
On the Nature and Relationship of
Individual and Collective Justification

Simon Graf

Submitted in accordance with the requirements for the degree of

Doctor of Philosophy

The University of Leeds

School of Philosophy, Religion and History of Science

May 2024

I confirm that the work submitted is my own and that appropriate credit has been given where reference has been made to the work of others.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

The right of Simon Graf to be identified as Author of this work has been asserted by Simon Graf in accordance with the Copyright, Designs and Patents Act 1988.

© 2024 The University of Leeds and Simon Graf.

Some parts of this thesis are based on the following published work, which will not be referenced in the text:

Graf, Simon (2023). Permissive Divergence. *Canadian Journal of Philosophy*, 53 (3), 240-255. <https://doi.org/10.1017/can.2024.4>.

Graf, Simon (2021). Review: The Epistemology of Groups by Jennifer Lackey. *Perspectives: International Postgraduate Journal of Philosophy*, 9 (1), 380-387.

Acknowledgements

A major chunk of this thesis was written on campus in Leeds, at the School of Philosophy, Religion and History of Science. More precisely, in the silent study room 02.08 on the 2nd floor of Botany house. A peculiar research environment, sometimes vacant and covered by an aura of contemplative solitude and sometimes a lively and energetic debate club. I developed some early thoughts and did some final polishing during my remote studies in Austria, either in ‘my office’ a.k.a. my old room in my brother’s flat or at various places I lived together with Melissa, my beloved partner.

During my time in Leeds, I was privileged to have the opportunity of being a postgraduate researcher at the ERC-funded research project *Group Thinking: New Foundations* (in short: *GROUNDS*). The group’s weekly meetings, informal post-seminar discussions and regular workshops with international speakers provided a vibrant research environment. I had the luck of working directly not only with the PI of the project Robbie Williams but also with all three of the post-docs who joined the project over the course of four years Haixin Dang, Andrew Peet, and Joshua-Habgood Coote.

Consequently, for helpful feedback and discussions, I am first and foremost grateful to my main supervisor Robbie Williams, who always provided very useful and thoughtful feedback on my half-baked chapter drafts. It was really Robbie who guided me through this academic journey from start to finish and who shaped my philosophical development in the most positive way imaginable. Thanks, Robbie, you are incredible. A special thanks also goes to Andrew Peet whose supervisions were immensely helpful. Many of the central ideas of this thesis emerged from discussions with him; especially the methodological interlude on epistemic continuity (chapter 4) is basically further developing an idea introduced in one of his unpublished manuscripts. I also want to thank Haixin Dang, who supervised me in my first year during the pandemic, an especially challenging and unusual time –and a stage of the PhD in which you are more lost and confused than anything else! Many thanks also to my final and fourth supervisor Joshua Habgood-Coote who helped me to substantially refine the thesis chapter by chapter and get it into the shape it currently is.

Much gratitude also goes to the other *GROUNDS* members Thomas Brouwer, Heather Logue, Ludovica Adamo, Daniel Elstein and Matt Clark, who always gave me constructive feedback when

I did some of my ‘conference presentation dry runs’ at the seminar (which I feel like I did way too often). I also want to mention my PhD colleagues, including Sam Mason, Radu Bumbacea, Jyl Schacher, Maddy Page, Bart Smith-Johnson, Katie Prosser, Keiran Bateman, Huaming Xu, Zach Gudmunsen, Miriam Bowen, Yihan Jiang, and Ben Davis who provided me useful feedback, during or after one of my countless PG seminar presentations, in the pub, or during one of our reading groups. Among my fellow PhD contemporaries, special appreciation goes to my colleague, friend, and PG seminar co-convenor Rodrigo Valencia whose support made the hardships of my PhD life in Leeds more bearable; as well as to Jakob Donskov, with whom I had many inspiring discussions and whose bright mind often managed to enlighten issues I previously was in the dark about. You guys are amazing.

I also want to mention my former colleague Raimund Pils who gave me feedback on various paper drafts which helped me to get them published, which was a long and painful process. As well as Sandy Goldberg who gave me valuable feedback on an early draft of chapter 7, in which I discuss his work on epistemic responsibility. Another big shout-out goes to Harald Wiltsche who supervised me during my undergraduate studies in Austria and who encouraged me to pursue a career in philosophy. Not to forget countless audiences at the University of Leeds’ Postgraduate and Thursday Philosophy seminar, the University of Salzburg’s Sophia conference, 2021, 2022, 2023, the ISOS conference at the University of Vienna 2022, and the University of San Diego 2021, the Centre for interdisciplinary research at the University of Bielefeld, the law department at the Jagiellonian University in Kraków, as well as the Irish Society for the Philosophy of Women in Maynooth. Among the people I met on one or multiple of these occasions, I remember giving me especially helpful feedback are Jesper Kallestrup, Brian Hedden, Chuizu Dawn Wang, Jessica Brown, Anne Schwenkenbecher, Christian List, Krzysztof Pośłajko, Nathan Biebel, Rik Peels, Konrad Werner, Edward Elliot, Gerald Lang, Nicholas Makins, Pekka Väyrynen, Adam Dyrda, Philipp Schoenegger, Lukas Schwengerer, Bill Wringe, Lily Tappe, Max Gab, Jessica Isserow, Daniel Garibay-Garcia, Heng Ying, Mike Schneider and probably a handful of people I forgot. Not to neglect the (too) many anonymous referees from multiple journals, such as Ergo, AJP, Synthese, CJP, Episteme and Philosophical Studies.

Many thanks also to my family, my mother Daniela Graf, my father Gregor Pokorny and my brother Lukas Graf, as well as my partner Melissa Mukeli who accompanied me on this voyage despite all of the difficulties that naturally come with having a long-distance relationship. Last but

definitely not least I want to express my love for my second family members, most of whom I have known almost as long as my family, Jonathan Schachenhofer, Daniel Schmalnauer, Jonathan Schmalnauer, Severin Gruber and Christoph Szalay. Without your constant support and warm-hearted community, I definitely wouldn't have made it this far.

The research leading to this thesis has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 818633).

Abstract

This thesis is an investigation into the nature of epistemic justification. It brings together themes from traditional, individual-centred epistemology, and collective, group-centred epistemology.

The first half of the thesis is concerned with the question of whether rationality is epistemically permissive; that is, whether one body of evidence can rationalise more than one doxastic attitude. In chapter 1, I argue that permissive cases are best understood as epistemic standard conflicts. Doing so provides us with a novel understanding of the arbitrariness objection against permissivism and enables us to reduce questions about epistemic permissibility to questions about the nature of incommensurability. In chapter 2, I show that the defended understanding of permissive cases generalises by defending it against an objection from self-fulfilling beliefs. In chapter 3, I demonstrate that we can use this view of epistemic rationality to generate so-called divergence arguments which show that the epistemic status of group-level attitudes and member-level attitudes can rationally diverge.

In the second half of the thesis, I develop a novel evidentialist theory of epistemic justification, called Continuous Evidentialism. Continuous Evidentialism is inspired by some general methodological reflections (chapter 4), which suggest that we should opt for a theory of epistemic justification that analyses the epistemic status of group-level attitudes and member-level attitudes continuously. According to Continuous Evidentialism, to have a justified belief is to possess sufficient evidence and utilise that evidence in an epistemically responsible way when forming the belief. While I argue that we can reduce epistemic responsibility to higher-order evidentialist requirements. In chapter 5, I develop a theory of evidence, evidence possession and epistemic basing. Chapters 6 - 8, discuss various complications of the proposed theory, having to do with the alleged defeasibility of justification (chapter 6), epistemic responsibility (chapter 7), and the proposed reduction of epistemic responsibility (chapter 8). In chapter 9, I compare Continuous Evidentialism to various extant accounts.

Contents

Acknowledgements.....	iii
Abstract.....	vii
Abbreviations.....	xi
General Introduction.....	1
<i>Part I: Epistemic Conflicts</i>	
1. Epistemic Conflicts.....	7
1.1 Introduction.....	7
1.2 Standard Uniqueness and Standard Permissivism.....	11
1.3 Value Conflicts and Standard Conflicts.....	13
1.4 Incommensurability: Three Features.....	16
1.4.1 Sweetening Insensitivity.....	17
1.4.2 Bindingness.....	19
1.4.3 Angst.....	20
1.5 Incommensurability and Permissibility.....	22
1.6 Arbitrariness, Bindingness, and Changes of Mind.....	25
1.6.1 Strong Bindingness.....	26
1.6.2 Weak Bindingness.....	27
1.6.3 Changes of Mind.....	31
1.7 Conclusion.....	33
2. Uniqueness, Permissivism and Self-Fulfilling Beliefs.....	34
2.1 Introduction.....	34
2.2 Self-Fulfilling Beliefs.....	36
2.3 Theoretical Deliberation and Self-Fulfilling Beliefs.....	40
2.4 Practical Deliberation and Self-Fulfilling Beliefs.....	43
2.5 The Impermissibility of Self-Fulfilling Beliefs.....	47
2.6 Genuine Indifference: A recipe.....	50
2.6.1 The Impossibility of Indifference?.....	52
2.6.2 Evidential Obligations.....	53
2.6.3 Doxastic Voluntarism Revisited.....	54
2.7 Picking Outcomes and Choosing Attitudes.....	55
2.8 Conclusion.....	57
3. Epistemic Divergence.....	58
3.1 Introduction.....	58
3.2 Summativism and Divergence.....	60
3.3 Types of Permissivism.....	64
3.4 Permissive Divergence: The Argument.....	68
3.5 Permissive Divergence: Discussion.....	72
3.6 Permissive Divergence: Comparison.....	76
3.7 Conclusion.....	79
Part I: Conclusion.....	82

Methodological Interlude: Epistemic Continuity

4. A Plea for a United Epistemology.....	85
4.1 Introduction.....	85
4.2 Collective and Individual Epistemology.....	86
4.3 Methodism and Particularism.....	89
4.4 The Method of Cases: Three Principles.....	91
4.5 Continuous Epistemology.....	94
4.5 Continuous Explanations.....	98
4.7 The Contemporary Landscape.....	101
4.8 Conclusion.....	106

Part II: Continuous Justification

Part II: Introduction.....	107
5. Rethinking Epistemic Justification.....	109
5.1 Introduction.....	109
5.2 Evidentialism.....	111
5.3 Evidence.....	114
5.4 Evidence Possession.....	119
5.5 Epistemic Support.....	121
5.5.1 Probabilism.....	123
5.5.2 Explanationism	125
5.5.3 A Hybrid View?.....	126
5.5.4 Subjectivism vs. Objectivism.....	128
5.6 Epistemic Basing.....	129
5.6.1 Proper Basing.....	131
5.7 Continuous Evidentialism.....	134
5.8 Conclusion.....	140
6. The No-Defeater Clause.....	144
6.1 Introduction.....	145
6.2 Types of Evidence.....	146
6.3 Defeaters.....	150
6.4 Evidentialism and Defeat.....	155
6.4.1 Evidentialism and Propositional Defeaters.....	157
6.4.2 Evidentialism and Mental State Defeaters.....	160
6.5 Responsibilism: Two Desiderata.....	165
6.6 Responsibility, Balance and Location.....	167
6.7 Propositional and Doxastic Justification: Revisited.....	168
6.8 Conclusion.....	173
7. Epistemic Responsibilism.....	174
7.1 Introduction.....	174
7.2 Epistemic Responsibility: The Entitlement Model.....	176
7.3 Sensitivity and Epistemic Strict Liability.....	178
7.3.1 Epistemic Ceiling 1: First-Order and Higher-Order Evidence.....	179

7.3.1 Epistemic Ceiling 2: Evidence You Should Have Had.....	182
7.3.1 Against Epistemic Strict Liability.....	183
7.4 Reducing Responsibility: Higher-Order Evidence.....	186
7.4.1 Responsibility and Higher-Order Evidence.....	186
7.4.2 A Higher-Order Evidence Clause	188
7.4.3 Why do we need Higher-Order Evidence?.....	190
7.5 Defeatism Defeated?.....	193
7.6 Conclusion.....	195
8. Higher-Order Requirements.....	196
8.1 Introduction.....	196
8.2 Defencism Defended.....	197
8.2.1 Higher-Order Requirements: Two Desiderata.....	198
8.2.2 Epistemic Basing and Higher-Order Requirements.....	202
8.2.3 Summary.....	203
8.3 Higher-Order Requirements: Overview.....	203
8.4 Higher-Order Requirements: Worries.....	206
8.4.1 Overintellectualisation.....	206
8.4.2 Epistemic Relevance.....	210
8.4.3 Vicious Regress.....	213
8.4.5 Summary.....	217
8.5 Evidential Requirements.....	218
8.6 Conclusion.....	224
9. Continuous Evidentialism.....	225
9.1 Introduction.....	225
9.2 Continuous Evidentialism: Comparison.....	227
9.3 Continuous Evidentialism: Application.....	229
9.3.1 Distributed Evidence.....	230
9.3.2 Collective Defeat.....	233
9.3.3 Collective Basing.....	239
9.3.4 Evidential Manipulation	241
9.4 The Big Picture.....	244
9.4.1 Epistemic Divergence: Revisited.....	245
9.4.1 The Best of Both Worlds.....	246
9.4.2 The Possibility of Continuous Epistemology.....	248
10. General Conclusion.....	249
10.1 Recapitulation.....	249
10.2 Outlook.....	255
Bibliography.....	257

Abbreviations

<Agency>	Epistemic Agency: There is a general way to characterise and identify epistemic agents. That is, there is one universal set of characteristics that enables us to identify potential bearers of epistemic states.
<E _{BASE} >:	The Evidential Base: The total evidence <i>possessed</i> by an epistemic agent S.
<E _{TOTAL} >	The total Evidence: The total evidence that is <i>relevant</i> to S's epistemic situation.
<ERG>	Evidentialist Responsibilism for Groups*
<Evidentialism>	Evidentialist No-Defeater Clause: S's <i>prima facie</i> justified belief that p is undefeated iff S's evidential base <i>on balance supports</i> p.
<Extended Evidentialism>	Extended Evidentialist No-Defeater Clause: S's <i>prima facie</i> justified belief that p is undefeated iff the <i>total evidence on balance supports</i> p.
<GEAA>	Group Epistemic Agent Account*
<HOR+>	Positive Higher-Order Requirement: If S's belief that p is justified then S <i>necessarily</i> has a <i>defender</i> Φ , that is a [belief/justified belief/evidence/...] <i>indicating</i> that p is justified.
<HOR+ dispositional>	Dispositional Positive Higher-Order Requirement: If S's belief that p is justified then S <i>necessarily</i> has Φ , a <i>disposition</i> to form beliefs in a way that <i>indicates</i> that p is justified, and this disposition caused S to believe that p.
<HOR+ doxastic>	Doxastic Positive Higher-Order Requirement: If S's belief that p is justified then S <i>necessarily</i> has Φ , a <i>doxastic</i> mental state that <i>indicates</i> that p is justified.
<HOR+ evidential>	Evidential Positive Higher-Order Requirement: If S's belief that p is justified then S <i>necessarily</i> has Φ , sufficient <i>evidence</i> that <i>indicates</i> that p is justified.
<HOR+ mental>	Mental Positive Higher-Order Requirement: If S's belief that p is justified then S <i>necessarily</i> has Φ , a <i>non-doxastic</i> mental state that <i>indicates</i> that p is justified.
<HOR->	Negative Higher-Order Requirement: S's <i>prima facie</i> justified belief that p is <i>ultima facie</i> justified iff S lacks a <i>defeater</i> ϕ [that S should have had], that is [a belief/justified belief/evidence...] <i>indicating</i> that the belief that p is <i>not prima facie</i> justified.

* See [S9.2] for a full definition of <ERG> and <GEAA>.

- <InterP> **Interpersonal Permissivism:** There are cases in which it is rationally permissible for one agent A to have the doxastic attitude D towards p, and for another agent A* to have a different doxastic attitude D* towards p instead, given the same body of evidence E.
- <IntraU> **Intrapersonal Uniqueness:** For any body of evidence E and any proposition p, there is a unique rationally permissible doxastic attitude D towards p for any epistemic agent A and their epistemic standards S, possessing that body of evidence E.
- <MESc> **Minimal Epistemic Summativist Commitment:** A group's doxastic attitude D towards p has the epistemic status S only if there is at least one member of G who has the doxastic attitude D towards p and for whom D has the epistemic status S.
- <MESc*> **Minimal Epistemic Summativist Commitment*:** A group's attitude D towards p has the epistemic status E only if there is at least one member of G for whom D has epistemic status E *qua* group member.
Positive Higher-Order Requirement:
- <MSC> **Minimal Summativist Commitment:** A group G has the doxastic attitude D towards p only if at least one individual m is both a member of G and has the doxastic attitude D towards p.
- <MSS> **Minimal Standard Summativism:** A group G has an epistemic standard S only if at least one member of G has the epistemic standard S.
- <non-Primacy> **Methodological non-Primacy:** Epistemological theorising has no preferred starting point. That is, all case judgements should *ceteris paribus* be weighted in the same way.
- <PDTemplate> **Permissive Divergence Template:** There are cases in which a group G is rationally required to have the doxastic attitude D towards p in relation to a body of evidence E and its permissible epistemic standards S. While all of G's members m_1 - m_n possess the same body of evidence E, they have different permissible epistemic standards S_1 - S_n which recommend different doxastic attitudes D_1 - D_n towards p.
- <Responsibilism> **Responsibilist No-Defeater Clause:** S's *prima facie* justified belief that p is undefeated iff S is *epistemically responsible* in believing that p.
- <Unity> **Conceptual Unity:** Epistemological concepts are general. That is, whenever we evaluate some type of epistemic state of different types of epistemic agents, we evaluate them under one unified concept.

General Introduction

Some of our beliefs are better than others. My belief that the academic job market is going to improve is undoubtedly a result of wishful thinking. In contrast, my belief that this thesis is among the last which have been written without support from AI-assisted language tools is probably not. We might say the latter belief is *rational* or *epistemically justified* while the former is not. This epistemic assessment of our beliefs is at the centre of the following investigations.

Traditionally, epistemic justification has received a lot of attention in epistemology. For one, (pre-Gettier) epistemic justification has often been thought to play a central role in the traditional analysis of knowledge, which understands knowledge as justified true belief.² Moreover, (pre-Williamson) epistemic justification was widely acknowledged to be the first or primary way to evaluate the normative standings of our beliefs.³ Despite these developments and open questions about the place of epistemic justification within our overall epistemic framework, epistemic justification has continued to play a central role in our epistemic theorising. Especially, during the so-called social turn in epistemology new questions about epistemic justification emerged having to do with socio-epistemic issues, such as testimony or peer disagreement.⁴ More recently, many have taken epistemic justification to the collective realm. That is, many have started to study the epistemic status of collective attitude ascriptions, such as ‘the colleague union believes that the industrial action is going to plan’ and developed theories of collective justification or group justified belief.⁵

This thesis aims to contribute to these endeavours in various ways. First, I agree with those who argued that we ought to overcome the restrictions of traditional, individual-centred epistemology and advocate for a social or collective understanding. In particular, I will analyse the conflicts we face as socio-epistemic agents and draw some general lessons about epistemic permissibility and the relationship between collective and individual justification from that.

² Gettier (1963).

³ Williamson (2000).

⁴ See, for example, Goldman (1999, 2001), Christensen (2007), Goldberg (2008), Lackey (2008, 2010), Elga (2007), Kelly (2010).

⁵ See, for example, Schmitt (1994), Hakli (2011), Goldman (2014), Lackey (2016; 2021), Hedden (2019) or Silva (2019).

Second, I will offer a novel theory of epistemic justification that can be used to analyse the epistemic status of individual and collective doxastic attitudes.

Accordingly, this thesis is divided into two parts and a methodological interlude in-between the two parts:

Part I: Epistemic Conflicts

Methodological Interlude: Epistemic Continuity

Part II: Continuous Justification

In the rest of this introduction, I will sketch the contents of each part chapter by chapter.

Part I: Epistemic Conflicts

I chose an unconventional entry point into the topic of epistemic justification. Namely, I begin with an observation that I take to be a brute fact of our epistemic life; the observation that we as epistemic agents sometimes face conflicts that demonstrate to us that there are alternative, seemingly equally rational ways to model the world.

The first chapter [ch.1] sets out to compare these epistemic conflicts to so-called value conflicts, which are cases in which the different values we subscribe to give us conflicting recommendations. In so doing, I will show that so-called permissive cases, which are cases in which one body of evidence seemingly rationalises multiple doxastic attitudes, are best understood as epistemic conflicts. I call this the Conflict View. The proposed resemblance arises from a similarity in the underlying conflicts displayed in these cases: the former involve conflicting values and the latter conflicting epistemic standards. In both instances, the agents are faced with incommensurable alternatives which are independent normative sources. By showing that permissive cases share the idiosyncratic features of comparisons under incommensurability, we gain a better understanding of numerous issues, such as the alleged arbitrariness of permissive attitudes. Furthermore, I will demonstrate that the proposed strategy is *prima facie* neutral regarding whether epistemic rationality is genuinely permissive. While some understandings of incommensurability support a permissive interpretation, others can be used to motivate impermissivism. This shifts the debate between permissivists and impermissivists by reducing

normative questions about epistemic rationality to more fundamental questions about the nature of incommensurability.

In the second chapter [ch.2], I will defend the proposed understanding of permissive cases against a possible objection from so-called self-fulfilling beliefs. Allegedly permissive cases based on self-fulfilling beliefs are problematic because they seem *not* to involve any conflicting epistemic standards. Despite their reception in the literature, I will show that self-fulfilling beliefs are not only governed by epistemic rationality but also compatible with impermissivism. Hence, they neither pose a threat to the conflict view nor do they support a permissive understanding of epistemic rationality. In particular, I will demonstrate, by carefully disentangling questions of practical and epistemic rationality, that for any self-fulfilling situation, there is not only a unique rational doxastic attitude to be in but also only at most one rational doxastic attitude to transition into. Doing so will also provide us with important insights into the nature of doxastic voluntarism.

The next chapter [ch.3] sets out to apply the conflict view to issues in collective epistemology. Within collective epistemology, there is a class of theories that understand ascriptions of collective attitudes, such as ‘the jury justifiedly believes that the suspect is guilty’, or ‘the IPCC knows that the currently observed climate change is predominantly anthropogenic’, as shorthand for saying that a relevant subset of group members has those attitudes. I will demonstrate that these summativist approaches to collective epistemology are incompatible with epistemic conflicts. In particular, I will provide a general recipe for generating so-called divergence arguments, which are motivated by divergence cases in which rationality requires group-level and member-level attitudes to diverge. I will call this class of cases *permissive divergence cases*. While other divergence arguments have been discussed in the literature, permissive divergence cases prove themselves to be less susceptible to many of the worries raised against its competitors, while being directly built on an often-defended epistemological thesis.

Methodological Interlude: Epistemic Continuity

This chapter [ch.4] is a methodological interlude into the nature and relationship of collective epistemology (concerned with collective epistemic agents) and individual epistemology (concerned with individual epistemic agents). It proposes and defends a novel approach to epistemology that considers individual epistemology and collective epistemology to be broadly continuous. The

approach is derived from an analysis of implicit methodological assumptions present in most contemporary approaches to the conceptual analysis of knowledge and justification. In particular, the proposed view has a conceptual and a methodological dimension: it aims not only (i) for general domain-independent epistemic concepts, but also (ii) treats insights made in the collective and the individual domain to be on *a par* or equally important. I will show how continuous concepts are more explanatory and powerful than their non-continuous competitors and that taking a continuous stance promises to give us a novel and fruitful way of handling various prevalent issues both in individual and collective epistemology.

Part II: Continuous Justification

In [ch.5], I will propose a continuous account of epistemic justification, that is built on evidentialist ideas discussed in individual and collective epistemology, called Continuous Evidentialism. Among other things, these investigations revealed that we need to embrace a view of collective justification that is sensitive to the overall evidence possessed by the group. This evidence can be (i) *distributed* among the group's members, and a group belief can be (ii) *defeated* by evidence possessed by a single member (even when that evidence wasn't taken into account when the respective group belief was formed). Furthermore, it is often argued that groups, like individuals, have certain (iii) *normative requirements* or epistemic obligations that affect how group members should gather, process and disclose available evidence. In trying to meet these challenges, I develop a continuous understanding of evidence, evidence possession and epistemic basing, which I will utilise in the next chapters to give a continuous account of epistemic justification.

Chapter six [ch.6] sets out to discuss different so-called no-defeater clauses. These clauses are motivated by Defeatism, the doctrine that we can have *prima facie* justified beliefs which are justified but defeasible. For instance, an otherwise justified belief based on visual loses its justifiedness if I acquire evidence that the underlying perceptual judgement is based on an optical illusion. I will survey various possible evidentialist as well as responsibility no-defeater clauses and develop a general taxonomy of defeater cases these clauses can be tested against. In so doing, I will reach a preliminary conclusion; namely, that only responsibility clauses, understood as being sensitive to the evidence one should have accessed and should have possessed, give us the right verdict with respect to the entire taxonomy of defeater cases.

In [ch.7] I will take a closer look at the concept of epistemic responsibility as understood by Sanford Goldberg. Goldberg (2018) argues that epistemic responsibility shall be understood via the socio-epistemic expectations we are entitled to have of epistemic agents. While I agree with Goldberg that epistemic responsibility is grounded in socio-epistemic expectations, I will argue that we shall not understand these expectations as making us strictly liable for the evidence we should have had. In contrast, we are entitled to expect from epistemic agents that they possess and utilise higher-order evidence when forming beliefs; and that acknowledging this enables us to reduce responsibility requirements of justification to higher-order evidentialist requirements. We can use these insights to revisit and dismiss Defeatism. This leaves us with a positive higher-order requirement of justification instead of a negative no-defeater clause.

In [ch.8], I discuss various positive higher-order requirements of justification and well-known problems faced by these requirements. I start by giving some additional motivations for thinking that we need to have some higher-order requirements built into our theory of epistemic justification. Next, I give an overview of different types of higher-order requirements, that have been proposed in the literature, including doxastic, non-doxastic, and dispositional requirements. In so doing, we can identify three worries often associated with such requirements having to do with over-intellectualisation, epistemic irrelevance, and infinite regresses. I conclude that neither the doxastic nor non-doxastic requirements proposed in the literature provide us with a satisfactory reply to all of these worries, while the defended higher-order evidential requirement does.

In the final chapter [ch.9], I bring together the ideas and concepts developed in the second half of the thesis to make use of the defended theory of justification Continuous Evidentialism. In particular, I compare it to some influential extant approaches in collective epistemology, such as Jennifer Lackey's Group Epistemic Agent Account (2016) and Paul Silva's Evidentialist Responsibility for Groups (2019). Doing so will help me to illustrate how Continuous Evidentialism explains case judgements about distributed evidence cases, collective defeater cases, collective basing cases, as well as evidential manipulation cases. This enables me not only to demonstrate the versatile nature of the proposed account but also allows me to highlight some interesting features of it.

But now, without further ado:

PART I:
EPISTEMIC CONFLICTS

1

Epistemic Conflicts:

Permissibility and Incommensurability

I will start my investigations into the nature of epistemic justification by taking a close look at what I call epistemic conflicts. These are cases in which we acknowledge that different modes of thinking or kinds of rationality are in conflict with each other. They are not the kinds of puzzles that we face in ordinary epistemic endeavours but the kinds of puzzles that arise when we (are forced to) take a step outside of our epistemic framework and reflect on our general way of thinking about the world. I chose this starting point, not only because I think that understanding these cases gives us general insights into the nature of rationality [ch.1] [ch.2], but also because I think that having a theory of epistemic conflicts helps us to understand the conflicts we as individuals face when participating as members of socio-epistemic entities [ch.3].

The following analysis is motivated by a comparison of these epistemic conflicts with so-called value conflicts, which are cases in which the different values we subscribe to give us conflicting recommendations. The proposed resemblance arises from a similarity in the underlying conflicts displayed in these cases: whereas the former involve conflicting values, the latter involve conflicting epistemic standards. In both instances, the agents are faced with incommensurable alternatives which are independent and non-directly comparable normative sources. By showing that these cases share the idiosyncratic features of comparisons under incommensurability, we gain a better understanding of numerous issues having to do with the alleged permissiveness of epistemic rationality. Most importantly, we can demonstrate that there is a way to analyse the belief dynamics that govern epistemic conflicts which is neutral with respect to the epistemic permissiveness of these conflicts.

1.1 Introduction

Sometimes we face choices in which the different values we subscribe to give us conflicting recommendations. Take, for example, this often-discussed case, found in Sartre (1966):

SARTRE'S STUDENT: Faced with the monumental decision of whether to care for his mother or to join the Free French Forces to revenge his brother's death, Sartre's student is forced to choose (or 'condemned to make a free choice') which person he wants himself to be.

Sartre describes his student as "hesitating between two kinds of morality; a morality motivated by sympathy and individual devotion and another morality of broader scope" (1966: 36). In other words, the student is torn between two alternatives, which are supported by conflicting but independent values. One value, let us call it the 'patriotic value' is in support of joining the Free French Forces, while the other, let us call it the 'filial value' is in support of caring for his dying mother.¹

Many have argued that these *value conflict* cases present us with distinct puzzles about rational choice since there is no direct way to compare or trade off alternatives supported by conflicting values against each other.² Decision theory, the predominant model of rational choice, works very well when we can assign numerical values to the possible outcomes and so calculate the relative choice-worthiness of these outcomes. However, value conflict cases are "cases where it's radically unclear how to assign values to the outcomes" (Williams 2016: 413). So, since there is no complete or even well-defined preference ranking with respect to the considered alternatives, value conflicts cannot be resolved by appealing to any standard means of decision-making.³

Compare such value-conflict cases to the following cases:

COMMUNITY:⁴ Ira has grown up in a religious community and believes in the existence of God. She has been given all sorts of arguments and reasons for this belief, which she has thought about at great length. One day she leaves that community and encounters arguments against her religious beliefs and learns that she only has the religious beliefs that she has, and only finds the reasoning that she engages in convincing, because of the influence of this community.

IDEALISM: When Phil is attending his first philosophy class, he is confronted with the idea that all his perceptual evidence equally supports two radically different hypotheses:

¹ This reading is inspired by Moss' (2014) and Williams' (2016) discussion of the case. Whether this is Sartre's own reading of the case is controversial and depends on the underlying understanding of the existentialist notion of radical freedom. For a contemporary analysis see, for example, Webber (2009: 44-73).

² See, for example, Levi (1986), Chang (2005), Ullmann-Margalit (2006), Schoenfield (2014a) or Williams (2016).

³ Note that this does *not* mean that we could not use choice functions which say that out of any menu of alternatives, it's permissible to choose any 'maximal element' (Sen 2017: 55).

⁴ A similar case is discussed by Schoenfield (2014b).

he could live in a mind-independent world composed of material objects, but he could also inhabit a Berkeleyan world composed solely of minds and their ideas.

Assuming that their beliefs have been rational so far, what is the appropriate epistemic reaction for Ira and Phil given their first encounter with this new information? Should Ira question her belief in God? Should Phil lower his credence in the existence of a mind-independent external world? While after carefully deliberating their new evidence, they need to form a new or update their old doxastic attitude towards the proposition in consideration p , it is unclear whether they should believe that p , disbelieve that p , or suspend judgement.

Some epistemologists, call this group the *impermissivists*, have argued that there needs to be a unique (positive) answer to these questions. For impermissivists, the evidence E_1 Ira possessed before encountering arguments against the existence of God and the evidence E_2 provided by those arguments together determine a unique doxastic attitude towards p . However, others, labelled *permissivists*, have argued that there does not need to be a unique answer. In other words, we might treat COMMUNITY and IDEALISM as ‘permissive cases’, cases where E_1 and E_2 together do *not* determine a unique doxastic attitude. Ira and Phil could (for reasons to be specified) rationally accept the encountered arguments and with it the implied conclusion, reject them and hold onto their beliefs, or suspend judgment on the respective proposition.⁵

This chapter will contribute to this debate, not by directly providing arguments in favour or against permissivism, but by arguing that cases such as COMMUNITY or IDEALISM are best understood as the epistemic equivalent of value conflict cases, such as SARTRE’S STUDENT. The proposed resemblance arises from a similarity in the underlying conflicts displayed in these cases: while Sartre’s student is in a case of value conflict, Ira and Phil are in a case of *epistemic standard conflict* (in short: epistemic conflict) [§1.2]. Both conflicts present the agents involved with alternatives which are supported by independent and non-directly comparable, i.e. *incommensurable*, normative sources. While incommensurability is often deployed when analysing cases of value conflict, epistemic conflicts are commonly understood as cases in which multiple doxastic attitudes are equally rational and, therefore, directly comparable [§1.3].⁶ In contrast, I will illustrate that epistemic conflicts have various idiosyncratic characteristics of comparisons under incommensurability, such as sweetening insensitivity [§1.4.1], normative bindingness

⁵ For an overview of the literature on epistemic permissivism see Kopec and Titelbaum (2016).

⁶ A notable exception is Moss’ (2014) treatment of credal dilemmas.

[§1.4.2], and angst [§1.4.3]. This will link the debate to numerous otherwise independently analysed phenomena, such as genealogical anxiety (Srinivasan 2015) and the perspective of doubt (Schoenfield 2023).

Despite the apparent permissive nature of epistemic conflicts, I will show that the proposed strategy is *prima facie* neutral with respect to the question of whether epistemic rationality is permissive [§1.5]. While some understandings of incommensurability support a permissive interpretation, others can be used to motivate impermissivism. This shifts the debate between permissivists and impermissivists by reducing normative questions about epistemic rationality to more fundamental questions about the nature of incommensurability. In so doing, we are able to tell different stories about how epistemic conflicts ought to be resolved conditional on different understandings of incommensurability.

In the final section [§1.6], I will show that viewing epistemic conflicts through the lens of incommensurability provides us with a novel understanding of deliberation and diachronic rational development in the face of such conflicts. This will provide us with a general solution to one of the most prominent objections against permissivism, the so-called arbitrariness objection (White 2005). Moreover, it will lead to novel insights into the overall dynamics of rational belief evolution in these conflicts. An evolution that is governed by various trade-offs between the idiosyncratic characteristics of comparisons under incommensurability identified in [§1.4].

Before I proceed, a few words of clarification are needed. The proposed analogy is different from recent comparisons between *epistemic dilemmas*, which are roughly cases in which no doxastic attitude, including suspension of judgement, is rationally permissible, and dilemmas in other normative domains.⁷ I am neither committed to the existence of genuine (epistemic) dilemmas, nor do I want to endorse this comparison, nor do I say that epistemic conflicts confront us with genuine epistemic dilemmas. My analysis is concerned with a comparison between permissive cases and cases of value conflict, which while confronting us with incommensurable alternatives, may have one or multiple normatively permissible resolutions.⁸

⁷ For an up-to-date overview of the literature on epistemic dilemmas see Hughes (2019), Hughes (forthcoming) or Priest (2022).

