem trades on the idea that one cannot come to have exclusively reflective knowledge of a proposition in cases where one already has an empirical basis for believing that proposition. Fortunately for us, it doesn’t. For notice that the challenge we are posing for the access problem is in fact far more specific. In the case just described where one moves from having empirical epistemic support to having (overwhelming) reflective epistemic support, and where one comes to know on this latter basis alone, what is key is that one is no longer basing one’s belief on the empirical epistemic support once the competent proof has been conducted. Indeed, it is only if this is so that the resultant knowledge is properly classed as exclusively reflective.”

(Pritchard, 2012: 48)

Pritchard’s point thus seems to be that when S comes to know that p in the conclusion of the Access Problem, S is still basing her belief on her empirical epistemic support once the deductive reasoning has been conducted. That is why, on his view, S’s belief that p in the conclusion cannot be classified as reflective knowledge. Crucially, it is really not clear why he would allow for reflective knowledge in the case of the student, but not in the case described by the Access Problem. What prevents the

subject of our Access Problem from believing that p on a mere reflective basis, once he came to know that p empirically? In a nutshell, Pritchard owes an explanation as to why the subject S in the Access Problem Argument cannot believe that p while no longer *basing* her belief that p on her empirical support R , while the student can do that.

Before moving on to the next section, I want to make a final remark. The considerations I have drawn so far, and the arguments I have provided, show that if Pritchard wants to save his Epistemological Disjunctivism he has some work to do. But what I have said so far also enables us to make a step forward towards understanding the source of dissatisfaction with Pritchard's response to the Access Problem. That is, the main reason why Pritchard's response is not a satisfying one derives from the fact that the notion of reflective access underpinning ED is unclear. In this Chapter, I have tried to fill in this gap by providing three ways of interpreting the notion of reflective access. Crucially, I have shown that none of them comes without serious problems. Before concluding the Chapter, I will consider an alternative diagnosis of Pritchard's response to the Access Problem, one recently put forward by Kraft (Kraft, 2015). I will argue that my diagnosis of what's problematic with Pritchard's response is to be preferred.

2.6 An Alternative Diagnosis: Tim Kraft's Source/Content Distinction

As I mentioned in the introduction of this Chapter, not much attention has been devoted to the Access Problem for Pritchard's Epistemological Disjunctivism. In this final section, I consider one of the few papers explicitly addressing this problem, namely, Tim Kraft's recent paper (Kraft, 2015), which provides an alternative diagnosis of where Pritchard's response to the Access Problem goes wrong. I argue that the source/content distinction underpinning his diagnosis is unmotivated.

Kraft's starting point in critically evaluating Pritchard's response is the same as mine. He quotes the relevant passage where Pritchard explicitly states that, contrary to first appearances, the conclusion of the argument fails to follow from the premises (Kraft quoting Pritchard, 2015: 318). Kraft then acknowledges that, by rejecting the entailment from the premises to the conclusion, Pritchard is rejecting a plausible closure principle for reflective knowledge. In fact, he takes rejection of closure to be the only palatable option available to the disjunctivist. He writes:

“[t]his leaves epistemological disjunctivism with only one option as a serious contender: the culprit must be closure of reflective knowledge. [. . .] [T]his is the option chosen by Pritchard.” (Kraft, 2015: 317)

However, Kraft points out that Pritchard does not provide an explanation as to why closure for reflective knowledge fails in the Access Problem argument. In his paper, Kraft attempts to fill in this gap. He argues that an explanation can be given by appealing to a distinction between the *source* and the *content* of the empirical support. Roughly put, the idea is that there are two ways in which our epistemic support can be empirical:

“According to the source criterion, R is empirical support for believing that p iff the source of one’s knowledge of, or of one’s access to, R is empirical. [. . .] According to the content criterion, R is empirical support for believing that p iff the content of R has the form $I \phi$ that p , with ϕ , being a perceptual or experiential verb” (Kraft, 2015: 319).

The advantage of appealing to the source/content distinction is, according to Kraft, twofold. First, the source/content distinction constitutes the conceptual tool the disjunctivist needs in order to account for the possibility of having reflective knowledge of an empirical support. Second, and more importantly, the source/content distinction would offer an explanation of why reflective knowledge fails to be closed under competent deduction. In order to see how it can do so, let us recall the original Access Problem argument:

(AP1) S can know by reflection alone that R.

(AP2) S can know by reflection alone that R entails p .

So,

(APC) S can know by reflection alone the specific empirical proposition p .

With respect to the first point, Kraft claims that, once the source/content distinction is available, we can classify S’s knowledge that R in the premise AP1 as reflective knowledge because, despite the fact that R is empirical in the *content* sense, R is nonetheless reflective in the *source* sense. Hence, what Kraft seems to be suggesting is that the nature of the belief’s *source* is what determines whether the target instance of knowledge should be classified as reflective rather than empirical.

But how is the source/content distinction supposed to explain why reflective knowledge is not closed under competent deduction? Here is what Kraft says with respect to this second point:

“If one deduces something from something known reflectively, one always ends up with more knowledge, but not necessarily with more *reflective* knowledge. *Reflective access is lost if the first belief is reflective in the source sense but not in the content sense.*” (Kraft, 2015: 319, italics is mine)

As I have just explained, Kraft takes S’s knowledge in AP1 to be reflective because R in AP1 is reflective in the *source* sense, despite being empirical in the *content* sense. Kraft does not explicitly say anything about S’s knowledge in AP2, but in the light of what we have just said, we can plausibly assume he would classify it as an instance of reflective knowledge as well. In fact, the rational support R in AP2 is empirical in the *content* sense, but reflective in the *source* sense. Finally – and here is the crucial bit - Kraft takes S’s knowledge in APC to be empirical knowledge because, as he says, “reflective access is lost if the first belief is reflective in the source sense but not in the content sense”. Why is this so? Here is Kraft’s motivation behind this:

“[w]hen I see that *p* and believe that *p* (*on that basis*), my belief that *p* is empirical in the *source* sense – no matter how I can know that I see that *p*” (Kraft, 2015: 319, italics added).

Remember that the source of our knowledge that R is what defines whether this target knowledge is empirical or reflective. To say that “my believe that *p* is empirical in the *source* sense” is thus to say, in Kraft’s jargon, that my belief (and hence my knowledge) that *p* is an instance of *empirical* knowledge. This is supposed to explain why APC fails to follow from AP1 and AP2. Nevertheless, in the second half of his paper, Kraft argues that the Access Problem for Pritchard’s Epistemological Disjunctivism cannot be solved anyway. This is because – Kraft says – the source/content distinction can be employed only if a so-called “independent requirement” is met, but, as it turns out, Pritchard’s disjunctivism does not meet such a requirement. Here I shall not discuss whether Pritchard’s disjunctivism does or does not meet this “independent requirement”. What I will focus on is the plausibility of Kraft’s diagnosis and his source/content distinction in the first place. In particular, in what follows, I show that Kraft’s source/content distinction is highly implausible, and that the way this distinction is exploited fails to capture the real reason why the Access Problem originally arises for Pritchard’s Epistemological Disjunctivism.

Before considering the source/content distinction, let me make an important clarification on the general dialectic within which Kraft develops his diagnosis of Pritchard’s response. Remember that Kraft takes the rejection of closure for reflective knowledge

to be the option chosen by Pritchard. However, this is far from being obviously the case. In fact, as I have stressed in the first half of this Chapter, Pritchard never mentions the closure principle for reflective knowledge when resisting the Access Problem. Kraft takes rejection of closure as the only plausible option available to the disjunctivist. Whether this is true or not, it is something Pritchard should not happily welcome. Instead, as I've shown in the first half of this Chapter, it is the undesirable consequence Pritchard has to pay if he wants to resist the Access Problem in a way that is in line with the overall spirit of his Epistemological Disjunctivism.

Let us now move onto considering Kraft's source/content distinction. First, Kraft's source/content distinction, and the way he exploits it in order to explain why APC fails to follow from AP1, is unmotivated. On one hand, Kraft seems to take the source criterion to be what determines the nature of one's knowledge. In fact, remember that it is exactly by appealing to S's knowledge being reflective in the *source* sense that Kraft explains out the possibility of having reflective knowledge in AP1. On the other hand, remember that Kraft takes S's knowledge that *p* in APC to be not reflective, because S's belief in APC is deduced from her belief that R in AP1, which, despite being reflective in the source sense, is empirical in the *content* sense. Again, as Kraft puts it: "reflective access is lost if the first belief is reflective in the source sense but not in the content sense". However, there is no reason why the content of S's knowledge in AP1 should have a role in determining the nature of S's belief in APC, namely, of the belief S infers from S's original instance of knowledge in AP1, together with S's reflective knowledge in AP2. That is, there is no reason why the content of S's belief in AP1 should have a role in determining the source of S's belief in APC. Furthermore, it is completely arbitrary to state, as Kraft seems to be doing, that the *content* of S's belief in AP1 is what determines the *source* of S's belief in APC, while leaving the *source* of the belief in AP1 untouched. I take this to be a very puzzling aspect of the way Kraft exploits the source/content distinction.

Second, if we follow Kraft's way of exploiting the source/content distinction, we have a very controversial scenario in which one starts off with reflective knowledge (as in (AP1)), and ends up with empirical knowledge (as in (APC)), without doing anything but bringing about a competent deduction. I believe this upshot is even more absurd than the one predicted by the original Access Problem.

Third, the source/content distinction, as Kraft uses it, has the undesirable result of classifying paradigmatic instances of reflective knowledge as cases of empirical knowledge. We have seen that, according to Kraft, one can have reflective knowledge

of R, insofar as one's knowledge is reflective in the *source* sense. However, if one's knowledge of R is empirical in the *content* sense, then any other knowledge one has after the first belief will lose its reflective status. Or at least, this is what Kraft maintains. Let's assume for a moment that he's right in saying that S's knowledge that R in AP1 can be classified as *reflective*, while S's knowledge that p in APC should be classified as *empirical*. From S's knowledge that p (in APC), S can easily infer, and thereby know, that $(p \vee \neg p)$. Crucially, if Kraft's diagnosis is correct, then we are forced to classify S's knowledge that $(p \vee \neg p)$ as empirical. But everyone would agree in taking S's knowledge that $(p \vee \neg p)$ to be non empirical. However, here Kraft might have an explanation accounting for the puzzling phenomenon I have just described, namely the fact that:

“[w]henever I see that p and believe that p (on that basis), my belief that p [in APC] is empirical in the source sense – *no matter how I can know that I see that p* ” (Kraft, 2015: 319, italics added).

If this were so, perhaps we should just accept this way of exploiting the source/content distinction – pace our pre-theoretical intuitions. Note, however, that the way he sets up the problem is mistaken in the first place. Kraft is assuming that S's belief that p in the conclusion APC is *based on S's seeing that p* . Crucially, this is exactly what is under discussion in the Access Problem. As I mentioned in the previous sections, the original problem is one in which S comes to know, by deductive reasoning alone, the conceptual implications of S's previous reflective knowledge. For the same reason, to say that it does not matter *how* I can know that I see that p (Kraft, 2015: 319), is to simply dismiss the Access Problem as it originally arises for Pritchard's view. Once we take seriously the challenge posed by the Access Problem, we are then back to consider a case in which S comes to believe and know that p in APC by deductive reasoning.

In the light of the above considerations, I take my diagnosis of where Pritchard's argument goes wrong to be preferred over Kraft's one. Besides avoiding the unacceptable results that Kraft's source/content distinction faces, my diagnosis better captures what is at the bottom of the Access Problem. That is, the diagnosis I've offered better highlights that what generates the Access Problem can eventually be traced back to a weakness of Pritchard's Epistemological Disjunctivism: a lack of a clear and plausible account of what he calls “reflective access”. Furthermore, while Kraft's diagnosis appeals to a controversial distinction, one that Pritchard does not explicitly embrace, the diagnosis I have provided relies on (and only on) claims that can be derived from Pritchard's own version of Epistemological Disjunctivism.

2.7 Conclusion

In this Chapter, I have addressed the Access Problem for Pritchard's Epistemological Disjunctivism (ED). The Access Problem, which mirrors the more notorious McKinsey-paradox, aims to show that the conjunction of accessibilism and factivity of rational support (as entailed by ED) leads to the unacceptable conclusion that one can have reflective knowledge of a specific empirical proposition. In his book, Pritchard offers a way of resisting the problem, namely, he argues that the conclusion fails to follow from the premises: S does not know the specific empirical proposition p by reflection. Pritchard points out that one's instance of knowledge can be classified as 'reflective' only if it does not essentially depend on any empirical reasons. Crucially, he does not provide us with a detailed account of what it takes for a doxastic attitude to essentially depend on empirical reasons. In this Chapter, I have argued that in order to assess the plausibility of Pritchard's response to the Access Problem, we first need to clarify what it takes for someone to have Reflective Access in the first place. In particular, a satisfying account of Reflective Access (and the related notion of essential dependency) has to meet the following desiderata: 1. it has to allow for reflective access of one's rational support; 2. it has to deny reflective access to the proposition one's rational support entails. I have thus provided three interpretations of the notion of essential dependency underlying Reflective Access: a metaphysical, a folk, and an epistemic interpretation. I have shown that none of the account is a satisfying one. The metaphysical account fails to meet the first desideratum. The folk interpretation fails to meet the second desideratum. Most of my attention has however been devoted to the epistemic interpretation insofar as I take this to be the most charitable interpretation of Pritchard's own words. Crucially, I have argued that the only way in which appealing to the epistemic interpretation enables us to resist the Access Problem, thereby meeting both the desiderata, comes at the high cost of rejecting a plausible closure principle for Reflective Knowledge. Finally, I have considered an alternative attempt to make sense of Pritchard's response to the Access Problem, namely, the one offered by Kraft (2015). I have argued that the source/content distinction Kraft appeals to is unmotivated and leads to undesirable consequences. As things stand, and without a clear and unproblematic account of Reflective Access, the Access Problem remains a challenge for Pritchard's Epistemological Disjunctivism.

Concluding Remark on Part I

In Part I, I have considered the so-called Access Problem for Evidential Externalism. More precisely, I have focused on the Access Problem as it seems to arise for Pritchard's Epistemological Disjunctivism (2012) and Williamson's E=K (2000).

We can now ask: does the Access Problem represent a challenge for Evidential Externalism? Given what I have argued for in Chapter 1 and Chapter 2, I believe we should answer this question as follows: it depends on the variety of Evidential Externalism we are considering. Chapter 1 and Chapter 2 show that, while the Access Problem represents a real threat for Pritchard's Epistemological Disjunctivism, it does not constitute a serious challenge for Williamson's E=K. More precisely, what Chapter 1 and Chapter 2 show is that, whether the Access Problem represents a serious threat for Evidential Externalism or not seems to depend on the specific variety of Evidential Externalism, and on what notion of 'reflective knowledge' the target Evidential Externalist view is committed to. Crucially, I have argued that no matter how we define this notion of 'reflective access', Pritchard's Disjunctivist view will be in trouble. That is, Chapter 2 shows that there does not seem to be a notion of 'reflective access' available that resists the access problem while remaining in the spirit of Pritchard's Disjunctivism (i.e., while remaining the alleged *holy grail* of epistemology). Williamson's E=K, however, is not committed to any specific thesis about 'reflective knowledge'. Chapter 1 shows that the account of reflective knowledge I develop, one on which 'reflective' knowledge is based on more than merely enabling experience, is perfectly compatible with E=K and it is able to successfully resist the Access Problem.

At the end of the day, I believe what Pritchard should do is to just give up on such a strong accessibility requirement and embrace a biting-the-bullet strategy. That is, he should just accept that we do have knowledge of contingent propositions about the external environment, and yet this 'reflective knowledge' is not absurd once we take seriously the idea that even 'reflective knowledge' is based on more than merely enabling experience, as Chapter 1 shows. Of course, as I have already stressed, this would come at the cost of not capturing the internalist intuition, and, furthermore, it would mean accepting that Epistemological Disjunctivism is a revisionist theory.

Part II

Evidential Externalism and The Infallibility Problem

Introduction

In Part I of this thesis, I have considered the so-called Access Problem for Evidential Externalism. In particular, I have focused on two prominent varieties of Evidential Externalism: Pritchard's Epistemological Disjunctivism and Williamson's E=K. I have argued that, while the Access Problem represents a challenge for Pritchard's view, it does not constitute a real threat to E=K.

In Part II I will address another argument that has been raised against Evidential Externalism, the so-called Infallibility Problem. Remember that by 'Evidential Externalism' I mean the mere denial of Evidential Internalism, the thesis on which one's evidence supervenes on one's non-factive mental states. As I have stressed in the Introduction of this thesis, the core idea underpinning Evidential Externalism is that features external to one's environment make a difference on what evidence one has. One worry one might have is that, if evidence is so strictly connected with how things are in the external environment, then Evidential Externalism seems to be committed to the idea that evidence gives you a sort of "infallible knowledge of the world". Assume for now that *propositionalism* is true (roughly put, the idea that evidence is propositional in its nature) (Cf. Williamson, 2000: Ch. 9). On Evidential Internalism, evidence is restricted to propositions about one's non-factive mental states. By denying Evidential Internalism, and assuming that all one's evidence is constituted by true propositions, Evidential Externalism will allow for one's "box of evidence" to contain propositions about the world. This seems to be a desirable feature of Evidential Externalism. However, as a consequence some propositions about the world are going to be certain²³ given our evidence, and this might clash with the intuition about the limits of our knowledge. Consider again Gary and Barry. Remember that Gary and Barry share the same non-factive mental states, but Barry is radically deceived, while Gary is in a paradigmatic good case. Both Gary and Barry have a seeming of two bananas in the bowl, but only Gary's environment actually includes two bananas in a bowl. Evidential Externalism entails that Gary and Barry differ in the evidence they have. In fact, the Evidential Externalist will say that Gary's evidence is incompatible with Gary being in a radically deceived situation as Barry is. How things are in the environment determine the evidence one has. In a sense, if Evidential Externalism is the correct view of evidence, and evidence is so strictly related to how in fact things are in the external

²³Note that it is not very clear how we should understand 'certain'. I will address this issue in the coming Chapter.

environment, then, given Garry is in a paradigmatic good case, he seems to have some sort of infallible knowledge. Crucially, it is overwhelmingly plausible to think that we can never get such an infallible knowledge. Plausibly, most of our beliefs are underdetermined by our evidence. In essence, one might worry that Evidential Externalism is committed to some sceptical variety of infallibilism. I will call this the Infallibility Problem for Evidential Externalism, and it will represent the topic of Part II.

This Part is constituted by Chapter 3, and it is followed by a general conclusion. In Chapter 3, I consider the Infallibility Problem for E=K as presented by Dylan Dodd (2005). I will argue that, although Williamson's E=K is committed to a form of infallibilism, this, by itself, does not give rise to any sceptical worry.

While in Part I I have considered the Access Problem for both E=K and Epistemological Disjunctivism respectively, in Part II, I will focus my attention on E=K. Remember that Pritchard's Epistemological Disjunctivism is a view that concerns paradigmatic cases of *perceptual* knowledge only. However, I take it that, at least in this debate, epistemologists are willing to grant that in paradigmatic cases of perceptual knowledge our beliefs are generally *not* underdetermined by our evidence. Nevertheless, one might feel the pull of the underdetermination thesis when it comes to beliefs about the future and, more generally, beliefs about what we cannot directly perceive. Given the restricted scope of Epistemological Disjunctivism, then, even if Epistemological Disjunctivism is committed to a version of infallibilism, it seems that it would have no straightforward sceptical implication. Contrary to Epistemological Disjunctivism, E=K does not restrict its scope to paradigmatic cases of perceptual knowledge and we should thus take the Infallibility Problem for E=K seriously.

Chapter 3

E=K and the Infallibility Problem

3.1 Introduction

Infallibilism has gained a very bad reputation amongst epistemologists, for it has been blamed as being a source of scepticism. In essence, the problem seems to be that the epistemic standards required by the infallibilist are too high to be realistically achieved. Hence, the sceptical worry arises. For this reason, as soon as a theory of knowledge, or epistemic justification, seemed to have an infallibilist vibe, philosophers have often not hesitated to discard the target theory as an implausible view, one that is incompatible with the assumption that - to use Lewis' words - "we know a lot" (Lewis, 1996). A perfect example of the idea that we should be ready to embrace scepticism if we want to be infallibilist can be found in Dodd (2005)¹. According to Dodd, Williamson's E=K is committed to infallibilism, thereby being the source of sceptical troubles. More precisely, Dodd argues that E=K leads to scepticism when jointly taken with the overwhelmingly plausible thesis that virtually all of our beliefs are underdetermined by our evidence. So, Williamson seems to face a dilemma: either he rejects the overwhelmingly plausible Underdetermination Thesis, or he accepts the sceptical conclusion. None of the options look very promising.

The aim of this Chapter is two-fold. First, by focusing on Dodd's argument against E=K, I will show that he provides us with a false dilemma: Williamson could both embrace E=K and a plausible Underdetermination Thesis without facing any sceptical problem. In order to show that Dodd provides us with a false dilemma, I will formulate varieties of infallibilism, and I will show that for each infallibilist thesis we can

¹Note that somewhere else Dodd argues in favour of infallibilism, while saying we should in fact be willing to accept scepticism (Dodd, 2011)

formulate a corresponding variety of underdetermination thesis. Developing this pluralism about infallibilism and Underdetermination theses will not only enable me to resist Dodd’s dilemma, but it will also enable me to draw some more general conclusion concerning the status of infallibilism. That is – and here is the second aim of the Chapter – it will enable me to show that the bad reputation infallibilism has gained is undeserved. This pluralism about infallibilism and underdetermination theses will show that infallibilism *per se* is not the source of any sceptical worry. Rather, each variety of infallibilism leads to scepticism when jointly taken with its corresponding variety of underdetermination thesis. One might think that this is far from a consolation for the infallibilist, insofar as this verdict concurs with Dodd’s conclusion regarding the incompatibility between infallibilism and underdetermination theses. Crucially, I argue that once we have all varieties of infallibilism and underdetermination theses on the table, we should take seriously the possibility of mismatching infallibilism and underdetermination across the taxonomy. This will be enough to show that Dodd’s dilemma is a false dilemma: E=K entails some varieties of infallibilism that, when combined with some non-corresponding varieties of underdetermination theses, avoid scepticism. Furthermore, this will show that not only there are various ways of being an infallibilist, but there are also different ways in which we can account for the intuition that most of our beliefs are underdetermined by our evidence (Section 3.3 and Section 3.4). After that, I will go back to evaluate Dodd’s argument with respect to the specific propositions he considers when formulating the dilemma, such as, lottery-style propositions and propositions about the future. After identifying the main problem with Dodd’s argument, I argue that by appealing to the pluralism I develop in the first half of the Chapter, we can account for our knowledge of (some) propositions about the future and the present, as well as for our ignorance of lottery-style propositions (Section 3.5). I conclude with some general considerations on infallibilism and, in particular, on the infallibilist nature of E=K (Section 3.6).

3.2 The Argument from E=K to Scepticism

In his “Why Williamson Should be a Skeptic” Dodd argues that Williamson’s E=K is committed to a variety of infallibilism that is incompatible with anti-scepticism (Dodd, 2007). More precisely, he argues that E=K is committed to the infallibilist thesis that if one knows that p , then the evidential probability of p is 1 (i.e., $P_s(p | E) = 1$). This is a straightforward consequence of explaining the notion evidential probability in terms

of knowledge in the way Williamson does (see Williamson, 2000, ch. 9 and 10). Normally, before knowing that p , $P_s(p | E) < 1$. But once p is known, p becomes part of one's evidence, thereby making $P_s(p | E) = 1$. Crucially, Dodd argues, this infallibilist thesis is incompatible with the allegedly overwhelmingly plausible underdetermination thesis on which for virtually all of our beliefs, $P_s(p | E) < 1$. Infallibilism and underdetermination theses lead to scepticism when jointly taken:

“For if anti-scepticism and the underdetermination view are both correct, then S can know that p even when S's evidence underdetermines S's belief that p , i.e., when p is only made probable and not certain by S's evidence (in other words, when $P_s(p | E) < 1$). I shall call the view that this is possible ‘fallibilism’, and the view that S can only know that p if $P_s(p | E) = 1$ ‘infallibilism’.”

(Dodd, 2007: 632)

Dodd puts the argument in a form of a dilemma. Either Williamson rejects the plausible underdetermination thesis, or he holds onto both his infallibilist E=K and underdetermination thesis, while embracing the skeptical conclusion that we know almost nothing at all. As anticipated in the Introduction of this Chapter, I believe Dodd is providing us with a false dilemma. But before explaining why it is so, let me first stress one more thing about Dodd's argument. That is, Dodd acknowledges that Williamson might have an explanation for why E=K does not lead to scepticism when it comes to cases of perceptual knowledge (Dodd, 2007: 637). In fact, on Williamson's Knowledge-first Epistemology, knowledge is the most general factive mental state to the extent that any factive mental state, e.g., *seeing* something, will be just a way of knowing something (Williamson, 2000: 33-35). Dodd is aware of this feature of Williamson's view, and he acknowledges that Williamson would be happy to say that, in these perceptual cases, it is just plainly false to claim that the target known proposition has evidential probability less than 1. Nevertheless, Dodd says, Williamson cannot appeal to the same explanation when it comes to other kinds of knowledge, such as, knowledge about the future. He says:

“I do not doubt that if Williamson's theory is true then we do know some things about the external world. [...] What about a sceptic who does not deny that I know what I can directly perceive and what I can remember, but who claims I can know very little about what it is like outside the region of the world with which I have direct perceptual contact, or about what will be the case in the future? Has Williamson a response to this sort of sceptic? I think not.”

(Dodd, 2007: 637)

This is sufficient, according to Dodd, to make Williamson a sceptic.

As I already mentioned, I believe there is nothing to worry about. Dodd's dilemma relies on a mischaracterisation of what a plausible underdetermination thesis looks like. In what follows, taking inspiration from and further exploring Littlejohn's response to Dylan Dodd (Littlejohn, 2008), I develop a taxonomy of infallibilist and underdetermination theses. This will enable me to: i) shed light on what the real source of scepticism is, thereby showing that infallibilism *per se* does not lead to any skeptical result; ii) argue that Williamson's E=K can escape the dilemma by embracing infallibilism, a plausible underdetermination thesis, as well as anti-scepticism. In Section 3.4, I will then further expand the taxonomy of infallibilism and underdetermination thesis in way that will enable me to iii) show that there is a combination of infallibilism and underdetermination thesis that is compatible with E=K and can account for knowledge of (some propositions about) the future, while capturing the intuition that there is a sense in which propositions about the future are underdetermined by our evidence. In order to show this last point, I will have to draw a distinction between propositions about the future, something that Dodd crucially fails to do (Section 3.5).

3.3 What is Infallibilism?

As mentioned in the Introduction, infallibilism has gained a very bad reputation, mainly because it has been thought of as leading to scepticism. But what is infallibilism? I believe there is no univocal way of answering this question. We should thus be wary of drawing any conclusions regarding the sceptical import of infallibilism, and we should clearly specify what we have in mind when we are talking about infallibilism. According to Dutant, "infallibilism is the claim that knowledge requires that one satisfies some infallibility condition" (Dutant, 2007: 59). In order to have an intelligible characterisation of infallibilism, we thus need to answer the question of what this infallibility condition is. We can start by considering one of the main intuitions that have traditionally been taken to underly infallibilism. That is, the idea that knowledge requires certainty. It is important to distinguish at least two notions of certainty, i.e., a *psychological* one and an *epistemic* one. To say that a subject S is certain in believing that *p*, is not the same as to say that the target proposition *p* is certain for S. We can thus sketch two different ways of understanding the infallibilist idea that knowledge requires certainty:

Psychological Certainty Infallibilism: If S knows that *p*, then S is certain

that p .

Epistemic Certainty Infallibilism: If S knows that p , then it is certain that p for S.

According to Psychological Certainty Infallibilism, knowledge that p requires one to be certain in believing that p , where certainty is understood as a psychological state, as oppose to the (psychological) state of doubt. A way of understanding Psychological Certainty Infallibilism is thus to say that being maximally confident about a proposition p is a necessary condition for knowing that p . By contrast Epistemic Certainty Infallibilism is silent on what one's degree of confidence should be². All it claims is that knowledge that p requires the proposition p being certain for S in some relevant *epistemic* sense. Even though there might be a plausible normative relation between the former and the latter, one according to which one should be maximally confident in believing that p only if p is epistemically certain for S, the two formulations do not entail each other. Think of someone who is affected by some kind of persecutory delusion and hence believes that everyone intends to harm her. In this scenario, the person is psychologically certain that everyone intends to harm her, despite the fact that it is not certain that that is actually the case (at least not in any plausible way of understanding *epistemic* certainty). Some epistemologists in the past have embraced something in the neighbourhood of Psychological Certainty Infallibilism (see Moore, 1959, Unger, 1971: 208, Unger, 1975: 95-6). However, this view is not very popular amongst epistemologists nowadays, as far as I am aware. And this is – I believe – rightly so. First, there seems to be cases in which we have knowledge despite being under-confident³. Second, if knowledge requires psychological certainty, then the question of whether we can have knowledge entails a question about whether there are propositions one can believe with maximal confidence. But this is clearly an empirical question, one psychologists should mainly be concerned about. Therefore, if anything, epistemologists who aim to offer an account of infallibilism in terms of certainty should be concerned with an epistemic notion of certainty. Crucially, things are not so simple even when we decide to understand infallibilism as Epistemic Certainty Infallibilism. For there seems to be different ways in which we can understand the notion of *epistemic* certainty. That is, there are different ways to spell out what it takes for a proposition p to be certain for a subject S in some relevant *epistemic* sense. A traditional way of doing it is to say

²See also Stanley, 2008.

³See Radford, 1966. Although Radford's notorious student example is originally supposed to show that there can be knowledge without belief, what I take it to show is that one can know that p without being confident that p .

that knowledge that p requires epistemic certainty in the sense that it “[demands] that the agent has eliminated all possibility of error associated with [p]” (Pritchard, 2005a: 18). Call this Ruling-Out Infallibilism:

Ruling-Out Infallibilism:

If S knows that p , then S has ruled out all the possibilities of error about p .

Notoriously, Lewis has made the link between infallibilism and ruling out error-possibilities explicit :

“[I]t seems as if knowledge must be by definition infallible. If you claim that S knows that p , and yet you grant that S cannot eliminate a certain possibility in which not- p , it certainly seems as if you have granted that S does not after all know that p . To speak of fallible knowledge, of knowledge despite uneliminated possibility of error, just *sounds* contradictory.” (Lewis, 1996: 549)

However, a good characterisation of ruling-out infallibilism needs to clearly explain what it means to “rule out” an error-possibility, and what counts as “error-possibility”. But this has revealed to be a very difficult task⁴. For instance, to say that knowledge requires ruling out all error possibilities has been often understood as the idea that one has to know that all error possibilities do not obtain. But if this is so, then knowledge is defined in terms of knowledge. This might not necessarily be a disastrous result by itself. However, as Dutant points out, if this is how we should understand what it takes for someone to ‘rule out’ error-possibilities, then Rule-Out Infallibilism is not an infallibilist position after all, yet it is the mere consequence of a plausible closure principle for knowledge (Dutant, 2007). Furthermore, it is important to stress that Lewis develops his infallibilist view within a contextualist framework, one that allows for ‘S knows that p ’ to be true even when some incompatible alternative possibilities are not ruled out, as long as these are not *relevant* alternatives. As he famously puts it:

“I say S knows that p iff p holds in every possibility left uneliminated by S’s evidence - Psst! - except for those possibilities that we are properly ignoring.” (Lewis, 1996: 561)

Overall, Lewis’ account of infallibilism does not seem very promising. One might think ‘contextualist’ infallibilism is unsatisfactory and does not really qualify as a genuine ‘infallibilist’ thesis. More generally, defining what it takes for someone to ‘rule

⁴See Hawthorne, 2003 and Dutant, 2007 for a relevant discussion.

out' an alternative, as well as defining what it takes for an alternative to be *relevant* might not be so easy⁵. Therefore, if we want to provide a satisfying characterisation of the notion of epistemic certainty, I believe we should avoid defining it in terms of ruling out error-possibilities. At the same time, it is worth stressing Lewis' important role within the infallibilism/fallibilism debate in epistemology, given that, in his paper, he provides us with an example of how infallibilism is not necessarily an implausible view. In a sense, this is exactly what I want to do in this Chapter.

For the reasons I have just mentioned, in this Chapter I will not be concerned with Ruling-Out Infallibilism, but I will nonetheless spell out in different ways the notion of *epistemic certainty* as it appears in Epistemic Certainty Infallibilism. More precisely, given that, as it stands, the notion of epistemic certainty is uninformative, I will treat it as nothing more than a mere placeholder standing for a general infallibilist position. In the next section, I define different varieties of infallibilism that can thus be understood as various ways of cashing out the infallibilist view on which knowledge requires (epistemic) certainty. This will hopefully shed light on what infallibilism is. It will also show that, in arguing that Williamson is a sceptic, Dodd sometimes interchangeably uses radically different notions of infallibilism. Furthermore, as I mentioned previously, for each variety of infallibilism, I will define corresponding varieties of underdetermination theses. This will enable us to detect what the real source of scepticism is. That is, it will enable us to see that infallibilism *per se* is not the source of any sceptical problem. Rather, epistemologists have blamed infallibilism as leading to scepticism because they have (more or less implicitly) assumed specific varieties of infallibilism together with corresponding underdetermination theses. In particular, this is the reason why Dodd takes E=K to be the source of sceptical worries. Finally, I will argue that, once varieties of infallibilism and underdetermination theses are on the table, we should take seriously the possibility of mismatching infallibilism and underdetermination across the taxonomy, thereby resisting scepticism.

3.3.1 Varieties of Infallibilism (and Underdetermination)

a. Probability 1 Infallibilism

The first variety of infallibilism I consider is what I call Propositional Probability 1 Infallibilism (cf. Brown, 2013):

⁵See also Dretske, 1981 and Nozick, 1981 for different ways of cashing out the notion of ruling-out error possibilities

Propositional Probability 1 Infallibilism (P1I):

If S knows that p , then S's evidential probability that p is 1 (i.e., $P_s(p | E) = 1$).⁶

How to best interpret the notion of probability is beyond the scope of this Chapter. However, for the sake of this Chapter the following clarification will be enough. That is, it is important to stress that *evidential* probability is different from probability understood as *frequency* (i.e., roughly put, the idea that probabilities are collections of events), and from probability understood as *subjective* probability as it usually underpins Bayesian theories (i.e., probability as representing the degree of confidence one has in a proposition).⁷ Furthermore, note that I am here assuming a reading of P1I in terms of the requirement of total evidence. According to P1I, knowledge requires that the evidential probability of the target proposition, namely, its probability given our total evidence, is 1. However, we can define another variety of infallibilism in terms of evidential probability, namely, what I call Doxastic Probability 1 Infallibilism:

Doxastic Probability 1 Infallibilism (D1I):

If S knows that p and believes that p on the basis of e , then $P_s(p | e) = 1$.⁸

The first thing to point out is that, as defined, neither Doxastic nor Propositional Probability 1 Infallibilism put a constraint on whether one's evidence has to be true or false. Of course, someone who's attracted to a more externalist account of evidence would probably not allow for evidence to be false. But my point is that P1I, as well as D1I, do not *by themselves* entail anything specific about what one's evidence can or cannot include. As stated, they are both silent on whether false propositions can be part of one's evidence. However, note that although both D1I and P1I are expressed in terms of evidential probability, the two definitions of infallibilism differ in some important respects. First, note that while D1I is a thesis about the specific instance of evidence e one bases one's belief on, P1I is a thesis about one's *total* evidence. Second, and related, note that D1I is cashed out in terms of the *basing* relation between one's belief and the evidence⁹ one is basing one's belief on (hence, *Doxastic Probabil-*

⁶As I mentioned in a previous footnote, Dodd himself defends this variety of infallibilism (Dodd, 2011). For interesting discussion on this view see Brueckner & Buford, 2012 and Brown, 2013.

⁷For an overview on philosophical interpretations of the notion of probability see Gillies, 2000, Hájek, 2008, and Hajek, 2008. See also Williamson, 1998 and Williamson, 2000, ch.9 and 10, for a discussion on the notion of evidential probability.

⁸From now on, I will use 'E' to designate one's total evidence, and I will use 'e' to refer to an instance of evidence.

⁹In this Chapter, I will assume 'evidence' and 'reasons' to be the same, although I am aware this is not completely uncontroversial.

ity 1 infallibilism)¹⁰. By contrast, PII does not put any constraint on what a subject should base her belief on. We can appreciate how different the two formulations are by considering the following scenario:

Unreliable Guru Testimony

S believes everything her unreliable guru says. One day it is sunny and S sees that it is sunny outside. However, she comes to believe that it is sunny outside for reasons that have nothing to do with how things look like. Instead, she believes that it is sunny outside on the basis of what the very unreliable guru told her.

What DII predicts is that S does not know that it is sunny outside: given she believes everything her guru says (and because her guru said so), the probability of the proposition that it is sunny outside given the evidence she bases her belief on, namely, the guru's testimony, is less than 1.¹¹ By contrast, PII is silent on what kind of reason you should base your belief on, in order to have knowledge. All PII says is that S's knowledge that p requires the probability of p given S's total evidence E to be 1. For all PII says, one's knowledge that p might also be compatible with basing one's belief that p on a specific instance of evidence e , where $P(p | e) < 1$.

We can now formulate the two corresponding Probability 1 Underdetermination Theses as follows:

Probability 1 Underdetermination (P1U):

For virtually¹² any proposition p , $P_s(p | E) < 1$.

Doxastic Probability 1 Underdetermination (D1U):

For virtually¹³ any proposition p and subject S, if S believes that p on the basis of e , then $P(p | e) < 1$.

As I mentioned previously in this Chapter, and as I will further stress later, Dodd has in mind PII and P1U when talking about infallibilism and the underdetermination thesis respectively.

¹⁰I am here using the notion of basing relationship as it is usually understood in the notion of doxastic justification. This is not an uncontroversial topic. More on doxastic justification in Chapter 4 of this Thesis.

¹¹Note that this is regardless of whether we are internalist or externalist about evidence, namely, regardless of whether we think evidence supervenes on our non-factive mental states or not.

¹²There are some exceptions. P1U is false when we let p be a logical truth, or when p belongs to E .

¹³There are some exceptions. For instance, when p is a logical truth or when $p = e$.

b. Entailing Infallibilism

A second variety of infallibilism I consider is what I call Entailing Infallibilism, roughly put the idea that knowledge requires truth-entailing evidence. Once again, we can distinguish a Doxastic and a Propositional variety of entailing infallibilism.

Propositional Entailing Infallibilism (PEI):

If S knows that p , then S's evidence E entails p (i.e., $E \models p$).

Doxastic Entailing Infallibilism (DEI):

If S knows that p and believes that p on the basis of e , then e entails p (i.e., $e \models p$)^{14, 15}.

Propositional Entailing and Doxastic Entailing Infallibilism differ in the same way in which D1I differs from P1I. DEI and PEI are both silent on whether one's evidence can or cannot include false propositions¹⁶. However, DEI and PEI differ in important ways. On one hand DEI (as D1I) concerns the specific instance of evidence one is basing one's belief on, while PEI (as P1I) concerns one's total evidence. Considering once again the *Unreliable Guru Testimony* case can shed light on the differences between PEI and DEI. DEI predicts that S does not know that it is sunny outside given that the evidence she bases her belief on (i.e., the guru testimony) does not entail that it is sunny outside. PEI, instead, does not put any constraint on what kind of evidence S has to base her belief on in order to gain knowledge. All it says is that knowledge requires one's *total* evidence to entail the target proposition.

We can now build the corresponding underdetermination theses as follows:

Propositional Entailing Underdetermination (PEU):

For virtually¹⁷ any proposition p and subject S, then S's evidence does not entail p .

Doxastic Entailing Underdetermination (DEU)

For virtually¹⁸ any proposition p and subject S, if S believes that p on the basis of e , then e does not entail p .

¹⁴Similar formulations can be found in Feldman, 1981, Cohen, 1988, Reed, 2002, Brueckner, 2005a, Stanley, 2005, Dutant, 2007.

¹⁵It might be already clear to the reader that, interestingly, and as I will stress in the Conclusion of Part II, Pritchard's Epistemological Disjunctivism seems to be committed to a version of DEI.

¹⁶Note that while DEI and PEI are silent on whether evidence is propositional, the same is not the case for P1I and D1I. Arguably, the only way in which the probabilistic relations expressed by $P(p | e) < 1$ and $P(p | e) = 1$ are intelligible is only if we assume evidence to be propositional. Cf. Williamson, 2000: ch. 9.

¹⁷There are some exceptions. DEU is false for p being a logical truth, and for p being e .

¹⁸There are some exceptions. As stated, PEU is false for p being a logical truth, and for p being e .

3.3.2 Comments on Varieties of Infallibilism

The taxonomy I have developed in the previous Section shows that, if we want to talk about infallibilism, we should clarify which infallibilist thesis we have in mind, and we should be very wary when drawing conclusions regarding the skeptical import of infallibilism. But let me now further analyse the relation between the varieties of infallibilism I have formulated above.

I think we can classify infallibilist theses into those defined in terms of the *basing* relationship between one's belief and the evidence one is basing one's belief on, and those that are not. Call Doxastic Infallibilism the former, and Propositional Infallibilism the latter. Doxastic Probability 1 Infallibilism and Doxastic Entailing Infallibilism are, not surprisingly, instances of Doxastic Infallibilism. By contrast, Propositional Probability 1 Infallibilism and Propositional Entailing Infallibilism can all be classified as varieties of Propositional Infallibilism.

We can now ask: how do these varieties of infallibilism relate to each other? By adding an extra (basing) requirement for knowledge, we should expect Doxastic varieties of Infallibilism to entail Propositional varieties of Infallibilism. Consider first the relation between Doxastic and Propositional Entailing Infallibilism. I have already pointed out that, according to Doxastic Entailing Infallibilism, knowledge requires one to base one's belief on one's evidence e , where e entails p , thereby putting a constraint on the kind of evidence one has to base one's belief on in order to have knowledge. By contrast, all Propositional Entailing Infallibilism claims is that knowledge that p requires one's overall evidence E to entail p . Propositional Entailing Infallibilism thus follows from Doxastic Entailing Infallibilism, but not vice versa. The same relation holds between Propositional Probability 1 Infallibilism and Doxastic Probability 1 Infallibilism: Doxastic Probability 1 Infallibilism entails Propositional Probability 1 but not vice versa. It should now be clear to the reader that while Doxastic varieties of Infallibilism are about the logical (or probabilistic) relation between the target proposition p and one's evidence e on which one bases one's belief that p (that is, they are about a *doxastic* relation), Propositional varieties of infallibilism are simply about the logical (or probabilistic) relation between one's total evidence E and the target proposition p . Let us now look at the relation between (Propositional and Doxastic) Entailing Infallibilism, and (Propositional and Doxastic) Probability 1 Infallibilism. Recall that P1I is an instance of Propositional Infallibilism insofar as it does not say anything about what I have called the "doxastic relation". According to P1I, if S knows that p ,

then the evidential probability of p given one's total evidence is 1 (i.e., $P(p | E) = 1$). But if $P(p | E) = 1$, then typically¹⁹ one's total evidence E entails p . Hence, Propositional Probability 1 Infallibilism entails Propositional Entailing Infallibilism and vice versa. For the very same reasons Doxastic Probability 1 Infallibilism entails Doxastic Entailing Infallibilism and vice versa. What about the relation between Doxastic Entailing Infallibilism and P1I? As Doxastic Probability 1 Infallibilism entails (but is not entailed by) Propositional Probability 1 Infallibilism, it is easy to see that, while Doxastic Entailing Infallibilism entails P1I, P1I does not necessarily entail Doxastic Entailing Infallibilism, for there can easily be cases in which one believes that p on the basis of evidence e , where e does not entail p , while the evidential probability of p given one's total evidence E is 1.

To sum up: the taxonomy of varieties of infallibilism I have developed in the previous section teaches us three important lessons. First, it teaches us that if we want to talk about infallibilism, we should clarify which infallibilist thesis we have in mind, for these varieties of infallibilism are not necessarily equivalent. In fact, Doxastic varieties of Infallibilism are the strongest infallibilist claims: they entail all instances of Propositional infallibilism, yet they are not entailed by any of them. Instead, the Propositional varieties of infallibilism, namely, PEI and P1I, are equivalent²⁰.

Second, it teaches us that not only we should be wary when talking about infallibilism, but we should also be careful when discussing the problem of underdetermination. We should acknowledge that there might be more than one way in which we can cash out the idea that our evidence underdetermines our beliefs.

Third, and finally, having formulated infallibilist theses together with their corresponding underdetermination theses helps us individuate what the actual source of scepticism is. That is, it is easy to see that each variety of infallibilism leads to scepticism only when taken together with its corresponding variety of underdetermination thesis. On one hand, each infallibilist thesis defines a specific standard for knowledge. On the other hand, their corresponding underdetermination theses claim that we are almost never in the position to meet that standard. Hence, the skeptical worry immedi-

¹⁹Note that there are some exceptions, such that despite one's evidential probability that p is 1, one's evidence does not entail that p . Thus P1I would fail to entail PEI. An example is given by lottery cases with an infinite number of lottery tickets. In these cases, a lottery proposition p has probability 0 (or, on some views, p has infinitesimal probability), yet p is compatible with the evidence one has. However, for the purpose of this Chapter, I will ignore these cases as they concern very particular and exceptional scenarios. Thanks to Martin Smith for pressing me on this point here. See Williamson, 2007a, Smith, 2010a and Williamson, 2014.

²⁰See previous footnotes for exceptions.

ately arises.

Let's now recall why Dodd thinks that E=K is in trouble. According to Dodd, given E=K is an infallibilist thesis, Williamson faces a dilemma. More precisely, Dodd says, E=K is committed to the infallibilist thesis that if S knows that p , then $P_s(p | E) = 1$. Crucially, this leads to skepticism when taken together with an (allegedly) overwhelmingly plausible underdetermination thesis, namely, the thesis that for almost everything we believe, $P_s(p | E) < 1$. So, according to Dodd, Williamson has two options: either he embraces both infallibilism and underdetermination, at the cost of being forced into scepticism, or he rejects the underdetermination thesis, at the cost of becoming too dogmatic. However, this should be hardly surprising at this point. For the infallibilist thesis and the underdetermination thesis Dodd is considering are cashed out in terms of the same epistemic standards (i.e., he is considering P1I and its corresponding P1U). But, as I have just pointed out, this is not a particular feature of P1I, rather it is exactly what we should expect anytime we combine any variety of infallibilism with its corresponding underdetermination thesis.

At this point one could point out that this is of little help for the infallibilist and, more specifically, for the defender of E=K. In fact, this seems to concur with what Dodd was arguing for, namely, that infallibilism is incompatible with the underdetermination thesis. Fortunately, I believe there is no reason to be seriously concerned. For once all these varieties of infallibilism and underdetermination theses are available, there is no principled reason why one should not embrace a specific formulation of infallibilism, while rejecting its corresponding underdetermination thesis. Instead, we should take seriously the possibility of mismatching infallibilism and underdetermination theses across the taxonomy. In the next Section, I will show how we can do so, thereby proving that Dodd's claim that "no infallibilist account of knowledge can be plausibly combined with anti-scepticism" is just plainly false.

3.3.3 E=K: what Infallibilism? What Underdetermination?

In the previous section, I have shed light on varieties of infallibilism and underdetermination theses. I will now appeal to the taxonomy of infallibilism and underdetermination to show that Dodd provides us with a false dilemma. If we focus our attention on E=K, we can now ask the following two questions:

1. What variety of Infallibilism is E=K committed to?

2. Is there a variety of Underdetermination thesis that, when combined with E=K, avoids skepticism?

Let's address 1. first. First of all, note that E=K entails (Propositional) Probability 1 Infallibilism. According to E=K, one's evidence is all and only one's knowledge. This means that, whenever one comes to know that p , p becomes part of one's total evidence. Given that once I come to know p , then p belongs to one's total evidence, it follows that $P_s(p | E) = P_s(p | p) = 1$. But as I have pointed out in the previous section, all Propositional Infallibilist theses are logically equivalent. Therefore given E=K is committed to P1I, then it is also committed to Propositional Entailing Infallibilism (i.e., if S knows that p , then $E \models p$). This comes as no surprise: if S knows that p , then given E=K, p is part of S's total evidence E, and, therefore, E entails p . We can thus say that E=K entails Propositional varieties of Infallibilism, and this is just what we expect to have once we use the concept of knowledge to define what the notion of evidential probability and evidence in the way Williamson does.

What about Doxastic varieties of Infallibilism? I have already stressed that, while Propositional varieties of Infallibilism are entailed by Doxastic varieties of Infallibilism, these in turn do not follow from Propositional Infallibilism. For, as I have just said, Propositional Infallibilism remains silent on what kind of evidence knowledge requires us to base our beliefs on. So, if E=K turns out to entail Doxastic Infallibilism, it is not going to be in virtue of its being committed to Propositional Infallibilism. In any case, it is easy to see why E=K does not *by itself* entail Doxastic Infallibilism. For as it stands E=K is merely a thesis about the *nature* of evidence, and it does not by itself answer the question of what kind of evidence one should base one's beliefs on in order to have knowledge, or whether you need to base one's beliefs on evidence at all (which is, instead, what Doxastic Infallibilism does). It is important to stress that I am not the first one to notice this. For instance, in his response to Dodd, Littlejohn similarly points out that, while E=K entails what I here call 'Propositional Probability 1 Infallibilism', Williamson is not committed to what I here call 'Doxastic Entailing Infallibilism' (Littlejohn, 2008: 682). Crucially, he says, insofar as Williamson's E=K is committed to Propositional Probability 1 Infallibilism but rejects Doxastic Entailing Infallibilism, then "there is no reason to think he is committed to skepticism (2008: 683)". And again, he writes: "it seems that while Williamson is an infallibilist in the sense that he [embraces Propositional Probability 1 Infallibilism], infallibilism so understood does not clearly entail skepticism" (2008: 682). Note, however, that the taxonomy I have developed in the previous section shows something even stronger

than that. It shows that *no* variety of infallibilism per se entails scepticism²¹. As there is nothing intrinsically sceptic about Propositional Probability 1 Infallibilism, there's also nothing intrinsically anti-sceptic about it. Once again, scepticism is what we get *only* when we combine infallibilism with corresponding underdetermination theses.

One might still object that, even if $E=K$ as stated does not entail any variety of Doxastic Infallibilism, there are good principled reasons why Williamson should embrace it anyway. I believe this is a mistake. If anything, I believe there are good reasons why the defender of $E=K$ should not embrace Doxastic Infallibilism. Consider for instance Doxastic Entailing Infallibilism. Remember that, on this version of infallibilism, in order for an agent S to know that p , then S has to believe that p on the basis of S 's (specific instance of) evidence e , such that e entails p . Note that if $E=K$ is true, then e must be something S knows. But if S knows e , and Doxastic Infallibilism is true, then it means S is believing that e on the basis of a further instance of evidence e_2 , such that e_2 entails e . It is easy to see how this leads to a *regressus ad infinitum*.

Let us recapitulate what I have done so far. An analysis of the relationship between Propositional and Doxastic Varieties of Infallibilism helped us to shed light on which infallibilist theses $E=K$ is committed to. I have shown that $E=K$ is committed to both varieties of – what I have called – Propositional Infallibilism (i.e., Propositional Entailing Infallibilism, and Propositional Probability 1 Infallibilism). This comes with no surprise, insofar as all versions of Propositional Infallibilism are trivial consequences of $E=K$. By contrast, $E=K$ is not committed to Doxastic varieties of Infallibilism (i.e., Doxastic Entailing Infallibilism, and Doxastic Probability 1 Infallibilism). Instead, I have shown that there are positive reasons why the defender of $E=K$ should reject Doxastic Infallibilism.²²

²¹It is worth stressing that given Doxastic Infallibilism is the strongest version of Infallibilism of the ones I have formulated so far, Doxastic Infallibilism leads to scepticism when taken together with any of the varieties of underdetermination theses I have formulated in the first half of this Chapter. However, my point is that *even* Doxastic Infallibilism does not *by itself* embody any kind of scepticism.

²²It is important to mention that Dodd believes there is no plausible view on which P1I is true while Doxastic varieties of infallibilism are false. The reason why he thinks this is the case is that, if Propositional varieties are true and Doxastic varieties are false, then you have a scenario in which the evidential probability of a proposition goes from less than 1 to 1 without having acquired any new evidence. I think this objection fails to understand what the real nature of $E=K$ is. In Chapter 4, I show further that $E=K$ by itself says nothing about what evidence one should take into account when coming to believe something. I believe there is a conflation between two claims: a claim about what it is appropriate to cite as reason for p ; a claim about what kind of relation is required between one's total evidence and p in order to have knowledge that p . I do not have time to address this issue here, for my aim in this Chapter is merely to show that Dodd's claim that no variety of infallibilism is compatible with anti-scepticism is just false. However, see Littlejohn, 2008 for a response to Dodd's worry. See also Brown, 2013 for a similar criticism to $E=K$.