⁸ Following Simion, I want to reserve the term epistemic dilemma to cases in which a subject “has only two available courses of action, both of which imply norm violation, where neither of the norms at stake takes precedence over the other” (2022: 112); which is different from how, e.g., Christensen (2022) understands epistemic dilemmas. For Christensen, epistemic dilemmas also include epistemic norm conflicts where one norm takes priority over another. As will be clear from the discussion later, the epistemic conflicts I am concerned with are quite different from either

1.2 Standard Permissivism and Standard Uniqueness

Permissivists and impermissivists disagree whether epistemic conflict cases, such as COMMUNITY and IDEALISM are *de facto* rationally permissive. However, they both usually agree that these cases initially appear permissive. After all, when discussing specific cases, permissivists want to argue for, while impermissivists often want to explain away, the apparent permissibility of these cases. This raises the question of what the source of this apparent permissibility is. Before I will answer that question in the next section [§1.3], some preliminaries on epistemic permissivism are needed.

While there are many different flavours of permissivism, most permissive understandings of rationality can be subsumed under the heading of *standard permissivism*; i.e., permissivism about epistemic standards. Standard permissivists think that given any body of evidence E, there could be independent non-evidential aspects of the overall epistemic situation that render different doxastic attitudes towards a given proposition rational. This conclusion is reached by pointing out that different (timeslices of) epistemic agents might, for example, have different preferred predicates (Titelbaum 2010), cognitive goals (Kelly 2013), Bayesian priors (Meacham 2013), or different standards they take to be truth-conducive (Schoenfield 2014b), which rationalise different doxastic attitudes with respect to E. If we subsume all these different *non-evidential but epistemic* aspects under the term *epistemic standards*, we get so-called Standard Permissivism:⁹

Standard Permissivism: For any body of evidence E and any combination of epistemic standards S_1 - S_n there is a unique rationally permissible doxastic attitude D towards any proposition p . There are cases in which one body of evidence E rationalises different doxastic attitudes D_1 - D_n relative to different sets of standards S_1 - S_n .

Here, a quick clarification about Standard Permissivism is needed. Defining Standard Permissivism this way gives us a rough understanding of what I take to be a general characteristic of epistemic standards (for a more detailed characterisation see [§5.5]). Epistemic standards fix the epistemic support relation; i.e. are the instructions that tell us what we ought to believe given

understanding of epistemic dilemmas. First, I think that epistemic standard conflicts are not conflicts between epistemic norms but at most conflicts about how to calibrate our norms by subscribing to different epistemic standards [§1.3]. Second, epistemic dilemmas usually are *not* understood to have the idiosyncratic properties of comparisons under incommensurability which I take to be the defining features of epistemic conflicts [§1.4.1]-[§1.4.3].

⁹ For a similar categorisation see Schoenfield (2014b), Ye (2019), Simpson (2017) or Jackson (2021).

our current evidential situation. As such, they have two dimensions. First, they define the degree of evidential support that is lent to a proposition p by a body of evidence E . Second, they have a doxastic dimension which specifies which doxastic attitude we ought to adopt towards the considered proposition given the evidential support we have for that proposition.¹⁰

Accordingly, Standard Permissivism relies on two assumptions. First, (i) that epistemic support is a three-place relation between the evidence, the epistemic standards, and a prescribed doxastic attitude (Decker 2012: 780-782; Kelly 2013: 307-308). And second, (ii) that epistemic standards are *variable*, e.g., across agents, contexts and/or times.

While standard permissivists disagree on how variable the epistemic standards are, they all share in common that fixing enough parameters of any epistemic situation mandates a unique rational response from any epistemic agent in that situation.¹¹ In other words, standard permissivists reject Evidential Uniqueness while endorsing Epistemic Uniqueness:

Evidential Uniqueness: For any body of evidence E and any proposition p there is *at most one* doxastic attitude that any epistemic agent A could rationally take.¹²

Epistemic Uniqueness: For any epistemic situation S (any body of evidence and any combination of epistemic standards) and any proposition p there is *at most one* doxastic attitude that any epistemic agent A could rationally take.

Impermissivists, on the other hand, either defend Evidential Uniqueness directly or Epistemic Uniqueness without allowing epistemic standards to be variable; i.e., they deny (ii). As such, impermissivists also often understand epistemic support as (i) a three-place relation between the evidence, a unique (set of) epistemic standard(s), and a prescribed doxastic attitude.¹³ I will call this version of impermissivism Standard Impermissivism:

¹⁰ We can now see that the things listed as examples of epistemic standards above concern both dimensions. For example, while Bayesian priors concern the evidential support dimension by telling us which evidence supports which proposition to which degree, Jamesian truth goals fix the doxastic dimension by telling us which doxastic attitude we ought to form. More on that in [§5.5]. I want to thank Edward Elliot for comments on this.

¹¹ Most standard permissivists are *interpersonal* permissivists (Kelly 2013), who claim that different standards might rationalise different doxastic attitudes for different agents given one body of evidence. For an *intrapersonal* version of standard permissivism see, e.g., Jackson (2021).

¹² Note that “at most one” signifies that there might be rational or epistemic dilemmas, i.e. cases in which no doxastic attitude is rationally permissible (Kopeck & Titlebaum 2016: 190-191).

¹³ Early defences of impermissivism, going back to White (2005), relied on a two-place relational understanding of epistemic support. However, more recent impermissivists’ views take epistemic standards to be necessary for the epistemic support relation. See, for example, Hedden (2015a, 2015b: 472-473).

Standard Impermissivism: For any epistemic situation *S* (any body of evidence *E* and the uniquely rational (set of) epistemic standard(s)) and any proposition *p* there is *at most one* doxastic attitude that any epistemic agent *A* could rationally take.

Given these definitions, we can now move on to analyse the nature of epistemic conflicts and the source of their apparent permissibility. In the next section [§1.3], I will demonstrate that the received view of epistemic conflicts as involving (apparently) *equally rational* doxastic attitudes is incompatible with both Standard Permissivism and Standard Impermissivism. Instead, those views suggest that we should think of these cases as confronting us with incommensurable alternatives, as value conflicts do. This analogy will then be further supported by showing that various defining features of comparisons under incommensurability are present in both types of conflicts [§1.4].

1.3 Value Conflicts and Standard Conflicts

A natural, and often implicitly presupposed, interpretation of epistemic conflicts, such as COMMUNITY or IDEALISM, is as cases in which we are faced with multiple (apparently) *equally rational* doxastic responses. For example, Greco and Hedden (2016: 387-390), in their defence of Evidential Uniqueness compare allegedly permissive cases to practical equilibria in which there are multiple equally preferred options.¹⁴ On this reading, the source of the apparent permissive character of these cases stems from an (apparent) rational indifference between the compared attitudes. In other words, Ira and Phil are like Buridan's Ass, not stuck between two equally delicious bales of hay but stuck between multiple seemingly equally rational doxastic attitudes. Since many have argued that practical equilibria such as Buridan's Ass present us with multiple permissible options (Ullmann & Morgenbesser 1977, Ullmann-Margalit 2006, Chang 2005, Blackburn 1996), this comparison is a natural one.¹⁵

However, neither Standard Permissivism nor Standard Impermissivism supports this reading. On both views, Ira and Phil are not in a situation analogous to Buridan's Ass, but in an (apparent) *epistemic standard conflict*.¹⁶ A conflict where different sets of (permissible) epistemic

¹⁴ See also Schoenfield (2014b: 200-202) or Kelly (2013: 302).

¹⁵ More on the relationship between practical permissivism and epistemic permissivism in [§2.7].

¹⁶ There is a cluster of allegedly permissive cases built around self-fulfilling beliefs that arguably cannot be understood via the proposed conflict reading since there are no conflicting standards involved in these cases. However, as I will

standards support different doxastic attitudes, rather than a case in which one set of standards supports multiple doxastic attitudes. Ira, faced with novel arguments against her strong conviction in the existence of God, is forced to re-deliberate while sticking to her belief or forming a new doxastic attitude forces her to subscribe to some underlying epistemic standards that rationalise the doxastic attitude given her evidence [§1.2].

While there might or might not be a unique set of epistemic standards that Ira should subscribe to from her perspective the case seems rationally permissive. She may either stay a theist, become an agnostic or an atheist, each of which comes with valuing different sources of evidence in a certain way, based on different (sets of) epistemic standards S_1 - S_n . Staying a theist, for example, may force her to give more credence to the testimonial evidence of her fellow believers than to the arguments presented by her new atheist friends and *vice versa*. Becoming an agnostic, on the other hand, might be the result of balancing the cognitive goals of ‘attaining truth’ and ‘avoiding error’ in different ways.¹⁷

Compare this interpretation of COMMUNITY with the following natural reading of SARTRE’S STUDENT: Sartre’s student is faced with a conflict between two opposing sets of values he subscribes to, the ‘filial value’ in support of taking care of his mother, and the ‘patriotic value’ in support of risking his life for the ‘greater good’. It is not his weakness of will or his uncertainty that fuels this conflict but his awareness that he faces a choice in which two or more values apply in a manner which yields prescriptions which cannot be implemented together.¹⁸ Since, for the student, there is no objective way to balance or trade-off these values against each other, from his point of view the two options are not directly comparable. Accordingly, the most natural interpretation of all considered cases, COMMUNITY, IDEALISM, and SARTRE’S STUDENT is as presenting us with two alternatives that are *incommensurable* due to an underlying standard or

argue in the next chapter [ch.2] I think these cases are best understood to be impermissible. As such, they neither pose any threat to Evidential Uniqueness nor the Conflict View (see below).

¹⁷ This also highlights the distinction between epistemic conflicts and epistemic dilemmas again. The latter are usually understood to be dilemmas arising from conflicting epistemic norms, such as first-order and second-order or substantive and structural norms (Hughes forthcoming: §2). The conflict I describe is different. It is not that one norm tells Ira to mistrust the testimonial evidence while another tells her to trust it. The conflict rather arises from the question of how much weight she should give to the testimonial evidence whereas she acknowledges that there are different permissible ways to answer that question.

¹⁸ For a systematic discussion of these so-called hard choices and other kinds of moral dilemmas see Levi (1986: 1-19).

value conflict, where incommensurability means that there is no common underlying measure which allows us to put the alternatives on an ordinal scale.¹⁹

In the practical domain, many have argued that not only alternatives that are equally good but also alternatives which are incommensurable are (apparently) permissible. For example, Ullmann and Morgenbesser (1977) have made a strong case that in practical equilibria rationality permits us to *pick* rather than choose among the available options; while ‘picking’ here is selection without preference and choosing is selection via preference. Furthermore, Ullmann and Morgenbesser argue that picking is not only present in micro-level practical equilibria but also in macro-level value-conflict cases. For them, “at the very deepest level of selection, involving the ultimate and most significant alternatives confronted by man, there can only be picking” (Ullmann & Morgenbesser 1977: 783). In other words, for Ullmann and Morgenbesser, Sartre’s student is rationally permitted to go either way, despite the incommensurability of the values involved, making his decision a matter of picking rather than choosing.²⁰

Importantly, while these transformations may be the result of voluntary processes, as it is in the case of picking, they do not need to be. Ullmann-Margalit (2006) speaks of *conversions* if we do not experience the transformation as voluntary, as is often the case in religious conversions. Furthermore, she distinguishes converting from another mode of picking, called *drifting*, which is an unrecognised gradual transition from one worldview into another.

For our purpose, these discussions reveal two important insights. First, apparent permissibility is a general feature not only of cases in which the compared alternatives are equally good but also of cases in which the alternatives are incommensurable. Second, apparent permissibility does not require voluntariness: while Sartre’s student’s decision is perceived to be voluntary –that’s the existentialist dilemma! – it does not have to be. He could also find himself on the convoy to join the Free French Forces, without perceiving it to be the result of a voluntary choice, while nonetheless, acknowledging that he would have been permitted to take care of his elderly mother instead. (He could be converted by one of his fellow comrades to become a

¹⁹ By an ordinal utility ranking I simply mean a preference ranking of any set of options S that satisfies Transitivity ($\forall A, B, C \in S: (A \leq B \wedge B \leq C) \rightarrow A \leq C$) and Completeness ($\forall A, B \in S: A \leq B \vee B \leq A$). This has to be distinguished from stronger cardinal preference rankings which also say “something about the desirabilistic ‘distance’ between options” (Steele & Stefánsson 2020: §2.2).

²⁰ Similar permissible notions of choosing under incommensurability, with slight variations in their psychological descriptions, have been defended in many places. See, for example, Nagel (1987), Blackburn (1996), Raz (1997), Broome (2001), Chang (2005), or Moss (2014).

righteous soldier or he could drift along with the *Zeitgeist* ignorant of the significance of his recent transformation.)

Similarly, epistemic conflicts such as *COMMUNITY* might sometimes be resolved by rational agents being simply dragged along by the current of their rational development. Ira may just end up becoming an atheist, without perceiving the change in her beliefs to be the result of conscious deliberation. This is not to say that there might be no initial struggle or deliberation but that the result is something that is reached not as a matter of choice but as a matter of (perceived) destiny. In other words, for those who are sceptical of the voluntariness of doxastic attitude acquisitions, epistemic standard conflicts might be best understood not as crossroads but as river bifurcations.

Together these two observations give some initial plausibility to the proposed understanding of epistemic conflicts in telling us that apparent permissibility neither requires being in an active choice scenario nor that one is in an equilibrium where multiple doxastic attitudes/alternatives are equally rational/good. This, together with Standard Permissivism or Standard Impermissivism, provides us with an argument to understand allegedly permissive cases, such as *COMMUNITY* or *IDEALISM* as epistemic standard conflicts rather than epistemic equilibria. Let us call this the Conflict View:

Conflict View: Apparently permissive cases are epistemic standard conflicts. That is, cases in which one body of evidence supports different doxastic attitudes relative to different incommensurable sets of (permissive) epistemic standards.

While the above-made observations are a good starting point, we can make a more general case for the Conflict View. To do so, we need to take a closer look at some features that distinguish comparisons under incommensurability from comparisons under equality and show that they are present not only in cases such as *SARTRE'S STUDENT* but also in cases such as *COMMUNITY* or *IDEALISM*.

1.4 Incommensurability: Three Features

To show that we should understand epistemic conflicts via incommensurability, I will spend this section demonstrating that various idiosyncratic characteristics of incommensurability are present in epistemic conflicts. In particular, I will highlight three often-mentioned features, which I will subsequently discuss in sections [§1.4.1] – [§1.4.3]:

- (i) *Sweetening Insensitivity*: improving either of the alternatives does not necessarily make the sweetened alternative rationally obligatory and the other rationally impermissible.
- (ii) *Bindingness*: ending up with one alternative binds one to appeal to similar resolutions in analogous future conflicts.
- (iii) *Angst*: resolving the conflict often leaves a persistent uncertainty about whether one has responded rightly.

1.4.1 Sweetening Insensitivity

One prominent feature distinctive of incommensurability is that the compared alternatives are *sweetening insensitive* or *resistant to small improvements* (Raz 1986: 332; Chang 2002; Schoenfield 2014a; Williams 2016). If you are indifferent concerning two equally good options A and B, adding additional incentives or sweeteners to one of the outcomes tips over the balance in favour of the sweetened option. For example, if you consider paying with cash or card equally good but then find out that if you are paying with card you have to pay an additional service fee, you *ought* to pay with cash.

In contrast, comparisons under incommensurability are not like this. If you are faced with the career choice between being an artist or a philosopher, neither of which you consider to be better than the other, adding a \$100 monthly salary to the future philosophy career does not necessarily break the tie (Chang 2002: 668; Raz 1986). The \$100 would simply not make enough difference to rule out adjusting your values in a way that aligns with your choice to become an artist.²¹ This also explains why Sartre's student may, after initially deciding to go to Paris to join the Free French Forces, change his mind and jump off the train to go home and care for his mother instead. Even considering the sunk cost of spending a few hours on the train and paying for the ticket did not make his decision definite.

Cases such as COMMUNITY provide us with the epistemic counterpart. If epistemic conflict cases were instances in which a body of evidence *equally supports* D_1 and D_2 , adding additional evidence in favour of D_1 would break the tie. However, this is evidently not the case, given the above-defended understanding of these cases as epistemic standard conflict cases. If we add yet another piece of testimonial evidence in support of theism on top of Ira's evidence, we have not

²¹ As such, these arguments are usually understood as arguments against Completeness. However, as pointed out by Raleigh (2023) these arguments are only effective if one already has some prior reason to prefer Transitivity to Completeness (see footnote 19).

broken the tie since it (apparently) would still be rational for Ira to become an atheist or an agnostic, even after adding this ‘sweetener’. While understanding the case as a genuine conflict proposes that she gives some credit to both sources of evidence, the assumed conflict leaves it open how she ought to balance her epistemic standards when interpreting them.

This sweetening insensitivity of epistemic conflicts is probably best illustrated by referring to philosophical discussions of scientific theory choice. In these discussions, it is commonly pointed out that different competing scientific theories are often supported by different theoretical virtues, such as adequacy, consistency, explanatory scope, simplicity, or fruitfulness.²² If we are choosing or picking between different theoretical virtues when we are opting for a specific theory, then theory choice cases are best understood as a special type of epistemic conflict:²³

THEORY UNDERDETERMINATION: Sarah is faced with two scientific theories A & B that are equally empirically adequate. A & B both explain 90% of the available data (though not the same 90%), which *ceteris paribus* makes it permissible for Sarah to believe either in A or in B. However, on top of being equally empirically adequate, both theories maximise different theoretical virtues. While A is more ontologically parsimonious than B, B is more mathematically elegant.

Arguably Sarah seems rationally permitted to believe in either theory, even when the theories are mutually exclusive. Furthermore, while Sarah values both parsimony and elegance she has no independent means to trade off these theoretical virtues against each other. This makes the theory choice, as well as the related belief formation, sweetening insensitive. Adding a little bit of evidence that can only be explained by A is surely not enough to break the tie between the permitted doxastic attitudes. If A can explain 91% of the data while being more parsimonious while B can only explain 89% of the data while being more elegant, Sarah could surely still value elegance more, so that she, despite this explanatory disadvantage, believes in B. Yet, this does not make the situation immune to any amount of sweetening: adding more and more evidence in favour of A would surely at some point break the tie, no matter how elegant B is in comparison to A.

²² This list of theoretical virtues is taken from Kuhn (1977: 320-339). For another canonical defence see also Longino (1996).

²³ An early comparison between moral dilemmas and scientific theory choice is found in Levi (1986: 36-46). For a comparison between underdetermination and permissivism see Jackson and Turnbull (forthcoming).

Interestingly, we could also reach the same result if we add the sweetener in the form of theoretical virtues. Take the original set-up of THEORY UNDERDETERMINATION. Adding a little bit of elegance to B, for example, by simplifying one of the central equations does not automatically make it the case that B is to be preferred to A. Sarah could still believe in A without being irrational. However, making the theory more and more elegant ultimately makes it the case that it is the only rational option. This is another similarity to value conflicts. In both cases, the form of the sweetener does not matter as long as it is something that is already valued by the agent. We could make Sartre's student's decision definite by either adding more and more patriotic duty (maximised by joining the Free French) or filial solicitude (maximised by caring for his elderly mother) to the respective options or by adding something else he already values, let's say money or social approval.

1.4.2 Bindingness

Another often-mentioned feature of incommensurability is that the resulting transformations have a *binding* nature, “resolving a case of conflicting value on one occasion commits one to similar resolutions of similar future situations” (Williams 2016: 424). Chang (2005) illustrates this bindingness via a money pump argument. If A and B are incommensurable, B and A+, where A+ is A plus \$100, might be incommensurable as well. In these cases, we are arguably permitted to choose B over A+. However, if we, after choosing B, are then faced with a choice between B and A, we are *not* permitted to choose A because this would leave us worse off since we could have had A+. Accordingly, Chang thinks that “the rational permissibility of choosing either of two items on a par, then, must be constrained by one's other choices [...] if one had not already chosen B over A+, it would have been rationally permissible to choose A over B” (2005: 347).²⁴

This again mirrors what happens in the epistemic case. Once Ira starts treating the encountered arguments against the existence of God in a certain way she commits herself to become a certain type of reasoner. As a consequence, she superimposes certain standards (rankings) onto her which changes the way she forms future attitudes and forces her to rethink old attitudes. For example, Ira, before encountering the arguments against theism might

²⁴ Chang uses ‘on a par’ in a technical sense as a fourth positive value relation (in addition to being better than, being worse than, and being equally good) that holds in cases of incommensurability. More on that below [§1.5].

subscribe to some testimonial epistemic standard, that tells her that if she receives testimony that *p*, then she is *ceteris paribus* rationally required to believe that *p*.²⁵

So, if Ira becomes an atheist, by lowering the credibility she gives to the testimonial evidence of her former fellow believers she is required to replace the standard regarding testimonial evidence with a weaker standard. Doing so means that she adjusts other beliefs based on testimony as well. Furthermore, it requires her to treat future evidence of a similar kind in the same manner independent of the proposition supported by that evidence. For example, if she receives testimonial evidence that atheists are on average more reasonable persons, she needs to dismiss that testimony in the same way she dismissed the testimony of her former fellows.²⁶

This is especially clear given the role that not only values but also doxastic attitudes play in practical decision-making since non-binding attitudes generate an unstable basis for successive decisions over time. In other words, non-binding doxastic attitudes make epistemic agents vulnerable to money pumps, such as the one described by Chang. Interestingly in the context of permissivism, similar worries have prominently been raised when discussing the alleged *arbitrariness* of permissive beliefs. The idea here is that permissivism makes doxastic attitudes unstable in a way that allows epistemic agents to arbitrarily switch between permissible doxastic attitudes, which in turn support different permissible actions (White 2005).²⁷ Before I go into this in more detail [§1.6], let me introduce a third property of comparisons under incommensurability that complicates the picture even further [§1.4.3].

1.4.3 Angst

These considerations about arbitrariness and bindingness directly lead us to the last characteristic I want to discuss, which has to do with the distinct phenomenology of acknowledging that one faces multiple incommensurable options. Due to the apparent permissibility of these conflicts, the agents who find themselves in such cases often experience *angst* in face of the choice or are

²⁵ For the discussion of such a principle see Elga (2007), Titelbaum (2015), or Bradley (2019).

²⁶ Note, that this treatment of epistemic standards assumes that there are general ways to carve up the space of possible epistemic standards. Otherwise, there might be a standard available to Ira that allows her to dismiss one type of testimony without dismissing the other. Otherwise we might face a generality problem about epistemic standards that is similar to the kind of generality problem faced by the reliabilist approaches to epistemic justification (Conee & Feldman 2004: 135-165). While, for the sake of narrative ease, I decided to not discuss this objection here I want to refer to reliabilist responses to this problem (Goldman & Beddor 2021: §2.3). I want to thank Pekka Väyrynen for raising a version of this objection.

²⁷ See also White (2013), Kelly (2013), Schoenfield (2014b), Simpson (2017), Ye (2019) or Jackson (2021).

left with a *resolutive remainder* after transitioning. This is true, regardless of whether one thinks that the faced conflict is genuinely permissive or not. Even if one knows that both possible routes are rationally permitted there is a “persisting perplexity about whether one has acted rightly” (Williams 2016: 425). This angst comes in varying degrees, and it is especially vivid in cases where the transformation has a deeper significance, where there is “no clash of arbitrary decisions but a substantive disagreement” (Chang 2002: 685). This is a key feature of macro-level value choices such as the one faced by Sartre’s student. Sitting on the train home to his mother, after having chosen to become the caring son, the student might feel the weight of the decision resting on his shoulders (irrespective of being aware of the alleged permissibility of the situation). Ullmann-Margalit calls this “the ghost of the Road Not Taken” (2006: 160).

Importantly, while angst may be present in equilibria cases when the choice is significant enough, the kind of angst characteristic of comparisons under incommensurability is different. We might experience angst, in “case of a moral tie, where we can e.g. save only one person out of five, and decide by lot who to save”; however, if there “are genuinely conflicting values, one feels angst over whether one has done the right thing” (Williamson 2016: 418). Here the former kind of angst solely concerns the fact that one had to make a difficult choice while one could know that one did as much as one could have done. In contrast, the latter kind of angst goes deeper. It is angst about the values that have led one to make the choice in the first place. The angst that leads one to question whether one has done the right thing, not knowing whether there has been one right thing to do.

In the epistemic domain, a similar phenomenon has been described as genealogical scepticism (Srinivasan 2015) or genealogical anxiety (Schoenfield 2023). For instance, finding out that one has been raised in a certain cultural or religious background may induce doubt about many beliefs that one otherwise takes for granted. One might believe in the analytic-synthetic distinction, the intrinsic value of free speech, or reincarnation just because genealogical coincidences have led one to adopt certain epistemic standards. The described anxiety then arises by acknowledging that certain belief sets might be the remnants of previous epistemic conflicts.

Based on similar observations, some (in addressing the so-called Arbitrariness Objection [§1.6]) have doubted the possibility of acknowledging that one is in a permissive situation (Cohen 2013; Smith 2020) or acknowledging that two sets of standards are equally truth-conducive (Schoenfield 2014b). However, as I will show in the next section, thinking of epistemic conflicts as involving incommensurability enables us to tell a different story: if Ira becomes an atheist by

balancing her standards in a certain way that enables her to reject the testimony of her fellow believers she is (sooner or later) bound to be a certain type of rational reasoner. Nonetheless, she might at some point look back with distressing puzzlement and ask herself whether that was the solely rational reaction. That is, she might acknowledge that she was – or still is – in a seemingly permissive situation.

One might go as far as to claim that this resolutive remainder is an intrinsic characteristic of our epistemic life. In philosophical reflection, we are often haunted by the malevolent sceptical demon in the very back of our minds that reminds us that the route of our epistemic development (even if rational) is rife with diverging paths. This insight plays a key role in contextualist epistemology and is often highlighted by those who want to tell a story about Cartesian reboots or the perspective of systematic doubt (Nguyen 2020; Schoenfield 2023).

In the final section, I will illustrate how the interplay between angst, sweetening sensitivity, and bindingness provides us with a novel understanding of these phenomena in the context of diachronic rationality [§1.6]. When engaging in systematic doubt angst mixes into the conflict, which, despite bindingness, can rationalise re-considerations of large belief sets and their underlying epistemic standards. This then tells us that switching between permissible standards and belief sets after being bound by them can be rational, despite the associated costs, if the underlying epistemic angst is strong enough. Before doing so, however, let us turn back to the question of whether epistemic conflicts are epistemically permissive or not [§1.5].

1.5 Incommensurability and Permissibility

What do these investigations tell us about the question of whether epistemic rationality is permissive or impermissive? Initially, the proposed analogy seems to support permissivism. If allegedly permissive cases are cases of epistemic standard conflict and there are different ways these conflicts can be rationally resolved there is some slack in how agents may rationally respond to a certain body of evidence. However, as I will show in this section, different underlying understandings of incommensurability present us with different stories of how this alleged permissibility comes about.

First, on some understandings of incommensurability, the failure of any (ordinary) value relation to obtain suggests that there are more value relations than previously believed (Griffin 1986; Parfit 1987; Chang 2002). For example, Chang argues that incommensurability cases are cases which are governed by a fourth positive “value relation of comparability that may hold

when “better than,” “worse than,” and “equally good” do not” (Chang 2002: 331). Here, Chang speaks of comparisons between alternatives that are *on a par*, whereby *parity* is a relation that obtains exactly then when a rational agent exhibits the behaviour described by the features discussed above [§1.4.1] – [§1.4.3]. If Chang is right, and the defended comparison to the epistemic domain holds, epistemic conflicts are genuinely or *strongly* permissive. From Ira’s perspective, there are multiple permissible doxastic attitudes which are rationally *on a par*, supported by the same evidence but different epistemic standards.

However, another influential understanding of incommensurability suggests or is at least compatible with epistemic impermissivism. On these views, not a fourth value relation but *indeterminacy* or *vagueness* is what explains incommensurability. For example, Williams (2016) argues, following Broome (2001), that decision-making under incommensurability is a special case of decision-making under indeterminacy. Drawing on his former work on decision-making under indeterminacy Williams (2014) reproduces the above-discussed features of incommensurability (sweetening insensitivity, bindingness, and angst) without introducing a fourth genuine value relation as Chang does. Interestingly, in his view, the alleged permissibility of decisions under indeterminacy is only *weak*, in the sense that there is no determinately correct criticism of the weakly permitted act (Williams 2016: 423-424):

“The Changian permissivist thesis is not true of decision making under indeterminacy under the strong reading of ‘permission’. The combination of indeterminacy and the strong notion of permission, however, generates a secondary ‘weak’ permission. Consider any option where it is not determinately the case that you are not permitted to perform it. Choosing such an option has a special status, because it is immune to neutral criticism and sanction. If your informed peers judge that you have done something morally impermissible, or impose sanctions, they are non-neutral in the sense that they take actions which are appropriate only if *p*, where *p* is indeterminate.”

Given the underlying understanding of incommensurability, this notion of weak permissibility can be used to explain the apparent permissibility of epistemic conflicts.²⁸ Compare this, for example, to Hedden’s (2015a; 2015b) objective Bayesian defence of Evidential Uniqueness. For

²⁸ Another way to understand weak permissibility is to note that the non-optimal attitudes and decisions are strictly impermissible but the respective agents are blameless when acting accordingly. So, permissibility and blame might come apart in these cases. For a recent investigation into the concept of epistemic blame see, for example, Brown (2018). Note, in chapter [ch.7] I will defend an understanding of epistemic justification in which epistemic justification and epistemic blamelessness go hand in hand. I want to thank Daniel Elstein for pointing this out.

Hedden, there is a unique rational Bayesian prior for every proposition determined by the uniquely correct set of epistemic standards which determines a unique rational doxastic attitude given any body of evidence. Nonetheless, there might be cases in which it is indeterminate which Bayesian prior and, therefore, which doxastic attitude is the uniquely rational one – while it “being indeterminate what you ought to believe is quite different from there determinately being multiple permissible doxastic states, given the same evidence” (Hedden 2015b: 469).²⁹ In William’s (2016) terminology, this is to say that while only one Bayesian prior for each proposition p may be strongly permissible, many Bayesian priors will be weakly permissible for p (i.e. there is no determinately correct criticism of them).

Compare this also to scientific theory choice cases [§1.4.1]. In these cases, various substantive rational requirements may definitely rule out certain epistemic virtue rankings and related interpretations of the evidence, while leaving others indeterminate. For example, one body of evidence might be in principle compatible with Cartesian scepticism or counter-inductivism, as well as the parsimonious theory A and the elegant theory B. While we might be able to rule out the former two by appealing to substantive rational requirements, whether we should believe in theory A or theory B could, nonetheless, be indeterminate.³⁰

In sum, competing explanations of incommensurability have diverging consequences with respect to the epistemic permissivism debate. If incommensurability is explained via a fourth genuine value relation, we end up with a permissive understanding of epistemic rationality. If, however, incommensurability is best understood via indeterminacy we end up with an impermissive view. Importantly, irrespective of the underlying explanation, the ‘epistemic life’ in the face of epistemic conflicts proceeds exactly the same. Whether those conflicts are grounded by indeterminacy, and to be resolved by indulging in weak permissions, or grounded in parity, and resolved by indulging genuine permissions, we get the same pattern of expected belief evolution.

While I will demonstrate in the final section that the exact pattern depends on the role one assigns to bindingness [§1.6], the discussion above shows that whether these situations are truly permissive or not is not crucial. This drains (some of) the apparent significance out of the

²⁹ See also Christensen (2007: 102, footnote 8), Hedden (2019: 584-585) or Greco and Hedden (2016).

³⁰ Note that this is compatible with saying that we should ‘go mushy’ over the standards and the associated credence functions that recommend A and B respectively (White 2013; Hedden 2015b), as well as saying that we should ‘identify with’ one of them, at least for the purpose of rational decision making (Moss 2014). Either way, we are not required to think that both attitudes are equally rationally permissible.

permissivism debate, and reduces questions about epistemic rationality to more fundamental questions about the nature of incommensurability.

1.6 Arbitrariness, Bindingness, and Changes of Mind

The considerations above show that we should think of competing doxastic attitudes in epistemic conflicts as incommensurable rather than equally rational. I demonstrated this by highlighting that various characteristics of incommensurability are present in these cases; namely, sweetening insensitivity [§1.4.1], bindingness [§1.4.2], and angst [§1.4.3]. Furthermore, I have shown, that conditional on different views of incommensurability, this leaves us with a permissive or an impermissive understanding of epistemic rationality [§1.5]. In this section, I will utilise these insights to give us a novel understanding of the so-called Arbitrariness Objection to permissivism, as well as related issues about the purported irrationality of *diachronically incoherent* belief evolutions. In particular, I will discuss three different understandings of bindingness [§1.6.1] – [§1.6.3], which will all present us with different resolutions to these worries.

When I discussed bindingness and Chang’s (2005) money pump argument above [§1.4.2], I pointed out that in the context of epistemic permissivism, similar worries have prominently been raised when discussing the alleged *arbitrariness* of permissive doxastic attitudes (White 2005, 2013; Kelly 2013; Schoenfield 2014b; Simpson 2017; Ye 2019). The overall idea of these worries is “that a permissive account of rationality introduces a kind of arbitrariness to our beliefs that can infect both practical and theoretical deliberations” (White 2005: 455). While there are different reasons why arbitrarily formed beliefs might be considered irrational, the most prominent is that such beliefs may lead to diachronically incoherent behaviour.³¹ In other words, permissivism seems to make doxastic attitudes unstable in a way that allows epistemic agents to arbitrarily switch between them.³²

³¹ White also considers an argument for the alleged arbitrariness of permissive beliefs based on a thought experiment involving magical belief-inducing pills (2005: 454). In this argument, White tries to show that while beliefs formed via popping magical pills are defective, permissivism cannot discriminate such beliefs from beliefs formed via proper belief formation processes, such as deliberation. For further discussion see White (2013: 315-317), Simpson (2017: 523-524) or Ye (2019).

³² Sometimes this specific aspect of the Arbitrariness Objection is called the toggling objection (Jackson 2021). For the sake of narrative ease, I will not discriminate this more specific worry from more general worries having to do with arbitrariness. For further discussions of the influence arbitrariness has on practical deliberation see White (2013: 318-322), Greco and Hedden (2016: 382-384), Simpson (2017: 525-526), Ye (2019: 669-672), or Jackson (2021: 322-324).