We can now address the second question:

2. Is there an Underdetermination thesis that, when combined with E=K, avoids skepticism?

I have already pointed out above that each variety of infallibilism leads to skepticism when jointly taken with its corresponding underdetermination thesis. Crucially, given I have shown E=K to be committed to Propositional Infallibilism, if Propositional Infallibilism and Propositional Underdetermination theses are both true, then we should be seriously worried about the skeptical threat. Fortunately, it is easy to see that, if E=K is true, then all versions of Propositional Underdetermination theses are simply false. Consider first infallibilism and underdetermination as understood by Dodd, namely, Propositional Probability 1 Infallibilism, and Probability 1 Underdetermination. According to (Propositional) Probability 1 Infallibilism, if S knows that p , then the evidential probability of p on one's total evidence is 1, i.e., $P_s(p | E) = 1$. According to P1U, however, for almost anything we believe the evidential probability of the believed proposition on our total evidence is less than one, i.e., $P_s(p | E) < 1$. But why should we accept underdetermination as stated here? If we assume E=K and the non-sceptical thesis, as Dodd himself is doing, then if P1I is true, then Underdetermination as understood by Dodd is just plainly false²³. So, if E=K is true, then what Dodd thought of as being an overwhelmingly plausible thesis actually turns out to be false. As things stand, there are no reasons why Williamson should be persuaded by Dodd into accepting P1U as true. The point can be generalised to any version of Propositional Underdetermination. If Propositional Entailing Infallibilism is true, and we assume the anti-sceptical thesis that we do have quite a lot of knowledge, then its corresponding underdetermination thesis turns out to be just plainly false.

Does this mean that Dodd was right after all, and that, if we want to save both E=K and anti-scepticism, then we need to give up on Underdetermination thesis? I don't think so. Williamson is not forced to embrace the first horn of Dodd's dilemma. In fact, he is not forced to embrace any horn of the dilemma. The defender of E=K does not have to choose between rejecting underdetermination and embracing skepticism. For once we have varieties of Underdetermination theses on the table, Williamson, and the defender of E=K more generally, can reject Propositional varieties of Underdetermination Theses, while embracing Doxastic varieties of Underdetermination Thesis. Interestingly, insofar as this is a mismatch across the taxonomy I have developed, the

²³Of course, as mentioned previously, P1U is also false for reasons that are independent of E=K, namely, for p being a logical truth and for p being E.

combination of Propositional Infallibilism and Doxastic Underdetermination does not lead to any skeptical consequences. As Littlejohn puts it: “Williamson would presumably accept [Doxastic Underdetermination] [...], but would insist that the Underdetermination thesis on this reading is compatible with [Propositional Infallibilism]” (2008: 684)

3.4 What is Underdetermination?

In the previous section, I formulated varieties of infallibilism and, for each of these, I formulated a corresponding underdetermination thesis. I have shown that, far from being an undesirable consequence peculiar to $E=K$, each variety of infallibilism has skeptical implications only when it is embraced together with its corresponding underdetermination thesis. At the same time, I have shown that while $E=K$ is committed to some varieties of infallibilism, namely, Propositional Infallibilism, it does not entail Doxastic Infallibilism. I have then argued that we should mismatch infallibilism and underdetermination theses across the taxonomy. Williamson, and the defender of $E=K$ more in general, could thus embrace Propositional Infallibilism, while accepting that our beliefs are underdetermined by our evidence in the sense expressed by Doxastic Underdetermination. That is, there seems to be a combination of (Propositional) Infallibilism and non-corresponding (Doxastic) Underdetermination available that Williamson could embrace, and that is not subjected to the skeptical threat. Furthermore, I have argued that, regardless of what the appearances might be, in formulating the dilemma, Dodd is not asking us to choose between infallibilism and a very plausible underdetermination thesis that everyone – including Williamson – should embrace at any cost.

However, one might worry that, even if I have identified formulations of Underdetermination thesis that are compatible with Williamson’s $E=K$ (and with its being committed to Propositional Infallibilism), these underdetermination theses, namely, Doxastic Entailing and Doxastic Probability 1 Underdetermination, fail to capture the general intuition behind the notion of underdetermination. It might be that, even though the Underdetermination Thesis presented by Dodd should be rejected (i.e., Probability 1 Underdetermination), the doxastic ones I proposed are not much better. That is, even if Doxastic varieties of Underdetermination are compatible with anti-scepticism and $E=K$, they might nevertheless fail to be fully satisfactory insofar as they might not seem to capture the way in which we usually talk about the problem of underdetermination.

Hence, we should look for a more promising formulation of the underdetermination thesis, one which is compatible with E=K and which does justice to the main intuitions behind the underdetermination thesis. The infallibilist thus still has some work to do.

We can thus ask: what is the main intuition behind the so-called underdetermination of theory by evidence? Once again, there is no agreement on how to clearly spell out the idea that most of our beliefs are underdetermined by our evidence. Within the context of epistemology, the problem of underdetermination is usually discussed in relation to the problem of radical skepticism. Two prominent formulations of the so-called underdetermination-based skeptical argument are provided by Bruckner and Pritchard, who claim that two beliefs are underdetermined by our evidence when our evidence does not favour one belief over the other. Although Bruckner's and Pritchard's formulations slightly differ in their presentation, what they both have in mind is that given two logically incompatible beliefs *b1* and *b2*, to say that our evidence underdetermines *b1* and *b2* is to say that our evidence does not enable us to regard *b1* as true rather than *b2* (Brueckner, 1994, Pritchard, 2005b, Pritchard, 2016a; see also Cohen, 1998, Byrne, 2004).

In the context of philosophy of science, the radical skeptical problem is not usually taken seriously. Nevertheless, I believe philosophers of science have in mind something similar when talking about underdetermination. Amongst philosophers of science, it is more common to define the notion of underdetermination in connection with the notion of empirical equivalence, so two theories are underdetermined when they are empirically equivalent. In their influential paper, Laudan and Lapin maintain that two hypotheses are empirically equivalent iff, were they true, they would have the same empirical consequences (Laudan & Leplin, 1991). Others, however, prefer to cash out the notion of empirical equivalence in terms of empirical indistinguishability or empirical congruence (see Earman, 1993, Kukla, 1998, Psillos, 1999, and Tulodziecki, 2012). However, both epistemologists and philosophers of science seem to refer, to some extent, to the idea that an hypothesis *hp1* is underdetermined insofar as the available evidence can be accommodated by *hp1* as well as by a different hypothesis *hp2*²⁴. I thus take the following Core Underdetermination thesis to capture the core idea shared by the different conceptions of underdetermination (both amongst philosophers of science and epistemologists):

²⁴It is worth noting that the problem of underdetermination of theory by data within philosophy of science is usually addressed in relation to the realism vs anti-realism debate concerning unobservable entities. I will not consider this issue here.

Core Underdetermination:

An hypothesis h_p is underdetermined for S , if the hypotheses h_p and $\text{not-}h_p$ are both compatible (in some relevant sense) with the available evidence.²⁵

Depending on what it takes for two hypotheses (or beliefs) to be compatible in some relevant sense with the available evidence, and depending on what we take evidence to be, we will have various ways of understanding the notion of underdetermination. (More on this later).

If we want an underdetermination thesis that does justice to the Core Underdetermination Thesis, then we need an underdetermination thesis that accounts for the idea that our belief that p is underdetermined when p and a logically incompatible proposition q are on par, in some relevant epistemic sense, with the available evidence. However, on Doxastic versions of Underdetermination the sense in which two beliefs can be underdetermined does not seem to be strictly related to the fact that they are both compatible with one's available evidence. Rather, Doxastic kinds of Underdetermination mainly focus on the logical (or probabilistic) relation between the specific evidence why one believes that p and the target believed proposition p . Therefore, although I believe both Doxastic Entailing and Doxastic Probability 1 Underdetermination provide us with an interesting way of cashing out the idea that our beliefs are often underdetermined, one might wonder whether there is a different characterisation of the underdetermination thesis, one that is more in line with how philosophers of science and epistemologists often talk about the problem of underdetermination.

Bearing this in mind, in what follows, I further develop my taxonomy of infallibilism and underdetermination thesis. While in the first part of the taxonomy I have started by considering plausible varieties of infallibilism, and I have then built corresponding underdetermination theses, I here work backwards. That is, first, I formulate underdetermination theses. Second, I formulate corresponding infallibilist theses. After developing these novel varieties of infallibilism and underdetermination, I will show that while the underdetermination theses I present here are plausible and are in line with the Core Underdetermination Thesis, their corresponding infallibilist theses are either extremely implausible or straightforwardly false. More importantly, none of

²⁵Note that what Laudan and Lapin seem to have in mind is something slightly stronger than this claim. That is, two hypotheses are empirically equivalent when they both *entail* the available evidence. However, what matters is that Core Underdetermination is broad enough to be something most epistemologists and philosophers of science will embrace as the core idea underlying the problem of underdetermination.

the corresponding infallibilist theses is something E=K is committed to. Again, once we take seriously the possibility of mismatching infallibilism and non-corresponding underdetermination across the taxonomy, it is easy to see that there is a plausible combination of infallibilism and underdetermination thesis available, one which resists skepticism, thereby showing Dodd's dilemma to be false.

3.4.1 Varieties of Underdetermination (and Infallibilism)

a. Objective Chance Underdetermination

The first formulation of the underdetermination thesis I propose is the following Objective-Chance Underdetermination Thesis:

Objective Chance Underdetermination (OCU):

For many²⁶ proposition p , time t and world w , the objective chance of p at time t in w is less than one and more than 0 (i.e., $0 < Ch_{tw}(p) < 1$).

Its corresponding infallibilist thesis goes as follows:

Objective-Chance Infallibilism (OCI):

If S knows that p at t in w , then the objective chance of p at t in w is 1 (i.e., $Ch_{tw}(p) = 1$)

Some qualifications are in order here. The notion of objective chance I have in mind is the one considered by Lewis (Lewis, 1980). Although it obeys the laws of probability, Objective Chance is different from probability understood as frequency, and it is different from a subjective interpretation of probability, one that takes probability of a proposition p to express the credence or degree of belief that a subject has towards p . It is also different from evidential probability: objective chance is not the probability of a proposition conditional on a subject's evidence. Objective Chance is a three-argument function. That is, it is the function that, to a given proposition p , time t , and world w , assigns a real number x , with $0 \leq x \leq 1$. If $Ch_{tw}(p) = 1$, then, on this interpretation, the event expressed by p is considered to be not a chancy event. Furthermore, the objective chance of p , where p is a proposition about a particular fact, is time-relative and it typically²⁷ depends on contingent features of the target world w . Propositions concerning events about the past always have objective chance 1, and are

²⁶As I will explain in a moment, arguably, (true) propositions about the past, (true) propositions about the present, and propositions about the future that are entailed by the laws of nature all have objective chance 1.

²⁷An example is constituted by contingent a priori truths. More on this later.

thus not chancy. By contrast, propositions about future events typically have objective chance less than 1. It is in this sense that, as Lewis says, there is a temporal asymmetry of chance. In more speculative terms, we can say that while the past is fixed, the future is wide open. In more technical terms, we can say that for any proposition p about the future, given a time t and a world w , there is a branching possibility in world w at time t where p is false. As Lewis says:

“The asymmetry of fixity and chance may be pictured by a tree. The single trunk is the one possible past that has any present chance of being actual. The many branches are the many possible futures that have some present chance of being actual.” (Lewis, 1980: 273)

Simply put, *branching possibilities* are possible ways in which the future could develop starting from a certain time t . That is, branching possibilities correspond to possible worlds that share their entire history with the actual world w up to a time t ²⁸. My proposal here is that there is one possible way of formulating the underdetermination thesis that is cashed out in terms of objective-chance so understood. A straightforward consequence of explaining the notion of underdetermination (and its corresponding infallibilist thesis) in terms of the notion of objective chance is that it makes the notion of underdetermination (and the notion of knowledge in the corresponding infallibilist claim) time and world relative. In this sense, a proposition p is underdetermined in a world w at time t if $0 < Ch_{tw}(p) < 1$ ²⁹. That is, a proposition p might be underdetermined in a world w at time t in the sense that there might be a possibility branching from w at time t in which the proposition p is false. Conversely, the corresponding Objective-Chance Infallibilism states that if a subject S knows a proposition p at time t in a world w , there is no branching possibility at time t in world w in which p is false. That is, according to Objective Chance Infallibilism knowledge of a proposition p in w at time t requires $Ch_{tw}(p) = 1$.

b. Appearance-Entailing Underdetermination

The next underdetermination thesis is what I call the Appearance-Entailing Underdetermination thesis:

²⁸See also Hawthorne & Lasonen-Aarnio, 2009

²⁹Note that while it is true that if $0 < Ch_{tw}(p) < 1$ then p is underdetermined, the other side of the conditional does not hold. As it will appear clear later, there might be another sense in which p is underdetermined by our evidence despite its objective chance being 1. For instance, it is consistent with my taxonomy that p is underdetermined in the sense expressed by what I'll call Appearance-Entailing Underdetermination, while having Objective chance 1. More on this later.

Appearance-Entailing Underdetermination (AEU):

For virtually³⁰ any proposition p , S's non-factive psychological states (i.e., S's impressions, appearances etc.) do not entail p .

Its corresponding Infallibilist thesis goes as follows:

Appearance-Entailing Infallibilism (AEI):

If S knows that p , then S's non-factive mental states (i.e., S's impressions, appearances etc.) entail p .

3.4.2 Comments on Varieties of Underdetermination

In Section 3.3 of this Chapter, I have defined varieties of infallibilism and corresponding underdetermination theses. By analysing the relationship between varieties of infallibilism and corresponding underdetermination thesis, I have drawn two important conclusions. First, infallibilism by itself does not have any skeptical consequences. Rather, it is only when taken together with its corresponding underdetermination thesis that infallibilism leads to scepticism. Second, and relatedly, mismatches of infallibilism and non-corresponding underdetermination thesis are perfectly compatible with anti-scepticism. In the previous Section, I have expanded my taxonomy by formulating two further varieties of underdetermination theses and corresponding infallibilist theses. Let me now make two brief remarks on these varieties of underdetermination theses and their corresponding infallibilist theses. The first remark concerns how and to what extent these two varieties of underdetermination are in line with the Core Underdetermination Thesis. The second remark concerns the plausibility of these Underdetermination theses as opposed to the highly implausibility of their corresponding infallibilist theses, thereby showing the legitimacy of embracing these novel varieties of Underdetermination thesis while rejecting their corresponding infallibilist theses.

To start with, remember that a motivation for further developing my taxonomy was that, despite Propositional Infallibilism being compatible with Doxastic Underdetermination, one might think that this version of underdetermination is not in line with the following Core Underdetermination thesis that seems to be widely shared by both philosophers of science and epistemologists:

³⁰Some of the few exceptions are when p is a logical truth, when p is the proposition that one has a seeming that p (i.e., the proposition that one has non-factive psychological states), when p is any other proposition that is entailed by the fact that it seems to one that p , and when p is any disjunction with a logical truth as one of its disjuncts.

Core Underdetermination:

An hypothesis h_p is underdetermined by evidence for S , if the hypotheses h_p and not h_p are both compatible (in some relevant sense) with the available evidence.

So, we want a variety of underdetermination that is both compatible with Propositional Infallibilism as well as being able to account for the Core Underdetermination thesis. At the same time, I have already mentioned that, depending on how we fill in the details in Core Underdetermination, we can have various interpretations of the notion of Underdetermination. Both Objective Chance and Appearance Entailing Underdetermination do justice to the Core Underdetermination thesis, whilst representing two radically different ways of understanding the notion of underdetermination. These two different interpretations of the notion of underdetermination rest on two different assumptions about what we take to be included in the subject's available evidence. Let us see why.

Let's start with Appearance-Entailing Underdetermination. What Appearance-Entailing Underdetermination says is that if neither a proposition p , nor the logically incompatible proposition $\neg p$ are entailed by S 's non-factive psychological states (e.g., it is seeming to S that p), then (trivially) both p and $\neg p$ are compatible with S 's non-factive psychological states. Appearance-Entailing Underdetermination captures the Core Underdetermination thesis when taken together with the assumption that the available evidence is constituted by facts about one's non-factive mental states³¹.

How and whether Objective-Chance Underdetermination captures the idea expressed in Core Underdetermination is less straightforward. As I said before, the notion of Objective Chance I have in mind in Objective-Chance Underdetermination is the one defended by David Lewis (1980). Crucially, while Lewis (and epistemologists more in general) have focused on the relation between objective-chance and credence, no connection has ever been drawn, as far as I am aware, between the notion of objective chance and the notion of underdetermination³². My proposal, as I mentioned above, is that we should use the notion of Objective-Chance to shed light on the notion of underdetermination in an interesting way. Remember that Objective Chance is time

³¹Note that this way of talking about a proposition p being underdetermined captures well the way in which philosophers of science and various of epistemologists talk about a proposition p being underdetermined for one in the sense that one is unable to subjectively distinguish a case in which p from a case in which not p .

³²For a relevant discussion on the relation between objective chance and credences see Lewis and his formulation of the "Principal Principle" in his 'A Subjectivist's Guide to Objective Chance'. See Lewis, 1980.

and world relative, and it obeys the laws of probability calculus. We can combine this notion of objective chance with the assumption on which the admissible evidence is constituted by all and only the facts about the present and the past. That is, on this view, the admissible evidence is constituted by all and only facts about p up to time t . To sum up: according to Objective-Chance Underdetermination, S 's belief that p is underdetermined for S at time t in world w , if the objective chance of p given all the admissible evidence about p is less than 1 and higher than 0, where the admissible evidence is all the facts about the present and the past³³.

Let me now move to the second remark I want to make. It must be clear by now that the Underdetermination theses have skeptical consequences only when taken together with their corresponding infallibilist theses. Once again, the suggestion is that we should embrace these varieties of Underdetermination theses while rejecting the corresponding infallibilist theses. In fact, both Underdetermination theses are plausible and they both capture different aspects of the idea that our beliefs are often underdetermined. At the same time, there are good reasons why we should reject their corresponding infallibilist theses. On one hand, Appearance-Entailing Infallibilism, as stated, is false, for it is seeming to one that p (i.e., one's non-factive psychological states), by definition, (typically) does not entail that p ³⁴. On the other hand, by stating that knowledge requires Objective-chance 1, Objective-Chance Infallibilism is highly questionable, and the connection between objective chance and knowledge arguably rejected. In fact, anyone who grants that we can have knowledge of (at least some propositions about) the future and, more generally, inferential knowledge, will reject Objective-chance Infallibilism. Given any proposition about the future has objective chance less than 1, Objective-chance Infallibilism entails that we are ignorant about most of the things we (plausibly) take ourself to know.³⁵ Overall, the varieties of in-

³³Requiring objective chance of the target proposition to be more than 0 allows me to avoid the very counterintuitive consequence that a false proposition about the past, one that has objective chance 0, is underdetermined by our evidence.

³⁴See footnote 30 for some exceptions.

³⁵Interestingly, Hawthorne and Lasonen-Aarnio show that knowledge does not require Objective Chance 1 by considering contingent a priori true propositions. As they point out, contingent a priori true propositions are propositions that have objective chance less than 1 at time t in a world w , and yet one can know these truths a priori at t in w . However, an important issue here is how objective chance relates the safety requirement for knowledge, namely, roughly put, the idea that knowledge of a proposition requires that proposition being true in all worlds close to the actual world in which the subject believes that proposition (more on Williamson's commitment to the safety requirement later). More precisely, the relevant question is whether every branching possibility is a *close* possibility. It seems that if one wants to hold onto both the safety requirement and the idea that it is possible to have knowledge of proposition with objective chance less than 1, one has to deny that every branching possibility is a close possibility. Hawthorne and Lasonen-Aarnio (2009) argue that rejecting this entailment is not enough

fallibilism I have formulated in this second half of the Chapter are not very promising.

3.4.3 E=K: What Underdetermination? What Infallibilism?

We can now put into practice my proposal of mismatching infallibilism and underdetermination across the taxonomy. Williamson and, more generally, the defender of E=K, can embrace the underdetermination theses provided in this second part of the taxonomy, while rejecting their corresponding infallibilist theses. Furthermore, these underdetermination theses are compatible with the varieties of infallibilism E=K is committed to, namely, what I have called *Propositional Infallibilism*. Propositional Probability 1 Infallibilism, Propositional Entailing Infallibilism, Doxastic Entailing Underdetermination, Doxastic Probability 1 Infallibilism, Objective-Chance Underdetermination, and Appearance-Entailing Underdetermination do not lead to skepticism when taken jointly. Once again, Dodd's dilemma is defeated.

3.5 Back to Dodd and Why He is Mistaken

In the previous sections, I have shed light on varieties of infallibilism and underdetermination theses. I have pointed out that no variety of infallibilism *per se* leads to skepticism: it is only when they are taken together with corresponding varieties of underdetermination theses that we have a sceptical result. I have clarified which varieties of infallibilism Williamson's E=K is committed to, and I have pointed out that there are at least four varieties of underdetermination theses that, when combined with E=K, avoid scepticism. With this in mind, I want to now go back to Dodd's argument and show why Dodd is wrong in taking Williamson to be a sceptic.

Remember that, as I mentioned above, Dodd acknowledges that Williamson can account for perceptual knowledge³⁶. However, he argues that if E=K is true, Williamson is forced to conclude that we know almost nothing "about what it is like outside the region of the world with which [we] have direct perceptual contact" (Dodd, 2007: 644). More specifically, Dodd considers the following three propositions: i) my lottery ticket will lose; ii) I will be alive tomorrow; iii) there are books on my shelves. For

for resisting skepticism. A discussion on this topic would lead too far afield. For the purpose of this Chapter it is enough to point out that there are good reasons why Williamson should reject Objective Chance Infallibilism.

³⁶It is worth noticing here that this is not a consequence of E=K, but rather it is a more general consequence of Williamson's commitment to the idea that knowledge is the most general factive mental state, thereby making seeing that *p* a way of knowing that *p*.

each proposition, Dodd argues that, if $E=K$ is true, we cannot have knowledge of the target proposition. Let p be either i), or ii), or iii). In each case, Dodd's reasoning runs - schematically - as follows:

1. There is a small non-zero chance that p is false
2. If $E = K$ is true, then knowledge that p is incompatible with there being a non-zero chance that p is false.
- C. Therefore, if $E = K$ is true we cannot know that p .

The main thought driving the argument from 1. to C. is that, in virtue of being committed to infallibilism, $E=K$ entails that knowledge of a target proposition is incompatible with a chance of that proposition being false. In fact, the argument from 1. to C. captures Dodd's general worry concerning the (allegedly) incompatibility between $E=K$'s commitment to infallibilism and the underdetermination thesis (supposedly captured by 1.). In this section, I will argue that Dodd is mistaken. I will evaluate the arguments for each of the three propositions, and, for each argument, I will show where Dodd's reasoning goes wrong. As I will point out, the foregoing argument relies on an ambiguity of the notion of chance in 1., and, hence, on a mischaracterisation of the underdetermination thesis. I will argue that whether we can or cannot know that p does not depend on the evidential probability of the target proposition. It follows that $E=K$ (and its commitment to P1I) typically does not by itself entail whether we *can* or *cannot* know a specific proposition. By contrast, whether a proposition is something we can know typically depends on factors related to the modal nature of knowledge³⁷.

Lottery Propositions

Let us start with the first proposition Dodd considers, namely, that my lottery ticket will lose. With respect to this proposition, Dodd says:

“According to Williamson one *cannot* know that one's lottery ticket will lose. This is an obvious consequence of his or any plausible version of infallibilism” (Dodd, 2007, footnote 11. Italic is mine)

Remember that the kind of infallibilism Dodd (rightly) attributes to Williamson's $E=K$ is what I have called P1I, namely, the thesis that if S knows that p , then the evidential probability of p is 1. P1I entails that, if I know that my lottery ticket will lose, the evidential probability of the proposition that my lottery ticket will lose is 1.

³⁷An exception concerns necessary false propositions.

However, Dodd says, the evidential probability of the proposition that my ticket will lose is less than 1. Hence, if $E=K$ is true, I *cannot* know that my lottery ticket will lose. Dodd is right in saying that Williamson takes lottery-style propositions to be a kind of propositions we just cannot know. Crucially, this is not a consequence of Williamson's commitment to infallibilism. As I have pointed out in the previous sections, P1I is just what trivially follows once we define the notion of evidence in the light of the notion of knowledge. In arguing that our ignorance of lottery-style propositions is a direct consequence of $E=K$, Dodd is committing a crucial mistake: he is taking the notion of evidential probability as what determines whether we *can* or *cannot* know a given proposition. By contrast, the evidential probability of a proposition is a mere indicator of whether, as things stand, we already know the target proposition or not³⁸. So, even if Williamson denies that we can have knowledge of the fact that my lottery ticket will lose, the explanation of why this is so is not to be found in $E=K$'s commitment to (Propositional) Probability 1 Infallibilism. Putting the matter in another way, my claim is that Dodd is reversing the order of explanation. Let "lottery" stand for the proposition that my lottery ticket will lose. The starting point in Dodd's reasoning is that $P(\textit{lottery} | E) < 1$. On his view, this is sufficient to lead to the sceptical conclusion that one cannot know "lottery" when taken together with $E=K$. Crucially, this gets the order of explanation wrong. Williamson agrees that we cannot know "lottery", but this is for reasons that, as I will explain shortly, are independent of $E=K$. We can agree with Dodd in saying that the evidential probability of "lottery" is less than 1 at t , but the evidential probability by itself does not explain, on Williamson's view, *why* I cannot know "lottery". If anything the fact that we cannot know "lottery" explains why the evidential probability of "lottery" remains less than 1. The fact that the evidential probability of a lottery-style proposition for a subject S remains less than 1 is thus a mere consequence of the fact that we *cannot* know lottery-style propositions in the first place, together with $E=K$ ^{39, 40}.

How can we explain our ignorance of lottery-style propositions if not by appealing to the evidential probability of "lottery"? The way we should explain it, or, at least,

³⁸Cf. also Littlejohn, 2008: 681 for a similar point.

³⁹It is important to point out that while, on Propositional Probability 1 Infallibilism, knowledge that p entails $P(p | E) = 1$, the other side of the entailment does not hold. We can thus have cases in which the evidential probability of p for S is 1 and yet one fails to know that p . It is for this reasons that I say that ignorance of a proposition p entails that the evidential probability of p *remains* less than 1. This won't matter for the purpose of my argument here.

⁴⁰Note that the idea that we cannot know that my lottery ticket will lose is widely embraced amongst contemporary epistemologists.

the way in which Williamson would explain it, is by appealing to the modal nature of knowledge, and to Williamson's commitment to the safety requirement for knowledge (Williamson, 2000, Williamson, 2009a). In a nutshell, this means that if S knows that p , then S's belief that p is safe from error. It is easy to see how the notion of safety is related to the notion of risk. For if S's belief is safe from error, then S is not at risk of falsely believing p (or a very similar proposition) in similar cases. Williamson thinks of the notions of safety and risk as modal notions. He says:

“[O]ne is safe in a possible world w at a time t from an eventuality if and only if that eventuality obtains in no world ‘close’ to w at t . Call this the ‘no-close risk’ conception of safety”. (Williamson, 2009a: 13)

Different formulations of the notion of ‘safety’ and ‘risk’ have been offered in the literature, but providing a detailed characterisation of the safety requirement would lead us too far afield⁴¹. Note, however, that my argument does not rely on any specific account of safety. The generality of Dodd's mistake in his arguing against E=K is such that any account of safety will do the job. For the purpose of this Chapter, it is thus enough to stress that, what all epistemic accounts of safety and risk have in common is that high probability of p being true is not enough for considering one's belief that p as safe⁴². That is, what ties together knowledge and risk of error is neither the objective chance of the target proposition, nor its evidential probability. Rather, it is the modal construal of the notion of safety⁴³.

The point is an important one, for it highlights the general problem with Dodd's argument. Remember that Dodd's argument concerns whether knowledge is compatible with the fact that virtually all of our beliefs are at risk of being in error (Underdetermination thesis). The way he cashes out the idea that virtually all of our beliefs are at risk of being in error is by assuming that evidential probability of virtually all our beliefs is less than 1. By doing so, he is wrongly assuming a notion of risk in terms of evidential probability, and he is illegitimately assuming that the evidential probability of a proposition is what determines whether at time t the target proposition is a good

⁴¹For formulations of safety see Sosa, 1999, Williamson, 2000: 123-130, Pritchard, 2005a. Recently Pritchard has also moved from talking about epistemic luck to talking about epistemic risk. See Pritchard, 2016b.

⁴²In fact, this is exactly what lottery propositions are supposed to show.

⁴³Williamson claims that “when epistemological risk is at issue, the explanation will be in epistemological terms” rather than in probabilistic terms (Williamson, 2009a, p.9). In fact, he argues that understanding risk in epistemological terms rather than probabilistic terms will allow us to avoid the notorious Reference Class problem that most theories of probability face (frequentists in particular). One worry I have is whether these modal accounts of risk face a similar reference class problem. This is an interesting question, but one that I am not going to address here.

candidate for being an instance of knowledge. As I have stressed above, however, the evidential probability of a proposition is merely an indicator of whether, as things stand, the agent knows the target proposition at time t . By contrast, it is by appealing to his modal construal of the notion of safety that Williamson can explain S's ignorance of the fact that S's ticket will lose: roughly put, S's belief in a world w at time t that her ticket is the losing one is, despite being highly likely, not safe, i.e., it is not true in all close possible worlds with respect to world w at time t .

Propositions About the Future

Let us now consider the proposition that I will be alive tomorrow. Dodd's argument here is analogous to the previous one. As with lottery-style propositions, the starting point in Dodd's reasoning is that there is a small chance that the proposition that I shall be alive tomorrow is false. From this, he infers that, if $E=K$ is true, then I cannot know that I will be alive tomorrow. Once again, Dodd is reversing the order of explanation. Now, let us assume that Williamson denies that we know that we will be alive tomorrow. Even if Williamson might grant that we cannot know that we will be alive tomorrow, he would make sense of our ignorance in this case by pointing out that the target belief does not meet the safety requirement for knowledge. It is only once we have established that we cannot have knowledge of the target proposition that, given $E=K$, we can conclude that the evidential probability of the target proposition remains less than 1.⁴⁴

Two points are worth stressing here. First, note that Dodd's reliance on the ambiguity of the notion of chance becomes particularly apparent when he explains why there is indeed a possibility that I will not be alive tomorrow. With this respect, Dodd points out the following two facts:

“Statistical facts about murder rates and automobile fatalities show that there is a small chance that I shall be killed in the near future.” (Dodd, 2007: 645)

“It is just a fact that my beliefs about what is going on to be the case in the future are *based on* information I have right now. And this information does not entail that I shall be alive next year, or that I shall not be able to afford an African safari.” (Dodd, 2007: 645. Italics are mine)

I have already pointed out that Dodd explicitly defines the underdetermination the-

⁴⁴As I will point out in a bit, the details of the case under consideration will matter in determining whether a target proposition is safe or not.

sis in terms of evidential probability⁴⁵. Crucially, in the first quotation, Dodd is appealing to *statistical* facts in order to explain *why* there is a chance that I shall not be alive tomorrow. By relying on an ambiguity in the notion of ‘chance’, he seems to be shifting to a notion of underdetermination that is different from the one he explicitly addresses in his paper. In fact, he seems to be appealing to a notion of underdetermination that is much closer to what I have here called Objective-Chance underdetermination. But, as I have shown above, Objective-Chance Underdetermination is perfectly compatible with the varieties of infallibilism E=K is committed to. In the second passage I have quoted, instead, we can see that Dodd is cashing out the idea that my belief that I shall be alive tomorrow is underdetermined by appealing to the fact that the statistical data about murders are part of my information, and this is what constitutes the *rational basis* on which I form my belief. This rational basis, however, does not entail that I will be alive tomorrow. Again, it is easy to see that Dodd is now moving far from his initial understanding of the underdetermination thesis, namely, one that is cashed out in terms of evidential probability. Instead, the formulation of the underdetermination thesis he is appealing to in the second quotation represents an instance of what I have defined as the ‘Doxastic Entailing Underdetermination Thesis’. Crucially, as I have shown in the first half of this Chapter, E=K, and the varieties of infallibilism E=K entails, are compatible with the Doxastic Underdetermination Thesis.

Second, from our ignorance of the proposition that my lottery ticket will lose, as well as the proposition that I will be alive tomorrow, Dodd generalises to all propositions about the future, thereby concluding that, at the end of the day, Williamson is a sceptic about almost all propositions about the future. However, making such a generalisation amounts to treating any proposition about the future as epistemically equal. It should be clear at this point why this is a crucial mistake. If we want to evaluate the overall epistemic status of a target proposition – whether it can be known – we need to ask whether the belief in the target proposition is safe from being mistaken. Crucially, given the modal construal of the notion of safety (one that does not cash out safety in terms of evidential probability), we will answer this question differently depending on which specific proposition about the future we are considering. For instance, while my belief that I will be alive tomorrow might fail to count as safe in w at t (for the proposition that I will be alive tomorrow is not true in all close possible worlds relevant to w at t), my belief that an alien will not abduct me might count as safe. Hence, while

⁴⁵In fact, as I have pointed out, this is exactly why the underdetermination thesis trivially leads to skepticism when matched with PII.

Williamson might be a sceptic about the proposition that I will be alive tomorrow, he is not necessarily a sceptic about *every* proposition about the future.

Propositions about the Present

Finally, Dodd argues that the case from E=K to scepticism can be generalised to propositions about the present that are not the content of any current perceptual experience. Consider a scenario in which you have just looked at your books on the shelves. Assume now that you are in a different room and you consider the proposition that there are still books on the shelves. Do you know that there are still books on the shelves? According to Dodd, if E=K is true, the answer to this question is negative. This is because, he says, according to some interpretation of quantum mechanics, there is some very small non-zero chance that the particles forming the books have just rearranged themselves in such a way that they do not count as books anymore⁴⁶. Again, Dodd's reasoning here is analogous to the previous two cases. The idea is that there is a small non-zero chance that the proposition that there are books on the shelves is false. That is, there is a sense in which my belief that there are books on the shelves is underdetermined by my evidence, where Dodd cashes out the notion of underdetermination in terms of evidential probability. Crucially, Dodd says, if E=K is true, then one's knowledge that there are books on the shelves is incompatible with there being a small non-zero chance that there are no books on the shelves. At this point, it should be clear to the reader where Dodd's argument goes wrong: once again, it reverses the order of explanation. Whether we *can* or *cannot* know that there are books on the shelves does not depend on what the probability of this proposition on our evidence is. Once again, by appealing to a modal construal of the related notions of safety and risk, Williamson can easily explain our knowledge of iii): S knows that there are books on the shelves in *w* at time *t* (assuming the proposition is true) because the proposition is true in all close possible worlds relative to the world *w* at *t*.

⁴⁶Dodd readapts this example from Hawthorne's *Knowledge and The Lotteries* (Hawthorne, 2003: 4-5).

3.5.1 Why we know propositions about the future and propositions about the (unperceived) present although they are underdetermined!

In the first half of this Chapter, I have put forward a taxonomy of varieties of infallibilism and underdetermination theses. I have argued that, once varieties of infallibilism and underdetermination theses are on the table, we should seriously consider mismatches of infallibilism and underdetermination theses. This enabled me to show that $E=K$ is indeed committed to a certain kind of infallibilism that I have called ‘Propositional Infallibilism’, and that this kind of infallibilism does not lead to scepticism when matched with non-corresponding Doxastic Underdetermination. Williamson can thus easily resist Dodd’s dilemma.

Furthermore, in the previous section, I have addressed the specific propositions Dodd considers when arguing that Williamson is a sceptic, and I have pointed out where Dodd’s argument goes wrong: by relying on an ambiguity in the notion of ‘chance’, he uses the notion of evidential probability as what determines whether we *can* or *cannot* know a given proposition. By contrast, I have stressed that the evidential probability of a proposition is an indicator of whether, as things stand, we do or do not know the target proposition. Once again, $E=K$ (and its commitment to Propositional Probability 1 Infallibilism) does not have by itself any sceptical implications.

However, I still have to provide an explanation of how Williamson can avoid Dodd’s dilemma with respect to the specific propositions Dodd considers, namely, that i) my lottery ticket will lose, that ii) I will be alive tomorrow, and that iii) there are books on the shelves.

I have argued that, for reasons that are independent of $E=K$, Williamson would deny knowledge of proposition i). A consequence of my ignorance of i) is that the probability of i) on my evidence remains indeed less than 1. It follows that there is no pressure to resist the dilemma anymore: we can happily embrace the ‘sceptical’ horn of the dilemma!

“[H]ow can it be anything but *painfully* obvious that I do not know that [I shall be able to go on an African safari next week] [...]? This is a good question.” Littlejohn says (2008: 684). As I have stressed above, whether Williamson allows for knowledge of propositions about the future depends on the details of the case and on the proposition in question. So, assuming we are talking about me, a philosophy PhD student, I can take the proposition that I shall be able to go on an African safari next week as

being epistemically on pair with a lottery proposition. Both propositions are not safe propositions, and we should thus accept that I in fact do not know that I will be able to go on an African safari next week.

But what about the proposition Dodd considers: ii) that I will be alive tomorrow? Depending on the details of the case, Williamson might deny knowledge of ii), thereby embracing the sceptical horn of the dilemma again! But what if the case is actually such that the proposition that ii) I will be alive tomorrow counts as safe on Williamson's framework? That is, what if Williamson can account for knowledge of ii)? Given $E=K$, when I know that ii) I will be alive tomorrow, it follows that the probability of ii) on my evidence is 1. We should thus take the dilemma seriously in these cases, and we should try to resist it. In order to successfully resist the dilemma, we have to explain how it is possible to have knowledge that ii) I will be alive tomorrow, without rejecting the overwhelmingly plausible thesis that there is a sense in which ii) is underdetermined by our evidence. After all, it just might be the case that I will have an unexpected heart attack tomorrow. Fortunately, appealing to the varieties of underdetermination theses I have developed enables us to resist the dilemma successfully. Besides being compatible with Doxastic varieties of Underdetermination, i.e., Doxastic Probability 1 Underdetermination and Doxastic Entailing Underdetermination, knowledge of propositions about the future is also compatible with both Objective-Chance and Appearance-Entailing Underdetermination.

Let us now move on to proposition iii). As explained above, the proposition iii) will count as safe on Williamson's view. Williamson can thus account for knowledge of iii). Again, we should take the dilemma seriously here. Given $E=K$, when I know that iii) there are books on the shelves, it follows that the probability of iii) on my evidence is 1. And yet, it might be the case that the particles of my book unfolded themselves in some weird ways so to make iii) false! Again, we can appeal to my taxonomy of varieties of underdetermination theses. In particular, we want to see if there are varieties of underdetermination theses that are compatible with $E=K$, *besides* the Doxastic varieties of Infallibilism. One way to go could be to appeal to the Objective-Chance Underdetermination thesis. This might allow us to say that, although the evidential probability of iii) for S at time t is 1, there is a non-zero objective chance at time t that iii) is false at t . Crucially, it is easy to see why this is not going to be a promising strategy. Remember that, by appealing to iii), Dodd is considering a proposition about the present. That is, the scenario is one in which S believes that there are books on the shelves at time t , whilst it might be that the particles of the books unfolded themselves

at t in such a way that they do not count as books anymore. What is important here is that the moment in which S forms her belief and the moment in which the particles of the book might unfold themselves in weird ways are the same. If this is so, then we have two possible scenarios. A first scenario is one in which S believes that there are books on the shelves in world w at time t , while the books indeed remain books in w at time t . This entails that $Ch_{tw}(iii) = 1$. A second scenario is one in which S believes that there are books on the shelves in world v at time t , while the books do indeed unfold themselves in a weird way so as to stop counting as books in v at time t . This entails that iii is false. Given factivity of knowledge, S would *not* know that there are books on the shelves in v at t . Note that, insofar as I wish to explain, *contra* Dodd, that *knowledge* of iii) is compatible with both E=K and a plausible underdetermination thesis, I am here concerned with the former scenario only. Crucially, if I assume that S knows that iii) at time t in the first scenario, iii) is not underdetermined in the sense expressed by Objective-Chance Underdetermination. For remember that we are here concerned with a true proposition about the present so that $Ch_{tw}(iii) = 1$. If we want to resist the dilemma, we thus need to find a different formulation of the underdetermination thesis. Fortunately, there is a variety of underdetermination thesis that is compatible with E=K as well as with the fact that S knows that iii) in w at t . That is, there is a sense in which iii) is underdetermined by our evidence in w at t , namely, in the sense expressed by what I have called the Appearance-Entailing Underdetermination Thesis:

Appearance-Entailing Underdetermination (AEU):

For any proposition p , S's non-factive psychological states (i.e., S's impressions, appearances etc.) do not entail p .

Appealing to AEU enables us to resist Dodd's dilemma with respect to the specific proposition iii). On one hand, S knows that there are books on the shelves in w at time t , and this is because – as I have explained – S's belief is safe from being mistaken in all close possible worlds relevant to w . Given E=K, once S knows that there are books on the shelves, S's evidential probability of the proposition that there are books on the shelves is 1. On the other hand, once we understand the underdetermination thesis in the sense expressed by Appearance Entailing Underdetermination, it is possible to account for the intuition that S's belief that there are books on the shelves is underdetermined by S's available evidence, when we restrict S's available evidence to facts about S's non-factive mental states: S's (non-factive) psychological states in w at t do not entail that there are books on the shelves. Putting the matter in another way: S's

non-factive psychological states in w at t , namely, S 's non-factive psychological states in the case in which S knows that there are books on the shelves, are compatible with S being in v at t , namely, with S being in the case in which S does not know that the books are on the shelves because their particles did unfold themselves in some funny way.

3.6 Conclusion

Infallibilism has traditionally been thought of as entailing a variety of scepticism. In this Chapter I have challenged this tradition. To do so, I have focused on Williamson's $E=K$ and Dodd's argument that allegedly shows that $E=K$ is committed to a sceptical version of infallibilism. More precisely, as I have stressed in this Chapter, Dodd presents his argument in the form of a dilemma. Given that $E=K$ entails a version of infallibilism that leads to scepticism when combined with an overwhelmingly plausible underdetermination thesis, Dodd argues that Williamson is forced to give up on something. The first horn of the dilemma involves holding onto both $E=K$ (and the entailed infallibilism) and Underdetermination thesis. This comes at the cost of embracing scepticism. The second horn of the dilemma involves saving $E=K$ while rejecting Underdetermination thesis. This comes at the cost of being too dogmatic. In this Chapter, I have shown that Dodd provides us with a false dilemma. More precisely, developing varieties of infallibilism and varieties of underdetermination theses enabled me to do the following two things: 1) identify the real source of scepticism, and 2) show how the defender of $E=K$ can have it both ways: she can retain both her infallibilist view and underdetermination thesis, while rejecting the sceptical conclusion. With respect to the first point, my taxonomy of infallibilism and underdetermination theses shows that there is nothing intrinsically sceptic about infallibilism. Rather infallibilism leads to scepticism only when combined with a corresponding variety of underdetermination theses. However, once we have clarified what the source of scepticism is, I show that we can combine varieties of infallibilism with non-corresponding varieties of underdetermination theses. This will allow us to retain our favourite infallibilist thesis as well as reject scepticism, while, at the same time, acknowledging that there is a sense in which (most of) our beliefs are underdetermined by our evidence. Focusing on $E=K$, my taxonomy of infallibilism and underdetermination theses allowed me to identify in which sense Williamson's $E=K$ is an infallibilist view, and which variety of underdetermination thesis avoids scepticism, when combined with $E=K$. More precisely, I

have shown that $E=K$ entails what I have called Propositional Entailing Infallibilism and Propositional Probability 1 Infallibilism. I have pointed out that these varieties of infallibilism do not lead to any sceptical result when combined with the following four varieties of underdetermination theses: Doxastic Probability 1 Underdetermination, Doxastic Entailing Underdetermination, Objective Chance Underdetermination, and Appearance-Entailing Underdetermination. Finally, I have reconstructed the details of Dodd's argument with respect to the specific three propositions he considers when arguing that Williamson should be a sceptic. I have pointed out that, besides overlooking varieties of infallibilism and underdetermination theses, his argument relies on an ambiguity of the notion of 'chance' and 'risk' underpinning the idea that our beliefs are underdetermined by our evidence. However, the lesson learnt from Dodd and $E=K$ can be generalised. What this Chapter shows is that we should be careful when talking about infallibilism, its sceptical import, and the problem of underdetermination. There is no univocal way of being an infallibilist (although, as my taxonomy shows, some ways are more plausible than others). Similarly, I believe there are different ways in which we can account for the intuition that (most of) our beliefs are underdetermined by our evidence. We should thus take seriously the possibility of having varieties of underdetermination theses. The main moral we can draw from this Chapter is thus the following: whether one thinks infallibilism is plausible or not, we should just stop blaming infallibilism for being the source of sceptical troubles.

Conclusion of Part II

In Chapter 3, I have considered the *Infallibility Problem* for $E=K$. That is, I have asked whether $E=K$ entails some sceptical variety of Infallibilism. By developing varieties of infallibilism and underdetermination theses, I have argued that $E=K$ entails different varieties of infallibilism that, when combined with at least four non-corresponding varieties of Underdetermination Thesis, do not lead to scepticism. I have thus focused on $E=K$ without considering Epistemological Disjunctivism. As mentioned previously in this Thesis, Epistemological Disjunctivism is a view that concerns paradigmatic cases of perceptual knowledge only, and yet epistemologists seem to take the Infallibility Problem more seriously when it comes to non-perceptual cases of knowledge (see Dodd (2007)). However, at this point one might nevertheless wonder whether Pritchard's Epistemological Disjunctivism entails any of the varieties of infallibilism I have formulated in Chapter 3 and, if so, whether these varieties are compatible with some Underdetermination Thesis. In order to answer these questions, let us recall what Epistemological Disjunctivism amounts to:

In paradigmatic cases of perceptual knowledge, an agent, S , has perceptual knowledge that P in virtue of being in possession of rational support, R [i.e., S 's seeing that], for her belief that P which is both factive (i.e., R 's obtaining entails P) and reflectively accessible. (Pritchard (2012): 13).

As I have argued in Chapter 2, the most plausible interpretation of the notion of reflective access is one which accounts for the fact that S knows that p in virtue of basing one's belief that p on one's rational support R . In other words, S has perceptual knowledge that p insofar as the reason *why* S believes that p is R . It thus follows that, on Epistemological Disjunctivism, what kind of reason one bases one's belief on are relevant for determining whether one has (perceptual) knowledge. Following the taxonomy of infallibilism I have developed in Chapter 3, we can define Epistemological Disjunctivism as an instance of, what I have called, Doxastic Entailing Infallibilism, namely, the thesis on which if S knows that p and believes that p on the basis of evidence e , then e entails p . However, as I have explained in Chapter 3, any view that is committed to Doxastic Entailing Infallibilism will also be committed to Doxastic Probability 1 Infallibilism. Furthermore, remember that, as I have pointed out in Chapter 3, Doxastic varieties of Infallibilism are the strongest versions of infallibilism on my taxonomy, and they entail all Propositional Varieties of Infallibilism. It turns out that Epistemological Disjunctivism is also committed to those varieties of

infallibilism. I have said above that, given that ED is restricted to paradigmatic cases of perceptual knowledge, we might be willing to accept that there is no available underdetermination thesis that, when combined with these kinds of infallibilism, avoid scepticism (that is we might just reject the idea that our beliefs in paradigmatic cases of perceptual knowledge are underdetermined at all). Nevertheless, the reader might still wonder if, in principle, there is a version of underdetermination that is compatible with ED. Once again, we can appeal to the varieties of underdetermination theses that I have developed in Chapter 3. Given that Disjunctivism is concerned with paradigmatic cases of perceptual knowledge, we can plausibly restrict the scope of Epistemological Disjunctivism to propositions about the present. For this reason, we cannot appeal to Objective Chance Underdetermination to account for the fact that there is a sense in which, although S knows that p at time t, S's belief that p is underdetermined. For if S knows that p at time t, then p is true, thereby making Objective Chance of p in world w at time t = 1. However, we can still say that S's belief that p is underdetermined in the sense expressed by Appearance-Entailing Underdetermination.

A final consideration is in order before concluding this Part. Note that, as mentioned in Chapter 2, Pritchard takes one of the advantages of his ED to be that it resists the sceptical paradox. In particular, he argues that his ED is able to resist the Underdetermination-based Sceptical Argument. While appealing to my taxonomy shows that there ED entails some versions of infallibilism that are, nevertheless, logically compatible with one variety of underdetermination theses, I think the question of whether Pritchard's ED is able to resist the radical sceptical paradox is still to be answered. However, this is something I will not address here⁴⁷

⁴⁷It is also important to point out that there is a distinction to be made between *local* and *global* sceptical challenge. Pritchard offers two different solutions to these problems. See for instance Pritchard, 2012 and Pritchard, 2016a. For a criticism to Pritchard's solution to global scepticism see Zalabardo, 2015. For an interesting recent criticism to Pritchard's solution to *local* scepticism see Smith, 2016b.

Part III

Evidential Externalism and Epistemic Justification

Introduction

I have considered Evidential Internalism as the thesis on which one's evidence supervenes on one's non-factive mental states. I have taken Evidential Externalism to be merely the denial of Evidential Internalism. As I have stressed in the Introduction of this thesis, Evidential Externalism, so defined, is silent on various issues concerning the notion of evidence. For instance, it does not specify what it takes for something to be evidence, what it takes for evidence to support a proposition, and what it takes for someone to possess evidence. Furthermore, Evidential Externalism defined as the denial of the supervenience thesis between evidence and non-factive mental states, remains silent on how we should characterise the notion of epistemic justification. However, although Evidential Externalism, as well as Evidential Internalism, *per se* do not entail anything about epistemic justification, I believe it is nevertheless plausible to consider Evidential Externalism and Evidential Internalism as two ways in which one can embrace an Evidentialist theories of epistemic justification. Following Conee and Feldman (2004), I will understand Evidentialism as the thesis on which one's justification supervenes on one's evidence. This thesis is also very broad and leaves various questions about evidence and the relation between evidence and justification unanswered. For all Evidentialism says is that the notions of evidence and justification are strictly intertwined, yet it does not specify how these notions are related to each other. Part III addressed the following question: how does Evidential Externalism understand the relation between evidence and epistemic justification? In other words, how can we combine Evidential Externalism with Evidentialism about epistemic justification?

This section is constituted of Chapter 4 and a general conclusion. In chapter 4, I consider two varieties of evidentialism, what I call, *Evidence-first Evidentialism* and *Knowledge-first Evidentialism*. I take Pritchard's Epistemological Disjunctivism (understood in evidentialist terms) to be a version of the former. I take Williamson's account of evidence and justification to be a version of the latter. After defining these evidentialist views, and after clarifying to what extent Knowledge-first Evidentialism differs from the more orthodox Evidence-first Evidentialism, I argue for novel evidentialist theory: Ecumenical Evidentialism. I show that my Ecumenical Evidentialism vindicates some of the intuitions underlying Evidence-first Evidentialism while remaining in the spirit of the Knowledge-first program.

Chapter 4

Evidence-first, Knowledge-first, and Ecumenical Evidentialism

4.1 Introduction

Let *Evidentialism* be the thesis on which one's justification supervenes on one's evidence (Conee & Feldman, 2004). On Evidentialism, the notions of evidence and justification are thus strictly intertwined. However, as defined, Evidentialism remains silent on lots of questions concerning the nature of evidence and justification, and how, precisely, these notions are related. It is important to point out that what Conee and Feldman have in mind when formulating Evidentialism is *propositional* justification (as opposed to *doxastic* justification), namely, evaluation of a *proposition* relative to a subject's evidence (as opposed to an evaluation of a *belief* relative to a subject's evidence). Nevertheless, I take it that a satisfying Evidentialist theory, namely, a theory that wants to cash out epistemic justification in terms of evidence, should provide us with an answer to both the following two questions:

- i) What is the relation between evidence and propositional justification?
- ii) What is the relation between evidence and doxastic justification?

Evidentialism, as defined, does not answer these questions. In this Chapter, I will take this to be a virtue of Evidentialism, for, depending on how we answer the foregoing questions, we will have varieties of Evidentialism. In particular, I will devote most of my attention to the second of the two questions above. This is because, as I will show later, while there is more agreement, amongst evidentialists, on how to spell out propositional justification in terms of evidence, there is general disagreement on how to answer the question about doxastic justification. An orthodox way of

answering the second question is to take doxastic justification to be a matter of appropriately responding to the evidence one has, thereby vindicating what I will call the ‘Responsiveness Intuition’. Call this view ‘Evidence-first Evidentialism’. However, I believe there’s another way of understanding Evidentialism, one that is in line with Williamson’s Knowledge-first Epistemology. As I will explain later, on Williamson’s view, doxastic justification is just a matter of complying with the Knowledge Norm of Belief (KNB), so that one is (doxastically) justified in believing that p iff one knows that p . By reducing justification to knowledge, this view doesn’t seem to vindicate the Responsiveness Intuition¹. Nevertheless, given Williamson’s E=K, his view of epistemic justification is also an instance of Evidentialism. Call this ‘Knowledge-first Evidentialism’.