Accordingly, we can put the Arbitrariness Objection into argument form as follows:

The Arbitrariness Objection:

- (P1) There are permissive cases in which one body of evidence permits multiple doxastic attitudes (D_1 - D_n). (*permissivism*)
- (P2) If you find yourself in a permissive case it is possible to acknowledge that multiple doxastic attitudes (D_1 - D_n) are permissible. (*acknowledged permissivism*)
- (P3) If you acknowledge that multiple doxastic attitudes are permissible you can arbitrarily switch from one permissible attitude to the other.
- (P4) Arbitrary switching between permissible attitudes is irrational.
- (C) Permissivism allows arbitrary switching and, therefore, irrational behaviour.

While there are manifold ways permissivists have responded to the Arbitrariness Objection, we can group most of them into two camps: those that reject (P2) and thereby say that acknowledged permissivism is impossible (Cohen 2013; Smith 2020), and those that accept acknowledged permissivism but reject (P3) by thinking that something else prevents us from switching arbitrarily between permissive attitudes (Schoenfield 2014b; Simpson 2017).

While I will not discuss any of these strategies here, it is easy to see that understanding allegedly permissive cases as conflicts between incommensurable epistemic alternatives makes these moves redundant, because we already have grounds for rejecting the Arbitrariness Objection. We just need to ask ourselves whether we think the above-described features of comparisons under incommensurability allow agents faced with these conflicts to arbitrarily switch between alternatives. Doing so allows us not only to work out a unified resolution that works for value conflicts and epistemic conflicts alike but also explains if, why, and when switching is (ir)rational. However, how this resolution precisely plays out depends on our exact views about the role of bindingness in conflicts. In the rest of this section, I will subsequently discuss three increasingly liberal understandings of bindingness [§1.6.1] – [§1.6.3].

1.6.1 Strong Bindingness

First, we can understand bindingness in a strong sense as described by Chang (2005) or Williams (2014). For them, once we opt for one alternative we are normatively bound and there is no going back. If this understanding of bindingness is right, we can easily see that acknowledging that we are in a conflict case does not allow us to switch. While this mirrors permissivist responses which reject (P3), the reasons why acknowledgement does not allow us to switch is different from

extant responses. For example, Schoenfield (2014b), thinks that while different epistemic standards might be acknowledged as being equally permissible, epistemic agents need to be *immodest*, that is, take their own standards to be more truth-conducive than any alternative permitted set of standards.³³ On the contrary, the proposed understanding of permissive cases as epistemic standard conflicts does not need to appeal to immodesty. We can reach a similar result by pointing out that unbinding standards would make epistemic agents vulnerable to diachronically irrational patterns of behaviour and, therefore, standards need to be binding. Here, bindingness is an intrinsic feature of incommensurability and not the consequence of some other independent principle (such as immodesty) but simply results from considerations similar to the ones motivating the Arbitrariness Objection.

If we subscribe to this strong reading of bindingness, we have a neat resolution to the arbitrariness worry. Taking on an epistemic attitude is normatively binding, at least until one encounters new reasons to re-deliberate. However, I think we can and *should* utilise the insights reached above to give a more nuanced view which tells us that arbitrary switching is not *necessarily* irrational (P4).³⁴ This requires us to employ a weaker reading of the bindingness of conflict resolutions.

1.6.2 Weak Bindingness

Let us take another look at SARTRE'S STUDENT. At first, while standing on the platform, Sartre's student is permitted to go home to care for his elderly mother (A) or join the Free French Forces (B). Moreover, since the situation is sweetening insensitive finding out (unexpectedly) that the train ticket home costs more than the train ticket to Paris he still faces a permissible choice: going home and buying a train ticket for £ 100 (A-) or joining the Free French Forces and buying a train ticket for £ 50 (B). Similarly, we may *mutatis mutandis* think that (if the tickets are equally expensive) he might also jump off the train halfway to Paris and go home instead (A-). If the additional cost of £ 50 did not rule out going home, the additional cost of spending a few hours on the train and paying for the extra ticket may as well not make his initial decision definite [§1.4.1]. However, this is incompatible with the alleged bindingness of conflict resolutions. If conflict resolutions were normatively binding, Sartre's student would not be permitted to jump off the train and go home instead (A-) since he could have gone home directly (A).

³³ The *locus classicus* on immodesty is Lewis (1971).

³⁴ A different resolution to the Arbitrariness Objection which also does away with (P4) is discussed by Ye (2019).

How do we resolve this apparent tension between sweetening insensitivity and bindingness? My preferred view is to loosen bindingness as follows. Standing on the platform at t_1 , Sartre's student is not only confronted with a choice between A and B but also with various non-optimal versions of A and B. For example, he could decide to board the train to Paris and jump off after ten minutes (A_{-}), an hour (A_{\sim}), or halfway through (A_{\rightsquigarrow}), to go home instead. Picking any of these non-optimal options at t_1 is clearly irrational since he could have had the optimal version of A. That is, he is not permitted to plan to be diachronically incoherent. However, once he is on the train at t_2 and he finds himself deliberating (for reasons to be specified) whether to jump off the train and go home to care for his mother instead, A is not on the table any more but only other incommensurable options, including A_{-} and B. Picking A_{-} at this point, would not be any less rational than picking A_{-} at t_1 would have been if he had faced a choice between A_{-} and B.

To admit this is to surrender to some diachronic incoherence. If Sartre's student is allowed to switch back at t_2 he might as well be allowed to switch back at t_3 , to just switch back at t_4 , and so on. In the worst case, he might find himself in a downward spiral, dithering between train platforms torn between increasingly worse incommensurable choices.³⁵ Nonetheless, this is not the kind of incoherence theorists of rationality are usually worried about. After all, we cannot abuse the situation and pump money from the student (or anything else he values throughout), since the student does not plan to behave in this indecisive way but just happens to find himself deliberating over and over again. Neither does the student ever have an incoherent set of attitudes with respect to the courses of action he thinks he ought to take. He never thinks that he should board one train and stay on it while he simultaneously thinks that he should jump off halfway.³⁶

One way to avoid this result is to give a theory of when it is rational to re-deliberate. The strict view that restores diachronic coherence is that re-deliberation or re-planning is only allowed if we acquire new relevant information; that is, information that could in principle change the nature of the choice. (This is the above-mentioned view defended by Chang (2005) and Williams (2014)). On the other hand, we might think that rationality only tells us how to resolve

³⁵ Chang (1997) uses a similar example in which someone is repeatedly trading between ever more lukewarm cups of coffee and tea. See also Moss (2014: 672).

³⁶ This is similar to the way some have depicted the argumentative role of Dutch Book arguments for Bayesian conditionalisation [§5.5]. Following Lewis (1999: ch.23) we might say that being vulnerable to Dutch Books *per se* is not irrational but to believe that one ought to take bet one (ϕ), take bet two (ψ) and take bet three (λ) but *not* ought to take all three bets (ϕ & ψ & λ) is. I want to thank Edward Elliot for comments on this.

deliberation and not when we ought to deliberate. On this later understanding, there is nothing irrational in deliberating on the same body of evidence multiple times even if it ends in switching endlessly between increasingly worse options. The act of deliberation might be an arational event that happens to Sartre's student, and as a rational agent, he is always permitted to resolve the conflict one way or another independently of his previous choices. This, however, seems to completely lose any sense in which rational choices are binding; i.e., we have lost any guarantee that Sartre's student is not switching arbitrarily over and over again.

I think that rather than going for one of the extremes we ought to aim for a middle ground and try to normatively evaluate some general habits and dispositions of re-deliberation. In his planning theory of intentions, Bratman (1999) spells out various conditions for when it is rational to reconsider an intention, which we could apply to our case of re-deliberation.³⁷ For Bratman, we need to investigate the question on two levels, the level of the general reasonableness of the habits of reconsiderations, as well as the particular reconsideration itself. On the general level, we should have some disposition to reconsider even if we do not bump into new evidence, since "it seems plausible to suppose that it is in the long-run interests of an agent occasionally to reconsider what he is up to" (1999: 75). In other words, there are some reasonable and some unreasonable habits and dispositions with respect to re-deliberation. Whereas Sartre's student disposition leading him to dither between train platforms is irrational, others that lead to one-off reconsiderations are not. Needless to say, judging whether a reasonable habit is manifested in a particular instance of re-deliberation is tricky, especially in conflict cases.

However, while I will not be able to offer a general systematic answer to the question of re-deliberation, I think the above-described angst or resolutional remainder [§1.4.3] can play a twofold part in the story. On the one hand, angst may induce re-deliberation. Acknowledging that we have just faced a (potentially permissive) conflict may induce angst that, in turn, may push us to re-deliberate. After all, the defining feature of the angst that governs these conflicts is that we cannot as in other difficult choices "take comfort that we did as much as we could" (Williams 2016: 418). This is the resolutional remainder that prevents such conflicts from ever being truly resolved. On the other hand, knowing that constant re-deliberation might trap us in

³⁷ Bratman thinks that, while the reconsideration of an intention might be initiated by deliberation to reconsider, most reconsiderations are non-reflective, "one typically does not deliberate about whether to reconsider, but just goes ahead and reconsiders" (1999: 67). This makes his analysis of non-reflective reconsiderations especially helpful when considering the question of re-deliberation, since, as in the case of reconsideration, one typically does not deliberate about whether to deliberate but just deliberates directly.

the above-described loop of uncertainty makes the faced conflict even more daunting, and the resolution opted for more binding. It is a common assumption of ordinary decision-making that our future selves will continue to be rational since otherwise we cannot trust them to act in accordance with our preferences and beliefs. Incommensurable choices are especially unnerving in telling us that even if our future selves are completely rational we may end up a plaything of cosmic forces, trapped between two incommensurable options. It is this angsty realisation that can have the opposite effect by giving us reasons to stick to our initial resolution. That is, angst has a stabilising as well as a destabilising force.

Nonetheless, all of this is *not* to say that Sartre's student is always permitted to turn back. Once, he is entangled enough with the subsequent consequences of his decision, say after he has arrived in Paris, at t_n , he ceases to be in a permissive case. At t_n he faces a choice between the strictly worse option of going home and a strictly better option of staying. In other words, once he has gone beyond the *sweetening-insensitive grey area* Sartre's student is bound by his decisions.³⁸

Let me summarise the observations made so far by applying them to the epistemic case. While sweetening sensitivity allows for some amount of switching, bindingness prevents epistemic agents from (a) planning to switch among permissible standards, as well as having a habit that (b) leads to constant re-deliberation or (c) switching once the sweetening-insensitive grey area has been surpassed. First, planning to be indecisive is clearly irrational. There is no point of view, neither the atheist nor the theist perspective from which switching between the permissible standards makes any sense. If Ira truly thinks that the testimony of her fellow believers justifies her belief in God, why should she think that she ought to be indecisive before settling on one attitude? Second, after recently becoming an atheist, Ira is still allowed to give the encountered arguments for atheism a different reading, adopt different standards and reacquire her belief in God. Even after having made some cognitive effort to become an atheist, turning back to theism is still a rational option. However, once she has surpassed the sweetening-insensitive grey area bindingness kicks in. Once, she is entangled enough with the subsequent consequences of her deliberation, once she has truly become an atheist, she has ceased to be permitted to switch by becoming a certain type of reasoner.

This weak understanding of bindingness gives us a novel view not only of epistemic conflicts and the Arbitrariness Objection but also of the normative evaluation of subsequent decision-

³⁸ Note that talking of a sweetening insensitive grey area does not imply that there is any determinate range within which sweetening insensitivity applies but only that there are clear cases in which this grey area has been surpassed.

making under conflicts. While I think that this understanding of bindingness captures the phenomenon better than the stronger view outlined above [§1.5.1], it is worth mentioning that some have argued that incommensurability allows for rational diachronic incoherence even on a deeper level. I will discuss one such view in the next section [§1.5.3], before summarising the proposed understanding of the bindingness of conflict resolutions.

1.6.3 Changes of Mind

Moss (2014) argues that *changing our minds* by changing our values can be rationally permissible despite sometimes resulting in situations in which we are accepting sure losses. In her treatment of credal dilemmas, she argues that agents can have imprecise credences among which they can choose or ‘identify with’ in case they are forced to act upon them, even when this would lead to diachronically incoherent behaviour.³⁹ She points out that while “in a decision situation, an agent must act to maximise expected value according to the precise mental state she identifies with [...] there is no rule of rationality saying that an agent cannot change which mental state she identifies with.” (2014: 673).⁴⁰ If Moss is right that loosens the alleged bindingness of conflict resolutions even further; i.e., the bindingness that arises after surpassing the sweetening insensitive grey area is conditional on us not changing our minds. Once we have truly changed our minds and acquired new values, we are simply not bound by our old values anymore. That is, while doing A and B simultaneously, or planning to do A and then B might be irrational, doing A, changing your mind and then doing B is not necessarily irrational. How these changes of mind occur might again be partly explained by the angst arising from acknowledging the alleged permissiveness of the situation one finds oneself in.

If we combine Moss’ insights concerning changes of minds with the above-introduced weak understanding of bindingness we get the following picture. First, sweetening sensitivity allows us to switch between alternatives. While we are not allowed to plan to switch, if we find ourselves deliberating, switching among incommensurable options is rationally permissible given the underlying dispositions are generally reasonable; but only as long as we have not surpassed the

³⁹ Imprecise credences are credences which are not represented by a single probability measure but by a set of probability measures. For early canonical defences of imprecise credences see Levi (1974) or van Fraassen (1990). For a critical discussion see, for example, Elga (2010).

⁴⁰ Note, that these changes of mind, while similar to the above-mentioned arational view of re-deliberation, take place on another level since re-deliberation as discussed above presupposes that our values remain constant throughout.

sweetening-insensitive grey area. At this point re-deliberating (without acquiring new evidence) is only rational if we truly change our minds. If we acquire new or change our old epistemic standards, we are not bound by our old attitudes and decisions anymore but are rationally required to change our beliefs in accordance with our new standards. This is true even if such a change would be considered irrational from our previous epistemic perspective. How such changes of mind can occur is partly mysterious. Sometimes, they can just happen. A bump on the head might lead to me valuing parsimony more than elegance. There is nothing intrinsically irrational about that. Nonetheless, we have identified one driving force of such epistemic transformations: angst [§1.4.3]. Angst, or genealogical anxiety, induced by acknowledging the permissiveness of our epistemic situation, may take over and exert enough psychological pressure to reconsider large belief clusters or even our entire worldview. While there is a cognitive cost to such Cartesian reboots, sometimes the experienced angst is simply strong enough that overcoming it by changing our minds outweighs this cost.⁴¹

The proposed understanding of the arbitrariness of doxastic attitudes and the bindingness of conflict resolutions rely on contentious assumptions about the overall role diachronic considerations have to play in rationality; a discussion of which I have only scratched the surface. Importantly, however, the proposed analogy between value conflicts and epistemic conflicts holds independently of these issues. Whether we take conflict resolution to be binding in the strong or in some weaker sense we are in possession of a neat resolution of the most prominent objection to permissivism: the Arbitrariness Objection. This, together with the observations made in the previous section [§1.5], gives us a novel understanding of the expected pattern of belief evolution in epistemic conflicts. Whether those conflicts are grounded by indeterminacy, and to be resolved by indulging in weak permissions, or grounded in parity, and resolved by indulging genuine permissions, we get the same pattern of expected belief evolution, a pattern which is explained by the trade-off between sweetening insensitivity, bindingness and epistemic angst, and which may sometimes be interrupted by a sudden change of mind.

⁴¹ Whether or not re-deliberation or changes of mind induced by angst are arational or rational is up for debate. But I think that acknowledging that angst is an intrinsic feature of the rational resolution of epistemic conflicts gives us some *prima facie* reason to think it is the latter.

1.7 Conclusion

We should think of doxastic attitudes in epistemic conflicts as incommensurable rather than equally rational. Allegedly permissive cases, as they are discussed in the recent literature, are best understood as cases of epistemic standard conflict and, thereby, analogous to cases of value conflict discussed in other normative domains. I have called this the Conflict View:

Conflict View: Apparently permissive cases are epistemic standard conflicts. That is, cases in which one body of evidence supports different doxastic attitudes relative to different incommensurable sets of (permissive) epistemic standards.

I have argued for the Conflict View by showing that various prominent features of incommensurability [§1.4.1] – [§1.4.3] are present in both value conflicts and epistemic conflict. As such, the Conflict View presents us with a novel understanding not only of the comparison of doxastic attitudes under conflict but also of issues linked to arbitrariness and belief evolution [§1.6]. Interestingly, viewing allegedly permissive cases as epistemic conflicts is in principle neutral on the question of whether epistemic rationality is permissive or not [§1.5]. Moreover, it shows that having a permissive or impermissive understanding of rationality has less impact on the expected belief-evolution of rational agents than is often assumed.

Overall, the Conflict View presents us with a general understanding of a wide range of allegedly permissive cases discussed in the literature that can be accepted by both permissivists and impermissivists. It covers all cases that involve any type of apparent epistemic standard conflicts and so all types of cases that have been used to motivate Standard Permissivism in one of its many forms. However, there is one cluster of cases having to do with self-fulfilling beliefs that have been argued to be permissive, but which seem to involve no conflicting epistemic standards. Accordingly, these types of cases pose a threat not only to Evidential Uniqueness but also to the Conflict View as a general explanation of apparently permissive cases. I will discuss these cases in the next chapter [ch.2] and argue that we should understand them to be epistemically impermissive. Hence, those cases are neither epistemically permissive nor do they display genuine epistemic conflicts.

2

Uniqueness, Permissivism, and Self-Fulfilling Beliefs

In the previous chapter [ch.1], I motivated and defended what I have called the Conflict View of allegedly permissive cases. The Conflict View rests on the assumption that permissive cases always involve a direct or indirect conflict of epistemic standards. This is promising since it provides us with novel insights into the overall belief-formation dynamics that govern these cases while remaining neutral with respect to their epistemic permissibility. However, there is one often-discussed type of allegedly permissive cases, so-called self-fulfilling belief cases, which appears to be immune to this kind of treatment. That is, despite their alleged permissibility these cases do not involve any conflicting epistemic standards.

While permissivists often utilise these cases to generate potential counterexamples to Evidential Uniqueness, impermissivists usually dismiss them as being beyond the scope of epistemic rationality. In this chapter, I will illustrate that self-fulfilling beliefs are not only governed by epistemic rationality but also compatible with Evidential Uniqueness. Hence they neither pose a threat to the Conflict View defended in the previous chapter [ch.1] nor do they support a permissive view of epistemic rationality. In particular, I will demonstrate, by carefully disentangling questions of practical and epistemic rationality, that for any self-fulfilling situation, there is not only a unique rational doxastic attitude to be in but also only a single rational doxastic attitude to transition into.

2.1 Introduction

Sometimes we are confronted with beliefs whose propositional content is more likely to be true in case we form the respective belief. Take, for example, the belief that you'll perform well in your upcoming job interview. Many of us are psychologically hardwired in a way that makes having this belief evidence for its own content: if you believe that you'll perform well, you're likely to do so, while if you believe that you'll not perform well, then your performance is likely to be mediocre. Let us call these kinds of cases self-fulfilling belief cases (hereafter: SFB-cases).

While SFB-cases are long-known sources of epistemological controversy they have recently been used to generate counterexamples to Evidential Uniqueness.¹ Here is the definition of Evidential Uniqueness I gave in [§1.2] again:

Evidential Uniqueness: For any body of evidence E and any proposition p there is at most one doxastic attitude that any epistemic agent A could rationally take.

Since Evidential Uniqueness is formulated as a general thesis a single convincing example of a permissive case would be enough to disprove it. While many alleged counterexamples have been discussed in the literature it has turned out to be very challenging to give clear-cut examples which withstand closer inspection (Kopec & Titlebaum 2016; Ross 2021). Among the best candidate cases proposed are *self-fulfilling-belief cases* such as the one described above (Raleigh 2015, 2017; Kopec 2015; Kopec & Titlebaum 2016; Drake 2017; Antill 2019; Dahlback forthcoming).

Furthermore, SFB-cases are challenging since they seem to make rationality epistemically permissive even in cases where there are no opposing epistemic standards. So, if successful, SFB-cases can not only be used as counterexamples to Evidential Uniqueness but also as counterexamples to some forms of permissivism such as Standard Permissivism, the kind of permissivism that motivated the Conflict View [§1.2]:

Standard Permissivism: For any body of evidence E and any combination of epistemic standards S_1 - S_n there is a unique rationally permissible doxastic attitude D towards any proposition p . There are cases in which one body of evidence E rationalises different doxastic attitudes D_1 - D_n relative to different sets of standards S_1 - S_n .

In other words, SFB-cases support a very strong and seldomly defended form of permissivism: strong synchronic *intrapersonal* permissivism.² This form of permissivism allows for possible evidential situations in which it is rationally permissible to believe that p , but it is also rationally permissible to believe that not- p for one epistemic agent at one moment in time. Consequently,

¹ There is a long tradition going back to James (1895) that uses SFB-cases to motivate pragmatist non-evidentialist understandings of epistemology. Others, most notably Velleman have used SFB-cases to argue for a framework that combines non-evidentialist, permissivist, and voluntarist ideas (1989). More recently, some have revived these discussions to argue for non-evidentialism (Reisner 2013; Antill 2019; Silva 2023), doxastic voluntarism (Reisner 2013, Peels 2015), or epistemic permissivism (Raleigh 2015, 2017; Kopec 2015; Kopec & Titlebaum 2016; Drake 2017; Antill 2019).

² This is pointed out by Kopec (2010). For a defense of diachronic *intrapersonal* permissivism see Jackson (2021). I will also discuss intrapersonal permissivism in [§3.3].

if SFB-cases are permissive they do not only present us with a counterexample to Evidential Uniqueness but also pose a threat to weaker and more promising *interpersonal* versions of permissivism, such as Standard Permissivism.

Despite their reception in the literature, I will illustrate that SFB-cases are compatible with Evidential Uniqueness. First, I will introduce one guiding SFB-case to illuminate the dialectical structure of the debate, identify common grounds between permissivists and impermissivists and discuss proposed impermissive defence strategies. I will conclude that none of these strategies provides us with a satisfactory resolution, especially when it comes to the question of how to think about doxastic attitude formation in SFB-cases [§2.2]. It has been argued that doxastic attitude formation in SFB-cases is problematic since the transparently self-fulfilling nature of these cases prohibits rational deliberation [§2.3]. Despite these worries, I will not only demonstrate that there is a way to rationally transition into a doxastic state [§2.4], but also that this transition supports an impermissive understanding of epistemic rationality [§2.5]. In SFB-cases, practical considerations about prospective behaviour together with general assumptions about the (in)voluntary nature of belief formation generate additional evidence that tips the evidential balance towards the preferred self-fulfilling doxastic attitude. This leaves us with genuinely indifferent SFB-cases, cases in which there is no preferred doxastic attitude or outcome [§2.6]. To resolve these cases, we need to appeal to practical permissivism, the view that we can *pick* rather than *choose* actions under preferential indifference. When picking among practically permissible alternatives epistemic agents generate evidence that again singles out a uniquely rational doxastic attitude [§2.7]. This not only refutes the alleged epistemic permissiveness of SFB-cases but also offers us important insights into the nature of doxastic voluntarism and ideal rational agency.

2.2 Self-Fulfilling Beliefs

Let me start with a specific SFB-case which will be used as a guiding example throughout the subsequent discussion. This will not only help me to reconstruct the dialectical structure of the current debate but also demonstrate the novelty of my proposal.

MILLIONAIRE FORECAST: A powerful magical Genie appears in front of an ideal rational agent, called Ira, and offers her the following proposal: If she believes that she will be a millionaire tomorrow, then the Genie will make sure that p ‘she will be a millionaire tomorrow’. But if Ira believes she won’t be a millionaire tomorrow, then the

Genie will make sure that not p ‘she won’t be a millionaire tomorrow’. After rationally convincing Ira of its magical powers and mind-reading abilities, it seems that Ira is rationally permitted to believe either, that she will be a millionaire tomorrow (henceforth: Bp) or that she won’t be a millionaire tomorrow (henceforth: $B\neg p$).³

This case is a prime example of a so-called *transparent* SFB-case. Those are cases in which not only do the doxastic attitudes of the epistemic agent influence their own propositional content but also in which the epistemic agent is perfectly aware of the self-fulfilling nature of their situation (Muralidharan 2021: 219).⁴ Many permissivists have argued that transparent SFB-cases are forthright counterexamples to Evidential Uniqueness (Raleigh 2015, 2017; Kopec 2015; Kopec & Titlebaum 2016; Drake 2017; Antill 2019; Dahlback forthcoming). When considering whether to believe that p , Ira seems rationally permitted to acquire either doxastic attitude Bp or $B\neg p$, since she possesses sufficient evidence that the respective belief, once formed, will become true. That is, both beliefs are rational since both beliefs are sufficiently supported by (and could be properly based on) the evidence possessed by Ira.

Furthermore, there are various reasons to think that suspension of judgement towards p is not a (maximally) rational option. First, since Ira is aware that whatever she believes will turn out to be true, suspending judgement means that she consciously dismisses a belief which is guaranteed to be true, once formed (Kopec 2015: 407; Antill 2019: 325; Muralidharan 2021: 218). Moreover, as Dahlback (forthcoming) has shown, it is possible to construct modified SFB-cases which rule out suspending judgement altogether. We just need to make sure that the Genie punishes Ira for suspending judgment by giving her sufficient evidence that she is guaranteed to have a false belief when doing so. (Without judging the legitimacy of these strategies, let us bracket this discussion for now, and assume that they successfully rule out suspension of judgment as a rational option. When I have spelt out my resolution of SFB-cases, I will be able to offer a general way of thinking about suspension of judgement in SFB-cases.)

So far this should be common ground: the evidential situation favours Bp as well as $B\neg p$ equally well (but disfavours suspension of judgement towards p). Therefore, conventional strategies used to dismiss other permissive cases, such as claiming that the underlying body of

³ A similar case is also discussed by Peels (2015) and Antill (2019).

⁴ For the sake of narrative ease, the following discussion is idealised in various ways. First, it focuses on the rational permissibility of full-on beliefs rather than fine-grained doxastic attitudes. Second, it is restricted to cases in which the self-fulfilling nature of the beliefs is not only completely transparent but also truth-guaranteeing. Despite these idealisations, I am confident that the observations made are general points which could be applied to non-idealised SFB-cases as well.

evidence, upon closer inspection, favours one of the doxastic attitudes or that the doxastic attitudes can only be held by being less than perfectly rational are not available to the impermissivist (Ross 2021: 206-212).

Accordingly, impermissivists proposed other strategies to save Evidential Uniqueness against SFB-cases. Ross in his systematic engagement with allegedly permissive cases, suggests a modification of the initial thesis which restricts Evidential Uniqueness “so that it only applies to act-state independent doxastic attitudes” (2021: 212). Other strategies argue that self-fulfilling beliefs are not governed by epistemic rationality but by practical rationality, or that they are not governed by rationality at all since they violate various rational principles (Antill 2020; Marxen 2021a, 2021b; Muralidharan 2021). Instead of discussing these strategies systematically, I will demonstrate that they are at least incomplete, for the simple reason that I am able to offer a purely epistemic resolution of SFB-cases that is compatible with Evidential Uniqueness [§2.5] [§2.7].

Another approach to deal with SFB-cases is to give Evidential Uniqueness a synchronic reading. Greco and Hedden have argued that Evidential Uniqueness is not about a uniquely rational state to transition into but about a uniquely rational state to be in right now (2016: 392-393):

“If you believe that you’ll give a great talk, you’re likely to do so, while if you believe that you’ll give a mediocre talk, then your talk is likely to be mediocre. [...] Relative to your initial evidential state, you ought to suspend judgment on whether you will give a great talk. But if you somehow then acquire the belief that you’ll give a great talk, then this belief is rational, relative to your new evidential state, for having the belief is evidence for its content. [...] All this is quite compatible with [Evidential] Uniqueness, for each body of total evidence uniquely fixes what beliefs you ought to have. It’s just that you can affect your evidential state by forming certain beliefs. Note that what has been said doesn’t commit us one way or another on the question of whether the transition–formation of such a self-fulfilling belief–would count as rational, irrational, or arational.”

I think Greco and Hedden make two important observations here. First, (i) having a self-fulfilling belief changes the body of evidence, and second, (ii) we need to distinguish between the permissibility of doxastic attitudes and the permissibility of doxastic attitude formations.

I take the former (i) to be uncontroversial. When a self-fulfilling belief is formed, we have effectively generated new evidence: once Ira believes that she will be a millionaire tomorrow she has good evidence that this belief is true. After having generated this additional evidence the

evidential balance shifts and the belief that p rationally dominates the belief that $\neg p$. So, having the belief makes it the only rationally permitted doxastic attitude.

This brings us directly to (ii). Greco and Hedden point out that we need to distinguish between two understandings of Evidential Uniqueness, where the former is concerned with the rationality of doxastic attitudes and the latter is concerned with rational ways of transitioning into or among doxastic attitudes. Although I agree with Greco and Hedden (2016) that this distinction is important, restricting our discussion to the former, synchronic version of Evidential Uniqueness is only satisfactory if we already defend a broadly synchronic understanding of rationality.⁵ While permissivists and impermissivists alike have often been unclear whether they are concerned with synchronic or diachronic matters about rationality, some have argued that we should pay more attention to diachronic questions (Podgorski 2016). Furthermore, many of the original discussions suggest a diachronic understanding of Evidential Uniqueness. For example, Schoenfield describes permissivism as the thesis that sometimes there is “more than one way to rationally respond to a given body of evidence” (2014b: 193). And Kopec and Titlebaum point out that most of the discussions have focused on *personal* Evidential Uniqueness, a thesis about how fully rational agents ought to handle a body of evidence, rather than *propositional* Evidential Uniqueness, a thesis about the relationship between bodies of evidence and propositions (Feldman 2007; Matheson 2011); where they understand the former as the thesis that “a body of evidence dictates a single attitude that a rational agent must settle upon after she reasons through the question at issue” (Kopec & Titlebaum 2016: 190).

In sum, the subsequent analysis of SFB-cases is built on two insights made in previous discussions. First, having one of the self-fulfilling doxastic attitudes changes the body of evidence in a way that makes it the only maximally rational option. Second, before forming any of the self-fulfilling attitudes it appears that in SFB-cases there can be multiple equally rational doxastic attitudes to transition into. Hence the task for those who defend a more general synchronic and diachronic understanding of Evidential Uniqueness is to refute the alleged permissibility of those doxastic attitude transitions. Accordingly, I will start out by taking a closer look at what could be meant by rational doxastic attitude transition both from a purely epistemic [§2.3] and from a practical [§2.4] point of view.

⁵ See, e.g., Hedden (2015b).

2.3 Theoretical Deliberation and Self-Fulfilling Beliefs

As pointed out above, when talking about the rationality of doxastic attitudes, we need to distinguish two often conflated epistemological endeavours: on the one hand, we may ask ourselves whether a doxastic attitude is rational given a certain body of evidence, and, on the other, we may ask ourselves what the appropriate response is to that body of evidence.

To analyse the rationality of a transition from one doxastic state D_1 into another D_2 we need to make some further clarifications. On the one hand, epistemology is concerned with (i) *theoretical deliberation*, the implicit or explicit process in which rational agents figure out what doxastic attitude they ought to have towards p given their evidence (or overall epistemic situation). On the other hand, we may think about (ii) the *doxastic attitude formation*, the process underlying the actual transition from D_1 to D_2 after having reached a conclusion in (i).

Ordinary epistemic agents are restricted with respect to both tasks: they often do not know what the most rational doxastic attitude is, but even if they do, a successful transition into this attitude is often prohibited. In other words, for any epistemic agent, two gaps need to be bridged when transitioning from D_1 to D_2 . The first gap is a consequence of having limited cognitive capacities. Ordinary epistemic agents can fail to settle on the most rational attitude because they fail to reach a conclusion or settle on a wrong attitude because they misinterpret the evidence. The other gap is related to the purportedly non-entirely voluntary nature of doxastic attitude formation. Ordinary epistemic agents can fail to settle on the most rational doxastic attitudes, even if they reach the conclusion that they should acquire D_2 . That is, they can after concluding remain in the akratic state of believing that they should transition into D_2 but fail to do so. In contrast, ideal rational agents, such as Ira, do not have any of these cognitive restrictions. They can directly derive the most rational doxastic attitude given their evidence and also transition immediately into the required attitude.⁶ In other words, they will always believe the most rational thing to believe given their evidence and their overall epistemic situation (that is, their evidence and their epistemic standards).

It is, however, important to point out that this does not suggest that ideal rational agents are doxastic voluntarists, who can acquire doxastic attitudes at will. First, doxastic voluntarism can mean a multitude of things, e.g., that epistemic agents can form beliefs, knowingly, intentionally, partly for practical reasons, or even independently of any truth-considerations (Peels 2015: 526-529). As I will illustrate in [§2.4], depending on different readings of doxastic voluntarism we can

⁶ See, e.g., Broome (2013: 154-155).

treat ideal rational agents as either voluntarist or involuntarist creatures. The above-made observations in contrast, just reveal that in any epistemic situation in which there is some proposition that some agent *A* ought to believe, if this agent is ideally rational they directly transition into that belief (upon discovering that they ought to believe it). This, however, does not settle the debate between permissivists and impermissivists. It does not tell us whether there are situations in which more than one doxastic attitude could be fully rational, but only that ideal rational agents after deliberation directly transition into one of the maximally rational doxastic attitudes.

We can now look at *theoretical deliberation* within SFB-cases. Theoretical deliberation is often depicted as a *functional process* taking the initial state of the epistemic agents as input and prescribing certain permitted doxastic attitudes as output. This way of thinking about rational deliberation goes back at least to Goldman (1980), who speaks of rational deliberation as a function whose inputs include “beliefs, perceptual fields and ostensible memories” and whose outputs are doxastic attitudes of all kinds, such as belief, disbelief, credences, or a suspension of judgement (Goldman 1980: 29). In short, in theoretical deliberation, we take all our currently (accessible) evidence and derive certain doxastic attitudes as appropriate responses to the evidence.⁷ By adopting this model, we treat the evidence possessed by the agent as given, fixed prior to and independently from the deliberation itself.