The aim of this Chapter is two-fold. It sheds light on how much and to what extent Knowledge-first Evidentialism differs from Evidence-first Evidentialism. It explores the prospects of a novel evidentialist view, one that is able to account for the so-called ‘Responsiveness Intuition’, while remaining in the spirit of Knowledge-first Epistemology. I shall call this ‘Ecumenical Evidentialism’².

This Chapter is structured as follows. In Section 4.2, I clarify the main commitments of Evidence-first Evidentialism. I then direct my attention to what I call Knowledge-first Evidentialism in Section 4.3. In Section 4.4, I provide a careful diagnosis of how much and to what extent Knowledge-first Evidentialism differs from Evidence-first Evidentialism. Finally, in Section 4.5, I develop the main theoretical commitments of my Ecumenical Evidentialism. I conclude in Section 4.6 with some remarks about possible future developments of Ecumenical Evidentialism.

4.2 Evidence-first Evidentialism

According to Evidentialism, the notions of evidence and justification are strictly intertwined. Evidence-first Evidentialism and Knowledge-first Evidentialism represent different ways of understanding this relation between evidence and justification. As I mentioned in the introduction, one aim of this chapter is to shed light on just how much and to what extent these evidentialist views differ from each other. In order to do so, however, we first need to have a better sense of what I mean by Evidence-first

¹I will point out that this is exactly what we should expect when we appreciate how radical the Knowledge-first Project is.

²Thanks for the suggestion of the name to Rachel Fraser.

Evidentialism and Knowledge-first Evidentialism. In this Section, I will clarify what I take to be the main commitments of Evidence-first Evidentialism.

Clarifying the relation between evidence and justification is not an easy task, especially given the disagreement over both the notion of evidence as well as the notion of justification. I will proceed as thus: first, I will look at the important distinction between propositional and doxastic justification and how this is traditionally understood. Second, after discussing the roles we expect evidence to play, I will suggest a way of understanding the relation between evidence and (propositional and doxastic) justification in a way that captures the orthodox intuition underpinning Evidence-first Evidentialism. Before proceeding, a disclaimer is in order. In what follows I will just assume that Evidentialism as stated above is the right view about epistemic justification. Furthermore, I will use the terms ‘reason’ and ‘evidence’ interchangeably. However, the reader who believes Evidentialism is wrong³ might still find the discussion in this chapter useful both because it clarifies a longstanding debate on evidence and justification by framing a new opposition, i.e., that one between Evidence-first and Knowledge-first Evidentialism, and because it explores a novel possibility in the debate that could bring together the two opposing views, i.e., Ecumenical Evidentialism.

4.2.1 Propositional and Doxastic Justification

As I already mentioned in the Introduction of this Chapter, I take propositional justification to be, roughly put, an evaluation of a proposition relative to a subject’s evidence. And I take doxastic justification to be, roughly put, an evaluation of a belief. I have also mentioned that a satisfying Evidentialist account is one that answers both the following questions:

- i) What is the relation between evidence and propositional justification?
- ii) What is the relation between evidence and doxastic justification?

Here is, in a nutshell, how I take propositional and doxastic justification to be traditionally understood. On one hand, one has *propositional* justification for believing a proposition p when, roughly put, one has good evidence for believing that p , where ‘good evidence’ is evidence that provides adequate epistemic support to p . On the other hand, one is *doxastically* justified in believing that p , when one justifiably believes that

³As mentioned in the Introduction of this thesis, Reliabilism represents a paradigmatic example of a non-evidentialist view about justification. However, see Comesana, 2010 for a recent discussion on how to combine reliabilism with evidentialism. For a more recent argument against evidentialism, one that does not rely on Reliabilism, see Littlejohn, 2018a

p , namely, one is believing that p for – roughly speaking – the right kind of reasons (i.e., in virtue of the evidence that makes the proposition p justified). Whether one has propositional justification for believing something thus depends on what counts as evidence and on how we spell out the evidential support relation. One way to understand this evidential support relation is in probabilistic terms. For instance, one has propositional justification for believing that p when one’s evidence makes p sufficiently likely. Of course, on this account, we need to set a threshold t that determines when p is sufficiently likely. Although other interpretations of the evidential-support relation have been offered⁴, this probabilistic understanding of the evidence-support relation, and thus its related notion of propositional justification, is widely embraced amongst epistemologists. Doxastic justification, instead, is traditionally understood as having propositional justification and *basing* one’s belief on one’s propositional justification. That is, on this view, one is doxastically justified in believing that p only if one has propositional justification for believing that p , and one bases one’s belief that p on one’s propositional justification (cf. Firth, 1978: 218, Kvanvig & Menzel, 1990, Pollock & Cruz, 1999: 35, Lasonen-Aarnio, 2010: 205-6).⁵ For the purpose of this section, it won’t matter what the right way of spelling out the details about the basing relationship is⁶. Remember that the aim of this section is to shed light on what a general orthodox Evidence-first Evidentialism looks like. Therefore, while assuming Evidentialism about justification, I will cash out propositional and doxastic justification as evaluations of propositions and beliefs respectively, thereby keeping a fairly neutral standpoint on which specific Evidence-first Evidentialist account is the right one.

Let us now go back to the two questions that every satisfying Evidentialist account has to answer:

⁴For instance, Conee and Feldman understand the evidential support relation in terms of inference to the best explanation (see Conee & Feldman, 2008), while Smith has recently offered a ‘normalcy’ account (see Smith, 2016a).

⁵While there is a general consensus (amongst evidentialists) on how to cash out the notion of propositional justification, there is less agreement on how to understand the notion of doxastic justification. For instance, Turri, 2010 has recently challenged this orthodox account by presenting cases in which someone believes that p on the basis of her propositional justification, yet she does so for the wrong reasons. The cases Turri presents are supposed to show that, given that one can base one’s belief on one’s propositional justification improperly, the orthodox account is mistaken. Instead, he argues that we should understand the notion of propositional justification in terms of doxastic justification (and not vice versa). For an interesting response to Turri, see Silva, 2015. Despite the disagreement, and despite the different ways in which one can fill in the details of this account, understanding doxastic justification in terms propositional justification plus basing remains the paradigmatic and orthodox way of understanding doxastic justification.

⁶For an excellent overview on different ways of understanding the notion of ‘basing’ see Korcz, 2010

- i) What is the relation between evidence and propositional justification?
- ii) What is the relation between evidence and doxastic justification?

Evidence-first Evidentialism will answer i) and ii) respectively as follows:

(Propositional Justification) (PJ): PJ is an evaluation of a proposition relative to a subject's evidence, so that a proposition p is propositionally justified for a subject S if and only if S 's evidence stands in a support-relation R towards p .

(Doxastic Justification) (DJ): DJ is an evaluation of a belief, so that S 's belief that p is doxastically justified for S only if S believes that p on the basis of one's evidence E that stands in a support-relation R towards p .

It is worth noticing that, Evidence-first Evidentialism, understood as a package of view embracing PJ and DJ, does not put any constraint on the nature of evidence. It is compatible with an externalist picture of evidence, as well as with evidential internalism. Similarly, it is compatible with the view on which evidence is propositional, as well as with the view on which evidence is constituted by one's mental states or experiences. Furthermore, note that, as stated, DJ is neutral on how to cash out the basing relation, and PJ and DJ are both neutral on how we should understand the evidential-support relation R .

Despite the disagreement on how to fill in the details in PJ and DJ, most epistemologists who embrace Evidentialism will agree on the following commitments (which is why I take this way of understanding doxastic and propositional justification to be the orthodox view):

1. Propositional Justification is an evaluation of a proposition relative to a subject and given the subject's total evidence⁷.
2. Doxastic Justification is an evaluation of a subject's belief relative to the subject's evidence.
3. While being doxastically justified in believing that p entails believing that p , one can have propositional justification for believing that p without believing p .
4. If one is doxastically justified in believing that p , then one has propositional justification for believing that p .

⁷It is important to note that the kind of propositional justification I am interested in here is what is sometimes called 'ordinary' propositional justification rather than 'objective' propositional justification. While the former refers to the evidence that are at one's disposal, the latter concerns the kind of evidence that is beyond the cognitive limits of a subject and that only idealised subjects have. Cf. Melis, 2018: 5.

Given PJ and DJ, Evidence-first Evidentialism comes together with the following two commitments:

PJ-DJ Relation: Propositional Justification is prior to Doxastic Justification and partly explains Doxastic Justification.

E-first: Evidence is prior to justification

There are different ways in which one can embrace Evidence-first Evidentialism.

As I said above, here I am not interested in what the exact way of understanding the notion of ‘basing’ is. Instead, I am interested in understanding what the intuition underlying this orthodox notion of doxastic justification is. More precisely, I’d like to ask: what has led epistemologists to put a ‘basing’ requirement in the definition of doxastic justification? The following scenario will help us seeing what the overall intuition motivating the basing requirement is. Imagine I believe that earth is not flat, and I have tons of scientific and testimony evidence that support my belief that the earth is not flat. However, assume the reason why I believe the earth is not flat is not constituted by all the scientific and testimony evidence I have. Instead, the reason why I believe the earth is not flat is that today it is Sunday and I am convinced that on Sundays the earth is not flat. In this scenario, I have tons of evidence supporting the proposition that the earth is not flat, thereby making it (propositionally) justified, and yet we wouldn’t judge my *belief* as being (doxastically) justified. The intuition driving our judgment in this scenario is that, in order for my belief to be doxastically justified, not only my belief that the earth is not flat has to be supported by my evidence, but I also have to believe that the earth is not flat on the basis of my evidence that supports the proposition that the earth is not flat. As Pollock and Cruz say:

“To be justified in believing something it is not sufficient merely to have a good reason for believing it. One could have a good reason at one’s disposal but never make the connection.”

(Pollock & Cruz, 1999: 35)

Similarly, when arguing that doxastic justification requires a belief to be “well founded”, Conee and Feldman say:

“The notion of a well-founded belief [...] serves [...] to accommodate the intuition that there is something epistemically defective about drawing justified conclusions for bad reasons. Such beliefs are ill-founded, in virtue of not being based on justifying evidence.”

(Conee & Feldman, 2004: 103)

Despite the different ways of spelling out the basing requirement, or whether one prefers to talk about “well-founded” belief, what all the traditional conceptions of doxastic justification seem to have in common is the following Responsiveness Intuition:

(Responsiveness Intuition): Doxastic justification is sensitive to whether one is evidence responsive. That is, doxastic justification, as evaluation of a belief, requires belief to be the result of appropriately taking into consideration one’s evidence in coming to form the target belief.

By understanding doxastic justification in terms of propositional justification plus basing, Evidence-first Evidentialism vindicates the foregoing Responsiveness Intuition. That is, on this view, doxastic justification is a matter of appropriately taking into account one’s evidence when forming one’s belief. That is, justification is a matter of being evidence-responsive.

Before moving on to the next section, let me note that, interestingly, Pritchard’s *Epistemological Disjunctivism* can be thought of as being an instance of Evidence-first Evidentialism in some relevant sense. To see why this is so, let us recall Pritchard’s formulation of Epistemological Disjunctivism:

“In paradigmatic cases of perceptual knowledge, an agent, S, has perceptual knowledge that *p* in virtue of being in possession of rational support, R [i.e., S’s seeing that], for her belief that *p* which is both factive (i.e., R’s obtaining entails *p*) and reflectively accessible.” (Pritchard, 2012: 13).

As I mentioned previously in this Thesis, and as it results clear from the above-mentioned passage, Pritchard does not talk about evidence. However, we can have an evidentialist formulation of Pritchard’s Disjunctivism, one on which S has perceptual knowledge in virtue of having evidence (i.e., S’s *seeing that*) for her belief that *p* that is both factive and reflectively accessible⁸. Assuming, knowledge entails doxastic justification, on this evidentialist formulation of Epistemological Disjunctivism it looks like one has doxastic justification (in paradigmatically good perceptual cases) in virtue of having evidence that one gets access to and in the light of which one thus believes that *p*. In a sense, by requiring one’s evidence to be reflectively accessible, Pritchard’s Epistemological Disjunctivism provides us with an Externalist version of the traditional Evidence-first view which takes doxastic justification as the result of believing that *p* for the right kind of (factive) reasons, thereby vindicating, the Responsiveness Intuition.

⁸This is a shift that Pritchard himself does in Pritchard, 2011b.

4.2.2 The Relation between Roles of Evidence and Justification

Recall the definition of Evidentialism as provided in the Introduction of this chapter: evidentialism is the thesis that one's justification supervenes on one's evidence. As I mentioned in the Introduction, Evidentialism, as stated, is relatively silent on how we should understand the relation between evidence and justification (both propositional and doxastic). In the previous section, I have evaluated how Evidence-first Evidentialism understands this relation. On Evidence-first Evidentialism, evidence provides one with propositional justification *for* believing a proposition p in virtue of appropriately supporting p (e.g., in virtue of making p sufficiently likely). Evidence provides one with doxastic justification *in* believing that p only if one believes that p on the basis of one's evidence that appropriately supports p (e.g., that makes p sufficiently likely). I have pointed out that one of the main features of Evidence-first Evidentialism is that it vindicates the so-called Responsiveness Intuition: (doxastic) justification is a matter of being *evidence-responsive*. In this Section, I suggest another way of understanding the relation between evidence and justification in an Evidence-first fashion, one that appeals to the roles we expect evidence to play.

Using the terminology familiar to the philosophers of reasons and action we can think of evidence as having the following main two roles: justifying and motivating (see Alvarez, 2008, Alvarez, 2009, Alvarez, 2010, Sylvan, 2016). Evidence plays a justifying role insofar as, roughly put, it is what favours having one doxastic attitude over another. Evidence plays a motivating role insofar as, roughly put, it is that in the light of which one believes something. To better see the contrast between the justifying and motivating role of evidence, we can think of an instance of evidence e as having a motivating role for S in believing a proposition p insofar S *takes* e as favouring believing p over not- p , thereby using e as a premise in deliberation for whether to believe that p . Instead, an instance of evidence e has a justifying role insofar as it is what, in fact, *favours* believing that p over not- p , thereby making believing that p a justified belief to have⁹. It is thus in virtue of its justifying role, that we are in fact justified when we appropriately believe in accordance to our evidence. Once again: depending on what view of evidence one subscribes to, and depending on what one

⁹Note that philosophers usually distinguish a third role reasons can have, namely, an explanatory role. A reason r (and thus evidence) plays an explanatory role for a subject S , for it is what explains why a subject has done a certain action, or has come to believe something. I believe Alvarez, 2009 is right in arguing that the explanatory and motivating reason can easily come apart. Here I won't be concerned with the explanatory role of evidence. Instead, I will focus on the motivating role for, as I will show later, I take this to play an important role in whether we should attribute doxastic justification to a subject.

takes the evidential-support relation to be, the details of how evidence plays a justifying rather than a motivating role will differ. However, the details won't matter for the purpose of this Chapter.

Let us go back to the two questions we expect an Evidentialist account to answer

- i) What is the relation between evidence and propositional justification?
- ii) What is the relation between evidence and doxastic justification?

Here is another way in which we can understand the relation between evidence and propositional justification and doxastic justification in an Evidence-first fashion:

PJ - Evidence Relation: Propositional justification is the “manifestation” of the justifying role of evidence.

DJ - Evidence Relation: Doxastic justification is the “manifestation” of both the justifying and the motivating roles of evidence, where these roles are manifested simultaneously and are in alignment, i.e., they are manifested by the same instance (or instances) of evidence.

We can ask: what notion of justification is involved when philosophers talk about the “justifying” role of evidence? Addressing this question will help us to shed light on what the notion of “manifestation” amounts to, and it will enable us to further understand the way in which Evidence-first Evidentialism understands the relation between justification and evidence.

To answer this question let us go back to the definitions of propositional and doxastic justification:

(Propositional Justification) (PJ): PJ is an evaluation of a proposition relative to a subject's evidence, so that a proposition p is propositionally justified for a subject S if and only if S 's evidence stands in a support-relation R towards p .

(Doxastic Justification) (DJ): DJ is an evaluation of a belief, so that one's belief that p is doxastically justified for S only if S believes that p in the light of one's evidence E that stands in a support-relation R towards p .

In the light of the notions of propositional and doxastic justification given here, the most plausible and intuitive way of accounting for the relation between the justifying role of evidence and epistemic justification is to take propositional justification to be the notion of justification underlying the justifying role of evidence. In order to see why this is the case, consider first the relation between the justifying role of evidence and the motivating role of evidence. As mentioned before, an instance of evidence

e plays a motivating role for one insofar as e is the reason in the light of which one believes that p . On the other hand, the justifying role of evidence is that role evidence plays when it supports a target proposition. However, evidence e plays its justifying role regardless of whether one in fact uses e as the reason for believing that p or not. In other words, the justifying and motivating roles of evidence may come apart. It is thus easy to see a similarity between the propositional kind of justification and the justifying role of evidence. As we have seen before, a proposition p is propositionally justified for one, when, given one's evidence, it is justified for one to believe that p . A proposition p can be propositionally justified regardless of whether one believes that p in virtue of p being propositionally justified by one's evidence. In fact, a proposition p can be propositionally justified for one regardless of whether one believes p at all. Doxastic justification, however, requires one to believe that p in the light of the evidential support provided by one's evidence. This suggests that, as it stands, doxastic justification is not a good candidate for the notion of justification that is in play when epistemologists talk about the justifying role of evidence. Our evidence does not provide us with doxastic justification merely in virtue of its justifying role. We thus have an alternative Evidence-first answer to the first question 1): what is the relation between evidence and propositional justification? I take the foregoing considerations to support my suggestion to conceptualise the notion of propositional justification as being the "manifestation", namely, the result, of evidence playing its justifying role.

This leaves us with the second question: what is the relation between evidence and doxastic justification? Bearing in mind what I have said so far about propositional and doxastic justification, and bearing in mind the relation between the justifying and the motivating roles of evidence, it should now be easy to see why understanding doxastic justification as the "manifestation" of both the justifying and motivating roles of evidence, where these roles are manifested simultaneously and by the same instance (or instances) of evidence, accounts for the orthodox way of understanding doxastic justification in evidentialist terms. That is, S's belief that p is doxastically justified as the result of both evidence e playing a justifying role and S "exploiting" the motivating role of evidence, i.e., as the result of S using e as the reason for believing that p . On this orthodox Evidence-first Evidentialist account, doxastic justification is thus a matter of whether one believes correctly given the evidence one has. This aligns with the intuition that, when 'praising' someone for believing accordingly to her evidence, we are not merely claiming that the *propositions* she believes are justified. What we normally care about – or, at least, what most epistemologists seem to care about – when

we evaluate the epistemic status of a subject is whether one appropriately believes in the light of one's evidence. It is in this sense that Evidence-first Evidentialism vindicates the Responsiveness Intuition according to which one is justified if one is appropriately evidence responsive:

(Responsiveness Intuition): Doxastic justification is sensitive to whether one is evidence responsive. That is, doxastic justification, as evaluation of a belief, requires a belief to be the result of appropriately taking into consideration one's evidence in coming to form the target belief.

4.3 Knowledge-first Evidentialism

In the previous section, I have considered 'Evidence-first Evidentialism', and I have elucidated how this variety of evidentialism understands the relation between evidence and justification. In this Section, I consider another way of cashing out Evidentialism, one that is in line with Williamson's Knowledge-first Epistemology. I call it 'Knowledge-first Evidentialism'.

First of all, let me point out the legitimacy of talking about Knowledge-first *Evidentialism* in the first place. It is worth noting that Williamson never explicitly defines himself as an evidentialist. However, he seems to be committed to something in the neighbourhood of Evidentialism in his *Knowledge and its Limits*, for instance when he writes:

"An epistemically justified belief which falls short of knowledge must be epistemically justified by something; whatever justifies it is evidence."
(Williamson, 2000: 208)

Furthermore, in *Knowledge and its Limits*, we can derive Evidentialism, from the following two theses that Williamson simultaneously embraces. First, the thesis that "in any possible situation in which one believes a proposition p , that belief is justified, if at all, by propositions $q_1 \dots q_n$ (usually other than p) which one knows" (2000: 203). Second, the thesis on which all and only propositions one knows are part of one's evidence ($E=K$) (2000, ch.9). I think we can thus safely define Williamson's theory of justification as being an instance of Evidentialism.

However, it might still be difficult to provide a unified Williamsonian account concerning the relation between evidence and justification, given that his view on this topic seems to have slightly changed over the years. For instance, while he seems to allow for justified false beliefs in his *Knowledge and its Limits* (as the above quote

shows), in his latest papers he reserves the status of justification to those beliefs that constitute knowledge (cf. Williamson, 2011, Williamson, in press). However, I think it is possible to define a coherent Knowledge-first Evidentialist view in a way that make sense of both his work in *Knowledge and Its Limits*, and what he argues for in his more recent papers. This is what I am going to do in the coming Section. That is, I want to shed light on how we can be evidentialist in a Knowledge-first way. In what follows, I will evaluate how Knowledge-first Evidentialism understands the relation between evidence and propositional justification on one hand, and the relation between evidence and doxastic justification on the other hand.

4.3.1 Knowledge-first Evidentialism: on Propositional Justification

In order to shed light on what Knowledge-first Evidentialism looks like, we will need to disentangle the main commitments underlying Williamson's Knowledge-first view about evidence and justification, and how and to what extent his Knowledge-first view of justification is informed by his theory of evidence. I will argue that, while his $E=K$ entails something about propositional justification, it is silent on what it takes for someone to be *doxastically* justified. It is by appealing to the Knowledge Norm of Belief (KNB) that Williamson provides us with an account of doxastic justification.

As already discussed in the previous chapters of this thesis, Williamson offers a theory of evidence in terms of knowledge (2000, ch. 9). On his view, one's evidence is all and only the propositions one knows ($E=K$). But what does $E=K$ entail about epistemic justification? More precisely, what does $E=K$ entail about propositional and doxastic justification? In order to have a charitable answer to this question, we need to consider $E=K$ as part of Williamson's broader picture of evidence. Call it the Evidence Package:

- i Propositionalism about evidence is true
- ii A proposition e is part of S 's evidence iff S knows that e
- iii A proposition e is evidence for q iff S 's evidence includes e and the probability of q conditional on e is higher than the unconditional probability of q .

Let us start with i). Take propositionalism to be the thesis on which, roughly put, the nature of evidence is propositional. In a nutshell, Williamson's argument for propositionalism is that only propositions can play the roles evidence is supposed to play,

e.g., in probabilistic reasoning, inference to the best explanation, and in ruling-out hypothesis (2000: ch.9)¹⁰. I will not rehearse Williamson’s argument for propositionalism here. I will assume that propositionalism about evidence is true, given that the aim of this section is merely clarificatory, and given that the nature of evidence does not have a straightforward implication on what DJ and PJ amount to. Let us move onto ii). If i) entails that something can work as evidence only if it is a proposition, commitment ii) specifies *which* propositions can work as evidence: all and only those propositions one knows can be evidence for one¹¹. Commitment iii) puts a further constraint on what counts as evidence. In particular, while commitment ii) is concerned with the *nature* of evidence, commitment iii) concerns the *evidence-for* relation, namely, what it takes for some instance of evidence *e* to provide evidential support to a proposition *p* (Williamson, 2000: 187-189).

We can see that although i), ii), and iii), by themselves, do not entail anything about epistemic justification, when we take these commitments together with the evidentialist claim that anything that is justified is justified by evidence (Williamson, 2000: 208), then we have an initial evidentialist picture of justification. But, we can ask, what kind of justification? I believe we should take the commitments of the Evidence Package to provide us with an evidentialist account of propositional justification, while remaining relatively¹² silent on what it takes for someone to be doxastically justified. In order to see why this is so, recall the core difference between propositional and doxastic justification: while propositional justification is an evaluation of a proposition, doxastic justification is an evaluation of a belief. Crucially, note that i) to iii), taken together, are either concerned with what counts as a *justifier* of a proposition (i, ii), or with how we should spell out the evidential-support relation R between the justifier and a target proposition (iii). None of i) - iii) involves an evaluation of a belief. Therefore, if i) to iii) entail anything at all about justification, they must entail something about propositional justification. In fact, by considering the evidence-for relation, and taken together with the evidentialist assumption on which “what justifies is evidence”, commitment iii) can be read as a way of spelling out the kind of evidential-support relation in virtue of which a proposition is justified. More precisely, it provides a probabilistic

¹⁰For an interesting argument in favour of propositionalism but against Williamson’s argument see Neta, 2008.

¹¹Note that i) clearly follows from ii). However, I think it is useful to make propositionalism explicit as one of the main commitments of Williamson’s theory of evidence.

¹²Of course the ‘Evidence Package’ is not completely silent on doxastic justification. In particular, E=K is going to be related to the notion of doxastic justification given that, as I will discuss in a bit, Knowledge-first Evidentialism equates doxastic justification to knowledge.

interpretation of such a relation. Assuming, further, that the evidence has to increase the probability of a target proposition over a specific threshold t , we have a first evidentialist account of propositional justification. The first conclusion we can draw here is that Williamson's account of evidence, and, more broadly, his Evidence Package, entails something about propositional justification, while remaining silent on what it takes for someone to be doxastically justified.

However, far from being uninterested in the notion of doxastic justification, Williamson has a way of accounting for when and why one's beliefs deserve the status of doxastic justification. In the next Section, I will evaluate his Evidentialist Account of Doxastic Justification. Once I have clarified the main commitments of Knowledge-first Evidentialism, I will evaluate to what extent Knowledge-first Evidentialism differs from Evidence-first Evidentialism.

4.3.2 Knowledge-first Evidentialism: on Doxastic Justification

In the previous Section, I have shown that the main commitments of Williamson's account of evidence (what I have called 'The Evidence Package') entails something about propositional justification while remaining virtually silent on what it takes for someone to be doxastically justified. Nevertheless, it would be uncharitable to conclude that Williamson is just not at all concerned with the notion of doxastic justification. Furthermore, while in his *Knowledge and Its Limits* he seemed to allow for justified beliefs that are not knowledge, in his more recent work he seems to reserve the notion of justification only to known propositions. To make sense of Knowledge-first Evidentialism in a coherent way, I think we should understand Williamson's words here (namely, when he equates justification and knowledge) as referring to doxastic justification. Let us see how.