When it comes to SFB-cases, some, including Grice (1972) and Antill (2019) have argued that this functional understanding of theoretical deliberation is problematic as it transforms any resolution of SFB-cases into a problematic act of bootstrapping. From the perspective of Ira finding herself deliberating whether to believe or to disbelieve that she will be a millionaire tomorrow, much of the relevant evidence accessible to her are current beliefs about her beliefs related to the possibility that she will be a millionaire tomorrow. This includes the conditional belief ‘if I believe that *p*, then *p*’ which appears vital in determining the appropriate doxastic attitude. However, believing the conditional is only useful to get us to the consequent if we already believe the antecedent. Since, in SFB-cases the antecedent just is the belief we are trying to acquire we cannot possibly use the conditional to reach a prescribed doxastic attitude (Antill 2019: 322). This shows that the fact that a belief is self-fulfilling has no bearing on whether to adopt a belief, it can only reinforce what we already believe (or give us a reason to disbelieve

⁷ Note that this presupposes either an abstractionist or mentalist understanding of evidence [§5.3].

things we have not considered so far).⁸ So, it seems like any way of acquiring one of the self-fulfilling beliefs cannot be a result of proper theoretical deliberation but must be the result of some form of bootstrapping oneself to the conclusion.

In response to these worries, Antill (2019) concludes that SFB-cases demonstrate that the function model of deliberation “is unable to model all the relevant factors that might go into a theoretical deliberation [which] shows not that it is impermissible to believe the proposition we are deliberating about, but that we have chosen a poor way of modelling deliberation” (2019: 324). To account for this deficiency, Antill supplements the function model with a transparent understanding of introspection. Following Evans (1982), many consider introspection to be *transparent* to the questions about the facts underlying the beliefs. If we introspectively consider whether we believe something we do not directly access facts about our mental lives, but facts which bear on the truth of the propositional content of the belief. If I ask myself whether I believe that *p*, I do not directly consider anything internal to me, but rather external facts that would determine my belief that *p*. In Evans’ words (1982: 225):

“ [...] in making a self-ascription of a belief, one’s eyes are, so to speak, or occasionally literally, directed outward –upon the world. If someone asks me ‘Do you think there is going to be a third world war?’ I must attend, in answering him, to precisely the same outward phenomena as I would attend to if I were answering the question ‘Will there be a third world war?’”

If this transparent understanding of introspection is correct this changes the inputs of our deliberation. Rather than taking our mental states to be the relevant inputs we directly take the facts that bear on *p* as inputs. But how does this affect theoretical deliberation in SFB-cases?

Here, Antill (2019) argues that even if we understand introspection to be transparent, the situation looks equally grim. In investigating whether to believe that *p* we need to consider the fact whether *p*. To do so we must direct our gaze outwards to determine whether *p* is the case. This, however, requires us not only to ask whether I will be a millionaire but to open up “a new sub-question: “whether I now believe that I will be a millionaire?” [...] this, in turn, requires us to determine whether we believe that we will be a millionaire, and so on *ad infinitum*” (Antill 2019: 324; *emphasis in the original*). In short, according to Antill (2019), SFB-cases force us into a chain

⁸ See also Reisner (2013).

of reasoning which keeps us deliberating indefinitely, without our evidence ever being able to terminate our deliberation and make up our minds.

In sum, in SFB-cases theoretical deliberation seems to have no epistemic bearing at all on our belief formation processes (if introspection is intransparent), or deliberating what to believe forces us into an infinite chain of deliberation (if introspection is transparent). Antill (2019) uses these observations to argue against evidentialism and for the permissiveness of SFB-cases. For Antill, SFB-cases demonstrate the failure of evidentialist understandings of deliberation “since in following our evidence, we’ve failed to come to any conclusion at all” (2019: 325). Instead, Antill argues that in SFB-cases we need epistemic-but-non-evidential considerations to break out of the deliberation loop to form a belief one way or another.

However, even if Antill is right, taking this observation to support permissivism is too quick. After all, Evidential Uniqueness is a thesis about the rationality of doxastic attitudes concerning a given body of evidence. Therefore, the need to consider non-evidential aspects of the situation in SFB-cases may also be understood to demonstrate that these cases are not (solely) governed by epistemic rationality but practical rationality.⁹

This seems like a stalemate. Either we think that non-evidential considerations can rationally influence theoretical deliberation or not, and in doing so we reach different conclusions with respect to the question of whether SFB-cases are epistemically permissive. Leaving this disagreement unresolved for now, I will proceed by answering the following question: given that practical considerations would enable us to transition into a doxastic attitude in SFB-cases, how does this transition exactly work [S2.4]? Only after we have answered this question, we have the required tools to reinvestigate the allegedly permissive nature of SFB-cases [S2.5]. Here, I will demonstrate that thinking about practical deliberation in SFB-cases enables us to demonstrate that we can rationally settle on a doxastic attitude in a way that is compatible with evidentialism and Evidential Uniqueness.

2.4 Practical Deliberation and Self-Fulfilling Beliefs

Let us start again by putting ourselves into Ira’s boots. Instead of introspectively investigating whether she already believes that p by looking at her mental states or by investigating the facts that would make her belief that p true, Ira might take a different approach. Instead of asking

⁹ This is similar to some of the above-mentioned impermissive defence strategies [S2.2].

what she ought to believe she could ask herself what she *ought to do*: —should she *try* to Bp (henceforth: φ), or should she *try* to $B\neg p$ (henceforth: ψ)?

In so doing, Ira needs to answer two subquestions: what do I *want* to believe? —and, what *can* I believe? That is, she needs to transform her situation into a practical decision situation in which she can assign *utilities* and *probabilities* to a set of possible outcomes or propositions: either she (a) is in the doxastic state Bp and p will (most likely) be the case, or (b) she is in state $B\neg p$ and p will (most likely) not be the case.¹⁰ Once she has assigned the relevant utilities and probabilities to propositions (a) & (b), she can then apply some formal framework, such as causal decision theory, to determine the most rational decision, whether to perform the subsequent actions φ or ψ .¹¹

To figure out what she *wants*, Ira needs some introspective insight into her inner mental life, to ascribe utilities and a preference ranking to the relevant propositions. That is, Ira engages in mind-reading or mentalising, a higher-order activity that allows her to conceptualise her own motivational states from a third-person perspective, based on the information accessible to her. While there is a certain set of problems associated with such introspective endeavours, so far this is nothing out of the ordinary, but something that epistemic agents must do in mundane (non-self-fulfilling) decision scenarios as well. However, since in SFB-cases certain facts are generated by having certain doxastic attitudes, there is also additional potential utility generated by having a true (and justified) belief. In other words, depending on her epistemological viewpoint, Ira may not only ascribe a certain utility to p but also to the state of the world in which she has a true belief Bp (based on her knowing that this belief is self-fulfilling).

Figuring out what she actually *can* believe is more challenging. To do so, Ira needs to settle on some theoretical assumptions with respect to the nature of doxastic attitudes and rational attitude formation [§2.3]. First, she needs to have some understanding of the ontology of beliefs that helps her answer the following questions: do mere dispositions to have a belief towards p already trigger the Genie's self-fulfilling powers, or does she need to be in an occurrent belief state? Ira, furthermore, needs to have some understanding of how ideal rational agents (voluntarily) form doxastic attitudes. Doxastic attitude formation might be completely voluntary, voluntary to a certain extent or completely involuntary. Settling on this question one way or another provides Ira with important information about the probabilities of the relevant propositions. While assuming a sufficiently strong doxastic voluntarism guarantees that φ ing

¹⁰ The probability considered here is *epistemic probability*, the probability given Ira's evidence.

¹¹ Note that here we directly ascribe probabilities and utilities to propositions (Jeffrey 1965).

(almost certainly) leads to her believing that p , a sufficiently strong doxastic involuntarism ascribes no such powers to Ira. In contrast, if she settles in the middle (indirect doxastic voluntarism) and thinks that she can only indirectly influence her beliefs, she needs to answer how successful belief induction is concerning the particular beliefs in consideration. While Ira as an ideally rational agent immediately forms maximally rational beliefs that are evidentially required it is not entirely clear how, if at all, she is able to transition into one of the self-fulfilling attitudes [§2.3]. For all she knows there are coherence constraints on inducing beliefs at will: she can certainly not induce beliefs which require her to perform radical reconsideration of many of her previous beliefs without changing her body of evidence. Having a belief that does not fit into her web of beliefs is certainly less than fully rational. Therefore, to consider how radical a particular belief shift would be, she again needs to model her inner mental life, this time her propositional mental states.

In sum, in deliberating what she *ought to do* Ira constructs a mental model of her own motivational and propositional mental states, as well as her agential capacities which are built on a cluster of theoretical background assumptions about the nature of doxastic attitudes, preferences and doxastic attitude formation processes. This enables Ira to assign both probabilities (P) and utilities (U) to the possible outcomes or propositions, which, in turn, allows her to decide whether to φ or ψ by applying decision theory. The resulting decision trees are depicted in [Fig.1] – [Fig.3].




Fig. 1: Doxastic Voluntarism




Fig. 3: Doxastic Involuntarism




Fig. 2: Indirect Doxastic Voluntarism

There are multiple things that need to be explained about [Fig. 1] – [Fig. 3]. First, Ira needs to assign utilities U to various propositions: being a millionaire (p), not being a millionaire ($\neg p$), but also, having a true belief Bp , and having a true belief $B\neg p$. Though the propositions, that p and that having the belief that p are strongly coupled in SFB-cases their utility could still be independent. So even if Ira is indifferent towards p the proposition ‘ $p \ \& \ Bp$ ’ could have a positive utility (U_1) if Ira ascribes a positive value to having the true belief Bp . Adding those utilities together then provides determinate utilities for the complex propositions ‘ $p \ \& \ Bp$ ’ (U_1) and ‘ $\neg p \ \& \ B\neg p$ ’ (U_2). The utility for suspension, however, is not as simple to determine as the other ones, since Ira does not know what happens if she suspends judgement towards p . The Genie could do nothing, punish Ira for not collaborating, or it may flip a coin to decide whether to make it the case that p or not p .

Once the utilities are assigned, we need to think about the probabilities P that those propositions become true conditional on Ira’s decision to φ (try to Bp) or ψ (try to $B\neg p$). These probabilities are fixed by the corresponding understanding of the nature of doxastic attitude formation, which can be broadly categorised as follows:

- (1) If *doxastic voluntarism* is true, then the probability that φ ing or is ψ ing successful in inducing a belief is (close to) 1. In this case, it becomes just a matter of choosing one of the outcomes since there are no further chances involved [Fig. 1].
- (2) If *indirect doxastic voluntarism* is true, we will assign conditional probabilities P to the following propositions: Bp given that Ira is φ ing ($P_{1.1}$), suspension given that Ira is φ ing ($P_{1.2}$), $B\neg p$ given that Ira is φ ing ($P_{1.3}$), Bp given that Ira is ψ ing ($P_{2.1}$), suspension given that Ira is ψ ing ($P_{2.2}$), $B\neg p$ given that Ira is ψ ing ($P_{2.3}$); whereby assuming indirect doxastic voluntarism $P_{1.1}$, and $P_{2.1}$ strictly need to be larger than $P_{1.2} \ \& \ P_{1.3}$ and $P_{2.2} \ \& \ P_{2.3}$ respectively [Fig. 2].
- (3) If *doxastic involuntarism* is true, then the probability that φ ing or ψ ing is successful in inducing the desired belief is the same as if she is not φ ing or ψ ing. In counterfactual terms, the probability that φ ing induces the belief (try to $Bp \ \Box \rightarrow Bp$) is just the probability of Bp itself.¹² In other words, Ira is not really in a decision scenario anymore since whatever the initial probabilities are they are not changed by Ira’s behaviour. [Fig. 3].

¹² See, e.g., Gibbard and Harper (1978).

After having settled on these important questions calculating what Ira ought to do given (1) - (3) should be straightforward. If (1), Ira just chooses the action with the highest expected utility, e.g., if she wants to be a millionaire, she should φ . Similarly, if Ira has some but not perfect control over her doxastic attitude formation (2), she should try to transition into the doxastic attitude which changes the world in a way that maximises expected utility (this should be obvious without plugging in any actual numbers). In other words, if she desires to be a millionaire, she *ought* to φ and if she doesn't desire to be a millionaire, she *ought* to ψ . If (3), Ira cannot do anything to influence the state of the world. Hence, the solely rational thing to do is to see what happens and hope that the most valuable doxastic state manifests in her mind. Importantly, this is only so if Ira is *absolutely certain* that involuntarism is true. If she has some evidence that indirect- or direct doxastic voluntarism might be true she should still do the action associated with the preferred proposition (I will talk about this in more detail in [§2.6.3]).

In sum, in contrast to epistemic deliberation, we have reached definite results when it comes to practical decision-making in SFB-cases: if we have any influence on our doxastic attitude formation capacities we ought to choose the outcome with the highest expected utility. While this result requires cases in which there is a definite preference ranking with one preferred outcome, I will investigate cases of preferential indifference in [§2.6] and [§2.7]. Before doing so, however, I will take a step back to use our findings to reinvestigate the initial question of epistemic permissibility in SFB-cases [§2.5].

2.5 The Impermissibility of Self-Fulfilling Beliefs

What do these considerations about practical decision-making in SFB-cases tell us about the *epistemic* permissibility of doxastic attitude transitions? I think they illustrate at least two things. First, the above-considered model can be utilised by epistemic agents to predict their own behaviour which, in turn, enables them to settle on a doxastic attitude even before the respective action is initiated. This demonstrates that SFB-cases such as MILLIONAIRE FORECAST are not only epistemically impermissible but also practically impermissible (independently of doxastic attitude formation being voluntary or not). Second, it gives us important insights into the nature of belief formation in ideal rational agents. If you assume some indirect doxastic control in SFB cases, those cases collapse into full-blown voluntarist scenarios.

First, let us assume that doxastic involuntarism is true, or at least that Ira believes that her belief formation is involuntary. If Ira has no influence on p being the case, there is no reason to

φ or ψ , no way to predict the outcome, and no doxastic attitude to transition into. Epistemic rationality requires Ira to remain in whatever state she was in before she started considering the proposition. Maybe she believed or had a disposition to believe that she won't be a millionaire tomorrow because there was no reason to think that she would encounter a magical Genie, or she was in some default state of suspension or 'having no doxastic attitude at all'.¹³ Whatever that default state is, given doxastic involuntarism, Ira is normatively bound to it (I'll come back to this in [§2.6.2]).

Second, let's assume that Ira desires to be a millionaire and, therefore, ascribes a higher expected utility to p than to $\neg p$, and, furthermore, that she (thinks that) she has some (indirect) voluntary control over what she believes; i.e. her actions positively influence doxastic attitude formation towards the desired outcome. Given these assumptions, as argued above, practical rationality tells her that she ought to φ . That is, the utilities and probabilities ascribed to the possible outcomes together single out φ ing as the uniquely rational action for Ira.

However, what if before φ ing, she hesitates and reflects on her situation once more? During this time, that potential pause, something crucial happens. At this point, Ira knows various things that she did not know before. First, she knows that φ ing rationally dominates ψ ing. Furthermore, she knows that φ ing increases the chance that she will end up believing that p , which in turn will make p true. Therefore, she has strong evidence that p will be the case; evidence, that is via reflection sufficient to support the doxastic state Bp . In sum, every bit of information about the nature of the required practical choice is transformed into evidence for or against p and, therefore, the overall evidential balance shifts towards Bp . Settling the practical question of what Ira ought to do has settled the purely epistemic question of what Ira ought to believe. This is true, irrespective of whether we understand theoretical deliberation to be transparent or not [§2.3]. Assuming the intransparency of introspection, Ira reasons as follows: I believe that I will φ , which will more likely make it the case that I Bp , which will make it the case that p (input), I ought to Bp (output). In contrast, transparent deliberation goes like this: p will probably be the case since I will φ which will make it more likely that I Bp , which will make it the case that p . Therefore, I ought to Bp .

Here is another slightly different way of depicting the situation Ira finds herself in. Ira takes a bird's-eye view at herself to predict her own behaviour as she would predict the behaviour of

¹³ On many views ideally rational agents either suspend judgement as a default attitude towards every proposition or at least necessarily form a rational attitude once the proposition is deliberated. For an influential critique of both assumptions see Friedman (2013).

another epistemic agent. This effectively turns the first decision node in the decision tree into a chance node [Fig. 3], whereby the preferred outcome is more likely to come about. Based on that she forms a belief about the future state of the world which will in turn make the belief true. This also means that Ira changed her evidential situation, which, given the mechanisms that govern belief formation in ideal rational agents [§2.3], means that Ira will form the evidentially required attitude B_p immediately. There is no trying to believe anymore, practical considerations have settled epistemic considerations, which have made ϕ ing superfluous.

Note, that this isn't dependent on considerations about the amount of voluntary control that Ira has over her beliefs. Some might worry that given a sufficiently weak version of indirect doxastic voluntarism, ϕ ing may not make it sufficiently probable that Ira will believe that p , since Ira won't have enough evidence to believe that she will believe that p and hence insufficient evidence to believe p outright.¹⁴ Instead, the evidence would only allow some positive credence (>0.5) towards p . However, If Ira believes that some positive credence in p is enough to trigger the Genie's powers, the situation is self-reaffirming: believing that p to some degree is strong evidence that the Genie will make it the case that p , which will justify an even higher degree in belief, which is even stronger evidence that the Genie will make it the case, and so on until the evidence justifies outright belief in p .¹⁵

While this appears to support doxastic voluntarism, this is not an argument for doxastic voluntarism *per se* since we started with the assumption that Ira has at least some voluntary control over her doxastic states. But it is an argument that indirect doxastic voluntarism collapses into direct doxastic voluntarism in SFB-cases. Furthermore, this also gives us an argument against the suspension of judgement in cases such as MILLIONAIRE FORECAST. If practical considerations rule out trying to suspend judgement towards p and practical considerations fix epistemic considerations, suspension of judgement is not a rational doxastic attitude to transition into.

What does all of this tell us about the allegedly permissive nature of SFB cases? At first, the whole situation seems permissive. Given that multiple preferences are possible and rationally permitted Ira may reach a different conclusion which would single out a unique rational attitude. Two ideal epistemic agents Ira_1 and Ira_2 in the same evidential situation, one desiring to be a millionaire and the other desiring to remain poor would be allowed to transition into different

¹⁴ I am grateful to an anonymous referee for raising this objection.

¹⁵ On the other hand, if Ira knows or is unsure that having some positive credence in p is insufficient for triggering the Genie's powers the case can be assimilated into the doxastic involuntarism case; i.e. she needs to remain in her default state, may it be suspension, having no doxastic attitude or disbelieving that she will be a millionaire tomorrow.

doxastic attitudes. This is a version of *interpersonal* permissivism, which is a subtype of Standard Permissivism. While *intrapersonal* permissivism allows for situations in which more than one doxastic state is fully rational for the very same agent, *interpersonal* permissivism is only permissive across individuals. As standard permissivists *interpersonal* permissivists rely on a 3-place relational understanding of evidential support (Decker 2012; Kelly 2013), where evidential support is a relation between the evidence, epistemic-but-non-evidential factors and doxastic attitudes. Analogously, permissivists may argue that in the given SFB-case the required doxastic attitudes are determined by the evidence and the agent-relative preferences of Ira_1 and Ira_2 .

However, this strategy would ignore the above-described synthesis of practical and epistemic rationality, which enabled us to resolve theoretical deliberation in SFB-cases in the first place. In SFB-cases, epistemic rationality requires rational agents to make predictions about them realising certain self-fulfilling events, based on *evidence* about their own motivational and propositional states. In other words, if Ira_1 and Ira_2 have different preferences they would also have different bodies of evidence. That is, not their evidence plus their epistemic-but-non-evidential features but their evidence alone singles out a uniquely rational doxastic attitude. After all, it is not about their actual preferences but about their respective mental model that takes self-ascribed preferences into account.

In sum, while SFB-cases seem to blur the line between epistemic rationality and practical rationality, deliberation in SFB-cases is nonetheless governed by epistemic rationality. Arguments bolstering a non-epistemic interpretation of these cases [S2.1] [S2.3] ignore the power of self-fulfilling beliefs to transform considerations about practical rationality into evidence. Based on predictions about practical decision-making [S2.4], epistemic agents can make predictions about their prospective behaviour which is vital evidence that determines a uniquely rationally permissible attitude.

2.6 Genuine Indifference: A Recipe

So far, I have demonstrated that SFB-cases with preferred outcomes are not only practically [S2.4] but also epistemically impermissible [S2.5]. By transforming practical considerations into evidence, they generate a body of evidence that rationalises a unique rational response. Yet, this leaves an important question unanswered: what about cases in which there is no preferred outcome? Let's put *Ira*, our fully rational guinea pig, into the following scenario:

ASTRONOMY ORACLE: A powerful magical Genie appears in front of Ira and makes the following proposal: If she believes that the number of stars in the universe is even, then the Genie will make sure that the number of stars in the universe is even. But if she believes that the number of stars in the universe is odd, then the Genie will make sure the number of stars in the universe is odd. After rationally convincing Ira of its magical powers and mind-reading abilities, it seems that Ira is rationally permitted to believe either, that the number of stars in the universe is even (henceforth: Bp) or that the number of stars in the universe is odd (henceforth: $B\neg p$).

The set-up of this case stipulates that there are no reasons, evidential, pragmatic or otherwise for Ira to form a doxastic attitude one way or another. Ira is simply indifferent in any meaningful sense with respect to the number of stars in the universe being even or odd. In such cases, it is hard to see how Ira could or whether she should form any doxastic attitude towards the considered proposition. Recall the way we modelled decision-making in SFB-cases in [§2.4] and how evidence about rational decision-making allowed doxastic attitude formation [§2.5]. Ira modelled her prospective behaviour which allowed her to make predictions about her future mental states and therefore states of the world, which in turn allowed her to form the respective doxastic attitude. Unfortunately, this strategy isn't available when it comes to these indifferent SFB-cases, since both doxastic attitudes, and the related states of the world, have no or the same expected utility. Ira's forecast simply ascribes equal chance to both φ ing (trying to Bp) and ψ ing (trying to $B\neg p$) which would realise the respective state of the world.

I will demonstrate, nonetheless, that these cases can be resolved. To do so, we again need to start with the practical question: what *ought* Ira to do? To adequately answer this question we first need to clarify various details omitted in the construction of ASTRONOMY ORACLE. In other words, need to stipulate additional ingredients to cook up a so-called *genuinely indifferent* SFB-case. In these cases, epistemic agents not only have a genuinely indifferent body of evidence with respect to multiple doxastic attitudes [§2.6.1] but also possess evidence suggesting that they *ought* [§2.6.2] and *can* [§2.6.3] transition into one of those attitudes. Adding these further ingredients to the recipe comes with additional epistemic and metaphysical assumptions some of which carry an additional theoretical burden. This may arguably make even the stipulation of genuinely indifferent SFB-cases controversial. However, I argue that even if we accept the underlying construction the best way to understand genuinely indifferent cases is to think of them as being practically permissive but epistemically impermissive [§2.7].

2.6.1 The (Im)possibility of Indifference?

A common initial response to cases such as ASTRONOMY ORACLE is to explain the indifference away. While Ira has no obvious preferred outcome some external facts or latent mental states will always determine which doxastic attitude she will form. Maybe she secretly prefers even over odd numbers or is habituated to go for the first mentioned possibility when it comes to indiscernible outcomes. In other words, by evoking ASTRONOMY ORACLE as a counterexample to Evidential Uniqueness we simply stipulate that *Buridan's ass-like* situations are metaphysically possible. But in apparently indifferent situations agents are not paralysed but *de facto* always choose an ever so slightly preferred outcome. This is a historically influential response to Buridan's ass is most prominently defended by Leibniz:¹⁶

“There is never any indifference of equipoise. [...] There will therefore always be many things in the ass and outside the ass, although they may not be apparent to us, which will determine him to go one side or the other. (Leibniz, theodicy: §46-49)”

While Leibniz's treatment of Buridan's ass has some intuitive appeal, suggesting his reply as an argument against the permissive nature of genuinely indifferent SFB-cases is mistaken in at least two respects. First, it begs the question. Claiming that there is no indifference by saying that there is always a determined outcome does not prove that one outcome is more rational than the other (Ullmann & Morgenbesser 1977: 759). Maybe there is always one outcome since it is rational to just 'pick' one of the permitted doxastic states when being in an equilibrium. I will return to this view in the next section [§2.7].

More important for our investigation is the insight that our current concern is not the metaphysical possibility of Buridan's-ass-like equilibria, but the *epistemological possibility* of having a body of evidence that suggests that one is completely indifferent towards a set of outcomes. Ira only needs to reasonably think that there is no uniquely optimal outcome for her, even if she is not genuinely indifferent with respect to the outcome. This insight gives us the first ingredient we need to cook up a genuinely indifferent SFB-case.

Ingredient (1): Ira's evidence needs to support that she is completely indifferent with respect to p.

¹⁶ See also Rescher (1967).

In other words, ASTRONOMY ORACLE is not stipulating that Ira is *de facto* completely indifferent but that for all she knows she is completely indifferent concerning the possible outcomes. This should be rather unproblematic since impermissivists and permissivists alike usually accept that there can be indifferent bodies of evidence. They only disagree about the rational requirements that govern these cases.

2.6.2 Evidential Obligations

After we have discovered that the indifference we are looking for is in the evidential body possessed by the epistemic agent, rather than the actual preferences, we need to return to one of the initial objections put forward by impermissivists against the allegedly permissive nature of SFB-cases. Again, we may ask: given indifference in the body of evidence why isn't then the most rational option to suspend judgement? Maybe genuine indifference leaves ideal rational agents, well, indifferent!

This would suggest introducing some form of *equal reasons agnosticism* which states that if no unique doxastic attitude is determined by the evidential or pragmatic factors one should suspend judgement. While stipulating equal reasons agnosticism in evidentially indifferent cases arguably presupposes Evidential Uniqueness¹⁷, in genuinely indifferent cases equal reasons agnosticism seems to be more plausible and is *de facto* directly defended by various permissivists. As illustrated above, *interpersonal* permissivists think that there is only ever one rational doxastic attitude to transition into for one epistemic agent (at one moment in time) determined by the overall evidence and their additional non-evidential features [§2.5]. In other words, to reach the result that Ira should suspend judgement we do not need to refer to equal reasons agnosticism but to 'equal evidence and non-evidential feature agnosticism', which is far less controversial.

Nonetheless, I think we can modify our case further, to rule out suspension of judgement as an option. This can be done in multiple ways. One would be to simply defend some form of epistemic consequentialism that renders having true beliefs more valuable than suspending judgement (or give Ira evidence that makes such consequentialism plausible). A probably less controversial way is to stipulate some kind of non-epistemic reward. For example, the Genie may promise Ira a million dollars if she either φ 's or ψ 's to generate some preferences. Doing so seems to make φ ing or ψ ing practically superior to *not* φ ing or ψ ing. This gives us the second ingredient we need for our recipe of genuine indifference cases:

¹⁷ See Kopec (2015: 407-408).

Ingredient (2): Ira's evidence needs to support that she prefers having a true belief with respect to p .

2.6.3 Doxastic Voluntarism Revisited

After adding ingredients (1) and (2), we have effectively cooked up a body of evidence that models Ira as preferring to have a true belief towards p while being indifferent whether p is the case. However, there is a final ingredient we need to finish our recipe. Ira might still not believe that she *can* form one of the required doxastic attitudes which would prevent her from engaging in φ ing or ψ ing. Maybe her body of evidence suggests that she necessarily needs to have either decisive evidence or preferences which enable her to transition into one of the doxastic attitudes; i.e. she might believe that she is a doxastic involuntarist (concerning indifference cases).

However, this only works if we stipulate that Ira is *absolutely certain* that involuntarism is true. If she has some evidence that indirect- or direct doxastic voluntarism might be true she should still try to φ or ψ , since she would still ascribe *some* probability to her action changing the world in a desired way. If we accept for a moment that Ira could have a body of evidence that indicates that involuntarism *must* be true, we have put Ira into a dilemma in which rationality demands her to do something she cannot do which forces her to do something that is not fully rational.¹⁸ If Ira's actions do not influence her doxastic attitude formation, then she cannot change the state of the world and there is no evidence that allows her to predict her future behaviour which would render transitioning into the doxastic attitude rational. She is normatively bound to the default state she was in before considering the proposition [§2.5].

To overcome this obstacle, we need a third and final ingredient for our recipe:

Ingredient (3): Ira's evidence needs to support that φ ing or ψ ing increases the probability of inducing the respective doxastic attitude.

In sum, to transform ASTRONOMY ORACLE into a *genuinely indifferent* case, Ira's overall evidence needs to support that (1) she really is indifferent with respect to the outcomes, (2) she prefers

¹⁸ Note that framing this as a dilemma does not rely on the existence of a general norm that 'epistemic ought-implies-can'. While I won't take sides on the debate of whether there is a general 'epistemic ought-implies-can' norm, the idiosyncratic nature of SFB-cases makes it the case that there is at least a local norm that 'ought-implies-can'. For a characterisation of epistemic dilemmas and some references to recent work see [§1.1].

having one of the self-fulfilling beliefs, and that (3) if she tries to transition into one of the respective attitudes this increases the chance that she ends up having that attitude.

2.7 Picking Outcomes and Choosing Attitudes

Having identified all the necessary ingredients for a genuinely indifferent case enables us to investigate the practical and epistemic requirements that govern these situations. While the final ingredient established that Ira has evidence that she can influence her doxastic attitude formation towards the respective proposition, we still need to think about *how* this attitude formation might come about. To do so we need some background in practical decision-making under preferential indifference.

It has been argued that in practical equilibria, such as Buridan's ass, it is rational for epistemic agents to *pick* rather than to *choose* one of the equally preferred outcomes, while the former is selection without preference and the latter is selection via preference (Ullmann & Morgenbesser 1977; see also [§1.3]). One illustrative example discussed by Ullman and Morgenbesser is the apparent choice situation in a supermarket between two superficially indiscernible cans of tomato soup (1977: 761):

“To be sure, given the variety of products on display and given your preference you may choose to get a can of soup. You may, further, choose to get tomato rather than mushroom soup, and you may, if you are particular about such matters, choose to get Campbell's tomato soup rather than Heinz's. But we hold that usually you cannot, and as a matter of fact do not, choose the can you end up throwing into the carriage: you pick it. That is, if [...] you are still facing at least two cans neither of which is discernibly superior to the other(s), then you are in a picking situation, willy-nilly.”

While upon closer inspection there might be some relevant difference between the two cans of tomato soup, rationality does not demand that we spend useless amounts of cognitive resources figuring out which can actually is the preferred one. Instead, it is perfectly rational to just stop our investigation and pick rather than choose among seemingly equal outcomes.

This notion of picking is helpful in resolving genuinely indifferent SFB-cases. As Buridan's ass, Ira is in a situation where there are exactly two positive outcomes that have the very same expected utility and are therefore equally preferred. Given that practical rationality requires her to pick among two actions in these situations Ira is in a practically permissive case. Yet, Ira is not in an epistemically permissive case, because before initiating any of the permitted actions no

attitude is epistemically permitted. As I have argued above, the only way to deliberate the evidence in SFB-cases is to take the birds-eye view and model prospective behaviour, which is not possible here. For all that Ira knows both outcomes are equally preferred and, therefore, equally likely to come about. Before picking one of the actions, Ira's prediction about her own behaviour is indecisive and, therefore, she cannot ascribe probabilities to either outcome. Before doing something, she *ought* to suspend judgement. However, after picking one of the actions predictions about her future behaviour are possible again. That is, in picking either φ or ψ , Ira effectively generated new evidence, which shifts the evidential balance one way or another and determines a unique doxastic attitude.

But what exactly is picked in these situations? For sure the end that is picked is one of the desired outcomes, but this end is picked via picking a mean, the action to either φ or ψ , which again is picked via picking an intention to either φ or ψ . Let me illustrate this via an analogy to direct and indirect picking mechanisms. Ira could pick φ or ψ by indirectly picking some external mechanism which decides for her which doxastic attitude to transition into.¹⁹ She may flip a coin, φ if it lands heads and ψ if it lands tails. While this seems reasonable, there is an air of paradox in this indirect way of reaching the desired result. By picking an external randomiser, Ira picks a certain way of utilising this external randomiser. In doing so, the state of the world and with it the overall evidence possessed by Ira has already changed, she generated evidence that if the coin lands head she will φ and if the coin lands tails she will ψ . This power to generate evidence in SFB-cases ironically makes it the case that picking an external randomiser is superfluous. The same result could be achieved by directly picking one of the actions. What is true for directly picking one of the actions is *mutatis mutandis* true for picking an intention to φ or ψ . Once one of these intentions has been formed evidence has been generated which makes the respective action redundant since Ira as an ideal rational agent directly forms the evidentially required doxastic attitude.²⁰

This has surprising implications for our understanding of ideal rational agency: in genuine indifference cases, rational agents not only *can* and *ought* to generate evidence to tip the evidential balance but also *cannot help* but do so as part of their practical decision process. Therefore,

¹⁹ For a discussion of internal and external picking procedures see Rescher (1967: 169-170) and Ullmann and Morgenbesser (1977: 769-770).

²⁰ Something similar could be said about belief pills (White 2005). If I pick to take a pill that induces (or likely induces) a belief, I already have generated evidence that I will have the belief, which paradoxically makes taking the belief pill redundant. See also [S1.6] (footnote 31).

acknowledging that picking actions is not only possible, but also rationally required if there are equally preferred outcomes, and recognising that in picking actions we generate decisive evidence about what to believe, is acknowledging that genuinely indifferent SFB-cases are epistemically impermissible.

2.8 Conclusion

In this chapter, I have demonstrated that even the most idealised versions of SFB-cases cannot be harvested to generate counterexamples to Evidential Uniqueness. The presented resolution is novel in various respects.

First, it applies to both synchronic considerations as well as diachronic attitude formation processes. The allegedly permissive nature of belief formation in SFB-cases is generated by confusing practical and epistemic aspects of doxastic attitude formation. In making a model of how one ought to act in these cases, a rational agent can make predictions of what they will believe, which tells them into which doxastic attitude they ought to transition. In other words, one idiosyncratic feature of SFB-cases is that they transform otherwise irrelevant beliefs about non-evidential features of the given situation into evidence.

Second, I have illustrated that if we accept practical permissivism we can understand even genuinely indifferent SFB-cases as being epistemically impermissible. This also told us something about the voluntarist nature of these cases: if ideal rational agents can indirectly influence their doxastic attitude formation processes in SFB-cases they can and also *de facto* will directly influence them, when deliberating what they ought to do. This result is important, not only for those who want to defend an impermissible understanding of epistemic rationality but also to motivate my general treatment of allegedly permissive cases as epistemic standard conflicts [ch.1].