In his [forthcoming], Williamson argues that justification is the result of complying with the following norm of belief: one ought to believe that p only if one knows that p . Call this the Knowledge Norm of Belief (henceforth, KNB). So one is justified in believing that p iff one knows that p . How does he get there? Williamson is not the only one who has argued for the thesis that being justified in believing that p is a matter of knowing that p (cf. Sutton, 2005, and Littlejohn, 2018b), but he is the only one, as far as I know, who argues for the KNB from considerations that are explicitly independent of epistemology. The idea in his paper is to "take a step back from the epistemological issue to make some [...] general normative distinctions, then return

with them to epistemology” (Williamson, forthcoming, *p.* 3). Consider, as Williamson does, the case of promise-keeping. If I promise to be in Italy for my dad’s birthday, the norm of keeping my promise has specific compliance conditions, namely, to be in Italy for my dad’s birthday. It follows that if I try my best to be in Italy for my dad’s birthday but my flight gets cancelled and I am stuck at the airport, I have failed to keep my promise. I fulfil my promise if and only if I end up being in Italy for my dad’s birthday. To put it in Williamson’s words:

“In general, if I promise to ϕ , then I keep my promise if and only if I ϕ , and I comply with the norm of promise-keeping with respect to that promise if and only if I do keep it.”

(Williamson, in press, *p.* 10)

However, any given norm, on Williamson’s view, generates other derivative norms. For any primary norm N we thus have, at least, a secondary and a tertiary norm, what he calls ‘DN’ and ‘ODN’ respectively. One complies with a secondary norm DN only if one has a general disposition to comply with the relevant primary norm N (D stands for disposition). The tertiary norm ODN, instead, requires one to do whatever someone who has a general disposition to comply with the primary norm N would do in the given circumstances (O stands for occurrent). Importantly, however, the secondary and tertiary norms do not inherit the full normative status of the primary norm. While the primary norm N has full normative strength, the secondary and tertiary norms are, so to speak, normatively weaker. Consider again the case of promise-keeping. Imagine I promise I will be in Italy for my dad’s birthday, but I am a compulsive liar and I am not that kind of person who makes a promise and keeps it. I have thus no intention whatsoever to end up in Italy for my dad’s birthday. Assume now that, unexpectedly, it turns out I really have to be in Italy for an important doctor’s appointment, and this turns out to be the same day as my dad’s birthday. I thus end up being in Italy for my dad’s birthday. What Williamson’s normative framework predicts in this case is that, although I fail to comply with the secondary norm DN, I have nevertheless complied with the relevant primary norm N of promise-keeping: I have kept my promise. This framework has important and interesting consequences when applied to the epistemological case. For, once we take this general normative framework seriously, and we go back to do epistemology, we have a framework on which we “restrict the term ‘justification’ to compliance with the relevant primary norm” (Williamson, forthcoming, *p.*10). On this view, justification is not something we can get from complying with DN or ODN, for their normative statuses are merely derivative from the primary norm N .

What we can get from complying with DN and ODN is, at most, an excuse¹³. Justification requires compliance with the primary norm of belief KNB¹⁴. At the same time, complying with the relevant primary norm of belief is sufficient for being justified. On this view, one is doxastically justified in believing that p if and only if one knows that p ¹⁵.

4.3.3 How to make sense of Knowledge-first Evidentialism

Let us recapitulate the main features of Knowledge-first Evidentialism and how it cashes out the notions of propositional and doxastic justification. On one hand, it seems to understand propositional justification in terms of evidence and the evidential support it offers to a proposition. This view of propositional justification is thus in line with the way in which Evidence-first Evidentialism understands propositional justification:

(Propositional Justification (PJ)): PJ is an evaluation of a proposition, so that a proposition p is propositionally justified for S only if S's evidence stands in a support-relation R towards p .

Similarly, we can understand propositional justification in terms of the justifying role of evidence as follows:

PJ - Evidence Relation: Propositional justification is the manifestation of the justifying role of evidence.

¹³It is worth noting that appealing to the excuse-justification distinction represents the typical externalist move to the New Evil Demon problem, an objection originally raised by Lehrer & Cohen, 1983 and again Cohen, 1984 to reliabilism. The problem goes, roughly, as follows. Consider a subject who's radically deceived by an evil demon and who doesn't know she's radically deceived. In fact, everything looks just as if she were not radically deceived. If Reliabilism about justification is right, then the subject is not justified in believing as she does because, given the evil demon, her beliefs are systematically false. Crucially, the internalist says, she is doing the best she can and she should thus be granted justification. Appealing to the excuse-justification distinction enables the externalist to account for the intuition that the radically deceived subject is doing well: she is not justified but she is nevertheless excused (or some would say blameless). For this line of response see also Littlejohn, in press. For a critique of this response see Gerken, 2011 who argues the excuses-justification distinction is an ad hoc move. I am not going to discuss the New Evil Demon Problem in this Chapter.

¹⁴The idea that Knowledge is the norm of belief already appears in Knowledge and Its Limits where Williamson first defends a Knowledge norm of Assertion and then he takes the belief to be the inner speech of assertion thereby suggesting an analogous knowledge norm of belief (see Williamson, 2000, ch. 10). For other arguments for the Knowledge norm of assertion and knowledge norm of action see Hawthorne, 2003: 30, Hawthorne & Stanley, 2008. For other arguments for Knowledge Norm of Belief see Littlejohn, 2018b. For another defence of how justification is a matter of complying with the relevant norm of belief see Littlejohn, 2018b.

¹⁵It is far from clear why compliance with the primary norm gives one epistemic justification and Williamson does not really provide an argument for that. In fact, as I will mention in the last section of this chapter, I think KNB does not set the standard for epistemic justification.

On the other hand, Knowledge-first Evidentialism understands the notion of doxastic justification in terms of compliance with the knowledge norm of belief:

K-first DJ: DJ is an evaluation of a belief, so that one is doxastically justified in believing that p iff one knows that p .

I've already mentioned at the beginning of Section 4.3 that, although Williamson does not explicitly define himself an evidentialist, his commitment to evidentialism seems to underlie his *Knowledge and its Limits*. Now that we have a better picture of what Williamson's account of doxastic justification amounts to, we can see more clearly why we can legitimately define this view an *evidentialist* view: if one is doxastically justified in believing that p if and only if one knows that p , and p is part of one's evidence if and only if one knows that p , then, on this view, one's justification supervenes on one's total evidence. Therefore, although Williamson's account of doxastic justification is importantly different from the orthodox Evidence-first way of understanding doxastic justification, it is nevertheless an instance of Evidentialism as defined it above.¹⁶

4.4 Evidence-first VS Knowledge-first Evidentialism

At the beginning of this Chapter, I offered a definition of Evidentialism as the thesis on which one's justification supervenes on one's evidence. I have noted that, as it stands, Evidentialism is silent on what evidence is and how we should understand the relation between evidence and epistemic justification, both in its propositional as well as in its doxastic form. I have thus considered the orthodox way of understanding Evidentialism, what I have called, 'Evidence-first Evidentialism'. I have then considered a different way of cashing out Evidentialism, namely, 'Knowledge-first Evidentialism'. On the way I understand Evidence-first Evidentialism, we should cash out the relation between evidence and propositional and doxastic justification in the following way:

(Propositional Justification (PJ)): PJ is an evaluation of a proposition, so that a proposition p is propositionally justified for S only if S's evidence stands in a support-relation R towards p .

(Doxastic Justification DJ): DJ is an evaluation of a belief, so that one's belief that p is doxastically justified for S only if S believes that p on the basis of one's evidence E that stands in a support-relation R towards p .

¹⁶In fact, note that Knowledge-first Evidentialism is even more 'hardcore' evidentialism than the Evidence-first Evidentialism. For remember, on this latter view, doxastic justification supervenes on evidence plus basing.

Alternatively, I have suggested that we can understand propositional and doxastic justification as being strictly intertwined with the justifying and motivating roles of evidence in the following way:

PJ - Evidence Relation: Propositional justification is the manifestation of the justifying role of evidence.

DJ - Evidence Relation: Doxastic justification is the manifestation of both the justifying and the motivating roles of evidence, where these roles are manifested simultaneously and are in alignment, i.e., they are manifested by the same instance (or instances) of evidence.

While Williamson's Knowledge-first Evidentialism understands propositional justification in a similar way, as seen in the previous Section, it embraces a different conception of doxastic justification. That is, Knowledge-first Evidentialism takes doxastic justification to be a matter of compliance with the knowledge norm of belief. I have also stressed that, despite the differences, both Evidence-first and Knowledge-first accounts can legitimately be considered as varieties of Evidentialism.

In this Section, I will examine how much and to what extent Knowledge-first Evidentialism differs from the more orthodox Evidence-first Evidentialism.

Justified False Beliefs

One of the main differences between Evidence-first and Knowledge-first Evidentialism is that Knowledge-first Evidentialism does not allow for (doxastically) justified false beliefs. This is not very surprising; rather it is just what trivially follows once we understand, as Williamson does, doxastic justification in terms of compliance with the knowledge norm of belief. I will thus move onto what I take to be more interesting features of Knowledge-first Evidentialism.

Responsiveness Intuition

Besides the fact that Knowledge-first does not allow for justified false beliefs, I take the core difference between Evidence-first and Knowledge-first Evidentialism to be the way in which they conceptualise the notion of doxastic justification. On one hand, as I've mentioned in the previous section, by understanding doxastic justification in terms of basing one's belief on one's (justifying) evidence, Evidence-first Evidentialism vindicates the intuition on which doxastic justification is a matter of whether one is appropriately responding to the evidence one has (Responsiveness Intuition). On

the other hand, by understanding doxastic justification in terms of compliance with the KNB, on Knowledge-first Evidentialism, whether one is doxastically justified in believing that p is not strictly related to whatever evidence one takes into consideration when coming to believe that p . In a nutshell, Knowledge-first Evidentialism does not capture the intuition that takes doxastic justification to be a matter of being *evidence-responsive* in the sense specified in the Responsiveness Intuition^{17, 18}.

Responsiveness Intuition and Epistemic Evaluations

Once we understand doxastic justification in terms of compliance with the Knowledge Norm of Belief, the Responsiveness Intuition seems to play no interesting function in our epistemological evaluations. In other words: if the notion of doxastic justification is separated from any basing requirement, namely, if the notion of doxastic justification is unrelated to what evidence one takes into consideration when forming the target belief, then the Responsiveness Intuition ends up playing no role when we evaluate one's doxastic attitude as being epistemically good. And yet, one might object that, when we evaluate one's doxastic attitude as being epistemically good or bad, epistemologists seem to be interested in whether one is appropriately responding to the evidence one has.

I can see two ways in which a defender of the KNB account of doxastic justification can respond here. One could either bite the bullet, thereby saying that, in fact, we should not care much about whether one is appropriately evidence-responsive when evaluating whether one's doxastic attitude is 'epistemically good'. That is, the Responsiveness Intuition should play no role in our epistemic evaluations. Alternatively, one could try to 'save' the Responsiveness Intuition by linking it to some other relevant epistemic concepts. I'm going to consider the latter response only.

One could point out that we can still appeal to the Responsiveness Intuition when making our epistemic evaluation, namely, when judging someone as being epistemically 'successful' or 'defective'. For instance, one could suggest that we can appeal to

¹⁷For those who prefer talking about justification in terms of the roles of evidence, it will be interesting to note that one consequence of taking doxastic justification to be a matter of complying with the KNB is that, by doing so, Knowledge-first Evidentialism separates the notion of doxastic justification from the motivating role of evidence.

¹⁸Brueckner has a similar concern when it comes to cases of non-inferential knowledge. As I will explain later, on a knowledge-first account of non-inferential knowledge one can have knowledge without having any previous evidence. As I will argue in Section 4.5, I do not think this is a problematic aspect of Knowledge-first Epistemology. On the contrary, I think this is something we should embrace as a nice feature of the view. See Brueckner, 2005b, Brueckner, 2009 and Williamson's response in Williamson, 2009b: pp.282-85.

the Responsiveness Intuition when evaluating whether one is personally justified in believing something. While propositional justification is an evaluation of a proposition, and doxastic justification is an evaluation of a belief, personal justification is generally taken to be an evaluation of the *subject*. Although the notions of doxastic and personal justification have been sometimes used interchangeably (see, for instance, Kvanvig & Menzel, 1990), there are good reasons to think these are two distinct notions. While doxastic justification has to do with whether one's *belief* that *p* is epistemically justified, personal justification has to do with whether the *subject* is justified in believing as she does, given the subject's perspective. Following Catherine Lowy, we can put things in terms of what we reasonably expect from a subject to do in a given situation with respect to figuring out whether one's believing in an epistemically appropriate way (e.g., one's belief is true, one's belief fits the evidence etc.) (see Lowy, 1978). A *subject* is thus personally justified if and only if she has done all we can expect her to do given the subject's perspective. Doxastic and personal justification can thus come apart, for one might do all one can given the circumstance, and yet nevertheless fails to be doxastically justified (which, within a Knowledge-first Evidentialist framework means that one fails to have knowledge). I think it is thus plausible to say that, while doxastic justification is sensitive to epistemic factors only, the standard for personal justification might change depending on epistemic as well as practical factors (see also Engel, 1992: 40).

One can thus think of 'saving' the Responsiveness Intuition in our epistemic evaluations by relating it to the notion of personal justification. On this view, whether one is personally justified has to do with what evidence one takes into consideration when believing something. A Knowledge-first Evidentialist could thus suggest the following evidentialist account of personal justification:

Personal Justification: One is personally justified in believing that *p* only if one believes that *p* on the basis on what one *takes* to be evidence that stands in an appropriate evidential-support relation R towards *p*.

We can see how personal and doxastic justification come apart on this account. For while whether one is doxastically justified depends on whether one's evidence stands in an appropriate evidential support relation toward the target believed proposition, personal justification embodies a more internalist vibe on which all that matters for the subject to be personally justified is whether, on the subject's perspective, the target believed proposition is well supported by what she *takes* to be one's evidence (regardless of whether this is actually the case). On this view, one might thus be personally

justified in believing something without actually being either propositionally or doxastically justified.

To see how this would be compatible within Knowledge-first Evidentialism, remember that, on the kind of Knowledge-first Evidentialism I am considering, a proposition p is propositionally justified for someone when one's evidence raises the probability of p over a certain threshold. Assume t to be whatever threshold we take to be relevant for assigning propositional justification. While doxastic justification, on his view, would thus be a case in which one knows that p , and thus, given $E=K$, one's evidence – trivially – fully supports p over any specific threshold, personal justification would be a case in which one believes that p on the basis of one's evidence E (that one *takes* to support p over t), and yet $P(p | E) < t$. In this sense, one might be personally justified in believing that p without being either doxastically or propositionally justified.

In this way, one might be able to hold onto a “norm-based” account of doxastic justification (like the one defended by Williamson), while maintaining that something in the neighbourhood of the Responsiveness Intuition still plays a relevant role in our epistemic evaluation. Whether one is appropriately *evidence-responsive* is not going to be relevant for attributing doxastic justification, but it is going to be relevant for attributing personal justification.

However, I do not think this is a very promising way to go, for at least two reasons. First, note that, as mentioned above, the notion of personal justification is meant to capture the idea that, despite failing to be doxastically justified, one might nonetheless be *excusable* in believing as one does. However, remember that Williamson already has a way of accounting for how one might be excusable despite not being doxastically justified. As he points out, while only compliance with the primary norm guarantees justification, compliance with secondary and tertiary norm can give one an excuse. There is already an easy way of making sense of when and how someone is excusable but not justified within Williamson's normative framework. Hence, there is already an easy way of accounting for what someone might call ‘personal justification’¹⁹. Besides ‘saving’ the Responsiveness Intuition, it is thus not clear why one might want to cash out the notion of personal justification in terms of being *evidence-responsive* while having a norm-based account of doxastic justification. If anything – and this brings me to the second point I want to raise – this view seems to make doxastic justification

¹⁹In fact, as I mentioned in a previous footnote, Williamson's normative framework is meant to make sense of the justification-excuse distinction often used by the externalists to account for the notorious New Evil Demon Intuition.

something that is easier to achieve than personal justification. On this view, personal justification seems something that one can get only with some sort of cognitive effort that involves considering the evidence one has and believing in the light of such evidence, and yet the notion of personal justification epistemologists seem to be interested in does not seem to require such an effort. Furthermore, note that, especially when we combine E=K with an account of personal justification understood in terms of being evidence-responsive, then we end up with a very restrictive picture of what it takes for someone to be ‘excusable’. For if one’s evidence is all and only one’s knowledge, an evidentialist account of personal justification entails that one is excusable (despite not being doxastically justified) only if one believes something in the light of one’s knowledge²⁰.

Propositional Justification and Evidence

A final implication of Knowledge-first Evidentialism and its way of understanding doxastic justification is that, not only does the Responsiveness Intuition seem not to be relevant in our epistemic evaluations anymore, but propositional justification seems also to play no substantial role in our evaluation of people’s beliefs. Despite the disagreement on how to best cash out the notion of doxastic justification, it is nevertheless widely accepted that propositional justification plays an important role in determining whether one is doxastically justified or not. That is, not only would virtually everyone agree that if one is doxastically justified, then one is also propositionally justified, but, as I mentioned in Section 4.2, epistemologists would also claim that the fact that one is propositionally justified partly explains why one is doxastically justified²¹. Note that, on Williamson’s norm-based account of doxastic justification, doxastic justification does trivially entail propositional justification. Given that one is doxastically justified in believing that *p* if and only if one knows that *p*, on this view, knowledge becomes a limiting case of propositional justification²². However, note that, on Knowledge-first

²⁰Of course this is not the only option available to save the Responsiveness Intuition. For instance, one could point out that Williamson could relate the secondary (derivative) norm to the Responsiveness Intuition. In this way he could keep his already existing normative framework while ‘saving’ the Responsiveness intuition. Considering this further option would lead us too far afield. Anyway, I suspect that linking the Responsiveness Intuition to the derivative norm will also be problematic. For instance, it wouldn’t be clear in which sense the derivative secondary norm would be in fact *derivative* from the primary norm given that the compliance with the two different norms would involve completely different “epistemic jobs”.

²¹For an interesting discussion on the conceptual and theoretical priority of propositional justification see Melis, 2018.

²²That knowledge is a limiting case of propositional justification is easily derived from the main commitments of Williamson’s Evidence Package. This trivially follows from ii). For assume S knows

Evidentialism, the fact that one has propositional justification for believing that p plays no role in explaining *why* one is doxastically justified. In principle, one can account for one's doxastic justification without appealing to the notion of propositional justification at all. Once again, we can see how Knowledge-first Evidentialism differs from Evidence-first Evidentialism: Knowledge-first Evidentialism seems to reject the PJ-DJ intuition. But if we don't need to appeal to the notion of propositional justification in order to account for doxastic justification, then it looks like the notion of propositional justification plays no substantial role either in our epistemic theorising or in our evaluation of whether one's beliefs are justified. This, by itself, is an interesting result, but it is not necessarily a bad result. Perhaps one could argue that this is actually a positive outcome for it would provide us with a simpler framework within which we could evaluate the epistemic situation of a subject. Perhaps, one could say, we should lighten our theoretical apparatus and we should quit using the notion of propositional justification while focusing on doxastic justification only. Although this is an available option, I am afraid this move would come at a very high cost. For remember that, as I have mentioned above, one of the main consequences of understanding doxastic justification in terms of compliance with the constitutive (knowledge) norm of belief is that this framework separates the notion of doxastic justification from any basing requirement and from the Responsiveness Intuition. That is, it detaches the notion of doxastic justification from whether one is responding to the evidence one has. However, even if we detach the notion of evidence from the notion of doxastic justification in this way, evidence could still be relevant in our epistemic evaluations: the evidence one has, and its 'evidential strength', are what we consider in order to determine whether one is *propositionally* justified or not. Therefore, getting rid of the notion of propositional justification from our theoretical apparatus would thus mean getting rid of the notion of evidence in our epistemic evaluations^{23, 24}. It is important to stress that, by this I

that p . Given ii), once S knows that p , p becomes part of S's evidence. It thus follows that the evidential probability of p conditional on one's evidence is 1, thereby meeting any specific threshold.

²³Of course given E=K, if doxastic justification is understood as compliance with the KNB, then we still have an evidence-based account of doxastic justification in a sense, for, on this view, if S is doxastically justified in believing that p , then p is part of S's evidence. However, my point is that, on this norm-based account of doxastic justification, appealing to the notion of evidence is not at all required in order to explain why one is doxastically justified. I think this by itself is an interesting result.

²⁴Once again, some similar and interesting consequences can be drawn when we understand propositional and doxastic justification in terms of the roles of evidence. Doxastic Justification understood in terms of compliance with the knowledge norm of belief deprives evidence of its motivating role. Evidence is thus reduced to its justifying function, namely, evidence is what provides us with propositional justification. Getting rid of the notion of propositional justification altogether thus means getting rid of the notion of evidence altogether from our epistemic evaluations.

do not mean to say that a defender of Knowledge-first Evidentialism is forced to get rid of the notion of evidence altogether. For, although evidence might end up playing no role in our epistemic evaluations, it might still fulfil an important role in other epistemic practises, such as, in ruling-out hypothesis, in grounding inference to the best explanation, as well as in doing probabilistic reasoning (cf Williamson, 2000, ch. 9). However, at the very least, the defender of Knowledge-first Evidentialism is forced to rethink the role evidence plays in our epistemic evaluations.

Let me recapitulate what I have done in this section. On one hand, Williamson's Knowledge-first Evidentialism understands propositional justification similarly to how Evidence-first Evidentialism does: one has propositional justification for a proposition p iff one's evidence appropriately supports the target proposition p (e.g., one's evidence makes p sufficiently likely). On the other hand, it understands doxastic justification as compliance with the Knowledge Norm of Belief, so that one's belief that p is doxastically justified iff one knows that p . In this section, I have clarified to what extent Knowledge-first Evidentialism differs from Evidence-first Evidentialism, and I have pointed out some interesting consequences of Knowledge-first Evidentialism. I've pointed out that, besides the (already noted) fact that Williamson's view does not allow for (doxastically) justified false beliefs, the main differences between Evidence-first and Knowledge-first Evidentialism concern the way in which the notion of doxastic justification is conceptualised. More precisely, I have pointed out that it is in virtue of understanding doxastic justification in terms of propositional justification plus basing requirement that Evidence-first Evidentialism accommodates the Responsiveness Intuition. By contrast, on Knowledge-first Evidentialism, the evidence one takes into account when coming to believe something is not relevant for determining whether one is doxastically justified. Doxastic justification is not a matter of being evidence-responsive (in the sense specified in Responsiveness Intuition). One consequence of this framework is that the Responsiveness Intuition seems to play no relevant role in our epistemic evaluation. Nevertheless, I have noted that, rejecting the Responsiveness Intuition from our epistemic evaluations does not necessarily mean giving up on the notion of evidence altogether. In principle one could reject the Responsiveness Intuition while still appealing to the notion of evidence in our epistemic evaluations: what evidence one has is what we consider when we want to determine whether one is propositionally justified. However, given that the reason why one might be interested in whether one has propositional justification is to determine whether one is doxastically justified, once we realise that, on Knowledge-first Evidentialism, propositional

justification does not explain why one is doxastically justified, it is not really clear what the need for these notions is. That is, it is not clear what the role of propositional justification (and thus the role of evidence) is in our epistemic evaluations.

4.5 Ecumenical Evidentialism

In the first half of this Chapter, I have considered two possible ways of understanding Evidentialism: Evidence-first Evidentialism and Knowledge-first Evidentialism. After shedding light on what these views amount to, I have considered to what extent the two views are different. In particular, I have pointed out that one important difference between the two views is that, while Evidence-first Evidentialism vindicates the so-called Responsiveness Intuition, the intuition that takes doxastic justification to be a matter of believing in the light of one's evidence, Knowledge-first Evidentialism does not vindicate the Responsiveness Intuition. On this view, being justified in believing that p is just a matter of knowing that p . In this section, I will suggest a novel Evidentialist view that accounts for the orthodox Responsiveness Intuition while remaining in the spirit of Knowledge-first Epistemology. I will call it Ecumenical Evidentialism.

4.5.1 A two-tiered account of knowledge

Ecumenical Evidentialism steers a course between Evidence-first Evidentialism and Knowledge-first Evidentialism. It retains the main commitments of Knowledge-first Epistemology, i.e., KNB , $E=K$, and it combines them with the intuition that doxastic justification is a matter of adequately responding to one's evidence. More precisely, Ecumenical Evidentialism embraces all the following Knowledge-first and Evidence-first commitments:

1. Doxastic Justification: DJ is an evaluation of a belief so that one's belief that p is doxastically justified for S only if S believes that p on the basis of one's evidence E that stands in a support-relation R towards p .
2. Responsiveness Intuition
3. PJ-DJ Intuition
4. JFB
5. $E=K$

6. KNB

How is it possible to bring together these Evidence-first and Knowledge-first commitments? The answer is to be found in the two-tiered account of knowledge that underpins my Ecumenical Evidentialism.

On one hand, Ecumenical Evidentialism understands inferentially-formed beliefs and inferential knowledge in an Evidence-first fashion. On the other hand, it embraces a Knowledge-first approach when it comes to cases of non-inferential knowledge. Let us first consider cases of inferential knowledge. Consider, for instance, my belief that my friend is in London right now. This belief is supported by my total evidence: I saw her standing next to the train from Edinburgh that goes to London this afternoon, every Tuesday she is in London and today is Tuesday, and so on. That is why I believe that my friend is in London. In fact, my friend is in London. What would Knowledge-first Evidentialism say about this case? If we understand Knowledge-first Evidentialism within Williamson's broader account of knowledge, then we will expect a defender of Knowledge-first Evidentialism to say that, in this scenario, I know that my friend is in London and this is because my belief is safe (and any other relevant necessary conditions for knowledge is met). Instead, Ecumenical Evidentialism would consider cases of inferential knowledge in an Evidence-first fashion. That is, given the orthodox notion of doxastic justification Ecumenical Evidentialism embraces, cases of inferential knowledge will be cases in which one knows that p in virtue of having doxastic justification for believing that p (while fulfilling whatever further condition is needed in order to have knowledge). For remember that on the orthodox Evidence-first view, doxastic justification is the result of basing one's belief on one's propositional justification. Putting the matter in terms of the roles of evidence, on Ecumenical Evidentialism doxastic justification is connected with the motivating role of evidence: one has doxastic justification as the result of taking into account one's evidence when coming to believe a specific proposition. Inferential cases thus represent exactly the context in which one can "exploit" the motivating role of evidence, insofar as one takes such evidence as a premise in one's reasoning.

Consider now a different case of inferentially-formed belief. Again, I believe my friend is in London right now. This belief is supported by my total evidence: I saw her standing next to the train from Edinburgh that goes to London this afternoon, every Tuesday she is in London and today is Tuesday, and so on. As it happens, however, my friend's train got delayed and she didn't manage to get to London. What would Ecumenical Evidentialism say about this case? While on Williamson's account, my

belief would be at most excusable, but it would fail to be justified, Ecumenical Evidentialism, instead, would classify my belief as being doxastically justified because, despite false, it is the result of appropriately believing on the basis of my supporting evidence. It is in this sense that Ecumenical Evidentialism accommodates at least two of the main commitments of Evidence-first Evidentialism: not only does it vindicate the Responsiveness Intuition, but it also allows for (doxastically) justified false belief.²⁵

Let us now move onto cases of non-inferential knowledge. While Ecumenical Evidentialism embraces an orthodox Evidence-first view about inferential knowledge (and inferentially-formed beliefs more in general), it retains a Knowledge-first view about non-inferential knowledge. Consider a paradigmatic case of non-inferential knowledge, such as, a case of perceptual knowledge. Assume Albert is in a paradigmatic good case and he sees that there is a table in front of him. Ecumenical Evidentialism understands this case in a Knowledge-first fashion; by seeing that there is a table in front of him, Albert comes to know that there is a table in front of him. *Seeing that there is a table* is just a *way of knowing* that there is a table. There is no interesting sense in which Albert has to previously consider the evidence he has in order to come to believe that there is a table²⁶. It is in this sense that Ecumenical Evidentialism (partly) captures the Knowledge-first commitment that takes knowledge to be the most general factive mental state.

It is important to note that epistemologists have sometimes taken this to be a problematic aspect of Knowledge-first Epistemology. Consider again the proposition that there is a table. Brueckner, for instance, has focused on perceptual knowledge and has argued against the Knowledge-first view, on which one perceptually knows that there is table merely by seeing that there is a table (Cf. Brueckner, 2005b, Brueckner, 2009 and Williamson's response in Williamson, 2009b: pp.282-85.). In particular, Brueckner has claimed that, on this Knowledge-first view of perception, it is hard to see what evidence one has and in virtue of which one knows that there is table. Given that, on E=K, once you know a proposition, then that proposition is part of your evidence, and given knowledge entails belief (thing that Williamson accepts as well), Brueckner

²⁵For other interesting views that aim to account for justified false belief while remaining Knowledge-first in some relevant sense see Bird, 2007, Ichikawa, 2014, and, in a sense, Smith, 2010b. For some objections to these views see McGlynn, 2012.

²⁶It is natural to think of cases of proprioception as similar to cases of non-inferential knowledge in this respect. Notoriously, Anscombe describes cases of bodily awareness as cases of 'knowledge without clues' (see Anscombe, 1962). However, I think the case of proprioception is a complicated one. For instance, I suspect that whether cases of proprioception will end up being cases of knowledge not based on any evidence in the same way as cases of perceptual knowledge are will depend a lot on the details of the case in question.

argues that the only option available is to say that cases of perceptual knowledge are cases in which one is justified in believing something (e.g., that there is a table) in virtue of one's belief that it is the case (e.g., in virtue of the belief that there is a table). I will not discuss Brueckner's arguments in this thesis, but I would like to briefly note three things. First, note that, as McGlynn has said, Brueckner's mistake here seems to be that "[he] assumes a picture according to which there must be independent evidence for the proposition that [there is a table], in virtue of which his belief in that proposition is justified and, given that the world is cooperating in the right way, knows" (2014: 65). Given what I have discussed in the first half of the Chapter, this does not come as a surprise. For Brueckner seems to assume what I have called Evidence-first Evidentialism. What Section 4.4. of this Chapter was meant to show is exactly that it is part of the Knowledge-first picture to reject the orthodox conception concerning the relation between evidence and justification, one that embraces the Responsiveness Intuition. Williamson is just offering us a completely new paradigm, one that is an alternative to the orthodoxy. Second, once we accept and realise that, the Knowledge-first picture of perceptual knowledge is such that perceptual knowledge is not based on any evidence, I do not think this is a negative result. Instead, as I will suggest later, I think it vindicates our intuitions concerning the role evidence plays in perceptual knowledge. Third, as I will argue in a moment, it does not follow from this Knowledge-first picture of non-inferential knowledge that one is justified in believing that there is a table in virtue of one's belief that there is a table. Considering my Ecumenical Evidentialism further will help us see why this is the case.

We can ask: what does Ecumenical Evidentialism say about the epistemic status of Albert's belief? At this point it is worth noticing that, virtually every epistemologist (both Knowledge-firsters and non) would grant that, in this scenario, Albert both knows and (thus) is doxastically justified in believing that there is a table in front of him. Those who are sympathetic to a traditional Knowledge-first approach would say that Albert is doxastically justified *because* he knows that there is a table in front of him: doxastic justification is just a matter of complying with the Knowledge Norm of Belief. By contrast, defenders of more orthodox Evidence-first view would say that he knows that there is a table because he is doxastically justified: seeing that there is a table provides him with the right kind of support necessary for knowing that there is a table²⁷. Whatever order of explanation one chooses, however, both Evidence-first and Knowledge-first defenders would grant that Albert is doxastically justified in believ-

²⁷This result especially clear in Pritchard's version of Evidence-first Evidentialism.

ing that there is a table before him. Knowledge entails (doxastic) justification. Here is where Ecumenical Evidentialism departs from both Evidence-first and Knowledge-first Evidentialism. Remember that Ecumenical Evidentialism combines the orthodox Evidence-first view of doxastic justification, on which one is doxastically justified as a result of taking into account one's evidence when coming to form one's belief, with a Knowledge-first view of non-inferential knowledge, on which perceptual (and other non-inferential) cases, are just cases in which one just comes to know something without taking into account any evidence when forming the target belief. But if we want to retain the orthodox intuition on which doxastic justification is a matter of believing on the basis of one's evidence, then we have to insist that cases of non-inferential knowledge are cases of knowledge without justification. This is what Ecumenical Evidentialism does. Let me show why. On Ecumenical Evidentialism, doxastic justification has to do with how much and whether one is appropriately *evidence-responsive* in a way that captures the Responsiveness Intuition. Putting things in terms of the roles of evidence, doxastic justification should be understood in an orthodox manner, namely, as the result of the justifying and, more importantly, the motivating role of evidence. Crucially, even if we grant, as Ecumenical Evidentialism does, that $E=K$, and thus we grant that cases of non-inferential knowledge are related to the justifying role of evidence (given that they are limiting cases of propositionally justified belief), cases of non-inferential knowledge are not cases in which the subject "exploits" the motivating role of evidence. As I mentioned above, when Albert sees a table in front of him, he just comes to know that there is a table in front of him. There is no interesting sense in which Albert had to previously consider any evidence at all in order to come to believe that there is a table in front of him.

To sum up: on Ecumenical Evidentialism, doxastic justification is the result of appropriately taking into account one's evidence when coming to believe something. In this way, Ecumenical Evidentialism captures the orthodox Responsiveness Intuition that sees doxastic justification as a matter of being evidence-responsive. Cases of inferential knowledge and inferentially-formed beliefs are understood in an Evidence-first way. Inferential knowledge is the result of having doxastic justification. At the same time, Ecumenical Evidentialism understands non-inferential cases, e.g., perceptual cases, in a Knowledge-first fashion. These are cases in which one knows that p without taking into account any evidence when coming to believe that p . On Ecumenical Evidentialism, cases of non-inferential knowledge are cases of knowledge without justification. Interestingly, by appealing to this two-tiered account of knowledge, not

only Ecumenical Evidentialism is in line with the Responsiveness Intuition, but it also seems to vindicate the Austinian intuition on which:

“if [a table] [...] stands plainly in view, there is no longer a question of collecting evidence; its [standing plainly in] view doesn’t provide me with more evidence that it is a table, I can now just see that it is, the question is settled.”

(Austin, 1962: 115)

For this, and other reasons that I will show in the remaining part of this chapter, I take the case of knowledge without justification to be a feature of my Ecumenical Evidentialism rather than a bug.

4.5.2 Meta-epistemological Pluralism

In the previous section, I have presented the core theoretical commitments of Ecumenical Evidentialism. As said above, this view steers a course between traditional Knowledge-first Evidentialism and orthodox Evidence-first intuitions about doxastic justification. It does so by appealing to a two-tiered account of knowledge, one that brings together a Knowledge-first view about non-inferential knowledge with an Evidence-first view about inferentially justified beliefs and inferential knowledge. Within this picture, Ecumenical Evidentialism restricts the notion of doxastic justification to those doxastic attitudes that are the result of appropriately taking into account one’s evidence, thereby taking cases of non-inferential knowledge to be cases of knowledge without (doxastic) justification. However, one might find my two-tiered account of knowledge unsatisfying with respect to what it predicts about Albert’s epistemic status. For wouldn’t everyone agree that Albert is successful in believing as he does? In particular, one might point out that my Ecumenical Evidentialism has a very counterintuitive result: on one hand it considers a (inferentially-formed) *false* belief as justified, while it deprives non-inferential knowledge of the status of justification. Drawing a distinction between *unjustified* and *non-justified* belief will be helpful here. Take an unjustified belief to be a belief that is based on evidence that does not support the target believed proposition (e.g., it does not make the target proposition sufficiently likely). Take a non-justified belief to be, instead, a belief that is not based on any evidence. While the case of non-inferential knowledge is a case of non-justified belief, we should not make the mistake to classify it as a case of unjustified belief. On my

view, cases of non-inferential knowledge are just cases of beliefs that are not justified. In fact, on my view, these beliefs are not even in the market for justification.²⁸

However, I can see why one might still find my view puzzling. For isn't an instance of knowledge always better than any false belief? But how can we account for Albert's success if, on Ecumenical Evidentialism, his belief can't be classified as being doxastically justified? In this section, I further develop and motivate Ecumenical Evidentialism by appealing to a form of meta-epistemological pluralism. I will show that the form of meta-epistemological pluralism I develop provides us with a framework that is sensitive to the different ways in which a belief can be epistemically successful as well as epistemically defective. I will show that this form of meta-epistemological pluralism best achieves the conjunction of Evidence-first and Knowledge-first features in a satisfying way.

We can ask: what, as epistemologists, do we evaluate when judging one's epistemic situation as 'epistemically good'?²⁹ That is, what are we interested in when we make epistemic evaluations of people's beliefs? When do we consider one's doxastic attitude as epistemically successful? Meta-epistemological Pluralism claims that there is no univocal answer to these questions. Instead, there are at least two Epistemic Success Criteria that we should consider in our epistemic evaluations: Epistemic Justification and Epistemic Excellence. We can further distinguish Epistemic Justification into Propositional and Doxastic Justification, thereby having three Epistemic Success Criteria: i) Propositional Justification, ii) Doxastic Justification, and iii) Epistemic Excellence. We already have an idea of what Propositional and Doxastic Justification amount to. When making epistemic evaluations, we can evaluate whether a *proposition* is justified for a subject given the subject's evidence (propositional justification); we can evaluate whether a subject's *belief* is justified given the evidence one has (doxastic justification). However, we can also evaluate whether a subject's belief is a good mental state to be in, regardless of whatever evidence one has (epistemic excellence). These three criteria enable us to appreciate that there are different ways in which one can be epistemically successful in a target situation. In particular, the criteria of Dox-

²⁸So far in this thesis, I have avoided talking about rationality. I will continue to do so. However, I am inclined to apply the unjustified-non justified distinction to the case of rationality. That is, I believe we should distinguish between *arational* beliefs and *irrational* beliefs. Arational beliefs are beliefs that are just not in the market for rationality. Given I believe rationality is a matter of appropriately responding to the evidence one has, I believe we should just accept that cases of perceptual knowledge are cases arational beliefs. However, this is not something I am going to defend here. For an interesting recent discussion on rationality see Williamson, 2017.

²⁹One's epistemic situation includes but is not restricted to, one's doxastic attitude. For instance, it includes the evidence one has.

astic Justification and Epistemic Excellence show that there are at least two different ways in which a *belief* can be epistemically successful. According to this form of meta-epistemological pluralism, on the Doxastic Justification Criterion, one's doxastic attitude is epistemically good insofar as it is the best *response* one can provide, given one's evidence. On the Epistemic Excellence Criterion, one's doxastic attitude is epistemically good insofar as it is an excellent mental state to be in. More precisely, on this form of pluralism, one's belief is Epistemically Excellent insofar as it complies with the constitutive norm of belief. To sum up, on this form of Meta-epistemological Pluralism, we should consider - at least - the following three *Epistemic Success Criteria* in our epistemic evaluations:

Propositional justification: an evaluation of a *proposition* relative to the evidence a subject has.

Doxastic justification: an evaluation of a subject's *belief* qua adequate response to the evidence she has.

Epistemic Excellence: an evaluation of a subject's *belief* qua belief that meets the standard set by the constitutive norm of belief.

How can this meta-epistemological pluralism support Ecumenical Evidentialism? That is, how can we make sense of this form of Meta-epistemological Pluralism within an Ecumenical Evidentialist framework?

My proposal is that, while we should cash out Epistemic Justification (both in its propositional and doxastic form) in an Evidence-first fashion, namely, as being related to the justifying and motivating roles of evidence, we should understand Epistemic Excellence within the normative framework developed by Williamson. Using Williamson's normative framework, and the KNB that underpins it, generates the claim that one's belief that *p* is epistemically excellent if and only if one knows that *p*. Ecumenical Evidentialism embraces the Knowledge-first commitment on which knowledge is the norm of belief. However, instead of setting knowledge as the standard for Doxastic justification, it takes knowledge to be the standard for *Epistemic Excellence*. Therefore, it is not by merely complying with the Knowledge Norm of Belief that one's belief can be judged as being doxastically justified. While doxastic justification is sensitive to how one responds to one's evidence, on Ecumenical Evidentialism, one's belief can be *epistemically excellent* even if no evidence is taken into account when coming to form the target belief. By capturing the epistemic successes as well as the epistemic failings of a target belief, my meta-epistemological pluralism helps us making sense of the way in which Ecumenical Evidentialism brings together an orthodox

way of understanding doxastic justification with some core features of Knowledge-first Epistemology. Let us see how.

Remember that, on Ecumenical Evidentialism, cases of non-inferential knowledge are understood within a Knowledge-first framework. Consider the case of Albert again. Albert's *seeing that* there is a table is just a *way of knowing* that there is a table. As I have noticed above, in these cases, there is no evidence that Albert previously considers in order to form the belief that there is a table before him. That is, his belief that there is a table in front of him is not the result of him responding to the evidence he had before forming the belief. When we take this Knowledge-first account of non-inferential knowledge together with an Evidence-first notion of doxastic justification, one on which whether one is doxastically justified depends on whether one believes appropriately in the light of the evidence one has, then we have to insist that cases of non-inferential knowledge are cases of knowledge without justification. Crucially, everyone would agree that Albert's belief is a good belief to have. That is, despite not being doxastically justified, Albert's belief is intuitively epistemically good. Appealing to my version of Meta-epistemological pluralism allows us to account for the sense in which Albert's belief is epistemically successful despite not being justified. This consideration can be generalised, of course, to all those cases of (doxastically) non-justified non-inferential knowledge. Albert's belief is successful because, by complying with the constitutive (Knowledge) Norm of belief, it is *epistemically excellent*. The criterion of Epistemic Excellence allows us to capture the idea that there is a sense in which one is doing well when one knows something, even if the belief is based on no evidence at all³⁰.

Let us now go back to the case of inferentially-formed belief. In particular, let's consider the case in which my belief that my friend is in London is false, despite being supported by my evidence. On Williamson's view, my belief that my friend is in London may be excusable, but it is not justified. One reason of dissatisfaction with this view is that, in this scenario, my belief seems to be more than merely excusable (cf. Boulton, 2017). My pluralist view accounts for the intuition that this belief is not merely excusable, but, in coming to believe in the light of my evidence, I'm doing epistemically well. My Meta-epistemological Pluralism captures the sense in which my belief

³⁰Of course, one might object that knowledge is not the constitutive norm of belief. I think there are good arguments in the literature for the KNB (see footnote 14), but assessing those arguments would lead me too far afield. In any case, this is not needed for the purpose of this Chapter. For remember the aim of this chapter is to explore the plausibility of a novel variety of evidentialism, one that brings together Evidence-first and Knowledge-first features. Here I am borrowing the KNB from Knowledge-first Evidentialism, thereby assuming it to be true.

is epistemically successful, namely, it is doxastically justified. At the same time, it captures the sense in which my belief is epistemically defective: is not epistemically excellent for it fails to meet the KNB.

What about the case in which my belief that my friend is in London is in fact true? Remember that, on Ecumenical Evidentialism, these cases are understood in a more orthodox way. Therefore, this would classify as an instance of inferential knowledge as being the result of having doxastic justification. My Meta-epistemological pluralism will thus predict that cases of inferential knowledge like the one described above are cases of successful beliefs, both in the sense of being doxastically justified as well as in the sense of being epistemically excellent beliefs. For remember that, on this view, a belief is epistemically excellent if and only if it complies with the Knowledge Norm of Belief. Therefore, cases of inferential as well as non-inferential knowledge will be classified as cases of epistemic excellence. By contrast, cases of epistemically excellent beliefs will not necessarily be cases of doxastically justified belief.

At this point one might worry that, by keeping an orthodox account of doxastic justification, and by understanding inferential knowledge in terms of more basic components (i.e., propositional and doxastic justification), my view does not maintain the spirit of a Knowledge-first account. I think this is a legitimate concern. However, once I have filled in the details of this account a bit more I will be able to show that, despite being revisionary, my Ecumenical Evidentialism can be developed in a way that retains some of the main features of traditional Knowledge-first Epistemology. I will show that not only can one be a “Knowledge-firster” when it comes to cases of non-inferential knowledge, but, on Ecumenical Evidentialism, knowledge also plays a fundamental role in explaining the how one is doxastically justified. In order to do so, I will first consider some potential objections for my view. Addressing (and resisting) these objections will enable me to provide further details on my account. Once I have done that, I will show that Ecumenical Evidentialism successfully retains the main orthodox intuitions about doxastic justification, while remaining in the spirit of the Knowledge-first program.

4.5.3 Non-traditional Foundationalism

As I have described in the previous section, Ecumenical Evidentialism rests on a two-tiered account of knowledge. On this two-tiered account of knowledge, we should understand knowledge differently, depending on whether it is an instance of non-

inferential knowledge, or inferential knowledge. It understands the former in a Knowledge-first fashion. It understands the latter in an Evidence-first fashion. In this section, I will address a potential problem for my view. Resisting this problem will enable me to provide more details and to further shed light on some important features and benefits of Ecumenical Evidentialism.

The problem I want to address is one that concerns the more orthodox part of Ecumenical Evidentialism. Remember that, on my view, if one is doxastically justified in believing that p , then one believes that p in the light of one's evidence that stands in an appropriate evidential support relation towards p . Furthermore, on my view, only cases of inferentially formed beliefs are cases of (potential) doxastic justification. Remember however, that Ecumenical Evidentialism brings together the orthodox notion of doxastic justification with the Knowledge-first thesis that one's evidence is all and only the propositions one knows, namely, $E=K$. Crucially, if one's doxastic justification depends on whether one believes appropriately on the basis of one's evidence, and one's evidence is one's knowledge, the worry of an infinite regress arises. Consider a chain of inferentially formed and justified beliefs. If a doxastically justified belief rests on one's knowledge k_1 , then this knowledge k_1 must also be the result of a belief based on evidence, namely, given $E=K$, on knowledge k_2 , which, in turn, must be inferentially justified in the light of further evidence, namely, given $E=K$, on further knowledge k_3 and so on.

There have been four main ways of dealing with the Regress Problem³¹. A first option is to provide a foundationalist response and argue that we have self-justifying beliefs at the end of the chain, and that these beliefs do not need further justification (BonJour, 1978, BonJour, 1985, BonJour, 1999, BonJour & Sosa, 2003). A second option is to embrace infinitism and to bite the bullet by accepting that the chain of justification continues *ad infinitum* (cf. Klein, 2005). A third option is to embrace coherentism. On this view, what it takes for a belief to be justified is to form a coherent system with a set of beliefs. Therefore, the regress stops because the chain of beliefs circles back on itself (cf. Lehrer, 1990). Finally, one can take a sceptical stance and argue that in fact we do not have justified beliefs.

One of the benefits of Ecumenical Evidentialism is that it provides a straightforward foundationalist solution to the Regress Problem. There are two ways in which one can provide a foundationalist solution to the Regress Problem. A traditional foun-

³¹For recent discussion on the Regress Problem as it arises for some forms of internalism see Bergmann, 2006.

dationalist solution resists the Regress Problem by claiming that we can find a self-justified belief at the end of the chain of beliefs. A non-traditional foundationalist solution resists the Regress Problem by claiming that what we find at the end of the chain is a non-justified belief. On this non-traditional form of foundationalism, inferential justification is provided eventually by beliefs that are not themselves justified. Ecumenical Evidentialism provides a non-traditional foundationalist solution to the Regress Problem. As it will be clear in a moment, considering this foundationalist solution enables us to see how the two parts of Ecumenical Evidentialism interact with each other. On this non-traditional foundationalist solution, the regress eventually stops because what we find at the end of the chain of beliefs is non-inferential knowledge, which, on Ecumenical Evidentialism, is not (doxastically) justified itself. Non-inferential knowledge stops the regress by being knowledge that is not based on evidence, namely, given $E=K$, on knowledge³². This foundationalist solution to the Regress Problem enables us to clarify in which sense my account is Knowledge-first: it shows that, on my account, knowledge is the most basic epistemic state. (Non-inferential) knowledge is first because it represents the grounds on which other beliefs are doxastically justified. It is in virtue of non-inferential knowledge that one can gain doxastic justification. It is in virtue of non-inferential knowledge that one can acquire further knowledge. On Ecumenical Evidentialism, (non-inferential) knowledge is more fundamental. (Non-inferential) knowledge is potential justification. It is at least in this sense that Ecumenical Evidentialism retains the spirit of Knowledge-first Epistemology³³.

³²Wittgenstein can be thought of as offering a similar non-traditional foundationalist solution. However, note that on his view hinge-propositions are not-justified and not-known (at least not in the traditional sense of knowing), yet they are *certain*. By just accepting the possibility of non-justified *knowledge*, my non-foundationalist solution avoids appealing to vague notion of certainty. (see Wittgenstein, 1969).

³³I have already mentioned above that the way in which Ecumenical Evidentialism resists the regress problem can be classified as a non-traditional foundationalist solution. That is, it claims that the regress is stopped by non-inferential knowledge that, nevertheless, is not itself justified. However, it is worth pointing out that there is at least another way in which the solution provided by Ecumenical Evidentialism differs from the one offered by traditional foundationalists. For a traditional foundationalist response would identify one's experience as the candidate for what lies at the end of the chain of inferentially justified beliefs. However, note that, on my view, everything that is known non-inferentially is suitable for being the foundation of our justificatory chain. Therefore, while a traditional foundationalist response would claim that empirical knowledge is what stops the regress, on my view, perceptual knowledge is only one of the various non-inferential candidate that stops the regress. On my view, any kind of non-inferential knowledge will be a good candidate for stopping the regress.

4.6 Conclusion

In this Chapter, I have considered the Evidentialist claim on which one's justification supervenes on one's evidence. However, Evidentialism so defined is silent on various issues, e.g., the nature of evidence and justification, and how to understand the relation between these two epistemic notions. All Evidentialism entails is that if two subjects have the same evidence, then they have the same justification. But, besides that, Evidentialism does not explain what relation there is between evidence and propositional justification, and between evidence and doxastic justification. Depending on how we fill in the details, we can have varieties of Evidentialism. The theories that I have considered in this Chapter, namely, what I have called, 'Evidence-first Evidentialism', and 'Knowledge-first Evidentialism', are just two ways in which one can be an Evidentialist. What I have called 'Evidence-first Evidentialism' constitutes the orthodox way of understanding propositional and doxastic justification in evidentialist terms. Roughly put, it understands propositional justification in terms of evidence and the support that evidence provides to a target proposition, and it understands doxastic justification in terms of propositional justification and a basing requirement. In particular, by understanding doxastic justification as being related to what kind of evidence one bases one's belief on, Evidence-first Evidentialism vindicates the orthodox intuition on which whether one is doxastically justified depends on whether one is appropriately evidence-responsive, namely, on whether one's belief is the result of appropriately taking into consideration one's evidence in coming to form the target belief. I have called this the 'Responsiveness Intuition'. Williamson (Williamson, 2011, Williamson, in press) has a different view about doxastic justification. He argues that doxastic justification is a matter of compliance with the Knowledge Norm of Belief. On this view, whether one is doxastically justified in believing that p is not strictly related to what evidence one takes into account in coming to believe that p . However, given his $E=K$, his view is also an instance of Evidentialism. I have called this view 'Knowledge-first Evidentialism'. In this Chapter, firstly, I have clarified the main commitments of both Evidence-first and Knowledge-first Evidentialism. Secondly, I have put forward a novel evidentialist view, one that retains some of the main Knowledge-first commitments, such as, KNB and $E=K$, whilst at the same time capturing the Evidence-first Responsiveness Intuition. I have called this view 'Ecumenical Evidentialism'. Ecumenical Evidentialism brings together Evidence-first and Knowledge-first features by appealing to a two-tiered account of knowledge. On one hand, it understands cases

of inferential knowledge in an orthodox Evidence-first way. This means that cases of inferential knowledge are cases in which one knows in virtue of having doxastic justification. On the other hand, it embraces a Knowledge-first picture of non-inferential knowledge. This means that cases of non-inferential knowledge are cases in which one's seeing something is just a way of knowing something. That is, one knows something by directly *seeing something* and without taking into account any evidence. However, I have pointed out that if we want to retain both an orthodox Evidence-first view of inferential knowledge (one that vindicates the Responsiveness Intuition), together with a Knowledge-first view of non-inferential knowledge, we have to insist that cases of non-inferential knowledge are cases of knowledge without doxastic justification. I have argued that this is a feature of my account rather than a bug. Furthermore, I have argued that a form of meta-epistemological pluralism best achieves the combination of Evidence-first and Knowledge-first features in a way that accounts for the fact that cases of non-inferential knowledge, despite being cases of non justified belief, are nevertheless importantly *epistemically* successful. Similarly, it accounts for the facts that, when based on supporting evidence, inferentially-formed false beliefs are more than merely excusable: they are justified. On the form of pluralism I defend, we should thus distinguish Doxastic Justification from Epistemic Excellence. While Doxastic Justification has to be understood along Evidence-first lines, we should understand Epistemic Excellence in terms of compliance with the constitutive (Knowledge) norm of belief. While one is *doxastically justified* in believing that p only if one believes that p on the basis of evidence that stands in a support relation towards p , one's belief that p is *epistemically excellent* if and only if one knows that p . Meta-epistemological Pluralism accounts for the virtues (and the vices) of cases of non-inferential knowledge as well as cases of inferentially justified but false beliefs.