3

Epistemic Divergence:[†]

When Member-Level and Group-Level Attitudes Ought to Diverge

The previous two chapters were concerned with the kinds of conflicts we face as individuals navigating our epistemic environment. In particular, I have motivated the Conflict View of permissive cases [ch.1] and defended it against a potential objection from self-fulfilling belief cases [ch.2]. In this chapter, I will take the idea of epistemic conflicts and apply it to issues in social and collective epistemology. That is, I will demonstrate that noticing that epistemic agents face epistemic conflicts when they navigate different socio-epistemic environments and act as members of different groups has interesting implications for debates in collective epistemology. More precisely, it motivates an inflationary understanding of the epistemic status of collective attitudes, such as group knowledge or collective justification.

Note, that for the sake of narrative ease, I will presuppose epistemic permissivism in this chapter. Nonetheless, I am confident that the presented argument generalises to my novel treatment of allegedly permissive cases as epistemic conflicts [ch.1]. I will illustrate how a generalised version of the presented argument might look like in the conclusion to this chapter [§3.7].

3.1 Introduction

What is the relationship between the epistemic status of group-level attitudes and member-level attitudes? It is natural to assume that when we talk of ‘collective justification’, or ‘group knowledge’ we are merely making generalisations about the epistemic status of the attitudes of the group’s members. For example, we might think that a group’s belief that *p* is justified iff all (or a significant percentage) of the group’s members justifiedly believe that *p*. This is a version of epistemic summativism: the view that, roughly, to ascribe an epistemic status to a group is to

[†] This chapter is based on the following published work:

Graf, Simon (2023). Permissive Divergence. *Canadian Journal of Philosophy* 53 (3), 240-255.
<https://doi.org/10.1017/can.2024.4>.

indirectly ascribe it to a sufficient portion of its relevant members (where debate exists about what proportion is sufficient and which members are relevant).

Many find summativism intuitive, and it is certainly both theoretically and ontologically conservative. Consequently, some of the most promising theories in collective epistemology are explicitly summative (Goldman 2014) or at least retain summativist elements (Lackey 2016; 2021; Tuomela 2011).¹ Yet, as I will argue, all of those accounts face a serious problem: they are incompatible with *permissivism* about epistemic rationality. According to permissivism, for some bodies of evidence, there is more than one rational doxastic attitude epistemic agents are permitted to hold. Analogously, permissivism allows a single body of evidence possessed by a group and its members to rationalise different doxastic states for different members of the group and the group itself. Even if all members have the same evidence, other non-evidential factors, such as the members' cognitive goals (Kelly 2013), or belief-formation mechanisms they take to be truth-conducive (Schoenfield 2014b), could rationally require them to form different doxastic attitudes. In such cases, the doxastic attitude that is rational for the group cannot be determined summatively as a function of the attitudes rationalised by its members, since the group could have different epistemic standards rationally requiring a divergent doxastic attitude. This might be the case if a group's judgements, but not those of its members, play institutional functions, have diverging normative requirements, or are governed by a specific form of charter. Call this class of cases *permissive divergence* cases.

These permissivism-based cases can be used to support so-called divergence arguments, a type of argument that purports to show that there can be a conflict between member-level and group-level attitudes (Lackey 2016, 2021: 56-58). Divergence arguments have been used to motivate and support non-summativist accounts of group belief (Gilbert 1987, 1989), group assertion (Lackey 2021: 158-163; Ludwig 2014), collective virtue ascriptions (Fricker 2010), collective justification (Schmitt 1994; Mathiensen 2011), or different kinds of collective knowledge, such as knowledge-how and knowledge-why (Bird 2010, 2014; Hutchins 1995a; 1995b; Habgood-Coote 2019). While the number of cases considered has accumulated, divergence arguments often rely on individual case judgements of differing intuitive appeal. Permissive divergence cases, on the other hand, are directly built on a position that, whilst controversial, has a great deal of antecedent support independent of the collective epistemology

¹ I will give a more extensive overview of summative approaches to collective epistemology in [S3.2].

debate. Using epistemic permissivism as a starting point, therefore, gives us a general recipe for generating epistemic divergence cases, which do not rely on intuitive case-by-case judgements but on an explicit theoretical commitment.

With the general idea in place, the remainder of my discussion focuses on working out the specifics. First, we need to clarify the general structure of divergence arguments and the summativist assumptions they are rejecting [§3.2]. Afterwards, I will introduce a very general version of epistemic permissivism and highlight some of the normative implications we can derive from it [§3.3]. This, together with the assumption that groups and their members may have different epistemic standards, enables me to provide a template for constructing epistemic divergence cases. These cases can then be used to undermine summative understandings of collective rationality, group justification and group knowledge. This will also render summative approaches to group belief implausible [§3.4]. Afterwards, I will address some possible objections and draw some general lessons about the nature of collective epistemic attitudes from my responses to them [§3.5]. In particular, I argue that any account of the epistemic status of group attitudes needs to account for the fact that groups have member-independent epistemic requirements. In the final section [§3.6], I will give a quick overview of other types of divergence arguments found in the literature and show that permissive divergence is immune to various shortcomings of these arguments.

3.2. Summativism and Divergence

The term ‘summativism’ subsumes a diverse set of theories about group-level attitudes that rest on the assumption that ascribing an attitude to a group is to indirectly ascribe it to (some of) its members. To figure out, for example, whether a group believes that p summativists think that we need to sum up or aggregate the beliefs of the respective members. Similar to belief-summativism, there are also summative views about the epistemic status of doxastic attitudes such as justification or knowledge, which are the main target of the argument presented in [§3.4].

Summativist views about attitudes are a subclass of deflationary views within social philosophy that think of group attitudes as being dependent on, or reducible to, rather than ‘over and above’ their member attitudes.² While summativism makes a very specific reductionist

² While deflationism and summativism are often used interchangeably in the literature, Habgood-Coote (among others) has pointed out the importance of this distinction (2019: 932).

claim, namely that the group attitudes are reducible to the respective member attitudes, there are also deflationary but non-summative views.³ For example, group-level beliefs may be reducible to member-level behaviour more broadly, and therefore depend on other member-level attitudes such as member-level acceptance.

Following Lackey (2021: 20-30), we can differentiate between conservative and liberal versions of summativism. The conservative summativist claims that a group *G* has the attitude *D* iff all the members of *G* have the attitude *D* (Lackey 2021: 21). On the other hand, liberal summativists only require some relevant subset of *G* to have the attitude *D*. The subset sufficient to ascribe the respective attitude to the group may be defined in general, e.g., the majority as in majoritarian summativism, or it might differ based on the type and organisational structure of the group. Sometimes the summativist needs every member to have a specific attitude, sometimes it's only the majority, a significant part, or some particularly influential or important member. These liberal, more flexible versions of summativism are especially promising since they can make sense of the fact that groups can be hierarchically structured in manifold ways. Groups may have operative and non-operative members, whereas only the attitudes of the former constitute the formation of the group's attitudes. Or, more precisely, operative members "are those who are responsible for the group belief having the content that it does which, in turn, is determined by the rules and regulations of the group in question" (2021: 27).⁴ For example, the fact that the president or the CEO of a group has the attitude *D* towards *p* could in some cases constitute or ground the fact that the group has attitude *D* towards *p*.

In light of these discrepancies between different versions of summativism it is useful to define a minimal summativist commitment (Lackey 2021: 21; Faria 2021: 85):

Minimal Summativist Commitment <MSC>: A group *G* has the doxastic attitude *D* towards *p* only if at least one individual *m* is both a member of *G* and has the doxastic attitude *D* towards *p*.

<MSC> is accepted by all summative (Quinton 1976; Cohen 1989; Faria 2021), as well as partially summative (Lackey 2021: 48-53) understandings of group belief and other doxastic

³ Nonetheless, some accounts which are understood to be non-deflationary defend a supervenience relation between member-level and group-level attitudes (List and Pettit 2011: 65-66). For further discussion see also Bird (2010) or Kallestrup (2022).

⁴ For the concept of an operative member see also Tuomela (2004: 113).

attitudes.⁵ Similar to <MSC>, we can define a minimal commitment for *epistemic summativism* as follows:

Minimal Epistemic Summativist Commitment <MESc>: A group's doxastic attitude D towards p has the epistemic status S only if there is at least one member of G who has the doxastic attitude D towards p and for whom D has the epistemic status S.

<MESc> is central to fully summative views of collective justification (Goldman 2014), group knowledge (Mokyr 2002; Tuomela 2011), as well as various other hybrid approaches to collective justification, which retain summativist elements (Lackey 2016, 2021).⁶

Consequently, any argument that renders <MSC> or <MESc> implausible or incompatible with other well-established views has a large impact on our theorising in social philosophy. Divergence arguments do exactly that: they illustrate possible cases in which the respective minimal summativist commitments are violated. The general structure of *psychological divergence* arguments against summative views of doxastic attitudes can be summarised as follows (Faria 2021: 85):

- (1) If <MSC>, then necessarily, G has the doxastic attitude D towards p only if at least one individual m is both a member of G and has the doxastic attitude D towards p.
- (2) It is possible that G has the doxastic attitude D towards p while no member m_1 - m_n of G has the doxastic attitude D towards p.
- (C) <MSC> is false.

The main job of the proponent of psychological divergence arguments is to give examples to support (2) and, therefore, demonstrate the possibility of divergence of group doxastic attitudes and member doxastic attitudes. In other words, psychological divergence arguments are usually accompanied by illustrative divergence cases that aim to serve as counterexamples to <MSC>.

⁵ Strictly Cohen's (1989) account, while implementing summativist elements rejects the existence of genuine group beliefs. Cohen argues that groups, while being capable of accepting certain propositions, can't have genuine beliefs since they lack certain mental capacities which he takes to be necessary for beliefs (Gilbert & Pilchman 2014: 190-197).

⁶ While Lackey thinks of her *Group Epistemic Agent Account* as being neither deflationary nor inflationary, for Lackey justified group belief still requires that "a significant percentage of the operative members of G [...] justifiedly believe that p" (2016: 381). For a discussion on the summative nature of Tuomela's (2011) account see Lackey (2021: 72). Another partly summative account, which is not directly affected by the presented argument is defended by Silva (2019). I will discuss Silva's account in more detail in [ch.9].

Here, the distinction between deflationism and summativism becomes important. By supporting (2) divergence arguments demonstrate that group-level attitudes are not solely dependent on the respective member attitudes, but they do *not* show that group-level attitudes are independent from other things happening on the member level. Consequently, attempts to support (2) via divergence arguments have not only been used to support inflationary but also deflationary non-summative accounts within social epistemology.⁷

Accordingly, we can capture the structure of *epistemic divergence* arguments as follows:

- (1') If <MESC>, then necessarily, G's doxastic attitude D towards p has the epistemic status S only if at least one member m of G has the doxastic attitude D towards p and m's attitude has the epistemic status S.
- (2') It is possible that G has the doxastic attitude D towards p with the epistemic status S while no member m_1 - m_n of G has the doxastic attitude D towards p with the epistemic status S.
- (C') <MESC> is false.

Epistemic divergence arguments are built on examples that support (2'), and these examples can work in two different ways. First, they can show that a group can have an attitude D with the epistemic status S while no member has the relevant attitude D (this would refute both <MSC> and <MESC>). Second, the divergence can arise because the epistemic status of the relevant member-level attitude diverges from the epistemic status of the group-level attitude (this would refute only <MESC>).

In what follows, I will introduce a new type of epistemic divergence argument of the second kind, called *permissive divergence*. This argument reveals that <MESC>, and thereby summativism about the epistemic status of doxastic attitudes, such as justification or knowledge, is incompatible with epistemic permissivism. While the main target of permissive divergence is summativism about epistemic states, I will also argue that retaining belief-summativism in light of my findings comes at a high theoretical cost [S3.4]. In contrast to other epistemic divergence arguments, permissive divergence will turn out, not only to be more appealing, by being built on a widely defended epistemological thesis [S3.3], but also to be less vulnerable to potential

⁷ For example, Habgood-Coote (2019) uses a variety of divergence arguments to motivate his deflationary but non-summative account of collective know-how.

objections [§3.5] [§3.6]. However, before I can move on to the proposed argument some preliminaries on epistemic permissivism are needed [§3.3].

3.3 Types of Permissivism

As pointed out in various places in the previous chapters [ch.1] - [ch.2], there are two opposing views about the relationship between evidence and doxastic attitudes central to much recent work in epistemology. While impermissivists defend Evidential Uniqueness, the thesis that there is just one rationally permissible doxastic attitude given one particular body of evidence, permissivists hold the view that one body of evidence might rationalise multiple doxastic attitudes.⁸ Let me take another dive into the subject matter to precisely work out the assumptions underlying Standard Permissivism, the view that motivated the Conflict View of permissive cases [ch.1]. This is important because it will help us to work out the precise commitments underlying the epistemic divergence argument presented below [§3.4].

One way to understand the debate is as, both permissivists, as well as impermissivists, making statements about the doxastic attitude(s) agents rationally *ought* to take given one body of evidence.⁹ While impermissivists and permissivists disagree about the precision of rational requirements, both agree that rationality requires us to be in at most one of the permissible states. Here is one way to formulate permissivism:

***Interpersonal Permissivism* <InterP>**: There are cases in which it is rationally permissible for one agent A to have the doxastic attitude D towards p, and for another agent A* to have a different doxastic attitude D* towards p instead, given the same body of evidence E.

This particular form of permissivism is called *interpersonal* because it relies on a specific reading of Uniqueness. As originally formulated, the quantifier within Uniqueness is ambiguous, it leaves open “whether different attitudes are permitted (or required) for *different* individuals on that same body of evidence” (Kopeck and Titlebaum 2016: 191; *emphasis in the original*). Based on two different readings of Uniqueness we can distinguish two kinds of permissivism. While

⁸ For an overview of the rich and vast literature on permissivism see Kopeck and Titlebaum (2016), or Jackson and Turnbull (2023).

⁹ See, e.g., Greco and Hedden (2016: 392-393).

intrapersonal permissivism allows for situations in which more than one doxastic state is fully rational for the very same epistemic agent, *interpersonal* permissivism is “only permissive across individuals” (Kelly 2013: 304).

Since this distinction has first been introduced, <InterP> has received a lot of attention in the literature. This is largely because <InterP> is less theoretically demanding and therefore less vulnerable to worries raised by impermissivists than its *intrapersonal* competitor. While any systematic comparison of different types of permissivism would exceed the scope of this chapter, there is one feature of <InterP> that helps us to construct divergence arguments based on permissivism: <InterP> can be understood as *explaining* the permissibility of multiple doxastic attitudes by appeal to the different *epistemic standards* had by different epistemic agents. In other words, <InterP> understands evidential support as a three-place relation between the evidence E, some epistemic standards S, and a prescribed doxastic attitude D. When criticising White (2005), whose defence of Uniqueness is explicitly built on a two-place relational understanding of evidential support, Decker points out that “evidence supports a proposition, not *simpliciter*, but rather relative to an interpretation” (2012: 780; *emphasis in the original*). Decker, therefore, concludes that a more reasonable way to understand Uniqueness would be, that “given any admissible total interpretation of one’s total evidence, there is a unique rational doxastic attitude that one can take to any proposition” (2012: 782). In that sense, most *interpersonal* permissivists defend *intrapersonal* Uniqueness:¹⁰

Intrapersonal Uniqueness <IntraU>: For any body of evidence E and any proposition p, there is a unique rationally permissible doxastic attitude D towards p for any epistemic agent A and their epistemic standards S, possessing that body of evidence E.¹¹

<IntraU> is compatible with <InterP> and with it, the possibility of situations in which two epistemic agents possess the same evidence but differ in their epistemic standards and, therefore, are rationally permitted to have diverging doxastic attitudes. That is, <InterP> rests on two assumptions:

¹⁰ Both Decker (2012) and Kelly (2013) independently point this out in response to White (2005, 2013) who understands evidential support to be a two-place relation.

¹¹ Note that this assumes that one epistemic agent cannot have multiple epistemic standards that recommend conflicting doxastic attitudes towards the same proposition simultaneously.

- (i) *epistemic support* is a 3-place relation between the evidence E, epistemic standards S and doxastic attitudes D.¹²
- (ii) epistemic standards are *variable*.

Some impermissivists, such as White (2005, 2013) deny (i). However, there are versions of impermissivism that are compatible with accepting (i) but rejecting (ii). Take, for instance, the objective Bayesian understanding of Uniqueness defended by Hedden (2015). For Hedden, the doxastic attitudes of epistemic agents should be the result of taking the “uniquely rational prior probability function and conditionalising it on your total evidence” (2015: 470). This could be read as understanding the evidential support relation to be a 3-place relation between the evidence, the unique rational prior and the doxastic attitude, whereby it is denied that there are multiple rationally permitted priors as claimed by subjective Bayesians. Hence, to end up with permissivism we need both assumptions.

These observations can be used to motivate a very general version of permissivism that results from combining (i) & (ii). In [ch.1] I called this version of permissivism Standard Permissivism:

Standard Permissivism: For any body of evidence E and any rationally permissible epistemic standards S_1 - S_n there is a unique rationally permissible doxastic attitude D towards any proposition p. There are cases in which one body of evidence E rationalises different doxastic attitudes D_1 - D_n relative to different rationally permissible epistemic standards S_1 - S_n respectively.

We can understand Standard Permissivism as subscribing to a very weak form of Uniqueness, which I called Epistemic Uniqueness in [§1.2]:

Epistemic Uniqueness: For any epistemic situation S (any body of evidence and any combination of epistemic standards) and any proposition p there is *at most one* doxastic attitude that any epistemic agent A could rationally take.

¹² Note that this is a metaphysical assumption about the nature of epistemic support and not solely about the number of syntactic places of the support relation. After all, an n-place relation can have many different grounding conditions that the syntax of an expression does not take into account. I want to thank an anonymous referee for pointing that out.

Given Standard Permissivism, for some bodies of evidence there are multiple permitted doxastic attitudes, while any epistemic situation (any particular body of evidence and set of epistemic standards) determines a uniquely rational doxastic attitude. Furthermore, Standard Permissivism does not make any assumptions about the nature of the variability: standards may vary across agents, time, contexts, or in any other evidence-independent way. As such, <InterP> is a version of Standard Permissivism, which restricts the variability of epistemic standards to the interpersonal dimension.

To illustrate how widespread Standard Permissivism is, I will spend the rest of this section providing examples of epistemological frameworks that imply Standard Permissivism.

Let me start by elucidating further how various versions of <InterP> rely both on (i) and (ii). As illustrated above, Decker (2012) thinks that we can only make sense of a certain body of evidence, given a certain (permissible) interpretation. While not all interpretations are born equal, for Decker, there is often a wide range of equally coherent, and similarly fruitful, simple, or elegant interpretations available—a fact that is well-known to philosophers of science.¹³ Similarly, Kelly (2013) contemplates a version of interpersonal permissivism built on the rivalling Jamesian cognitive goals of ‘attaining truth’ on the one hand and ‘avoiding error’ on the other. He thinks that “subtly different ways of responding to the same body of evidence seem equally reasonable, given corresponding differences in the weights that we give to our shared cognitive goals” (Kelly 2013: 302). Other proposals for defending interpersonal permissivism, which are compatible with (i) and (ii), are found in Titlebaum (2010), who thinks of evidential support as being relative to a preferred set of predicates, or Meacham (2013), who points out that subjective Bayesianism implies that evidential support is relative to an agent’s prior probability functions. Schoenfield (2014) is defending an intrapersonal permissivism that arises from a divergence in epistemic standards that are taken to be truth-conducive by the epistemic agent.

Apart from these interpersonal understandings, other theoretical frameworks that can be utilised to generate versions of Standard Permissivism. Take, for example, inductive risk. It is a long-known conundrum in the philosophy of science that any belief or disbelief of an empirical hypothesis *H* involves an inductive leap that risks accepting “*H* while *H* is in fact false, and, conversely, [...] rejecting *H* when *H* is in fact true” (Hempel 1965: 91–92). Since there is no

¹³ For *loci classici* on the subject matter see Kuhn (1977: ch.13) or Longino (1996). For a comparison between arguments from the underdetermination of scientific theories and permissivism see Jackson and Turnbull (forthcoming).

general trade-off rule for inductive risk, we may argue that there are different rationally permissible ways to handle the evidence given different inductive standards (Wilholt 2009: 94). Similar things can also be said about stances (Van Fraassen 2002), cognitive resources (Jackson 2021), or stakes (Stanley 2005); all of which not only, (i) sometimes rationalise different doxastic attitudes with respect to a single body of evidence, but also, (ii) are variable in the above-considered way.

While this illustrative catalogue of epistemological frameworks that accept both (i) and (ii), isn't intended to be exhaustive, it gives a good impression of the wide range of epistemic positions that could and have been used to support Standard Permissivism. This is an important observation that gives the upcoming argument against epistemic summativism [§3.4] additional theoretical appeal since it is directly built on Standard Permissivism.

3.4 Permissive Divergence: The Argument

What does permissivism tell us about the relation between the epistemic status of member-level and group-level attitudes? The two assumptions identified above that Standard Permissivism is based on, namely that, (i) evidential support is a 3-place relation, and that (ii) epistemic standards are rationally *variable*, establish that there could be cases in which the group-level and member-level attitudes rationally *ought* to diverge. Given one body of evidence but diverging epistemic standards, different doxastic attitudes might be rationalised for the group and its members. This provides us with a general template for constructing epistemic divergence cases:

Permissive Divergence Template <PDTemplate>: There are cases in which a group G is rationally required to have the doxastic attitude D towards p in relation to a body of evidence E and its permissible epistemic standards S . While all of G 's members m_1 - m_n possess the same body of evidence E , they have different permissible epistemic standards S_1 - S_i which recommend different doxastic attitudes D_1 - D_n towards p .

<PDTemplate> is reliant on understanding evidential support as a 3-place relation. It acknowledges that evidence *simpliciter* cannot be understood as a signpost that designates a particular doxastic attitude without any appeal to some epistemic standards (Decker 2012; Kelly 2013). Accordingly, if we think of groups as being able to rationally acquire and possess doxastic attitudes, we need to think that they possess, or are governed by some epistemic standard(s)

which make(s) sense of the evidence possessed by the group. Otherwise, no doxastic attitude can ever be rationalised for G.

In real-life scenarios, the group's epistemic standards S will often be represented by an idiosyncratic *modus operandi* of the group. Groups such as juries or investigative panels have charters or rules of conduct. Scientific groups often have explicit or implicit norms and are shaped by their inner-disciplinary and wider scientific culture, defining proper and improper inquiry, methodology and research questions, within a paradigmatic framework. While some of these *modi operandi* govern non-epistemic aspects of the group's functioning, others will influence the group's belief formation processes in a way that is analogous to how epistemic standards govern the belief formation processes of individuals. While what exactly those processes are will depend on the exact nature of group beliefs, there are many plausible candidates compatible with a wide range of understandings of group-level attitudes. For example, in scientific research groups, there are often norms influencing how much weight is given to research articles published in journals of varying qualities or different test results from specific types of experiments. Other examples include norms governing representative samples, sample sizes, or statistical thresholds for labelling something a discovery, such as the famous '5 sigma rule' in particle physics.¹⁴

These observations, together with the provided template <PDTemplate> give us an instruction on how we can cook up scenarios in which member-level and group-level attitudes rationally *ought* to diverge. The only thing left to argue for is that there *actually* are such cases. This should be an easy task given the long and diverse list of possible epistemic frameworks that could be used to construct such cases [§3.3]. Take, for instance, this illustrative example:

DIVERGING STANDARDS: A group of scientists form a scientific collaboration G to investigate a particular empirical question regarding the anthropogenic impact on climate change. After several years of careful empirical studies and processing of enormous amounts of data, G possesses a body of evidence E. During one of their biweekly meetings, the relevant body of evidence E is carefully displayed and disclosed among the members and based on that all members m_1 - m_n form the belief that p "cutting greenhouse gas emissions by 50% by 2030 would keep the global mean temperature below 1.5 °C above pre-industrial levels".¹⁵ Nonetheless, none of the members of the scientific

¹⁴ See, for example, Lamb (2012).

¹⁵ This statement is found in the *Intergovernmental Panel on Climate Change (IPCC) special report on the impacts of global warming*.

collaboration publicly endorse that p , since they know that their epistemic standards S_1 - S_n which rationally require them to believe that p are not shared by G . G 's standards are set up in a way that G values avoiding errors more than attaining true beliefs about p . That is, G has a high but not unreasonable epistemic threshold for believing empirical propositions about prospective climate developments supported by its epistemic standards S^* which rationally requires G to suspend judgement about p .

Consider, for example, Ambitious Amy who is one of the theoretical physicists of the scientific collaboration G . While Ambitious Amy desires to avoid errors her primary epistemic goal is to attain true beliefs. As such, she has a low but not unreasonable liberal epistemic threshold for believing empirical propositions about prospective climate developments. This allows her to mistrust the scientific methodology underlying the group's suspension of judgement about p .¹⁶ Simultaneously, however, Amy understands that due to the inherently complex nature of the research, her conclusion that p would be regarded as being overzealous given G 's epistemic standards S^* (in this case their epistemic goal of avoiding error). In other words, Amy is perfectly capable of processing the data and inferring certain results, in accordance with S^* . However, while she believes that the collaboration would not be justified to believe that p , given S^* , she is simply not convinced that G is right in being so cautious with respect to p . Given that being epistemically ambitious in Amy's way is an *epistemically* permissible endeavour, she is rationally permitted to believe that p based on the body of evidence E possessed by the group. Since equally all other members of G could be in the same situation, this directly refutes <MESC>.

As constructed, DIVERGING STANDARDS directly violates any summative understanding of rational group belief: while G rationally suspends judgement about p all of G 's members rationally believe that p . Equally, any summative understanding of group justification is rendered implausible. The fact that the doxastic attitudes of both G and all members m_1 - m_n are based on E , as well as S_1 - S_n and S^* respectively, ensures that these attitudes are rational. Standard Permissivism, as a thesis about how doxastic attitudes are rationalised by a body of evidence E

¹⁶ One way to think about these cases is as cases in which pragmatic considerations of the group and the group members influence which of the permissible epistemic standards they subscribe to. That is, pragmatic factors, such as public pressure or the potentially disastrous consequences of getting things wrong might influence how the group is set up and, therefore, explain its risk aversiveness. In so doing, we treat questions about the way agents ended up with a particular (permissible) cognitive structure as independent from questions of whether they form beliefs in an epistemically permissive way. For a related discussion on the relationship between pragmatic encroachment and epistemic permissivism see Quanbeck and Worsnip (forthcoming). For a recent paper on pragmatic encroachment and collective justification see (Biebel 2023). I want to thank an anonymous referee from the Canadian Journal of Philosophy for comments on this issue.

makes restrictions on propositional justification, the kind of justification that arises from the mere fact of having sufficient evidence for a proposition p (Matheson 2011: 360-361). Given Standard Permissivism, multiple doxastic attitudes towards p may be propositionally justified for different agents given one body of evidence E , but at most one attitude is propositionally justified for one agent having a distinct set of epistemic standards S . That is, suspension of judgement on p is propositionally justified for G independent of the members' doxastic attitudes and their justificatory status. Furthermore, as described in DIVERGING STANDARDS, m_1 - m_n properly base their beliefs that p on E , while G also properly bases its suspension of judgement on E , which makes their attitudes not only propositionally but also doxastically justified (Turri 2010b: 312-314).

Given the traditional analysis of knowledge as (non-gettierised) justified true belief, this similarly proves summative understandings of group knowledge to be wrong. If both the group G and all members m_1 - m_n can have diverging but nonetheless justified beliefs, at most one, either G or m_1 - m_n , could have a true belief. In other words, by stipulating either that p is true or that p is false, we can generate a divergence case for knowledge.¹⁷

Interestingly, this also has consequences for how we think about group beliefs, despite permissive divergence being about the epistemic status of doxastic attitudes rather than doxastic attitudes per se. When confronted with permissive divergence cases the belief-summativist is in a dilemma: if they insist that the group cannot have a doxastic attitude that diverges from all its member's attitudes, they are forced to accept that there are cases in which the epistemic status of the member-level beliefs and group-level beliefs diverge. So presented with DIVERGING STANDARDS the summativist needs to say that the group *irrationally* believes that p (since G 's standards S^* require suspending judgement about p) in virtue of all members m_1 - m_n *rationally* believing that p based on E and S . For the summativist, the only way to ensure that both the group and the members are rational is to ensure that the group standards and individual standards align (which leads to all kinds of problems which I will discuss in the next section [§3.5]). The non-summativist about belief, by contrast, by allowing diverging belief states as well as diverging justificatory statuses, makes room for both the group and its members to hold rational beliefs in DIVERGING STANDARDS. Thus the best way to understand cases such as

¹⁷ Another, slightly more controversial, but more interesting way to generate divergence cases for knowledge is to treat the cautiousness in not forming a belief about p as a sign that G has a different stakes sensitivity than its members with respect to p . See, for example, Stanley (2005) or Simion (2021).

DIVERGING STANDARDS is as cases in which both the doxastic attitudes and the epistemic standards of the members and the group diverge.

In sum, Standard Permissivism, along with the observation that groups and their members may have divergent epistemic standards, gives us a recipe for crafting counterexamples to <MESC> and consequently summativism in general. This recipe is encapsulated in the template <PDTemplate> which inspired the concrete example DIVERGING STANDARDS. While this illustrative example enabled us to demonstrate a direct conflict between epistemic permissivism and summativism about the epistemic status of doxastic attitudes, such as justification or knowledge, it also rendered summative understandings of group belief implausible.

This simple method makes permissive divergence theoretically appealing since it is directly built on (a very general version of) permissivism, a thesis that has a great deal of antecedent support independent of the collective epistemology debate. In the remainder of the chapter, I will demonstrate that in contrast to other divergence arguments found in the literature, permissive divergence is not only theoretically more appealing, but also *less vulnerable* to potential objections. I will do so by first introducing and discussing potential summativist defence strategies [§3.5]. Afterwards, I will compare permissive divergence to other divergence arguments [§3.6] and illustrate that it is immune to multiple concerns raised against these competitors.

3.5 Permissive Divergence: Discussion

The most straightforward objection to the above-presented argument is to deny that groups can have epistemic standards in the same sense as individuals do. Based on the discussion above [§3.4], I can resist such kind of scepticism in two ways. First, if we think that epistemic support is a three-place relation (as implied by Standard Permissivism) we need to think that groups have epistemic standards which tell us which attitudes are justified given the group's evidence. Otherwise, we cannot make sense of the epistemic status of group attitudes at all. That is, any scepticism against the idea of group-level epistemic standards would lead to general scepticism with respect to collective epistemology. Second, if we take epistemic standards to be mechanisms that have a systematic influence on the formation of doxastic attitudes, given a body of evidence E, there are numerous plausible examples of group-level epistemic standards. In particular, I have mentioned various scientific standards governing the weighing of different types of evidence, the ways experiments are designed and conducted, and how scientific findings are reported. Since these mechanisms have a direct influence on the overall behaviour of the group, given minimal

assumptions about the nature of group beliefs, I think these mechanisms are best understood as group-level epistemic standards.¹⁸

Another way to try to resist permissive divergence arguments is to challenge the assumption that epistemic standards can diverge among group members and the group itself. Let's call the view that group-level standards cannot diverge from member-level standards, *standard summativism*. As with other types of summativism, we can define a minimal commitment as follows:

Minimal Standard Summativism <MSS>: A group G has an epistemic standard S only if at least one member of G has the epistemic standard S.

Unfortunately, <MSS> does not protect the summativist from permissive divergence arguments. First, even if <MSS> is true, to avoid permissive divergence cases, the relevant epistemic standards would need to be aligned with the group's doxastic attitudes in the right way. Otherwise, we can still have situations in which G has the doxastic attitude D with the epistemic status E while no member m_1 - m_n of G has an attitude D with the epistemic status E. This happens precisely in cases in which G's attitude D is rationalised by standard S, which G has due to m_1 - m_n having S, while m_1 - m_n do not have the doxastic attitude D (while any member that has D doesn't have the relevant epistemic standard).¹⁹

Second, the fact that groups are composed of independent epistemic agents, each with their own independent epistemic life makes standard summativism very implausible. So for <MSS> to be viable the summativist not only needs to deny that the norms identified above [§3.4] are group-level epistemic standards, but also that there is some mechanism that ensures that member-level and group-level epistemic standards do not diverge.

In so doing, standard summativists might claim that epistemic agents can perhaps endorse one set of beliefs privately based on S but need to temporarily adopt diverging epistemic standards S* (together with the relevant doxastic attitude) when acting as a member of G. For example, in DIVERGING STANDARDS Ambitious Amy adopts S* when she is in 'scientist mode' but exchanges S* with S when she is at home in 'ambitious mode'. Assuming that, we then treat

¹⁸ I want to thank an anonymous referee from the Canadian Journal of Philosophy for helpful comments and suggestions on these issues.

¹⁹ Note that these kinds of cases are possible given most versions of belief summativism.

Amy_[scientist] and Amy_[ambitious] as distinct epistemic agents, who can permissibly believe different things.

The problem with this framing of the case, however, is that Amy seems to be able to voluntarily switch her epistemic standards from *S* to *S**. This would transform Amy's doxastic attitudes in a way that they provide an unstable basis for action and practical deliberation. This concern is discussed as the so-called arbitrariness objection in the literature on epistemic permissivism (White 2005, 2013; Kelly 2013). While I discussed this worry extensively in [§1.6], we can summarise the basic idea behind this worry by referring to White (2005: 455):

“The arguments [...] suggest that a permissive account of rationality introduces a kind of arbitrariness to our beliefs that can infect both practical and theoretical deliberations. [...] If my current beliefs are not rationally obligatory for me, why should I take propositions that I actually believe as a basis for action and reasoning, rather than some others that I don't believe, but would be rational in believing? Why should my beliefs be privileged in my practical and theoretical deliberations, over equally rational alternative beliefs?”

Accordingly, the most prominent versions of permissivism defend themselves against the arbitrariness worry by denying that voluntary switching of epistemic standards is allowed. (The Conflict View allowed us to reach a similar result based on the (weak) bindingness of conflict resolutions [§1.6]). Take, for example, Schoenfield's (2014) interpersonal permissivism. While for Schoenfield there are potentially many different sets of permissible standards rationalising different attitudes regarding the given evidence *E*, epistemic agents cannot switch them voluntarily. From their individual perspective, rational agents must be *immodest*, that is, believe that their way of forming doxastic attitudes is truth-conducive, which, in turn rationally demands them to hold onto their standards (Schoenfield 2014b: 201-202).²⁰ In other words, treating Amy_[scientist] and Amy_[ambitious] as different epistemic agents conflicts with <InterP> (as well as most *intrapersonal* versions of Standard Permissivism), as traditionally conceived.²¹

²⁰ As pointed out in [§3.3], Schoenfield understands epistemic standards in a narrower sense, as the “function from bodies of evidence to doxastic states which the agent takes to be truth conducive” (2014: 199). For more discussions on immodesty see also Lewis (1971) or Schoenfield (2023: 279-285).

²¹ Likewise, on the Conflict View presented in [ch.1], conflict resolutions can be binding to various extents. I will reflect on how the presented epistemic divergence argument works out in the context of the Conflict View in the conclusion to this chapter [§3.7].

However, regardless of whether we think that rational agents ought to be *immodest* or not, or whether conflict resolutions are binding, the burden to prove that members switch their epistemic standards in permissive divergence cases is, after all, on the summativist. In other words, the summativist not only needs to show that Amy and all other scientists *can* and *ought* to switch back and forth between standards (from S_1 - S_n to S^* and vice versa), but also that this always and *inevitably* happens. Otherwise, we would still have permissive divergence cases.

One way to do this would be to insist that when acting as a member of a group, individuals necessarily form attitudes *qua* group member and *not qua* (independent) individual.²² Hence, the minimal epistemic summativist commitment could be updated accordingly:

Minimal Epistemic Summativist Commitment* <MESCS*>: A group's attitude D towards p has the epistemic status E only if there is at least one member of G for whom D has epistemic status E *qua* group member.

While <MESCS*> illustrates a way in which epistemic agents could and do *voluntarily* but *not arbitrarily* switch between different sets of epistemic standards, it is already a significant deviation from the initial summativist position. First, it admits that groups have an epistemic life of their own that forces certain attitudes and 'modes of thinking' onto their members. Second, even if we think that this voluntary switching of standards is unproblematic, it is far from clear that the attitudes adopted *qua* group member are genuine doxastic attitudes. As a group member, you can be perfectly aware of the epistemic standards of the group and how you ought to act according to them without temporarily forming a diverging doxastic attitude towards p. In other words, none of the scientists in DIVERGING STANDARDS needs to actually suspend judgement towards p to display the described behaviour and to understand that suspension is the rationally required attitude relative to G's standards. (The scientists can even be immodest in not acknowledging that G's and their own standards are rationally *on a par*).

This becomes especially clear when we think of instances in which agents are members of multiple groups simultaneously. If m is a member of two groups with different epistemic standards (but the same evidence), one which mandates believing that p, and another which mandates suspending on p, m would be committed to suspending judgement towards p and believing that p at the same time. The only way to make sense of these cases, without requiring

²² I am grateful to an anonymous referee for pressing a version of this objection.

m to have an incoherent set of doxastic attitudes, is then, to understand qua-group member attitudes *not* as genuine doxastic attitudes.

In sum, the best explanation of permissive divergence cases is that there are divergent standards in place which rationalise the group's attitudes. This demonstrates that minimal standard summativism <MSS> is false. While it is an open question how group-level epistemic standards relate to the member's behaviour and the member's qua attitudes (assuming that there are such qua attitudes), there is no reason to think that individuals form novel doxastic attitudes when acting as members of a group. But for (epistemic) summativism to be true the attitudes must be genuinely doxastic.

These considerations provide us with some general lessons about the epistemic status of collective attitudes, as well as a recipe for how to modify existing views to account for these insights. First, every understanding of the epistemic status of group-level attitudes needs some specification of the relevant agent *for whom* the evidence provides sufficient reason for belief. Here the presented argument shows that we cannot take the relevant agent to be (some) member(s) of the group –rather, the relevant agent must be the group itself. Second, we need to tell a substantial non-summative story about what it means that something is rational for the group to believe; i.e. how the group's epistemic standards come about.

3.6 Permissive Divergence: Comparison

Having demonstrated the resilience of permissive divergence arguments against potential objections enables me to systematically compare them to other epistemic divergence arguments found in the literature. One major virtue of permissive divergence is that it is immune to multiple concerns raised against other divergence arguments. I will reconstruct some of these worries here and demonstrate how and why permissive divergence does not succumb to them.

The first set of problems faced by other epistemic divergence arguments, particularly those proposed by so-called distributed cognition accounts (Bird 2010, 2014; Hutchins 1995a; 1995b) has to do with the connection of group-level and member-level cognition. According to distributed cognition accounts socio-epistemic entities can form doxastic attitudes by directly processing and deliberating evidence, which is spread among the group's members or integrated into the group's structure. Since this can happen without the awareness of most or all of the group's members, the doxastic states of the group produced in this way may differ from any of

the doxastic member states. So, distributed cognition accounts allow “an organic group to know that *p* without individual belief that *p*, joint acceptance that *p*, commitment that *p*; indeed, without a single individual person—group member or not—even being aware of that *p*” (Lackey 2021: 113).

This raises various problems, which I will mention here and discuss in more detail in the second half of the thesis.²³ First, for Lackey, this severs the vital connection between group knowledge, group action, and collective responsibility (2021: 115-123). Second, it is not entirely clear how to understand the collective evidence possessed by the group in these cases. On the one hand, distributed cognition accounts have been criticised for leading to unnecessary and implausible inflation of the group’s evidential base: any evidence possessed by any group member that could play a functional role within the group would count as part of the group’s evidence (Habgood-Coote 2019: 948). This leads to problems given the plausible assumption that evidence possessed by some members could serve as a defeater for the group-level attitude (Lackey 2020: 123-126). Accordingly, other promising accounts of group justification often demand the group to partly disclose their evidence (Silva 2019), or at least that disclosure wouldn’t undermine the group attitude (Lackey 2016, 2021).

Without judging the legitimacy of these worries, we can immediately demonstrate that permissive divergence is immune to them.²⁴ First, situations of permissive divergence can arise even if the relevant body of evidence is fully disclosed, e.g., via deliberation within the group. With the group and all of its members sharing the exact same body of evidence any worries about defeaters or wilfully excluded evidence are effectively ruled out.

Second, the permissive divergence case described above [§3.4] assumes mutual awareness not only of the evidential disclosure but also of the inner cognitive workings and the *modus operandi* of the group. This means that permissive divergence does not threaten the connection between justified group belief/knowledge and group action as distributed cognition does. In permissive divergence cases, members are aware of what is rational for the group to believe, and this awareness makes it rational for the members to act on the group’s behalf even though they don’t themselves (rationally) believe that *p* [§3.5].

²³ For some discussion of group evidence see chapter [ch.5], for extensive discussions of defeater cases see [ch.6] and [ch.7]. For a comparison of various accounts of collective justification see [ch.9].

²⁴ As I will argue in chapter [ch.9], many of the worries raised against distributed cognition do not survive closer inspection.

Another worry with extant epistemic divergence arguments is that even if we presuppose evidential disclosure and mutual awareness in the supporting example, we may indirectly manipulate the evidential base of the group by having different notions of what counts as evidence on the individual and the collective level. One often-discussed case, concerning a jury deliberating the guilt of a defendant, is found in Schmitt (1994).²⁵ In this case, all jury members possess an epistemically reliable piece of hearsay evidence supporting the guilt of the defendant, which is ruled out as being inadmissible by the judge. Schmitt thinks that while the jury as a chartered group justifiably believes that the defendant is innocent, none of the jurors justifiably believe this proposition because their justification is defeated by the relevant reliable hearsay evidence. This is possible since the group's charter prohibits it from considering the hearsay evidence and, therefore, excludes it from its evidential base. In its "legal capacity, the court rightly excludes hearsay, and its legal capacity is the only capacity in which it operates" (Schmitt 1994: 274).

However, if the evidential base can be changed in this way, we face similar worries to the ones concerned with distributed cognition. Namely, epistemically relevant pieces of evidence or defeaters possessed by the members may be left unconsidered. After all, as pointed out by Lackey, the exclusion of hearsay is based on "practical or procedural concerns but not necessarily epistemic ones" (2016: 355). That is, the court deeming the evidence as being inadmissible is not based on its epistemic properties, such as its reliability or truth-conduciveness (Lackey 2021: 68).²⁶ Treating these cases otherwise might open the door for instances in which groups wilfully manipulate their evidential base (Lackey 2021: 68-69). In other words, if declaring the hearsay evidence inadmissible would be an epistemic rather than a juristic-pragmatic decision, we could also imagine groups being chartered in a way that allows them to manipulate their evidential base inappropriately by, e.g., cherry-picking evidence that fits their practical goals.

While I agree that any cherry-picking of evidence is unacceptable, permissive divergence *only* relies on different evaluations of the *same* body of evidence via different rationally permitted methods based on different epistemic standards. We do not need to allow for any cases in which

²⁵ Following Gilbert (1989, 2004), Schmitt (1994) defends a so-called joint commitment account of collective justification.

²⁶ Here Lackey refers to Wigmore (1904), who points out that, the inadmissibility of hearsay evidence is vindicated since it doesn't allow the opposing side to confront the source of information. Instead, a witness needs to be "brought to testify in court on the stand, where he may be probed and cross-examined as to the grounds of his assertion and of his qualifications to make it" (Wigmore 1904: 437).

there are defeaters possessed by any of the group members –the evidential base itself is not manipulated or changed in any way.

This also shows that permissive divergence arguments are compatible with a wide range of views about group evidence, including various kinds of summativism about group evidence, which understand the group’s evidence as being the evidence possessed by some (or all) group members. After all, permissive divergence cases are possible even if the group’s evidential base is the same as the evidence possessed by all of the group’s members.²⁷

Furthermore, any divergence arising from diverging epistemic standards is still epistemically rational. While the juristic standards governing a court may sometimes lead to epistemically irrational group attitudes, in cases like DIVERGING STANDARDS the divergence is purely epistemic.²⁸

In sum, permissive divergence distinguishes itself from its competitors in allowing for divergence even under very stringent conditions, namely: (a) full disclosure of the evidence within the group, (b) mutual awareness of the disclosure and the *modus operandi* of the group, even if, (c) the group G itself, as well as all members m_1 - m_n are fully rational.

3.7 Conclusion

In this chapter, I have demonstrated that a very modest form of epistemic permissivism, called Standard Permissivism, is in direct conflict with any minimal summativist understanding of epistemic states such as justified beliefs or knowledge. In particular, I developed a new kind of epistemic divergence argument, called *permissive divergence*, that gives us a recipe to generate situations in which group-level and member-level attitudes rationally *ought* to diverge. This divergence argument differs from its competitors in allowing for a divergence between member-level and group-level attitudes even under very stringent conditions, namely: (a) full disclosure of

²⁷ For a summative understanding of group evidence see Silva (2019). For a discussion of various summative and non-summative understandings of group evidence see Brown (2022b; forthcoming: ch.2). I will discuss those views in [§5.7].

²⁸ Another case, also discussed by Lackey (2016; 2021) is presented by Mathiesen (2011). In this case, a hiring committee reaches a conclusion that is not supported by any committee member because of differences in their epistemic risk aversions. Lackey points out that the epistemic risk settings in the case (as Mathiesen describes it) are directly influenced by pragmatic considerations (Lackey 2021: 70-71). Depending on the exact reading of this case and the underlying understanding of the epistemic role of pragmatic considerations it may be interpreted as an exemplar of permissive divergence. See footnote 16.

the evidence within the group, (b) mutual awareness of the disclosure and the *modus operandi* of the group, even if, (c) the group G itself, as well as all members m_1 - m_n are fully rational. In so doing, the presented argument demonstrates that divergence among member-level and group-level attitudes is sometimes rationally required, without relying on any particular understanding of collective attitudes. This makes permissive divergence, in contrast to other divergence arguments found in the literature, not only theoretically more appealing but also less vulnerable to potential objections.

However, for the sake of brevity, I decided to frame the argument presented under the assumption of a permissive interpretation of epistemic conflicts. While this made the discussion more streamlined and more independent from the rest of the thesis it also raised some questions about how we can extend the presented argument to impermissible understandings of epistemic conflicts. Let me use this opportunity to reflect on how an impermissivist version of the presented epistemic divergence argument might look like.

It seems like that even if we subscribe to Evidential Uniqueness, as long as we analyse allegedly permissive cases as epistemic conflicts, the dialectic of the above-presented argument will play out very similarly. That is, all the members of a group may resolve an epistemic conflict over epistemic standards in one way, while the group resolves it in a different way, and all of this might be not determinately irrational (and all of the members can be aware that this is so). This gives us an epistemic divergence case under similarly stringent conditions to the ones identified above, namely: (a) full disclosure of the evidence within the group, (b) mutual awareness of the disclosure and the *modus operandi* of the group, even if, (c)* the group G itself, as well as all members m_1 - m_n are not determinately irrational. (Here only condition (c) has been weakened to (c)*, to make it compatible with the impermissible interpretation of epistemic conflicts.)

Now, as a matter of philosophical substance, we might wonder whether that will really present us with epistemic divergence cases. After all, on the indeterminacy-theoretic treatment of Evidential Uniqueness, there is no precisification on which both the members m_1 - m_n and the group G itself count as justified. So, strictly speaking, we cannot say that the epistemic status of the group-level attitude and the member-level attitude can diverge. At most, we can say that there are cases in which the epistemic status of their attitudes diverges according to their epistemic standards (of which at most one is correct) while there is no determinately correct criticism of neither G nor m_1 - m_n . I have called this immunity to criticism, following Williams (2016), weak permissibility. But what do these weakly permissible divergence cases tell us about the

relationship between the epistemic status of member-level and group-level attitudes? First, as argued in [§1.6], nothing changes in the overall epistemic development of members opting in and out of groups if those cases are only weakly permissive. Since after all, even if the members acknowledge that the group has different epistemic standards, they cannot know whether their standards, the group-level standards, or neither of them are strictly permitted or not. So, they are in a situation in which they cannot determinately criticise the group or any member (including themselves) who acts according to the group standards.

Second, the members might be aware of the conflicting epistemic standards but not be in an epistemic conflict. That is, they might when acting as a member of the group contribute to the group's overall behaviour in a way that makes the group form a doxastic attitude that aligns with the group's epistemic standards. This enables the members to act freely without facing any epistemic conflicts when opting in and out of their member roles. This means that they themselves when acting as a member do not adopt the epistemic standards of the group and, therefore, do not face any of the discussed problems associated with standard-switching [§3.5]. However, considerations about bindingness and the interplay of sweetening sensitivity and angst might allow for cases in which members do switch their standards when opting in and out of groups. I have discussed several understandings of bindingness that rationally allow for such behaviour in [§1.6], where I took a weak understanding of bindingness to be the most plausible candidate.

In sum, I think we can utilise the Conflict View to generate divergence arguments independently of whether we take epistemic conflicts to be genuinely permissive or not. This makes the presented divergence argument more general than the permissivist version above since it makes it more independent from prior theoretical commitments. This leads us closer towards an inflationary approach to collective epistemology in general. In the next chapter [ch. 4], I will explore some methodological principles underlying contemporary approaches to epistemology and draw some general lessons about how we ought to do conceptual analysis in collective and individual epistemology.

Part I: Conclusion

The first half of the thesis was concerned with the question of whether rationality is epistemically permissive; that is, whether one body of evidence can rationalise more than one doxastic attitude. In the first chapter [ch.1], I attempted to answer this question by comparing allegedly permissive cases to value conflict cases from other normative domains. In so doing, I have argued that allegedly permissive cases are best understood as epistemic standard conflicts, conflicts that confront us with different doxastic attitudes supported by competing but incommensurable epistemic standards. I called this the Conflict View.

Interestingly, the Conflict View is compatible with both permissive and impermissive views of epistemic rationality. More precisely, it allows for a permissive as well as an impermissive interpretation of epistemic conflicts depending on the underlying understanding of incommensurability. If comparisons under incommensurability are best understood to be governed by a fourth positive value relation, these cases are genuinely permissive; if they are, however, best understood to be a result of indeterminacy they are not genuinely permissive. Nonetheless, independent of the underlying understanding of incommensurability, the epistemic conflict resolutions exemplify the same overall pattern of expected belief evolution. While the exact details of this pattern depend on our understanding of the bindingness of epistemic conflict resolutions.

In the second chapter, I defended the Conflict View against an objection from self-fulfilling beliefs [ch.2]. This was important, since if self-fulfilling belief cases were genuinely permissive, they would present us with counterexamples not only to Evidential Uniqueness but also to the Conflict View because these cases do not involve any conflicting epistemic standards. In particular, I demonstrated that if we carefully disentangle questions of practical and epistemic rationality, for any self-fulfilling situation, there is not only a unique rational doxastic attitude to be in but also only a unique rational doxastic attitude to transition into.

In contrast, chapter three [ch.3] can be understood as an application of the Conflict View. In particular, I used the observation that there are epistemic conflicts to generate so-called epistemic divergence cases, which are cases in which the epistemic statuses of group-level and member-level attitudes diverge. While, for the sake of narrative ease, I decided to frame the argument presented under the assumption of a permissive interpretation of epistemic conflicts, I illustrated how we can utilise the Conflict View to generate a non-permissive version of the

given epistemic divergence argument. This makes the presented divergence argument more general since it makes it independent from prior theoretical commitments concerning epistemic permissivism.

This leads us closer towards an inflationary approach to collective epistemology since it suggests that the epistemic status of group-level attitudes and member-level attitudes is independent to the extent that these statuses can sometimes diverge. In the next chapter [ch.4], I will explore some methodological principles underlying a popular way to do conceptual analysis in epistemology. In so doing, I will draw some general lessons about how we ought to do conceptual analysis in collective and individual epistemology. Using these insights I will then propose and defend a novel theory of epistemic justification in the second half of the thesis [ch.5] - [ch.9] that can be used to analyse the epistemic statuses of doxastic attitudes of individual and collective epistemic agents.

**METHODOLOGICAL INTERLUDE:
EPISTEMIC CONTINUITY**

4

Epistemic Continuity

A Plea for a United Epistemology

This chapter is a methodological interlude. It examines the underlying methodology of contemporary approaches to epistemology and proposes a way to compare and analyse epistemic attitude ascriptions in various domains, such as collective and individual epistemology. In so doing, I arrive at a novel proposal of how we ought to do conceptual analysis and/or conceptual engineering which unifies debates about collective knowledge and justification and individual knowledge and justification. I call this approach Continuous Epistemology. While I will present some arguments in favour of Continuous Epistemology in the later sections of this chapter, these arguments will be schematic and at best establish the *ceteris paribus* superiority of Continuous Epistemology. So, the main function of this methodological interlude, in the context of the thesis can be understood as articulating the methodology used in the chapters to come.

4.1 Introduction

We frequently ascribe knowledge, as well as rational or justified beliefs not only to individual epistemic agents but also to collectives, such as groups, communities, or other social entities. We say that ‘the college union believes that the industrial action is going to plan’ or that ‘the IPCC knows that the currently observed climate change is predominantly anthropogenic’. Among theorists of collective intentionality, it is nowadays common to take some of those ascriptions to be non-metaphorical, genuine attributions of epistemic states, which has led to a plethora of accounts of group knowledge, collective justification or collective rationality.

This raises the question of how we should think of the relationship between, e.g., the concept of ‘knowledge’ employed in ‘collective knowledge’ and the concept of ‘knowledge’ employed in ‘individual knowledge’. Are these concepts the same? And if they are not, can or should these concepts be the same? And, can or should we follow the same methodology when we analyse them?

While these questions are indirectly addressed by many social philosophers in making methodological and ontological assumptions when approaching the subject matter, there is

surprisingly little explicit engagement with these methodological questions in the literature.¹ In what follows, I will advocate for an understanding of epistemology that seeks to unify individual and collective epistemic theorising, as well as further potential epistemic endeavours having other (non-human) entities at the centre of their study. I will call this approach *Continuous Epistemology*. The basic principles of Continuous Epistemology are simple. First, epistemic concepts such as justification, rationality, knowledge, or different kinds of epistemic virtues are independent of the instantiating entity in question. That is, there is no *sui generis* ‘collective knowledge’ or ‘individual knowledge’, but just ‘knowledge’. Second, epistemology has no preferred methodological starting point. That is, insights made in different domains of inquiry, such as individual or collective epistemology bear equal weight in our analysis of epistemological concepts.

I will start with some terminological preliminaries [§4.2]. In particular, I will distinguish three different kinds of epistemology, individual epistemology, social epistemology and collective epistemology. Afterwards, I will discuss the general methodological framework which is at the centre of most contemporary epistemological endeavours. Namely, a specific version of the so-called method of cases, the view that we start with intuitive judgments about certain cases and then derive a general method from those judgments [§4.3]. Doing so enables me to identify some of the implicit working assumptions the method of cases is reliant on [§4.4] and demonstrate that Continuous Epistemology results from a natural extrapolation of these assumptions [§4.5]. The second half of the chapter locates Continuous Epistemology within the present landscape as shaped by extant views and provides arguments for a continuous approach. In particular, I will argue that continuous concepts possess more explanatory power than their competitors [§4.6]. Furthermore, I will use the previously identified working assumptions of the method of cases and demonstrate how they are often misapplied by extant views [§4.7].

4.2 Collective and Individual Epistemology

Traditionally, epistemology has been concerned with individual epistemic agents and the normative evaluation of their doxastic attitudes, such as credences, beliefs, disbeliefs or suspensions of judgment. It is the study of the character of individual epistemic agents, how they evaluate their epistemic situation, respond to their evidence, and form and update their doxastic

¹ A notable exception is Gilbert and Pilchman (2014). I will discuss their analysis of the relationship between individual and collective epistemology later [§4.7].

attitudes. As such, traditionally epistemologists have focused on “the study of epistemic evaluation or normativity, represented by such evaluative concepts as justifiedness, rationality, and knowledge” and asked “such questions as how individuals can acquire knowledge and maintain justified or rational credal states” (Goldman 2010: 2). Within the last decades, however, there has been a trend to transcend this traditional individualistic conception of epistemology in favour of a more *social* understanding.

Yet, the meaning of ‘social’ in *social epistemology* is ambiguous. First, social epistemic theorising may depart from traditional, individual epistemology in situating epistemic agents within a socio-epistemic environment. In so doing, epistemology becomes concerned with the shared aspects of beliefs, examines relations and interactions among agents and discusses issues related to testimony (Goldberg 2008, Lackey 2008), peer disagreement (Christensen 2007; Elga 2007; Kelly 2010; Lackey 2010) or expertise (Goldman 2001). While the epistemological questions that arise from these endeavours are truly social, they preserve the individualistic spirit of traditional epistemology since they investigate how *individual epistemic agents* ought to navigate their socio-epistemic environment. Accordingly, we can think of research projects in this vein as *preservationist* social epistemology (Goldman 2010), or *individual* social epistemology (Bird 2010).

However, there is another understanding of epistemology, sometimes called *expansionist* social epistemology (Goldman 2010), *social* social epistemology (Bird 2010), or simply *collective epistemology* (Gilbert & Pilchman 2014; Kallestrup 2022). Collective epistemology is concerned with whether and how social entities, such as groups can have epistemic properties. Whether groups can be rational, epistemically virtuous or genuine bearers of knowledge or justified beliefs. In this respect, social epistemology is often understood as an extension of individual epistemology. It takes well-researched concepts from individual epistemology in an attempt to extrapolate them to all kinds of socio-epistemic entities. This kind of epistemic theorising has led to a plethora of accounts of collective rationality and justification (Schmitt 1994; Mathiesen 2011; Hakli 2011; Goldman 2014; Lackey 2016; 2021; Hedden 2019; Silva 2019, Schwengerer 2021), different kinds of collective knowledge, such as knowledge-how and knowledge-that (Tuomela 2004; Bird 2010, 2014; Hutchins 1995a; Palermos & Tollefsen 2018; Habgood-Coote 2019; Kallestrup 2022), and collective epistemic virtues (Fricker 2010; Kallestrup 2016; Harris 2021).

With this rough categorisation at hand, we can start to investigate the relationship between individual epistemology, social epistemology and collective epistemology.² First, I follow Goldman (2010) and Bird (2010) in taking social epistemology to be just an extension of traditional epistemology that “retains the focus on individual knowers while socialising their epistemic context” (Bird 2010: 23). Analogously, collective epistemology can be socialised as well by examining collectives within their social environment. Research projects in this vein may be called *social-collective* epistemology. These kinds of investigations ask questions about the epistemic warrant of group testimony (Tollefsen 2007) or investigate disagreement within and between collectives (Miragoli a& Simion 2020).

However, answering how individual epistemology and collective epistemology hang together is less straightforward. As a first approximation, we can think of collective epistemology and individual epistemology as two parallel research projects that are methodologically similar but concerned with different kinds of epistemic agents. While the former takes (a certain type of) human individual as the focus of their epistemic investigations the latter centres around (a certain type) of social entity.

To make this more precise let me introduce some terminology. First, collective and individual epistemology are concerned with different *domains of inquiry* that are inhabited by different types of *epistemic agents*; where the term ‘epistemic agent’ shall be understood as a placeholder for entities that are taken to be ‘proper bearers of *epistemic states*’; and ‘epistemic state’ (for the lack of a better term) is another placeholder, for rational or justified doxastic attitudes, knowledge, or epistemic virtues—in short, every property that an epistemic agent can have, that is a valid target of epistemic evaluation.

Given these terminological preliminaries, we can contrast collective and individual epistemology as follows:

Individual Epistemology: The study of *individual epistemic agents* and their *epistemic states*.

Collective Epistemology: The study of *collective epistemic agents* and their *epistemic states*.

This rough categorisation is a helpful heuristic since it illustrates how individual epistemology and collective epistemology are usually contrasted. However, it leaves the initial guiding questions

² While I think that this rough characterisation is a useful heuristic, there are additional epistemological projects which understand themselves to be social but do not clearly fall into any of those camps. For a more extensive overview of different types of social epistemology see Goldman (2010) or Longino (2022).

unanswered. Namely, it neither tells us whether we (ought to) use the same concepts to analyse epistemic states in collective and individual epistemology nor does it tell us how we (ought to) analyse and refine the used concepts [§4.1]. In other words, it leaves open whether we should analyse individual and collective epistemic states differently.

To be able to provide general answers to these questions [§4.5] we need to take a step back and investigate how epistemologists usually analyse (and/or engineer) the concepts they employ in their normative evaluations of epistemic states [§4.3] [§4.4].³

4.3 Methodism and Particularism

Following Chisholm's (1977) meta-epistemological investigations of the conceptual analysis of knowledge, we can distinguish two general methodological strategies within epistemology, *particularism* and *methodism*.⁴ Chisholm points out that when trying to get a grip on the nature of knowledge we typically ask two distinct but interdependent questions, 'what do we know?', and 'how do we know?', which require two distinct answers; whereby particularists start with the first question, and methodists start with the latter.

In other words, we either start with some uncontroversial examples or instances of knowledge and then deduce a general method for acquiring knowledge from them, or we start with a general method and then use it to analyse specific cases. Traditional epistemology, such as empiricism or rationalism is often said to be methodist in spirit (Adam & Sosa 2022: §3.1). These frameworks start with a list of certain methods of how we acquire knowledge, such as via empirical investigations, deductive inference, or rational insight (withstanding systematic doubt), and then make judgments about specific cases. In contrast, nowadays most approaches (both in individual epistemology and collective epistemology) are versions of *particularism*. They start with some representative cases and then try to derive a general method or analysis from them. Think, for example, about the *conceptual analysis* of knowledge and justification emerging from the discussion of Gettier cases (Gettier 1963) as counterexamples to the traditional analysis of knowledge as justified true belief (Shope 1983; Ichikawa & Steup 2018; Horvath & Koch 2020). Here the standard procedure is to make intuitive judgements about potential instances of

³ For the sake of narrative ease my investigations focus primarily on conceptual analysis. However, I am confident that the proposed continuous approach to conceptual analysis can be extended to conceptual engineering as well.

⁴ See also Amico (1993), Fumerton (2008), or Adam and Sosa (2022).

knowledge, derive a theory from them, and test it against hypothetical counterexamples.⁵ Furthermore, this method applies not only to contemporary analysis of knowledge but “*mutatis mutandis*, to projects in epistemology that investigate other epistemic phenomena of interest, including, e.g., epistemic justification, understanding, intellectual virtues, and the like” (Adam & Sosa 2022: § 3.1).

If we look at currently influential accounts of *collective* knowledge, rationality or justification this procedure seems to be more or less the same.⁶ We start with intuitive judgments about specific cases, such as ‘NASA knew how to build a Saturn5 rocket’, or ‘the jury justifiedly believes that the suspect is guilty’. Afterwards, we derive a general method from these intuitive judgements, e.g., a group *G* justifiedly believes that *p* iff “*G* has a good reason *r* to believe that *p* [...] and only if all members of *G* would properly express openly a willingness to accept *r* jointly as the group’s reason to believe that *p*” (Schmitt 1994: 265). Then we introduce counterexamples in which a group seems unjustified despite it having the alleged justification-conferring property. For example, we show that Phillip Morris does not seem justified in believing that secondhand smoke is harmless despite having a good reason *r* to hold this belief and despite all of Phillip Morris members jointly accepting *r* as the group’s reason to believe (Lackey 2016). Then based on these counterexamples, we derive a new or modified theory and test it against new hypothetical counterexamples until we reach a relatively stable ground.

As often pointed out, there are some central working assumptions to get the particularist methodology off the ground. Namely, particularists need to presuppose that there are good cases (and bad cases), to begin with, and that we can, in principle, distinguish (via intuitive judgements) between instances of knowledge, justification, rational belief and so on, and lack thereof. In other words, the particularist rules out scepticism from the outset not only by assuming that there are at least some instances of epistemic success but also that we have some fallible capacities to identify them (Chisholm 1982: 67-70). This also resembles current trends in collective epistemology. Many simply start with the assumption that at least some social entities, such as chartered groups or structured collectives, are in certain circumstances proper bearers of epistemic states. That is, they rule out eliminativism or rejectionism concerning collective epistemic states from the outset.

⁵ This procedure is also sometimes called the method of cases (Horvath & Koch 2022). I want to reserve this term for a more specific methodological approach which I will specify below.

⁶ See, for example, Schmitt (1994), Hakli (2010), Goldman (2014), Lackey (2016; 2021) Silva (2019) or Habgood-Coote (2019).

While this rejection of scepticism (about individual epistemic states) or eliminativism (about collective epistemic states) is central to particularist approaches to conceptual analysis there are different ways to motivate this move. Some start with central, uncontentious case judgements that seem so fundamental that we can never give way in the face of a conflict with other cases or general principles (Greco 2005a; Sankey 2010). Maybe judgements such as ‘I know that I have hands’, or ‘I know that identical objects are indiscernible’ cannot be refuted for the sake of maintaining an otherwise successful method. However, it is not clear that the particularist has to accept this in-principle irrefutability of central cases (Sankey 2010: 7). Others understand particularism to be compatible with the method of *reflective equilibrium* (Goodman 1955), on which judgements and methods are brought into balance by a process of mutual adjustment (Lemos 2001). On this later understanding of particularism, every case judgement is in principle revisable, regardless of how central it is to our understanding.

We do not need to take a stance on whether particularism needs or can be entirely reconceptualised in terms of a reflective equilibrium approach or not. Yet, it is important to point out that particularism always relies on the method of reflective equilibrium in a looser sense. Not all judgements can be uncontentious. So, to effectively derive a method from intuitive but contentious judgments, we need to trade off these judgements and the derived method against each other. Contentious case judgements can and shall be adjusted in favour of an otherwise extremely successful method, that can accommodate a large number of cases. I will call this way of doing conceptual analysis that combines particularist ideas with a reflective equilibrium approach the *method of cases* (Ichikawa & Steup 2018; Horvath & Koch 2022).

In the next section [§4.4], I am going to take a closer look at the method of cases to identify some of its working assumptions. We can then use these working assumptions to define my novel approach to epistemological conceptual analysis called Continuous Epistemology [§4.5].

4.4. The Method of Cases: Three Principles

We can summarise the above-made characterisation of the method of cases as a stepwise procedure (i) - (v):

- (i) First, we specify our *domain of inquiry* D.⁷

⁷ This is often done implicitly by focusing on a specific type of case judgements, such as judgements involving human agents or judgements involving human collectives.

- (ii) Second, we make judgements about a representative set of cases within D.
- (iii) Third, we come up with a (preliminary) *method* that can categorise and explain as many cases as possible in accordance with our judgements.
- (iv) Fourth, we test our *method* against potential counterexamples within D.
- (v) Fifth, If we manage to refute the method, we need to try to come up with a new one, if not, we temporarily accept it. Alternatively, we might also adjust some of our initial judgements or redefine our domain of inquiry D for the sake of reaching a *reflective equilibrium*.

This stepwise procedure (i)-(v) gives us a useful heuristic that encapsulates the overall strategy of the method of cases. This procedure, however, should not be understood as a recipe that needs to be followed step by step. Especially, the decision about when to try to introduce counterexamples (iv), how to weigh these and how to react to alleged refutations (v) are things that do not follow clear guidelines.

However, something that is highlighted by having a stepwise instruction such as this is that the method of cases as just described does not start from the cases without having any initial methodology built into it (something that is already suggested by speaking of the ‘method’ of cases). In particular, it relies on some assumptions about how to categorise, compare and value intuitive case judgements.⁸

This later observation is important since we are trying to compare contemporary conceptual analysis in collective epistemology with traditional conceptual analysis in individual epistemology. In particular, it enables us to highlight some methodological assumptions implicit in the method of cases, that can be captured via the following three principles:

Conceptual Unity <Unity>: Epistemological concepts are general. That is, whenever we evaluate some type of epistemic state of different types of epistemic agents, we evaluate them under one unified concept.

⁸ This mirrors something that is already pointed out by Chisholm (1973) as well as Amico (1988). They both emphasise that methodists start with the implausible assumption that we can grasp the understanding of epistemic concepts *a priori* without making any judgements about specific cases in advance. A similar worry, however, also applies to particularists. It is implausible that particularists can make any considered judgements about cases without implicitly relying on some methodology based on which they make their judgements. This mutual dependence demonstrates that the initial distinction between methodism and particularism should not be understood as a sharp dichotomy but rather as contrasting methodologies that utilise case judgements to different degrees and in different ways. In other words, the methodists need some implicit case judgement to come up with a method while the particularists are committed to some initial minimal methodology.

Methodological non-Primacy <non-Primacy>: Epistemological theorising has no preferred starting point. That is, all case judgements should *ceteris paribus* be weighted in the same way.

Epistemic Agency <Agency>: There is a general way to characterise and identify epistemic agents. That is, there is one universal set of characteristics that enables us to identify potential bearers of epistemic states.