One might wonder why we should take Ecumenical Evidentialism seriously. Here are what I take to be the following two main benefits of Ecumenical Evidentialism:

First, Ecumenical Evidentialism brings together Evidence-first commitments while remaining in the spirit of Knowledge-first Epistemology. In particular, Ecumenical Evidentialism is Evidence-first in the following sense:

- It understands Doxastic Justification in terms of propositional justification plus basing, thereby vindicating the Responsiveness Intuition.
- It allows for (doxastically) justified false beliefs
- It understands the relation between doxastic and propositional justification in a

traditional way: propositional justification partly explains doxastic justification

It is Knowledge-first in the following sense:

- It embraces E=K
- It embraces KNB
- It takes knowledge (in its non-inferential form) to be more basic and fundamental: inferential knowledge and doxastic justification are explained in terms of non-inferential knowledge.

Second, as mentioned above, by appealing to two-tiered account of knowledge and a form of Meta-Epistemological Pluralism, my Ecumenical Evidentialism is sensitive to the different ways in which a doxastic attitude can be epistemically successful and/or epistemically defective.

Let me conclude with some more speculative remarks on possible future developments of my view. I believe there are other ways in which, if developed further, Ecumenical Evidentialism will show to be a promising view. For instance, I believe the two-tiered account of knowledge underpinning Ecumenical Evidentialism might provide us with a solution to the so-called ‘Basis Problem’. The Basis Problem is a problem that arises for views that seem to embrace the following two theses: i) *seeing that* p entails knowing that p , and ii) *seeing that* p is the reason in virtue of which one believes (and knows) that p . Pritchard’s Epistemological Disjunctivism represents an instance of such a view. The problem arises because a logical gap seems to be required between *believing* that p and the *reason* on the basis of which one believes that p (cf. Pritchard, 2012: 19). Interestingly, the Basis Problem does not seem to arise once we just reject, as Ecumenical Evidentialism does, the thesis that *seeing that* p is the reason in the light of which one knows that p . Furthermore, if developed further, my Ecumenical Evidentialism might provide us with an anti-sceptical solution that is alternative to (and in between) the solutions offered by the *conservative* on one hand and the *dogmatist* on the other hand (cf. Wright, 2004 for the former, and Pryor, 2000 for the latter). Imagine that I see that I have hands. Like the dogmatist, Ecumenical Evidentialism rejects the idea that knowledge that I have hands requires antecedent justification for the belief that I am not a brain in a vat. Unlike the dogmatist, and like the conservative, Ecumenical Evidentialism predicts that I do *not* have justification for the belief that I have hands. Nevertheless, by appealing to a knowledge-first view of non-inferential knowledge, Ecumenical Evidentialism can explain how I can

have knowledge of my hands despite not having a justified belief. Finally, I believe that by developing further this view, we can appreciate how Ecumenical Evidentialism can account for how a target doxastic attitude can be evaluated as being epistemically successful or defective within a novel framework than the one provided by traditional virtue-epistemologists (cf. Sosa, 2005). Of course, a complete defence of Ecumenical Evidentialism would involve addressing and resisting arguments that potentially threaten my view. In particular, I think there are two main ways in which one could reject Ecumenical Evidentialism. One could argue that, contrary to the appearances, cases of non-inferential knowledge are cases in which the subject gains knowledge as a result of taking into account one's evidence (as a traditional Evidence-first view would do). Alternatively, one could reject the Responsiveness Intuition and argue that doxastic justification is not a matter of being appropriately *evidence-responsive* (as traditional Knowledge-first Evidentialism would do). Note, however, that anyone who opts for the former strategy will have to explain exactly how that is possible while resisting the above-mentioned Basis Problem. On the other hand, anyone who wants to say that non-inferential knowledge is both non evidence-based as well as a case of justified belief, will have to come up with an account of justification that is sensitive to the different ways in which a belief can be epistemically successful. However, the aim of this Chapter was more modest than providing a full defence of Ecumenical Evidentialism. All I wanted to do was to explore a novel evidentialist view, one that brings together Evidence-first and Knowledge-first elements, and evaluate the prospects of such a view. I believe the work in this Chapter shows that, overall, Ecumenical Evidentialism, when developed further, could represent a view that we should take seriously.

Concluding Remarks Part III

In Part III, I have addressed the following question:

How does Evidential Externalism understand the relation between evidence and justification?

Previously in this Thesis, I have pointed out how Evidential Externalism, defined as the denial of the supervenience thesis between evidence and non-factive mental states, is silent on what justification is. Nevertheless, it is plausible to take Evidential Externalism as a way of embracing *Evidentialism*. Following Conee and Feldman (2004), I define *Evidentialism* as the thesis on which one's justification supervenes on one's evidence. However, as I have already stressed, *Evidentialism* so defined does not specify which kind of relation there is between evidence and justification. In Part III, I have formulated various ways in which one can be an evidentialist: Evidence-first Evidentialism, Knowledge-first Evidentialism, and Ecumenical Evidentialism.

Conclusion

In this Thesis, I have focused on Evidential Externalism and I have addressed the following three questions:

1. Does the Access Problem represent a real threat to Evidential Externalism?
2. Is Evidential Externalism committed to a sceptical variety of infallibilism?
3. How does Evidential Externalism understand the relation between evidence and justification?

The arguments I have developed in the three Parts of this Thesis answer the questions respectively as follows:

1. The Access Problem represents a serious challenge for Pritchard's Epistemological Disjunctivism, but it does not pose a real threat to Williamson's E=K.
2. The Infallibility Problem does not represent a serious threat to Williamson's E=K. Given that Pritchard's Epistemological Disjunctivism is a claim restricted to paradigmatic cases of perceptual knowledge, it seems it is not committed to any problematic variety of infallibilism.³⁴
3. There are different ways in which an Evidential Externalist can cash out the relation between evidence and justification. I have argued that we can understand Pritchard's Epistemological Disjunctivism and Williamson's account of evidence as two instances of, what I have called, Evidence-first and Knowledge-first Evidentialism respectively. Furthermore I have defined a novel Evidential Externalist view which I have called *Ecumenical Evidentialism*.

³⁴In fact, following Dodd, I have pointed out that we should take seriously the threat of the Infallibility Problem only when it comes to cases of knowledge of "what it is like outside the region of the world with which [we] have direct perceptual contact" (Dodd, 2005: 644). However, I have shown that even when focusing on these cases, we can realise infallibilism does not lead to scepticism.

Where does this leave us? I believe we can summarise the main conclusion of this thesis in the following way. On one hand, I have noticed that Pritchard's Epistemological Disjunctivism can be classified as an instance of *Evidence-first Evidentialism*. In fact, despite restricting its scope to paradigmatic cases of perceptual knowledge, Epistemological Disjunctivism vindicates the *Responsiveness Intuition* underpinning Evidence-first Evidentialism, namely, the intuition that takes doxastic justification to be a matter of responding to the evidence one has (or, to use Pritchard's terminology, a matter of responding to the 'rational support' one has). Furthermore, I have noticed that it seems to be in virtue of its accessibility requirement that Epistemological Disjunctivism can capture the orthodox Evidence-first view of doxastic justification. Crucially, the accessibility requirement is exactly what gives rise to the Access Problem. On the other hand, Williamson's E=K successfully resists the Access Problem, insofar as it is not committed to such a strong accessibility requirement. At the same time, by reducing doxastic justification to knowledge, his account of justification, which can be defined as an instance of *Knowledge-first Evidentialism*, does not vindicate the *Responsiveness Intuition*. Interestingly, the *Ecumenical Evidentialism* I have developed in the last Chapter, is a Knowledge-first version of Evidential Externalism that vindicates the Responsiveness Intuition. It vindicates the *Responsiveness Intuition* because it takes doxastic justification to be a matter of being *evidence-responsive*, and it can do so without posing any strong accessibility requirement. It is an instance of *Evidential Externalism* because, by embracing E=K, it denies the thesis that evidence supervenes on non-factive mental states. It is *Knowledge-first* because, besides understanding cases of non-inferential knowledge in a Knowledge-first fashion, it explains evidence and justification in terms of knowledge.

References

- Alston, W. P. (1988). The deontological conception of epistemic justification. *Philosophical Perspectives*, 2(n/a), 257–299.
- Alvarez, M. (2008). Reasons and the ambiguity of 'belief'. *Philosophical Explorations*, 11(1), 53–65.
- Alvarez, M. (2009). How many kinds of reasons? *Philosophical Explorations*, 12(2), 181–193.
- Alvarez, M. (2010). *Kinds of reasons: An essay in the philosophy of action*. Oxford University Press.
- Anscombe, G. E. M. (1962). On sensations of position. *Analysis*, 22(3), 55–58.
- Ashton, N. A. (2015). Undercutting underdeterminationbased scepticism. *Theoria*, 81(4), 333–354.
- Audi, R. (2001). An internalist theory of normative grounds. *Philosophical Topics*, 29(1/2), 19–46.
- Austin, J. L. (1962). *Sense and sensibilia*. Oxford University Press.
- Ayer, A. J. (1936). *Language, truth and logic*. London: V. Gollancz.
- Bergmann, M. (2006). *Justification without awareness: A defense of epistemic externalism*. Oxford University Press.
- Bird, A. (2007). Justified judging. *Philosophy and Phenomenological Research*, 74(1), 81–110.
- Boghossian, P. A. (1998). What the externalist can know "a priori". *Philosophical Issues*, 9(2), 197–211.

- BonJour, L. (1978). Can empirical knowledge have a foundation? *American Philosophical Quarterly*, 15(1), 1–14.
- BonJour, L. (1985). *The structure of empirical knowledge*. Harvard University Press.
- BonJour, L. (1999). Foundationalism and the external world. *Philosophical Perspectives*, 13(s13), 229–249.
- BonJour, L., & Sosa, E. (2003). *Epistemic justification: Internalism vs. externalism, foundations vs. virtues*. Wiley-Blackwell.
- Boult, C. (2017). Epistemic normativity and the justification-excuse distinction. *Synthese*, 194(10), 4065–4081.
- Brown, J. (2004). *Anti-individualism and knowledge*. MIT Press.
- Brown, J. (2013). Infallibilism, evidence and pragmatics. *Analysis*, 73(4), 626–635.
- Brueckner, A. (1994). The structure of the skeptical argument. *Philosophy and Phenomenological Research*, 54(4), 827–835.
- Brueckner, A. (2005a). Fallibilism, underdetermination, and skepticism. *Philosophy and Phenomenological Research*, 71(2), 384–391.
- Brueckner, A. (2005b). Knowledge, evidence, and skepticism according to williamson. *Philosophy and Phenomenological Research*, 70(2), 436–443.
- Brueckner, A. (2009). E = k and perceptual knowledge. In P. Greenough & D. Pritchard (Eds.), *Williamson on knowledge*. Oxford University Press.
- Brueckner, A., & Buford, C. T. (2012). A tale of two fallibilists: On an argument for infallibilism. *Thought: A Journal of Philosophy*, 1(3), 195–199.
- Burge, T. (1998). Computer proof, apriori knowledge, and other minds. *Noûs*, 32(S12), 1–37.
- Byrne, A. (2004). How hard are the sceptical paradoxes? *Noûs*, 38(2), 299–325.
- Chisholm, R. M. (1988). The indispensability of internal justification. *Synthese*, 74(3), 285–96.
- Cohen, S. (1984). Justification and truth. *Philosophical Studies*, 46(3), 279–95.

- Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives*, 2(n/a), 91–123.
- Cohen, S. (1998). Two kinds of skeptical argument. *Philosophical and Phenomenological Research*, 58(1), 143–159.
- Comesana, J. (2010). Evidentialist reliabilism. *Noûs*, 44(4), 571–600.
- Conee, E., & Feldman, R. (1998). The generality problem for reliabilism. *Philosophical Studies*, 89(1), 1–29.
- Conee, E., & Feldman, R. (2004). *Evidentialism: Essays in epistemology*. Oxford University Press.
- Conee, E., & Feldman, R. (2008). Evidence. In Q. Smith (Ed.), *Epistemology: New essays*. Oxford University Press.
- Dennis, P. (2014). Criteria for indefeasible knowledge: John McDowell and 'epistemological disjunctivism'. *Synthese*, 191(17), 4099–4113.
- Dodd, D. (2007). Why Williamson should be a sceptic. *Philosophical Quarterly*, 57(229), 635–649.
- Dodd, D. (2011). Against fallibilism. *Australasian Journal of Philosophy*, 89(4), 665–685.
- Dretske, F. (1981). *Knowledge and the flow of information*. MIT Press.
- Dunn, J. (2012). Evidential externalism. *Philosophical Studies*, 158(3), 435–455.
- Dutant, J. (2007). The case for infallibilism. In C. Penco, M. Vignolo, V. Ottonelli, & C. Amoretti (Eds.), *Proceedings of the 4th latin meeting in analytic philosophy* (pp. 59–84). Genoa: University of Genoa.
- Earman, J. (1993). Underdetermination, realism, and reason. *Midwest Studies in Philosophy*, 18(1), 19–38.
- Engel, M. (1992). Personal and doxastic justification in epistemology. *Philosophical Studies*, 67(2), 133–150.
- Farkas, K. (2008). *The subject's point of view*. Oxford University Press.
- Feldman, R. (1981). Fallibilism and knowing that one knows. *Philosophical Review*, 90(2), 266–282.

- Feldman, R. (1985). Reliability and justification. *The Monist*, 68(2), 159–174.
- Firth, R. (1978). Are epistemic concepts reducible to ethical concepts? In A. Goldman & J. Kim (Eds.), *Values and morals: Essays in honor of william frankena, charles stevenson, and richard brandt* (pp. 215–229). Dordrecht: Kluwer Academic Publishers.
- Fratantonio, G., & McGlynn, A. (2018). Reassessing the case against evidential externalism. In V. Mitova (Ed.), *The factive turn in epistemology*. Cambridge: Cambridge University Press.
- Gerken, M. (2011). Warrant and action. *Synthese*, 178(3), 529–547.
- Ghijzen, H. (2015). The basis problem for epistemological disjunctivism revisited. *Erkenntnis*, 80(6), 1147–1156.
- Gillies, D. (2000). *Philosophical theories of probability*. Routledge.
- Ginet, C. (1975). Knowledge, perception and memory. *Philosophical Quarterly*, 26(104), 279–280.
- Goldman, A. I. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge* (pp. 1–25). Boston: D. Reidel.
- Goldman, A. I. (1986). *Epistemology and cognition*. Harvard University Press.
- Goldman, A. I. (1989). Interpretation psychologized. *Mind and Language*, 4(3), 161–185.
- Goldman, A. I. (2009). Internalism, externalism, and the architecture of justification. *Journal of Philosophy*, 106(6), 309–338.
- Goldman, A. L. (2008). *Simulating minds: The philosophy, psychology, and neuroscience of mindreading*. Oup Usa.
- Gordon, R. M. (1986). Folk psychology as simulation. *Mind and Language*, 1(2), 158–71.
- Greco, D. (2014). Could kk be ok? *Journal of Philosophy*, 111(4), 169–197.
- Greco, J. (2014). Pritchard's epistemological disjunctivism: How right? how radical? how satisfying? *Philosophical Quarterly*, 64(254), 115–122.

- Hájek, A. (2008). Interpretations of probability. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy*. [Urlhttp://Plato.Stanford.Edu/Archives/Spr2010/Entries/Probability-Interpret/](http://Plato.Stanford.Edu/Archives/Spr2010/Entries/Probability-Interpret/).
- Hájek, A. (2008). A philosopher's guide to probability. In G. Bammer & M. Smithson (Eds.), *Uncertainty and risk: Multidisciplinary perspectives*. Routledge.
- Hawthorne, J. (2003). *Knowledge and lotteries*. Oxford University Press.
- Hawthorne, J. (2005). The case for closure. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology* (pp. 26–43). Blackwell.
- Hawthorne, J., & Lasonen-Aarnio, M. (2009). Knowledge and objective chance. In P. Greenough & D. Pritchard (Eds.), *Williamson on knowledge* (pp. 92–108). Oxford University Press.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105(10), 571–590.
- Heal, J. (1986). Replication and functionalism. In J. Butterfield (Ed.), *Language, mind, and logic* (pp. 135–150). Cambridge University Press.
- Hicks, R. D., & Aristotle. (1907). *Aristotle, de anima*. Cambridge University.
- Hintikka, J. (1962). *Knowledge and belief*. Ithaca: Cornell University Press.
- Hyman, J. (1999). How knowledge works. *Philosophical Quarterly*, 49(197), 433–451.
- Hyman, J. (2006). Knowledge and evidence. *Mind*, 115(460), 891–916.
- Ichikawa, J. J. (2014). Justification is potential knowledge. *Canadian Journal of Philosophy*, 44(2), 184–206.
- Jackson, M. B. (2016). Knowing by imagining. In A. Kind & P. Kung (Eds.), *Knowledge through imagination*. Oxford University Press.
- Jackson, M. B. (2018). Justification by imagination.
- Jenkins, C. S. (2008). *Grounding concepts: An empirical basis for arithmetical knowledge*. Oxford University Press.
- Kallestrup, J. (2011a). Recent work on mckinsey's paradox. *Analysis*, 71(1), 157–171.

- Kallestrup, J. (2011b). *Semantic externalism*. Routledge.
- Kelly, T. (2008). Evidence: Fundamental concepts and the phenomenal conception. *Philosophy Compass*, 3(5), 933–955.
- Kennedy, M. (2010). Naive realism and experiential evidence. *Proceedings of the Aristotelian Society*, 110(1pt1), 77–109.
- Klein, P. D. (2005). Infitism is the solution to the epistemic regress problem. In S. Matthias & S. Ernest (Eds.), *Contemporary debates in epistemology*. Blackwell.
- Korcz, K. A. (2010). *The epistemic basic relation*. Stanford Encyclopedia of Philosophy.
- Kraft, T. (2015). Epistemological disjunctivism's genuine access problem. *Theoria*, 81(4), 311–332.
- Kukla, A. (1998). *Studies in scientific realism*. Oxford University Press.
- Kvanvig, J. L., & Menzel, C. (1990). The basic notion of justification. *Philosophical Studies*, 59(3), 235–261.
- Lasonen-Aarnio, M. (2010). Is there a viable account of well-founded belief? *Erkenntnis*, 72(2), 205–231.
- Laudan, L., & Leplin, J. (1991). Empirical equivalence and underdetermination. *Journal of Philosophy*, 88(9), 449–472.
- Lehrer, K. (1990). *Theory of knowledge*. Routledge.
- Lehrer, K., & Cohen, S. (1983). Justification, truth, and coherence. *Synthese*, 55(2), 191–207.
- Lewis, D. (1980). A subjectivist's guide to objective chance. In R. C. Jeffrey (Ed.), *Philosophy of probability: Contemporary readings* (pp. 83–132). University of California Press.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Littlejohn, C. (2008). From $e = k$ to scepticism? *Philosophical Quarterly*, 58(233), 679–684.

- Littlejohn, C. (2011). Evidence and armchair access. *Synthese*, 179(3), 479–500.
- Littlejohn, C. (2012). *Justification and the truth-connection*. Cambridge University Press.
- Littlejohn, C. (2018a). Evidence and its limits. In C. M. J. Way & D. Whiting (Eds.), *Normativity: Epistemic and practical*. Oxford University Press.
- Littlejohn, C. (2018b). How and why knowledge is first. In A. Carter, E. Gordon, & B. Jarvis (Eds.), *Knowledge first*. Oxford University Press.
- Littlejohn, C. (in press). A plea for epistemic excuses. In F. D. J. Dutant (Ed.), *The new evil demon problem*. Oxford University Press.
- Lowy, C. (1978). Gettier's notion of justification. *Mind*, 87(345), 105–108.
- McDowell, J. (1995). Knowledge and the internal. *Philosophy and Phenomenological Research*, 55(4), 877–93.
- McGlynn, A. (2012). Justification as 'would-be' knowledge. *Episteme*, 9(4), 361–376.
- McGlynn, A. (2014). *Knowledge first?* Palgrave Macmillan.
- McKinsey, M. (1991). Anti-individualism and privileged access. *Analysis*, 51(1), 9–16.
- Melis, G. (2018). The intertwining of propositional and doxastic justification. *Australasian Journal of Philosophy*, 96(2), 367–379.
- Mitova, V. (2015). Truthy psychologism about evidence. *Philosophical Studies*, 172(4), 1105–1126.
- Moore, G. E. (1959). *Philosophical papers*. New York: Macmillan.
- Neta, R. (2008). What evidence do you have? *British Journal for the Philosophy of Science*, 59(1), 89–119.
- Nozick, R. (1981). *Philosophical explanations*. Harvard University Press.
- Nuccetelli, S. (1999). What anti-individualists cannot know a priori. *Analysis*, 59(1), 48–51.

- Nuccetelli, S. (2003). *New essays on semantic externalism and self-knowledge*. MIT Press.
- Pollock, J., & Cruz, J. (1999). *Contemporary theories of knowledge*. Lanham, MD: Rowman and Littlefield.
- Pritchard, D. (2005a). *Epistemic luck*. Oxford University Press UK.
- Pritchard, D. (2005b). The structure of sceptical arguments. *Philosophical Quarterly*, 55(218), 37–52.
- Pritchard, D. (2011a). Epistemological disjunctivism and the basis problem. *Philosophical Issues*, 21(1), 434–455.
- Pritchard, D. (2011b). Evidentialism, internalism, disjunctivism. In T. Dougherty (Ed.), *Evidentialism and its discontents*. Oxford University Press.
- Pritchard, D. (2012). *Epistemological disjunctivism*. Oxford University Press.
- Pritchard, D. (2015). Epistemic dependence. *Philosophical Perspectives*, 29(1), 305–324.
- Pritchard, D. (2016a). *Epistemic angst: Radical skepticism and the groundlessness of our believing*. Princeton University Press.
- Pritchard, D. (2016b). Epistemic risk. *Journal of Philosophy*, 113(11), 550–571.
- Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*, 34(4), 517–549.
- Psillos, S. (1999). *Scientific realism: How science tracks truth*. Routledge.
- Putnam, H. (1975). The meaning of 'meaning'. *Minnesota Studies in the Philosophy of Science*, 7, 131–193.
- Radford, C. (1966). Knowledge—by examples. *Analysis*, 27(1), 1–11.
- Ranalli, C. (2014). Luck, propositional perception, and the entailment thesis. *Synthese*, 191(6), 1223–1247.
- Reed, B. (2002). How to think about fallibilism. *Philosophical Studies*, 107(2), 143–157.

- Sawyer, S. (1998). Privileged access to the world. *Australasian Journal of Philosophy*, 76(4), 523–533.
- Silins, N. (2005). Deception and evidence. *Philosophical Perspectives*, 19(1), 375–404.
- Silva, P. (2015). On doxastic justification and properly basing one's beliefs. *Erkenntnis*, 80(5), 945–955.
- Smith, M. (2010a). A generalised lottery paradox for infinite probability spaces. *British Journal for the Philosophy of Science*, 61(4), 821–831.
- Smith, M. (2010b). What else justification could be. *Noûs*, 44(1), 10–31.
- Smith, M. (2016a). *Between probability and certainty: What justifies belief*. Oxford University Press UK.
- Smith, M. (2016b). Scepticism by a thousand cuts. *International Journal for the Study of Skepticism*, 6(1), 44–52.
- Sosa, E. (1999). How to defeat opposition to moore. *Philosophical Perspectives*, 13(s13), 137–49.
- Sosa, E. (2005). Epistemic competence and judgment. In M. A. F. Vargas (Ed.), *Performance epistemology: Foundations and applications*.
- Stanley, J. (2005). Fallibilism and concessive knowledge attributions. *Analysis*, 65(2), 126–131.
- Stanley, J. (2008). Knowledge and certainty. *Philosophical Issues*, 18(1), 35–57.
- Sutton, J. (2005). Stick to what you know. *Noûs*, 39(3), 359–396.
- Sylvan, K. (2016). Epistemic reasons i: Normativity. *Philosophy Compass*, 11(7), 364–376.
- Tulodziecki, D. (2012). Epistemic equivalence and epistemic incapacitation. *British Journal for the Philosophy of Science*, 63(2), 313–328.
- Turri, J. (2010). On the relationship between propositional and doxastic justification. *Philosophy and Phenomenological Research*, 80(2), 312–326.
- Unger, P. (1971). A defense of skepticism. *Philosophical Review*, 80(2), 198–219.

- Unger, P. (1975). *Ignorance: A case for scepticism*. Oxford University Press.
- Williamson, T. (1998). Conditionalizing on knowledge. *British Journal for the Philosophy of Science*, 49(1), 89–121.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.
- Williamson, T. (2007a). How probable is an infinite sequence of heads? *Analysis*, 67(3), 173–180.
- Williamson, T. (2007b). *The philosophy of philosophy*. Wiley-Blackwell.
- Williamson, T. (2009a). Probability and danger. In *Amherst lecture in philosophy* (pp. 1–35).
- Williamson, T. (2009b). Replies to critics. In D. Pritchard & P. Greenough (Eds.), *Williamson on knowledge* (pp. 279–384). Oxford: Oxford University Press.
- Williamson, T. (2011). Knowledge first epistemology. In S. Bernecker & D. Pritchard (Eds.), *The routledge companion to epistemology* (pp. 208–218). Routledge.
- Williamson, T. (2013). How deep is the distinction between a priori and a posteriori knowledge? In A. Casullo & J. C. Thurow (Eds.), *The a priori in philosophy* (pp. 291–312). Oxford University Press.
- Williamson, T. (2014). Very improbable knowing. *Erkenntnis*, 79(5), 971–999.
- Williamson, T. (2016). Knowing by imagining. In A. Kind & P. Kung (Eds.), *Knowledge through imagination* (p. 113–123). Oxford University Press.
- Williamson, T. (2017). Ambiguous rationality. *Episteme*, 14, 263–274.
- Williamson, T. (in press). Justification, excuses, and sceptical scenarios. In J. Dutant (Ed.), *The new evil demon: New essays on justification, knowledge, and rationality*. Oxford University Press.
- Wittgenstein, L. (1969). *On certainty* (ed. anscombe and von wright). Harper Torchbooks.
- Wright, C. (2004). Warrant for nothing (and foundations for free)? *Aristotelian Society Supplementary Volume*, 78(1), 167–212.
- Zalabardo, J. (2015). Epistemic disjunctivism and the evidential problem. *Analysis*, 75(4), 615–627.