These three principles will turn out to be important not only in spelling out *Continuous Epistemology* [§4.5], but also when distinguishing the proposed research program from extant approaches [§4.7]. But, let me start by highlighting the reasons why <Unity>, <non-Primacy> and <Agency> are important principles that need to be followed when applying the method of cases.

<Unity> is important since it stops us from fracturing our domain of inquiry unnecessarily. It enables us to assume that we are studying the same concepts when thinking through different cases, and, therefore, prevents our concepts from becoming increasingly ambiguous and/or fractured the more cases we consider. It tells us that when considering instances of knowledge involving philosophers and laypersons we are not studying two *sui generis* kinds of knowledge, ‘philosopher knowledge’ and ‘layperson knowledge’ but just knowledge. As such, <Unity> is one of the underlying principles behind the analytic project of trying to disambiguate and generalise the loose character of our natural language.⁹ This, however, does *not* rule out *any* fracturing of our epistemic concepts. Sometimes, for the sake of reaching or maintaining a reflective equilibrium, we can fracture our domain of inquiry. For example, when we become convinced that there are different types of knowledge, such as knowledge-how or knowledge-that. <Unity> just tells us that this fracturing can only be part of a trade-off that comes at a theoretical cost.¹⁰

<non-Primacy> ensures democracy within the given domain. It, together with <Unity> enables us to introduce supporting cases as well as counterexamples to any given theory. It prevents us from discounting intuitions about certain cases in favour of others, without having further theoretical reasons to do so. Note that <non-Primacy>, does *not* say that all intuitive judgements should be given equal weight but instead that the way we use our intuitions in

⁹ This is reminiscent of Carnapian conceptual engineering (Carnap 1962: 3-8). For recent discussions see Brun (2016) or Cappelen (2020).

¹⁰ This raises the question of whether and to which degree <Unity> is compatible with epistemic contextualism. While I can’t hope to answer this question here, certain contextualist claims are compatible with <Unity>. For example, even if knowledge attributions are stakes sensitive, <Unity> holds if the way knowledge attributions are stakes sensitive is uniform across the entire domain of inquiry. Thanks to Raimund Pils for helpful discussions on this. For a related discussion see Chalmers (2011: 540-543).

weighting these cases should be the same. For example, we might think that there is a distinction between initial intuitive judgements and considered judgements and that we shall give the latter more weight than the former. Likewise, <non-Primacy> is compatible with the assumption that the strength of intuitions matters. In other words, very strong intuitive judgements about central uncontested cases might weigh much more because we have stronger intuitions about them (again whether there are cases which are in principle uncontested is left open).

<Agency> fulfils a similar role as <Unity>. It prevents us from unnecessarily fracturing our domain of inquiry by distinguishing between different types of epistemic agents. That is, it enables us to assume that we are studying the same type of epistemic agent when thinking through different cases. To use the above-given example, there are no *sui generis* types of epistemic agents called philosophers and laypersons. However, this does not mean that to save intuitive case judgements, sometimes reflective considerations can restrict or expand our domain of inquiry by re-considering our understanding of epistemic agency. For example, if we encounter too many counterexamples to the claim that infants can have justification or knowledge through testimony we might stop thinking about infants as being genuine epistemic agents. This is even more clear in the case of collective epistemology, where much of the effort goes into identifying the types of social entities that ought to be considered genuine epistemic agents. While some are more liberal with their epistemic state ascriptions others think that only structured collectives, chartered groups, or collectives with proper agential structure can be bearers of epistemic states.¹¹

In sum, most collective and individual epistemologists follow roughly the method of cases. As such, they make various implicit methodological assumptions, most notably they presuppose <Unity>, <non-Primacy> and <Agency> as underlying principles about how to categorise, utilise and weigh case judgements. In the next section [§4.5], I will demonstrate that taking these assumptions seriously leads us to a new understanding of epistemology, called Continuous Epistemology.

4.5 Continuous Epistemology

Having identified the methodological principles <Unity>, <non-Primacy>, and <Agency> enables us to reinvestigate the relationship between individual and collective epistemology. In particular, I will demonstrate that these methodological principles show that we ought to do collective and individual epistemology continuously. That is, we start with judgements about cases involving

¹¹ See, for example, List and Pettit (2011: 59-79) or Hess (2018).

individuals and cases involving collectives, treat them equally (following <non-Primacy>) and under one unified concept of epistemic agency (following <Agency>), and then derive a general method from them (following <Unity>). In other words, we start with one extended but unified domain of inquiry encompassing case judgements such as:

‘S knows that she has hands’,
 ‘I am justified in believing that London is the capital of the UK’,
 ‘NASA knew how to build a Saturn5 rocket’,
 ‘the IPCC rationally believes in anthropogenic climate change’,
 ‘Smith doesn’t know that Jones will get the job’,
 ‘I don’t know that I am looking at a Zebra rather than a cleverly disguised mule’,
 ‘Phillip Morris is not justified in believing that second-hand smoking is harmless’
 ...

When we have made enough judgements about cases from this extended domain we follow the recipe further, propose a method, and then test it against potential counterexamples [S4.4].

Once we operate in this extended domain, it is only a small conceptual step to extend our domain of inquiry even further to include other potential bearers of epistemic states. Remaining at least open to the possibility that not only human individuals and collectives can be proper bearers of epistemic states, but that there might be other epistemic agents such as animals, animal collectives, AIs, aliens, etc. suggests that there is a general domain of inquiry which encompasses all possible types of epistemic agents. In so doing, we take the method of cases and its underlying principles <Unity>, <non-Primacy> and <Agency> seriously. That makes our epistemological analysis *continuous* in the desired sense:

Continuous Epistemology: Doing epistemology is doing epistemology *per se*. That is, when we engage in epistemological conceptual analysis we analyse general domain-overarching concepts that enable us to capture and uniformly evaluate the epistemic states of (every kind of) epistemic agent.

I think the best way to understand Continuous Epistemology is as an *epistemic stance*, an idiosyncratic way to handle epistemic conceptual analysis and engage with epistemological puzzles.¹² Accordingly, as with other stances, we might characterise Continuous Epistemology as

¹² While thinking of Continuous Epistemology as a stance is a helpful heuristic, I do think that there are convincing arguments for it [S4.6]. This conflicts with a popular understanding of stances as voluntary non-epistemic positions, which cannot be argued for (see for example van Fraassen 2002).

a cluster of orientations, commitments, strategies and attitudes underlying our mode of epistemic inquiry (van Fraassen 2002; Chakravartty 2018). If we take a continuous stance we integrate <Unity>, <non-Primacy> and <Agency> deeply into our way of thinking about the world. We reject any approach that treats human epistemic agents or any other kind of epistemic agent to be superior or the primary object of epistemological analysis. As various philosophers of mind dismiss certain theories of consciousness as species-chauvinistic (Block 1978), the continuous epistemologist dismisses non-continuous approaches to epistemology to be equally prejudiced. After all, why should there be anything special about human epistemological inquiry?¹³

While Continuous Epistemology appears to be a radical stance, interestingly, many epistemologists, implicitly presuppose a continuous understanding of some concepts, while explicitly defending non-continuous understandings of others. For example, many accounts of collective justification incorporate notions of evidence, defeaters, normative obligations, and other issues related to evidential support (Goldman 2014; Lackey 2016; Silva 2019; Hedden 2019; Brown 2022a, forthcoming), that are widely acknowledged to play a role in individual justification. Furthermore, many have argued that reliability plays a key role both in collective and individual justification (Tollefsen 2002a; Goldman 2014; Dunn 2021). Often, these underlying concepts (evidence, reliability,...) are understood continuously, while the defended understanding of knowledge or justification is non-continuous (see [S4.7]).

Overall, theorists typically try to preserve their favourite theory of epistemic states when extrapolating it to the collective domain; e.g. when they are reliabilists or evidentialists about individual justification they are also often reliabilists or evidentialists concerning collective justification. Here the implicit assumption is that when we have learned that, for example, reliability has something to do with justification, we have learned something about justification in general (across all domains). In other words, being continuous with respect to some concepts is quite natural and indeed a common practice, and, that is, for good theoretical reasons. These

¹³ Recent epistemological discussions, inspired by work on extended cognition and the extended mind (Clark & Chalmers 1998) have raised the question of how we should demarcate the epistemic boundaries of epistemic agents (Clark 2008; Carter & Kallestrup 2019). Without discussing any of these proposals, it is easy to see that considerations like these are crucial for doing Continuous Epistemology. If we always need to make some assumptions about the boundaries of epistemic agents when considering and categorising cases, these assumptions can be traded off for the sake of approximating a reflective equilibrium as well. Accordingly, we should understand not only the defining characteristics of epistemic agents generally (following <Agency>) but also have a continuous understanding of the boundaries of epistemic agents. I think these issues, while interesting, are complex and in order to maintain the length and focus of the chapter, I decided not to include this discussion in the main text. I want to thank Jesper Kallestrup for helpful discussions on these issues.

reasons are the same as the reasons to presume <Unity>, <non-Primacy> and <Agency> in the first place. These principles shield our domain of inquiry from fracturing, help us to organise case judgements and prevent our concepts from becoming increasingly ambiguous [S4.4].

At this point, it is important to make some clarifications about the limitations and commitments of Continuous Epistemology. Continuous Epistemology should *not* be understood as absolute in the sense that it needs to apply to all possible epistemological investigations. Maybe some concepts are per design only relevant for describing certain types of phenomena involving certain types of epistemic agents. It may be reasonable to talk of, e.g., ‘common knowledge’, in the sense that it is ‘mutually shared knowledge within a certain collective’, or theorise about ‘individualistic rational insights’ in the sense that the notion incorporates distinct human phenomenology. We might call these kinds of concepts domain-specific or *local* concepts.

Since I think that analysing local concepts can sometimes lead to interesting and important results, we need to spell out the commitments of Continuous Epistemology more precisely to avoid the risk of making the proposed project vacuous. First, the continuous epistemologist can acknowledge that local, domain-restricted investigations can be of genuine importance. However, Continuous Epistemology rests on the assumption that there are global concepts enabling us to investigate epistemic agents and their epistemic states in general. As such, global concepts are the main targets of epistemic conceptual analysis, while local epistemic concepts just describe emergent epistemic phenomena arising from idiosyncratic characteristics of the type of epistemic agent in consideration. Second, Continuous Epistemology rests on the assumption that taking the continuous stance enables us to gain important insights into the properties of epistemic agents across all, or most, domains that we would miss otherwise. That is, following <non-Primacy>, the continuous epistemologist predicts that analysing the general epistemological nature of epistemic agents and their epistemic states is only possible if we rest our investigations on a vast range of case judgements from different domains.¹⁴ Those assumptions can be understood as the theoretical predictions of Continuous Epistemology, and, therefore, provide us with standards according to which we can measure the success and fruitfulness of a continuous approach to epistemology.

Having outlined these central features of Continuous Epistemology enables me to spend the rest of the chapter further motivating the continuous approach. First, I will argue that, if

¹⁴ At least we have to test our concepts against case judgements from all domains to be allowed to treat them as global. How many of our currently existing concepts, such as ‘evidence’, ‘knowledge’, ‘justification’ or ‘rationality’ shall be regarded as local rather than global, then becomes an open question, which I will not discuss here.

successful, the more continuous our approach to epistemology is the more explanatory power is provided by the concepts it analyses and employs and, therefore, the more explanatory knowledge is generated [§4.6]. I will end with a comparison to extant approaches in collective epistemology. This will demonstrate that despite being a natural extension of the method of cases continuous epistemology is practically unexplored [§4.7].

4.6 Continuous Explanations

Having spelt out the underlying principles of Continuous Epistemology [§4.5], as well as the methodological origin of those principles [§4.3] [§4.4], enables us to address the question of why we should take a continuous stance in the first place. While I have partly answered this question above, by arguing that Continuous Epistemology is a natural extension of the method of cases that comes with all of its theoretical merits, in this section I will provide a more general argument for Continuous Epistemology. Particularly, I will demonstrate that if Continuous Epistemology is successful, it will offer us more explanatory powerful concepts. This will give us a general argument for doing epistemology the continuous way.

Let me start by making a little detour into the philosophy of explanations. <Unity>, one of the central theses of Continuous Epistemology, is a thesis about epistemic concepts, the kind of concepts used in epistemic explanations. Explanations are often said to provide us with a special kind of knowledge, explanatory knowledge, or knowledge-why, in contrast, to merely descriptive knowledge or knowledge-that (Salmon 1984; Kim 1994).¹⁵ The method of cases rests on the assumption that case judgements provide us with (prima facie) descriptive knowledge; i.e., knowledge that an epistemic agent S has or lacks a certain epistemic state D in the given case, without giving us an explanation of why S has or lacks D. For example, I know that “I know that I have hands” via considering this specific case and making a judgement about it, without knowing why or how I have this knowledge. To gain explanatory knowledge about these judgements we need to appeal to epistemic explanations featuring epistemic concepts. For example, I know that I have hands because my belief is not only true but justified since it is formed via a reliable process: visual perception in bright daylight. Once I have identified this process, I have identified a method to generate justified beliefs: looking at objects in bright

¹⁵ Note that this distinction is different from the often-discussed distinction between knowledge-how and knowledge-that. See for example Stanley and Williamson (2001).

daylight.¹⁶ And, in so doing, I have gained explanatory knowledge: I am justified in believing that *p* because I have reached that belief via a reliable method. In a sense, this highlights the distinction between methodism and particularism again. The methodist starts with the knowledge-why, ‘I know everything that I cannot possibly doubt’, and generates knowledge-that in a second step, ‘Therefore, I know that I exist’; while the particularist starts with instances of knowledge-that, ‘I know that identicals are indiscernible’, or that ‘every natural number has a successor’, ... therefore, ‘I can know certain propositions via rational insight’.

With these remarks about explanatory knowledge in mind, we can proceed more systematically. So, how do we explain why ‘S knows that *p*’? First, we explain judgements about cases, by analysing the concept of knowledge and how it is realised in the given case. In other words, we engage in conceptual analysis, break down the complex concept of ‘knowledge’ further (if possible) into, e.g., ‘justified true belief’ and see how its component parts are realised. Afterwards, we break down simpler concepts such as ‘justification’ further, e.g., ‘properly based on sufficient evidence’, and see how this is realised, and so on. (Until at some point, we hit explanatory bedrock, the point at which we cannot break down the concepts any further.)¹⁷ This tells us something about epistemic explanations (as well as explanations in general). Explanations can have different amounts of complexity, some go deeper and reveal more structure while others remain fairly superficial. Let us call this feature of explanations, *explanatory depth*.

Judging explanations as having different amounts of depth rests on the assumption that explanations primarily generate explanatory knowledge by generating unifying patterns or structures by linking different phenomena in novel ways (Kitcher 1981; Friedman 1974; Kim 1994). This is a widely acknowledged feature of explanations that is independent of discussions about the nature of explanations. That is, all of the observations made so far leave us agnostic

¹⁶ One often-mentioned feature of explanations, that is important here, is that the amount of understanding they generate is context- and audience-sensitive. What counts as a good explanation largely depends on the audience, their prior knowledge, their cognitive set-up, and their interests (Lewis 1986). However, while the amount of understanding that particular acts of explanations generate is subjective or audience-relative, the explanations themselves aim to be objective. In other words, I can have explanatory knowledge without understanding why or how it explains the respective case in mind. So, telling you that S is justified in believing that *p* because her belief is formed via a reliable process, generates understanding only if you already believe that reliability has something to do with justification.

¹⁷ Importantly, normative explanations probably need further normative facts as part of their explanation to be complete while some of these normative facts would need to be considered brute (Väyrynen 2013: 158-160). For example, to fully explain why the belief that *p* is justified we might evoke causal explanations such as ‘the belief was formed via a reliable process’ or ‘is non-deviantly caused by some evidence sufficient to support that *p*’. However, these explanations cannot be full normative explanations of why it is good to form beliefs in that way unless they rely on further facts about the normative force of evidence or reliability.

concerning the debate between explanatory realism and antirealism.¹⁸ While explanatory realists think that explanatory knowledge is generated by giving us an understanding of the objective metaphysical relation between the explanans and the explanandum, anti-realists think that explanations only generate understanding by linking different ideas, concepts, and beliefs within a body of knowledge (Kim 1994).¹⁹ However, while realists and anti-realists differ in their interpretation of the origin or nature of the generated structure or pattern, they both often emphasise the structure-generating property of explanations as central (Kitcher 1981; Friedman 1974; Kim 1994; Dennett 1991). So, independently of whether we have a realist or an anti-realist view of explanations, we can (partly) measure the explanatory power of a particular explanation by referring to the amount of structure that is generated and/or revealed by it.

This allows us to formulate a simple but effective argument for Continuous Epistemology. Continuous Epistemology, if successful, would provide us with more powerful explanations because continuous explanations would generate more explanatory knowledge by revealing more *structure* or *patterns* across the individual domain, collective domain and beyond. That is, continuous explanations would be more powerful because they relate vastly different phenomena such as instances of various types of epistemic agents having various types of epistemic attitudes (realism) or link vastly different concepts and ideas in novel ways (anti-realism). Let us call this feature of explanations *conceptual range*.

Note that there are some trade-offs between *explanatory depth* and *explanatory range*, the two highlighted structure-generating properties of explanations. Acknowledging that explanations can have different degrees of depth means that there are different layers on which we could maximise their explanatory range. Maximising the range on a very high level is cheap. It only requires that there is some level of description that unifies intuitive judgements about epistemic agents in all domains. For example, we might say on a very general level that knowledge is justified true belief (in all domains); while an example of a very deep but wide-ranging explanation would be:

¹⁸ Kim (1994) illustrates this by pointing out that when thinking about explanations we not only raise epistemological questions, such as ‘what is gained by explanatory knowledge?’, but we also often raise metaphysical questions, such as ‘is there an objective metaphysical relation between the explanans and the explanandum?’. So-called explanatory realists often answer the first question by answering the second question. For them, explanations generate understanding by telling us something about the world. They reveal causal structures (Salmon 1984; Lewis 1986), dependence relations (Kim 1994), or constitution relations (Shafer-Landau 2003). Famous proponents include Hempel (1965) whose deductive-nomological model suggests that the value of explanations is grasping the nomic expectability of the explained event, or Kitcher (1981) and Friedman (1974) who think that explanations generate understanding solely by unifying our epistemic system(s).

¹⁹ For the relationship between explanations and understanding see footnote 16.

knowledge is justified true belief (in all domains) and justification is being properly based on sufficient evidence (in all domains), where evidence is all the propositions one justifiably believes (in all domains), and proper basing is a non-deviant causal relation (in all domains), and so on.

As pointed out above [§4.5], however, Continuous Epistemology rests on the assumption that the most explanatory powerful concepts are domain-independent. That is, Continuous Epistemology assumes that we do not need to trade off explanatory depth for explanatory range or vice versa. It simply predicts that epistemology is most fruitful when looking for patterns that are maximally deep and maximally wide-ranging. This is a prediction that can only be verified (or falsified) once we engage in Continuous Epistemology.

In sum, Continuous Epistemology promises to deliver more explanatory powerful concepts while being neutral with respect to the nature of explanations. This rests on the general observation that explanations generate structures or patterns by unifying different phenomena (realism) and/or linking different ideas (anti-realism). As such, Continuous Epistemology is promising since it has the capacity to provide maximally wide-ranging explanations (governing a vast range of cases involving different types of epistemic agents) while reaching as deep as the concepts provided by extant, non-continuous approaches; i.e., to use the introduced terminology, continuous explanations maximise explanatory range without trading off explanatory depth.

4.7 The Contemporary Landscape

In the final section, I will highlight the novelty of Continuous Epistemology by comparing it to extant research programs in collective epistemology. In particular, I will demonstrate that many extant understandings of the relationship between individual and collective epistemology are non-continuous by violating some or all of the principles underlying Continuous Epistemology [§4.5]. Those who are not interested in such comparative analysis can safely omit this section.

While, as noted above [§4.1], there is almost no direct engagement with questions about the relationship between collective and individual epistemology, a notable exception is Gilbert and Pilchman's (2014) comparison of different notions of group belief and group acceptance. Gilbert and Pilchman diagnose a stalemate between rejectionists about group belief, who deny the existence of genuine group beliefs, and non-rejectionists who advocate for the existence of group beliefs. Though Gilbert and Pilchman acknowledge arguments made on both sides, they think that the participating parties are talking past each other due to an implicit disagreement about the relationship between individual and collective epistemology (2014: 211):

“For, although the collective case is at issue, the debate operates with accounts and distinctions tailored specifically to the case of the individual. This is surely the wrong procedure. Though individual and collective epistemology will doubtless have related results, neither one should rely on accounts and distinctions tailored specifically for the other. Rather, those working on either one of these projects should develop concepts [...] appropriate to the particular project at hand, without being constrained by the results of the other, however helpful these results may be from a heuristic point of view.”

While I largely agree with Gilbert and Pilchman’s diagnosis that we are often too overzealous in generalising findings from the individual case to the collective case, I disagree with their pluralist conclusion that collective epistemology and individual epistemology should be more or less independent. In other words, Gilbert and Pilchman’s proposal is discontinuous since while they accept <non-Primacy> they reject <Unity> and <Agency>.

Similarly, these principles are often implicitly affirmed or rejected by epistemologists by engaging with socio-philosophical questions in a specific way. Taking a closer look at how those theorists operate we can categorise various extant approaches to collective epistemology concerning their stance on these principles. For the sake of narrative ease, I will do so only using <Unity> and <non-Primacy> since <Agency> often directly correlates with <Unity>.²⁰ In contrast, both <Unity> and <non-Primacy> are (relatively) independent of each other and can therefore be accepted or rejected autonomously. This means we are faced with the following space of possibilities: (1) \neg <Unity>, +<non-Primacy>; (2) +<Unity>, \neg <non-Primacy>; (3) \neg <Unity>, \neg <non-Primacy>; (4) +<Unity>, +<non-Primacy>.²¹

I will now briefly walk through these four types of accounts, which I call Pluralism, Extrapolationism, Parasitism, and Continuous Epistemology respectively [Fig.1].²²

²⁰ Usually, the proposal to use the same concepts to analyse the epistemic states of collective and individual epistemic agents comes with acknowledging that these are sufficiently similar types of epistemic agents. Nonetheless, I acknowledge that the relationship between <Agency> and <Unity> is more complex and is worth studying independently.

²¹ “+” stands for affirmation and “ \neg ” for rejection.

²² Note, that these categorisations do not need to apply to every aspect of the discussed theory. Some theories may satisfy <Unity> with respect to one concept while violating it concerning another, or take a methodological approach compatible with <non-Primacy> in some respect while being strictly individualistic or collectivist in another.

	+ <Unity>	¬ <Unity>
+ <non-Primacy>	Continuous Epistemology	Pluralism
¬ <non-Primacy>	Extrapolationism	Parasitism ²³

Fig. 1

(1) **Pluralism** (¬<Unity>, +<non-Primacy>): *Slogan: there are at least two genuine but (relatively) independent epistemological domains of inquiry.* For the pluralist, collective epistemic states and individual epistemic states come apart and our theories about them should be built on distinct insights made in the respective domains. The above-described approach by Gilbert and Pilchman (2014) is an exemplar of this. Gilbert and Pilchman argue that “one should not assume that accounts and distinctions arrived at within individual epistemology are appropriately applied within collective epistemology, however central they are to individual epistemology” (2014: 190). And, therefore, we should “develop concepts [...] appropriate to the particular project at hand, without being constrained by the results of the other” (2014: 211). Accordingly, one class of accounts that falls into this category are joint commitment models of collective epistemic states (Schmitt 1994; Schwengerer 2021), which are directly inspired by joint commitment accounts of group attitudes (Gilbert 1987).²⁴

(2) **Extrapolationism** (+<Unity>, ¬<non-Primacy>): *Slogan: There is one unified domain of inquiry but a preferred methodological starting point.* Accounts within the second category assume that epistemic concepts shall be understood uniformly while taking a preferred methodological starting point. Depending on the alleged success of this approach we can distinguish two subtypes: *eliminativist* and *extrapolation* accounts. Extrapolation accounts can go in one of two directions, they either start at the individual level and extrapolate their concepts to the collective level or *vice versa*. One set of accounts that are of the former type are reliabilist understandings of collective

²³ I want to thank Joshua-Habgood Coote for suggesting this name.

²⁴ Arguably, so-called epistemic agent accounts of collective epistemic states fall into this camp (List & Pettit 2011; Lackey 2016, 2021). They treat groups as genuine epistemic agents but analyse the epistemic status of their attitudes differently than they would analyse the epistemic status of the attitudes of individuals. An interesting question is whether these accounts rest on a unified understanding of epistemic agency or not.

justification, which take reliability as understood in individual epistemology and apply it to groups (Goldman 2014; Dunn 2019). On the other hand, we have radically social approaches to epistemology such as Longino's contextual empiricist analysis of science (2001; 2022). For Longino the extrapolation goes in the other direction, she thinks that "communities are the focal epistemic agents and individual knowledge is derivative, depending on participations in interactions" (2022: 176).

Likewise, there are two directions eliminativists can go. The most common type of eliminativist starts with the assumption that individual epistemology is primary to then reach the conclusion that collectives lack the requirements or properties to be genuine bearers of (certain) epistemic states. One such example is Cohen's analysis of collective belief ascriptions, which concludes that while groups may be capable of accepting propositions, they are not able to have genuine beliefs, due to their lack of "a disposition to have a certain kind of mental feeling" (1995: 11). On the other hand, some standpoint epistemologists endorsing the community model of knowledge can be understood as eliminativists in the reverse sense. Those theorists sometimes take collective epistemology as primary in the sense that individual agents do not live up to the standards of knowing, or even justifiedly believing.²⁵

(3) Parasitism (\neg <Unity>, \neg <non-Primacy>): *Slogan: There are two genuine domains of inquiry, but one domain (usually the collective) is merely derivative of the other.* This third category is inhabited by theories that take collective epistemic states to be genuine but distinct from individual epistemic states while advocating that we base our understanding of epistemic states primarily on insights made in individual epistemology. One group of theories that falls into this category are the summative understandings of collective epistemology (Quinton 1976; Goldman 2014; Moky 2002), which I have discussed in [ch.3]. Summativist theories are unified through the assumption that to ascribe an epistemic state to a group is to indirectly ascribe it to (some of) its members.²⁶ To figure out, for example, whether a group G has a justified belief that p we would need to sum up or aggregate the respective justified beliefs of the members m_1 - m_n . In so doing, summativists take an analogous approach to extrapolation accounts by predominantly relying on insights made

²⁵ As Longino's contextual empiricism, some standpoint epistemologies might be regarded as extrapolation rather than elimination accounts with respect to individual knowledge. For an overview see Grasswick (2018: § 2.3).

²⁶ Others, such as Lackey's (2016) understanding of collective justification, or Tuomela's (2004) understanding of collective knowledge are not completely summative but retain summativist elements. Equally, different types of epistemic agent accounts which rely on judgement aggregation procedures, such as premise-wise majority aggregation are partly summative. See, for example, List and Pettit (2011) or Bright, Dang and Heesen (2018). While those accounts can be considered non-continuous it is hard to put them into one of the four camps. See also footnote 24.

in individual epistemology. However, for summativists, there cannot be any conceptual unity since they embed the individual epistemological concept into the collective one. Take, for example, summative understandings of collective justification. The justification conferring property P^* for the collective cannot be the same as the justification conferring property P for individuals, if P^* is something in the vein of ‘some sufficient subpart of the collective entity has the justification conferring property P ’. As a result, for summativists, the concepts employed when analysing group-level epistemic states and the concepts employed when analysing member-level epistemic states are necessarily distinct.

(4) Continuous Epistemology (+<Unity>, +<non-Primacy>): Any approach that accepts both <Unity>, <non-Primacy> (and <Agency>) is continuous. The closest example of such an account in the literature is Bird’s so-called distributed cognition account of social knowing. For Bird, collective epistemology, or social-social epistemology in his terminology, “is liberated more fully from linkage to individualistic perspectives by being seen as strictly analogous to individual epistemology” (2010: 24). Furthermore, Bird argues for something along the lines of <non-Primacy>, when discussing two potential strategies to make collective epistemology analogous with individual epistemology. Here, he dismisses, extrapolationist approaches which take their “favourite analysis of (individual) knowledge and seek social analogues for the various components” (2010: 39). Instead, Bird proposes the following functionalist approach, which is heavily inspired by works of sociologists such as Durkheim (2010: 40):

“it would be preferable to seek to understand the nature of social knowing in a way that does not presuppose any particular theory of knowledge [...] I suggest that we can [...] make progress in understanding the social case by locating the analogy not in the analyses of individual and social knowing but in their role or function. The approach I shall take is motivated by the conclusion that sophisticated societies are organised organically and by the sociological methodology that this gives rise to.”

However, while broadly continuous, Bird’s account falls short of being fully continuous in some respects. First, while talking about individual knowledge being analogous to collective knowledge Bird often says that both individual and collective knowledge and beliefs are subtypes of the same more general type (2019: 282). This, however, might not be the right way to think about the relationship between collective and individual knowledge since the continuous stance leaves us agnostic about whether there are *sui generis* (sub)types of knowledge which we might call individual and collective knowledge. The only thing the continuous epistemologist studies is the

general type called knowledge. Second, Bird often argues that certain insights in individual epistemology should be treated as evidence for his understanding of social knowing. For example, he thinks that the compatibility of social knowing with knowledge-first views (as opposed to joint commitment models) is strong evidence in favour of his model (2010: 24; 2014, 2019: 282). This, however, is problematic if we take the initial motivation for knowledge-first epistemology to purely stem from the study of individual epistemic agents.

In sum, many contemporary approaches to collective epistemology assume a non-continuous understanding of epistemic states that violates either <Unity> or <non-Primacy> or both. This observation, together with the above-presented arguments in favour of Continuous Epistemology, makes the proposed research program a novel and promising enterprise.

4.8 Conclusion

Most contemporary approaches to individual and collective epistemology rely on the method of cases [§4.3]. In so doing, they rely on three principles, <non-Primacy>, <Unity>, and <Agency> [§4.4]. Continuous Epistemology is a novel and unexplored [§4.7] extension of the method of cases, that arises from a consequent application of those principles [§4.5]. In so doing, Continuous Epistemology unifies collective and individual epistemology, as well as other potential endeavours investigating other types of epistemic agents and their epistemic states. Furthermore, I have illustrated that continuous explanations of epistemic phenomena promise to be more explanatory powerful than their non-continuous competitors [§4.6].

PART II:
CONTINUOUS JUSTIFICATION

Part II: Introduction

Part II of this thesis will be dedicated to the development of a continuous theory of epistemic justification; that is, a theory of justification that enables us to capture the epistemic status of doxastic attitudes held by epistemic agents of all kinds, including collective and individual epistemic agents.

Accordingly, the following chapters form a compact block that has more the form of a monograph rather than the paper collection feel that the first part of the thesis had. Moreover, in contrast to the ideas and arguments developed in *Part I*, much of what I say in the following chapters [ch.5] – [ch.8] will be programmatic. That is, in my search for a continuous theory of justification, I will heavily rely on ideas developed and discussed elsewhere while elaborating the details of the underlying frameworks will sometimes not be possible. After all, any development of a complete theory of epistemic justification would come close to developing a complete epistemology, which would go beyond what is reasonably expected to be achievable in a few chapters. Nonetheless, I am confident that if we fill in the right details in the account provided, we will arrive at an understanding of epistemic justification that surpasses most of the current competitor theories in various respects. Illustrating this will be the job of the last chapter [ch.9].

5

Rethinking Epistemic Justification:

Continuous Evidentialism

Let us start our quest to formulate a continuous theory of epistemic justification by taking another look at collective attitude ascriptions, such as ‘the college union believes the industrial actions are going to plan’, or ‘the mathematical community regards the proof of the Poincare conjecture to be the most important achievement in recent mathematical history’.

Observations made in the previous chapters show that we should think of at least some of these collective attitude ascriptions as genuine non-metaphorical attributions of doxastic attitudes to social entities. In particular, I argued that various considerations about the (apparent) permissibility of rationality recommend an inflationary understanding of groups as independent epistemic agents having epistemic states over and above the epistemic states of the group members [ch.3]. In the following chapters [ch.5] - [ch.9], I will further analyse the epistemic status of the doxastic attitudes ascribed, by offering and discussing a *continuous* theory [ch.4] of epistemic justification. That is, instead of solely inquiring, e.g., whether ‘Jones believes that Smith will get the job’ or ‘British Petroleum believed that the Deepwater Horizon drill was safe’, I will ask what it means for Jones or British Petroleum to be justified or rational in holding their beliefs.

5.1 Introduction

Discussions of collective justification have only very recently received focused attention in the literature (Schmitt 1994; Goldman 2014; Lackey 2016, 2021; Silva 2019; Dunn 2021; Hedden 2019; Kallestrup 2022, Brown 2022a, forthcoming). Among other things, these investigations revealed that we need to embrace a view of collective justification that is sensitive to the *evidence possessed* by the group. This evidence can both be (i) *distributed* among the group’s members, and a group belief can be (ii) *defeated* by evidence possessed by a single member (even when that evidence was not taken into account when the respective group belief was formed). Furthermore, it is often argued that groups, like individuals, have certain (iii) *normative requirements* or epistemic obligations that affect how group members should gather, process and disclose available evidence. Unfortunately, most, if not all, approaches to collective justification fail with respect to one or

more of these desiderata (i) – (iii) (even those accounts that are directly designed to address them).¹

Many accounts of collective justification rest on a questionable *discontinuous* understanding of collective epistemology. They either divorce collective epistemic theorising from individual epistemic theorising by thinking of collective justification as being distinct from individual justification, or they treat collective justification as being just a special case of individual justification, by straightforwardly extrapolating findings from individual epistemology to the social domain. I will propose an account of justification, that is more theoretically appealing than current competitor accounts by accommodating the above-mentioned desiderata (i) – (iii) and by being epistemically continuous. Here is an evidentialist account of justification that is well-equipped to do just that:

Continuous Evidentialism:

S justifiably believes that p *iff*:

- (0) S believes that p (call this proposition “Bp”).
- (1) (a) S possesses some evidence E, which is *sufficient to support* Bp, and (b) Bp is *properly based* on E, and
- (2) (a) S possesses some higher-order evidence E_H, which is *sufficient to support* q, the proposition that the total evidence E_{TOTAL} on balance supports p and (b) S’s belief that p is *properly based* on E_H.

We need to unpack several things. First (0), we need to have a continuous understanding of belief that can be applied to collective and individual epistemic agents. I will not provide any account of group belief here but I am happy to rely on the interpretationist (Tollefsen 2002b, 2015; Brouwer, Ferrario & Porello 2021) or functionalist frameworks (List & Pettit 2011; Brown 2023a, forthcoming: ch.4) proposed by others.² These frameworks are promising since they are able to provide us with continuous, agent-type independent understandings of doxastic attitudes [ch.4].³ Furthermore, they can be motivated by the divergence arguments against summativism discussed in [ch.3].

¹ Both Lackey (2016; 2021: ch.4) and Silva (2019), who acknowledge these desiderata ultimately fail to live up to them. I will discuss their accounts in [§9.2].

² Note that List and Pettit’s (2011) theory of group belief combines interpretationist and functionalist ideas.

³ My preferred account of group belief roughly follows a Dennett-style (1981) version of interpretationism. On this version of interpretationism, making attitude attributions is an explanatory project that relies on taking the so-called intentional stance towards a system. If we take the intentional stance, we attribute beliefs and desires to an entity to predict how it will rationally behave given those beliefs and desires. The basic idea of using this strategy to analyse collective belief attributions is then that groups are genuine believers because they are (i) reliably interpretable from

Second, we need to have an understanding of evidentialism [§5.2] that is epistemically continuous in the desired way. Third, we need to think about the *nature of evidence* [§5.3], the ways in which epistemic agents *possess* evidence [§5.4], as well as the distinctive obstacles we face in generalising these concepts to collective epistemic agents. After that, we need to clarify the concepts embedded in clause (1) subclause by subclause. In particular, we need to analyse what it means for some evidence E to be *sufficient to support* the respective belief that p (1a) [§5.5]. Next, we need to disentangle (1b), which will require an examination of the notion of an *epistemic basing* relation [§5.6]. In [§5.7] I will show that the provided framework is epistemically continuous.

In [§5.8], I will introduce a no-defeater clause (2), which ensures that the epistemic agent is *not ignorant* of the entire possessed and unpossessed evidence relevant to the epistemic situation when basing the belief that p on E. Since justifying (2) has turned out to be an intricate task, I will spend the following chapters discussing and dismissing various alternative no-defeater clauses [ch.6] [ch.7] and defending higher-order evidential requirements against potential objections [ch.8]. After having spelt out the details of Continuous Evidentialism, the final chapter situates the proposed understanding of justification within the wider literature on (group) justification [ch.9]. In particular, I will illustrate how Continuous Evidentialism manages to resolve a wide variety of problematic cases discussed in the literature, such as distributed cognition and evidential manipulation cases [§9.3]. In so doing, I will also compare Continuous Evidentialism to other promising accounts of collective justification, that served as a major influence in the development of Continuous Evidentialism [§9.2]; namely the *Group Epistemic Agent Account* (Lackey 2016; 2021), as well as the *Evidentialist Responsibility for Groups* (Silva 2019).

5.2 Evidentialism

The general framework Continuous Evidentialism rests upon is an *evidentialist* understanding of epistemic justification. Evidentialism takes evidence and its relationship with doxastic attitudes to be the centrepiece of epistemic justification.⁴ Following Conee and Feldman, the leading proponents of evidentialism, we can portray evidentialism as “the view that the epistemic

the intentional stance and (ii) taking an intentional stance towards a group can yield new predictive powers that we did not antecedently have. For such views of group belief see Tollefson (2002b; 2015). For a critical discussion of those views see Backes (2021).

⁴ See, for example, Kelly (2016: §1) or Graham and Lyons (2023).

justification of a belief is determined by the quality of the believer's evidence for the belief" (2004: 83). Conee and Feldman (2004) defend the following thesis:

Evidentialist Justification: A doxastic attitude *D* towards proposition *p* is epistemically justified for *S* at *t* if and only if having *D* towards *p* fits the evidence *S* has at *t*.

Sometimes the core commitment of evidentialism is also expressed as a supervenience thesis, which states that normative facts about the doxastic attitudes of the epistemic agent directly supervene on facts about their evidence. That is, any two epistemic agents possessing exactly the same evidence would be exactly alike concerning what they are justified in believing about any given issue (Kelly 2016: §1; Fratantonio forthcoming: §1). Conee and Feldman express this supervenience thesis as follows (2004: 101):

Evidential Supervenience: The epistemic justification of anyone's doxastic attitude toward any proposition at any time strongly supervenes on the evidence that the person has at the time.

Many evidentialists take Evidentialist Justification as expressing only a necessary condition of epistemic justification rather than a sufficient condition, because even if the evidence possessed by an epistemic agent overall supports a given proposition, the way in which the respective attitude is formed can influence its normative status. The idea here is that if a belief is arrived at in an epistemically defective way, e.g., via wishful thinking, it is not fully justified even if it is supported by the possessed evidence. To this end, we need to distinguish between *propositional* and *doxastic* justification, where evidence alone only determines whether a belief is propositionally justified, for it to be doxastically justified, the belief also needs to be *properly based* on the evidence.⁵

We can modify Evidentialist Justification and Evidential Supervenience accordingly:

Evidential Doxastic Supervenience: The justification of anyone's doxastic attitude toward any proposition at any time strongly supervenes on the evidence that the person has at the time as well as the facts about how the person bases their doxastic attitude on that evidence.

⁵ Sometimes the distinction is put in terms of "justifiable" and "justified" belief, or "being in a position to justifiedly believe" versus "justifiedly believing" (Turri, 2010b: 312). Other epistemologists, including Conee and Feldman (2004: ch. 4), Feldman (2002: 46) and Pryor (2004: 365), instead distinguish between well-founded and non-well-founded justified beliefs. More on those distinctions at the end of the next chapter [§6.7].

Evidentialist Doxastic Justification: A doxastic attitude D towards proposition p is justified for S at t if and only if having D towards p fits the evidence S has at t and S properly bases D on her evidence.

If we look at clause (1) of Continuous Evidentialism we can see that it takes the form of Evidentialist Doxastic Justification. As long as we interpret “a doxastic attitude fits the evidence S has” as being equivalent to “ S possesses some evidence which is *sufficient to support* Bp ”.

(1) (a) S possesses some evidence E , which is *sufficient to support* Bp , and (b) Bp is *properly based* on E , and

However, as it stands (1) is very schematic. As Conee and Feldman note “among the things needed to develop the theory more fully are accounts of what evidence is, what it is for a person to have something as evidence, when a body of evidence supports a proposition, and what the basing relation is” (2008: 84). Those are what Fratantonio calls the four fundamental questions that any complete evidentialist theory needs to answer (forthcoming: § 2):

Nature of Evidence: What is evidence?

Evidence Possession: What does it mean to possess evidence?

Evidential Support: What does it take for evidence to support a proposition?

Evidential Basing: What does it mean for one to base her belief on one’s evidence?

I will dedicate each of the following sections [§5.3] – [§5.6] to one of these questions. In so doing, I will commit myself to a particular package of answers that result in an internalist version of evidentialism, which is committed to the following supervenience thesis:

Mentalist Doxastic Supervenience: The justification of any epistemic agent’s doxastic attitude toward any proposition at any time strongly supervenes on the mental states that the person has at the time as well as the causal relations that obtain between the person’s mental states.

At this point, it is important to point out that I take Mentalist Doxastic Supervenience to be compatible with my understanding of epistemic support [§5.5] as a 3-place relation between

evidence, epistemic standards and doxastic attitudes [ch.1], since on my understanding the epistemic standards possessed by an epistemic agent supervene on their mental states.

Note further, that while I will provide a particular package of views that accords with Mentalist Doxastic Supervenience, there are other possible ways to fill out the details of (1) that are compatible with a wide variety of internalist as well as externalist views of epistemic justification (more on the distinction below [§5.3]).⁶

5.3 Evidence

The term evidence is used in many ways. For example, in juristic or scientific contexts, we speak of evidence as ‘pieces of evidence’, as discrete *publicly available* objects such as findings, models or test results. This also accords with the way we provide justification for a doxastic attitude in ordinary contexts. For example, when providing justification for the belief that Aristotle was the teacher of Alexander the Great, we might refer to the manuscripts of Plutarch.⁷ This kind of publicly available evidence, however, is at best potential evidence in the epistemological sense, since the investigator needs to be able to access and process pieces of evidence accordingly for them to play any epistemological role (Conee & Feldman 2008: 101, 2010: 84-85; McCain 2014: 10). This highlights one specific role the concept of evidence plays in our epistemic theories: evidence, whatever it is, needs to be able to affect the normative status of our beliefs.⁸

There are multiple views of evidence that are suited for this job. For example, one suggestion is that when justifying our belief that Aristotle was the teacher of Alexander the Great we do not refer to Plutarch’s manuscript but *our belief* that it is written in Plutarch’s Parallel Lives that Aristotle was the teacher of Alexander the Great. Accordingly, for Conee and Feldman evidence is some type of mental state, such that given two epistemic agents that “are exactly alike mentally, then they are alike justificationaly, e.g., the same beliefs are justified for them to the same extent” (2004: 56). Call this the *mentalist* understanding of evidence.⁹

⁶ For the distinction between internalist and externalist versions of evidentialism see Bergman (2018), Berghofer (2022: 74-76), or Fratantonio (forthcoming: §1).

⁷ Following Turri (2009) we might want to call such a view of evidence factualism. However, I want to reserve the term factualism for a different view.

⁸ Note that acknowledging this does *not* commit us to evidentialism since other views about justification can accept that the normative status of beliefs can at least be influenced by the evidence possessed.

⁹ For mentalist understandings of evidence see Conee and Feldman (2004), Kelly (2008), Turri (2009) or McCain (2014).

The mentalist understanding of evidence is in stark contrast with the recently influential *propositional* understanding of evidence which originated in Williamson (2000). Williamson (2000) argues that although we also cite physical objects and experiences as evidence only propositions are part of one's evidence because only propositions can play the central functions of evidence, such as supporting hypotheses in probabilistic reasoning or by inferences to the best explanation.¹⁰

Orthogonal to debates between mentalist and propositional understandings of evidence is the question of whether evidence is *factive* or not. That is, does E need to be true for it to be evidence?¹¹ So we have two theses:¹²

Propositionality [P]: E is a piece of evidence only if it is a proposition.

Factivity [F]: E is a piece of evidence only if it is true.

Accepting both theses, none, or only Propositionality gives us the following three views:¹³

mentalism [-P/-F]: E is a piece of evidence only if it is a non-factive mental state M.

abstractionism [+P/-F]: E is a piece of evidence only if it is propositional content p of a mental state M.

factualism [+P/+F]: E is a piece of evidence only if it is propositional content p of a mental state M and p is true.

While my preferred view of evidence can be framed in a way that is compatible with Propositionality (more on that below), I want to start by providing some reasoning for why I prefer a *non-factive* understanding of evidence.

Historically the non-factive understanding of evidence has emerged from troubles arising from sceptical worries, and the urge to differentiate rational beliefs from irrational ones even in

¹⁰ For Williamson “knowledge, and only knowledge, justifies belief” and “what justifies belief is evidence” (2000:184). Williamson summarises this equivalence of evidence and knowledge in the equation E=K. Since knowledge is propositional, for Williamson evidence must equally be propositional, that is any piece of evidence can be identified with a (set of) proposition(s) (2000: 194-200). For a critical discussion of Williamson’s view see (McCain 2014: 12-18).

¹¹ For defences of Factivity see Littlejohn (2013) and Williamson (2000). For objections to Factivity see Schroeder (2008), Comesana and Kantin (2010), Comesana and McGrath (2014) or McCain (2014: 12-27).

¹² For similar distinctions see Turri (2009), or McCain (2014: 11).

¹³ Here “+” is meant to indicate affirmation and “-” rejection of the respective theses.

global sceptical scenarios (Kelly 2008). The underlying intuition is that one's beliefs about the external facts would be no less justified, or rational "even if one were a recently envatted brain or the plaything of a Cartesian evil demon" (Kelly 2008: 952; see also McCain 2014: 10-11, or Sillins 2005). The envatted brain has no access to ancient manuscripts, but it may possess mental representations of alleged ancient manuscripts and forms its beliefs about Aristotle based on them. So, we want to distinguish a careful and judicious but unlucky reasoner such as your brain in a vat duplicate from improper reasoners, which form beliefs on a whim, for no reason at all, or via wishful thinking. Let us call this the *Internalist Intuition*:

Internalist Intuition: Epistemic agents can have justified beliefs independently of whether they are in a lucky or an unlucky case.

The scenario described above that fuelled the Internalist Intuition has come to be known as the new evil demon problem (Cohen & Lehrer 1983). But there are other often-discussed anti-externalist thought experiments such as Bonjour's (1980) reliable clairvoyant or Lehrer's (1990) Mr. Truetemp that can be brought to bear to further support the Internalist Intuition. These latter arguments support the Internalist Intuition by showing that there are cases in which agents are lucky (to fulfil the required externalist conditions) but do not properly form their beliefs and hence are not justified. That is, while the new evil demon is a powerful argument against the *necessity* of externalist criteria those arguments seek to demonstrate that externalist criteria are also not *sufficient* for justification.¹⁴ While I will not expand on these foundational issues here, I will come back to them later, when I defend the idea of higher-order requirements on justification [ch.8].

The Internalist Intuition fits into an evidentialist framework, so long as we have the right understanding of evidence. Here we can turn either to mentalism or abstractionism. However, while mentalism or abstractionism tell us what kind of entity evidence is, they do not tell us anything further about which type of mental state we need to consider evidence (or turn to in order to pick out the propositions that are evidence). As a first approximation, we may distinguish factive from what Conee and Feldman (2008), as well as McCain (2014), call non-factive mental states. Those are the states that I and my demon-world counterpart share. For example, seeing that there is a coffee cup on my table is a factive mental state because I can only be in the mental

¹⁴ See, for example, Berghofer (2022: 224-226) or Goldberg (2018: 42, 50-52).

state of seeing that there is a coffee cup on my table when there actually is a coffee cup. In contrast, seeming to see that there is a coffee cup is a non-factive mental state. I can seem to see that there is a coffee cup even though I am in fact tricked by an evil demon.

So far so good. However, not every non-factive mental state can be considered evidence. As McCain (2014: 10-11) notes:

“Of course, not every non-factive mental state is part of one’s evidence. Imagining a tree and desiring a tree are non-factive mental states directed toward there being a tree, but they are not evidence that there is a tree. The non-factive mental states that [...] counts as evidence are ones that represent the world as being a certain way—things such as beliefs, introspective experiences, perceptual experiences, memorial experiences, and perhaps others such as intuitions and rational insights.”

So, for McCain (2014), evidence is a representational non-factive mental state, where the most plausible candidate of such states are all kinds of experiences, intuitions, rational insights as well as beliefs. This is compatible with Conee and Feldman’s doctrine that we should distinguish between “ultimate” and “intermediate” evidence, whereas “all ultimate evidence is experiential” (2008: 88).¹⁵ So, beliefs can only be evidence for anything other than the fact that we have them if they are themselves justified. To be justified it must be a belief for which we have sufficient evidence. This evidence consists “of other justified beliefs or experiences, but if we trace back the evidence far enough, it is reasonable to think the evidence will bottom out in experiences of some sort” (McCain 2014: 19).

In sum, if justification is determined by the evidence and the evidence for the envatted and the embodied brain are the same, then the representational non-factive mental states (mentalism) or their propositional content (abstractionism) are the only plausible candidates for evidence. Furthermore, on either of those views ultimate evidence is experiential, so experiential non-factive mental states or their propositional content. While there could be further representational non-factive mental states, such as justified beliefs which can be considered evidence, these states are only intermediate evidence since they ultimately rely on other evidence to gain their justificatory status.

¹⁵ In this context Conee and Feldman (2008) make it explicit that epistemological evidence differs from what they call scientific evidence. For them “scientific evidence is something like the publicly accessible reliable indicators that are recognised by science. Again, we think this has an understandable connection to the justifying evidence. The scientific evidence is particularly useful in acquiring justifying evidence. But the two are clearly different.” (2008: 101).

Note that, strictly speaking, both of these views are compatible with Propositionality. That is, on both of these views, pieces of evidence are propositions. So, for instance, when I am forming a belief that there is a laptop in front of me via visual perception my evidence for that belief is either the proposition that I am having a non-factive experiential mental state of a laptop in front of me (mentalism) or the proposition that it seems like there is a laptop in front of me which is the content of a non-factive experiential mental state of mine (abstractionism).¹⁶ Furthermore, in some sense, both views also are compatible with Factivity since both propositions are true. In the former case, it is the true proposition that I have the mental state of seeming to see a laptop in front of me. In the latter case, it is the true appearance proposition that I seem to see a laptop in front of me, which is the content of my experiential mental state.¹⁷

In sum, taken together evidentialism and the Internalist Intuition require a specific version of the mentalist or abstractionist understanding of evidence. Both of these views are live options for the internalist evidentialist and we do not need to choose between them here (McCain 2014: 20; Lutz 2020). Instead, we can postpone this choice until we have discussed epistemic support in [§5.5], where I will argue that the abstractionist version of these two available understandings of evidence fares better with respect to my preferred understanding of epistemic support. Let me call this version of abstractionism, Statism:¹⁸

Statism: E is an ultimate piece of evidence only if it is the propositional content of a representational non-factive mental state M.

Note that Statism is compatible with internalist mentalism since it allows us to subscribe to the above-mentioned mentalist version of the supervenience thesis [§5.2]:

Mentalist Doxastic Supervenience: The justification of any epistemic agent's doxastic attitude toward any proposition at any time strongly supervenes on the mental states that the person has at the time as well as the causal relations that obtain between the person's mental states.

¹⁶ Note that this implies that the propositional content of my experiences that there is a laptop in front of me is not the proposition that there is a laptop in front of me but the proposition that it seems to me that there is a laptop in front of me.

¹⁷ For an understanding of evidence that comes close to this version of mentalism see Lewis (1996).

¹⁸ The term 'statism' is taken from Turri (2009), who uses it for a non-propositional mentalist view.

With these considerations about the nature of evidence in mind, we can now analyse what it means for an epistemic agent to *possess* evidence.

5.4 Evidence Possession

Having an answer to what it means to possess evidence is crucial for internalist evidentialists, such as Conee and Feldman or McCain. After all, for them doxastic justification depends on propositional justification which depends on the *entire* evidence a person has or *possesses* at any given time. While I will argue that this is mistaken, it is nonetheless helpful to give an overview of the general dialectic of debates concerning evidence possession; especially since it will help us to reflect on issues arising in the context of collective epistemology [§5.7].¹⁹

Let us call the total evidence an epistemic agent possesses at a given time their evidential base (in short: $\langle E_{\text{BASE}} \rangle$).

Evidential Base $\langle E_{\text{BASE}} \rangle$: The total evidence *possessed* by an epistemic agent S.

It seems like that independently of which of the above-mentioned ontologies of evidence we subscribe to we already have a simple answer to the question of how it is that we possess evidence. We possess evidence by having mental states that either are evidence themselves or that have propositions that constitute our evidence as part of their content. So, in some important sense all of the evidence an agent *possesses* $\langle E_{\text{BASE}} \rangle$ is or is picked out by all of the relevant mental states this agent has.

However, many have argued that for something to be part of one's evidence it needs to be part of one's epistemic perspective. Let us call this *perspectivalism* about evidence.²⁰ One often-made perspectival constrain is doxastic in nature (Brown 2022b: 165):

Doxastic Constraint: E is part of S's evidence if and only if S bears doxastic relation, D, to E.

Proponents of doxastic constraints hold that the evidence possessed by an epistemic agent is the evidence that they have a certain doxastic relationship with, such as knowledge or justified belief

¹⁹ In particular, I will argue that cases in which some possessed (in the intuitive understanding of the term) evidence conflicts with the evidence one uses as a basis for one's belief are similar to defeater cases [ch.6].

²⁰ The label is taken from Brown (2022b: 164).

(Brown 2022b: 164-165). That is, for instance, the set of propositions that are known (Williamson 2000; Bird 2007; Hedden 2019), believed and true (Mitova 2014), or non-inferentially justifiedly believed (Goldman 2009; Littlejohn 2013) by the respective agent.

However, there are some problems with proposing such doxastic constraints in our internalist evidentialist framework. First, the required doxastic relation cannot be knowledge or true belief since this would require a factive understanding of evidence which is in conflict with the Internalist Intuition. Second, even if the doxastic constraint is non-factive we cannot treat it as a general constraint. Even if we count some justified beliefs to be part of one's evidence those beliefs can only be evidence in the intermediate sense [§5.3]. This is also pointed out by Beddor (2016) who argues that if the evidentialist puts a justificatory constraint on evidence possession they run into a circularity problem. After all, evidentialism is a theory that aims to *explain* justification in terms of evidence and the relation doxastic states bear towards it. So, by putting a justification requirement on evidence possession, we are appealing to the notion of justification to explain what counts as evidence for someone, which amounts to circular reasoning.²¹

Another, often-discussed constraint compatible with internalist evidentialism has to do with the *accessibility* of mental states. Here some argue that possessing some evidence E requires being presently conscious of E or being able to recall E in a suitable way. This gives us an accessibility constraint of the following form:

Accessibility Constraint: E is part of S's evidence if and only if S bears accessibility relation, A, to E.

This constraint could take more or less restrictive forms. Feldman, for example, has argued that "S has p as available evidence at t iff S is currently thinking of p" (2004: 232), while McCain has argued that "S has p available as evidence at t iff S is currently aware of p or S could easily access a memory of p" (2014: 33). Independently, of which accessibility constraint we prefer we face the hard question of spelling out what it means to be able to access some evidence. Are simple implications of your beliefs that you are not currently thinking about accessible in the right kind of sense? –and, what about deeply buried memories that could be recovered only through years of extensive psychoanalysis?

²¹ See also Fratantonio (forthcoming: §2.2).

Luckily, we do not need to enter this debate here. Instead, in my view, getting a grip on what it means to possess some evidence is only of secondary importance. The reasons for this are twofold. First, we can refine the possession relation indirectly by thinking about what it means to base one's belief on some evidence; where for one to be able to base one's belief on some evidence one needs to possess it and *vice versa* [§5.6]. Second, as we will see in the next chapter [ch.6], my view cuts the orthodox understanding between propositional and doxastic justification; which takes the former to be a necessary precondition for the latter. So demarcating possessed from unpossessed evidence that is currently not used when forming our beliefs is not important. All that matters is whether some, to be spelt out [ch.7], no-defeater condition obtains towards evidence relevant to a situation (which might be possessed or not).

In sum, we can move on with the intuitive understanding of the evidential base $\langle E_{\text{BASE}} \rangle$ as all of the propositions picked out by the relevant non-factive representational mental states had by S. Nonetheless, in so doing, we have left various questions unanswered that arise specifically when thinking about the evidential base of collective epistemic agents. Since collective epistemic agents are (partly) composed of multiple individual epistemic agents, we need to specify whose evidential base we have in mind when talking about the group's evidence. I will address these issues below [§5.7] after I have answered the remaining fundamental questions for evidentialism; that is, after I have sketched my preferred theory of epistemic support [§5.5] and epistemic basing [§5.6].

5.5 Epistemic Support

Having settled on an understanding of evidence [§5.3] and evidence possession [§5.4] enables us to turn to the question of what it means for a piece of evidence to support a given proposition or belief [§5.5] and what it means for an epistemic agent S to base their belief on that evidence [§5.6]. These issues are represented in clauses (1a) and (1b) of Continuous Evidentialism.

- (1) S possesses some evidence E, which is (a) *sufficient to support* Bp, and (b) Bp is *properly based* on E.

E is the segment of $\langle E_{\text{BASE}} \rangle$ that serves as *the basis* for the respective belief; i.e., it is the part that is used when forming the belief. This *operational part* of the evidence is sometimes called the “relevant evidence” or “source” of the belief that p (Swain 1979: 25; Korcz 2000: 525).

Operational Evidence: Some evidence E is *operational* concerning a doxastic attitude D iff it is used as the basis for D when D is formed.

While I will turn to the question of epistemic basing in the next section [§5.6], any operational piece of evidence E needs to be *sufficient to support* Bp (1a) to have the potential to make Bp justified. So, for example, my belief that there is a copy of Lackey’s *The Epistemology of Groups* on my bookshelf seems to be sufficiently supported by my belief that I put it there yesterday. But why and how does this memory support my belief? In Conee and Feldman’s words, “if perceptual and memorial experiences are justifying (with the proper background in place), then, there is something about them that makes this the case. [...] There must be a more illuminating truth about why the experiences are justifying” (Conee & Feldman 2008: 97). That is, any complete evidentialist theory of justification needs to be supplemented by a theory of epistemic support. A theory of the epistemic support relation between evidence and doxastic attitudes.

There are various theories of epistemic support compatible with evidentialist theories of justification, including probabilistic, normic, explanationist, subjective, phenomenological, as well as hybrid views.²² Which view one subscribes to partly depends on one’s understanding of evidence. For example, probabilistic accounts are preferred by those who understand evidence to be propositional, while explanationist accounts are often said to be compatible with non-propositional mentalist understandings of evidence. While I will discuss both probabilistic [§5.5.1] and explanationist [§5.5.2] views of support, as well as a hybrid approach [§5.5.3] below, let me start with some general observations.

The first thing to note is that epistemic support is governed by what I have called epistemic standards [ch.1]. As I have pointed out epistemic standards govern the *epistemic support relation*, where epistemic support has two dimensions, the *evidential support relation* and the *doxastic support relation*. More precisely, epistemic standards tell you what your evidence supports and then once you have a fix on the evidential support relation, they tell you what you ought or are permitted to believe given your evidence. Hence, to clarify clause (1a) we need to get a grip on the evidential and the doxastic support relation.

²² For an overview see Fratantonio (forthcoming: §2.3) or Conee and Feldman (2008: 94-98). For explanationist views see Lipton (2004), Conee and Feldman (2008), McCain (2014), Conee (2020), or Lutz (2020). For a phenomenological account see Berghofer (2022). Hybrid probabilist-explanationist accounts, similar to the one I consider below, are discussed by Lutz (2020) and defended by Hedden (2015a).

This makes epistemic support a 3-place relation between the evidence, the supported doxastic attitude and the epistemic standards. Understanding epistemic support this way comes with potential problems for those who want to defend an internalist version of evidentialism that is committed to the Mentalist Doxastic Supervenience thesis:

Mentalist Doxastic Supervenience: The justification of any epistemic agent's doxastic attitude toward any proposition at any time strongly supervenes on the mental states that the person has at the time as well as the causal relations that obtain between the person's mental states.

For the Mentalist Doxastic Supervenience thesis to be true on this 3-place relational understanding of support epistemic standards need to be something that can be possessed by epistemic agents similarly to how evidence is possessed by epistemic agents. Or, more precisely, not only evidence but also epistemic standards need to supervene on the mental lives of epistemic agents. So, epistemic standards cannot be something external to the epistemic agent or something internal but non-mental.²³

5.5.1 Probabilism

Most people think that rational agents respect their evidence because evidence indicates truth or likelihood. Seeing that there is a wet coffee cup stain on my desk is evidence E for the proposition p , that someone recently placed a coffee cup there. E is evidence for p because it makes it more likely that p , or at least makes p more likely than it would have been without E . This is encapsulated in the following principle:

Probability Raising Requirement: E is evidence for p iff $P(p|E) > P(p)$.

In a Bayesian framework, this Probability Raising Requirement can be understood as governing the relation between the conditional and unconditional degrees of belief. So, E supports p if the conditional credence of p given E $P(p|E)$ is higher than the unconditional credence $P(p)$. While

²³ Note that this understanding of epistemic standards fares well with the divergence argument discussed in [ch.3], where I argued that groups and their members could have different rational requirements because they *possess* the same evidence but different epistemic standards.

Bayes Theorem gives us the relationship between the conditional probabilities $P(p|E)$ and $P(E|p)$ and unconditional probabilities $P(p)$ and $P(E)$:²⁴

$$\text{Conditional credence: } P(p|E) = \frac{P(E \& p)}{P(E)} \quad \text{Bayes Theorem: } P(p|E) = P(p) \frac{P(E|p)}{P(E)}$$

However, this does not tell us when and how E supports having a specific *doxastic attitude* with respect to p. After all, the posterior probability, even if higher than the prior probability, might still be very low. To do so we need some assumptions about doxastic support. So, to make E strong enough to make B_p justified, we need to define some threshold (Lutz 2020: 2639; Fratantonio, forthcoming: §2.3.1):

Threshold Requirement: E is sufficient to support S's belief that p, only if E is evidence for p and $P(p|E) > t$.

Whatever this threshold t is, the Threshold Requirement captures the idea that there needs to be some threshold (while any plausible candidate is probably >0.5). (Note that, a threshold might not be required if we are probabilists about beliefs, i.e., model beliefs as credences).

Note that this probabilist understanding of epistemic support is committed to a propositional understanding of evidence since what we are updating on when conditionalising are propositions rather than mental states. We ascribe prior probabilities to propositions so only when we learn that some proposition is true we can draw consequences for the likelihood of some other propositions. Further, given prior internalist commitments the only propositions we can update on are appearance propositions, the kinds of propositions that are described by non-factive representational mental states. This is my preferred understanding of evidence which I call Statism [§5.3]. For instance, when I have an experience of a coffee cup stain on my desk we can only come to justifiably believe that someone has recently placed a coffee cup there by conditionalising on the proposition that it appears that there is a coffee cup stain, which –let's stipulate– raises the probability that there is a coffee cup stain, which raises the probability that someone has recently placed a coffee cup there.²⁵

²⁴ For an introduction into a Bayesian account of evidential support see, e.g., Bradley (2015: 5-13).

²⁵ This might avoid various problems propositionalists about evidence are often said to face. Namely the circularity problem (Turri 2009: 497-499, McCain 2014:19), as well as problems associated with cases involving experiences which are different considering their vividness but in which the propositional content is allegedly the same (Dougherty 2011; McCain 2014: 18-19).

Note that as I pointed out above [§5.3], I think that there is a way to spell out a mentalist understanding of evidence which is compatible with a propositional understanding of evidence. On this version of mentalism, the proposition that I have the non-factive mental state that there is a coffee cup stain on my desk is the evidence for my belief rather than the mental state itself.²⁶ That means that on both Statism, which is a version of abstractionism, or this version of mentalism evidential support is a relation between propositions, the considered proposition and the supporting proposition. Hence both views are compatible with a probabilist theory of support. However, there is a crucial difference between those two views which motivates my preference for Statism; namely that according to Statism evidential support is a content-to-content affair, while on the mentalist alternative, it is not. Instead of having to say that the true proposition that I am in the non-factive mental state M is what supports my belief I can say that the *content* of the non-factive representational mental state M supports the *content* of my belief.

5.5.2 Explanationism

Another popular understanding of evidential support is explanationism (Lipton 2004; Conee & Feldman 2008; McCain 2014; Conee 2020; Lutz 2020).²⁷ Explanationism is the view that evidence supports propositions whenever the best, or some sufficiently good, explanation for E includes p. In the case of perceptual experiences Conee and Feldman write, “perceptual experiences can contribute toward the justification of propositions about the world when the propositions are part of the best explanation of those experiences that is available to the person” (Conee & Feldman 2008: 97). That is, the proposition p, that someone recently placed a coffee cup on my desk, is supported by E, me having an experience of a fresh coffee cup stain on my desk, because p is sufficiently explained by E.²⁸

While explanationism tends to go under the name ‘inference to the best explanation’, Lipton (2004) has pointed out that this name might misleadingly suggest that E can only support p if p is the best possible explanation of E. But this support can also go *vice versa* since if p is true E

²⁶ Lewis (1996) seems to defend such a mentalist view of evidence.

²⁷ In so doing, I take explanationism to be not only compatible but also to be more fundamental than probabilism. See Lutz (2020: 2640-2643).

²⁸ This is to reject a primitivist view about evidential support (Hedden 2015a). For the primitivist facts about evidential support are brute facts that are not determined by other, explanatorily prior facts. As Hedden points out, on these views, we “would still be able to give rules of thumb or *ceteris paribus* principles (e.g. simpler hypotheses tend to be better supported by the evidence) for getting a grip on the evidential support relation, but there aren’t further facts that ground the evidential support relation” (2015a: 724; *emphasis added*).

would be explained by *p*. Lipton calls this ‘inference *from* the best explanation’. More generally, while explanations look asymmetric, evidential support is symmetric since propositions if supportive are always mutually supportive (Hedden 2015a: 732). Notably, this is true, independently of whether we think that evidence is some kind of mental state or the propositions picked out by them. Accordingly, some speak of explanatory coherence (Conee & Feldman 2008: 98; see also Lutz 2020: 2632; McCain 2014: 63):

“The best available explanation of one’s evidence is a body of propositions about the world and one’s place in it that make best sense of the existence of one’s evidence. This [...] can be equally well described as fitting the presence of the evidence into a coherent view of one’s situation. The coherence that justifies holds among propositions that assert the existence of the non-doxastic states that constitute one’s ultimate evidence and propositions that offer an optimal available explanation of the existence of that evidence.”

But what does it mean that *p* is (part of the) best explanation for *E*? Here the traditional answer is that taking *E* to explain *p* (and *vice versa*) would maximise various explanatory virtues such as adequacy, coherence, unification, parsimony, elegance, etc. (McCain 2014: 66; Hedden 2015a: 735; Berghofer 2022: 79; Fratantonio forthcoming: §2.3.2). That is, the fresh coffee cup stain *E* supports that *p*, someone has recently put a coffee cup there rather than *q*, a conspirator who is trying to trick me has broken into my office and painted a coffee cup stain on my table, because the former would be a better explanation than the latter. Here *p* is a better explanation than *q* since, while both equally explain the data *p* is more parsimonious (no need to postulate a conspirator) and more coherent (no need to change any of my background beliefs about coffee cup stains) as well as more unificatory (it gives me a general explanation of the phenomenon of coffee cup stains) and so on.

5.5.3 A Hybrid View?

My favoured answer would be a combination of the explanationist approach [§5.5.2] with the probabilistic constraints formulated above [§5.5.1]. Here the symmetrical understanding of explanatory coherence fits nicely with the symmetrical understanding of support argued for by the probabilist; whereas *E* supports *p* only if *p* supports *E* as well.²⁹ So we can interpret the Bayesian formalism by supplementing it with a more fundamental explanationist understanding

²⁹ See, for example, Bradley (2015: 10).

of evidential support. According to explanationism p is likely, given E , when E tends to explain p , and p receives greater support when the evidence that supports it is surprising (Poston 2014: 151-152). That is, the posterior probability of a proposition will be positively correlated with how surprising the new evidence is and how strong the explanatory connection between the evidence and the proposition is (Lutz 2020: 2642). This highlights another attractive feature of explanationism. Namely, we can make straightforward sense of degrees of justification since explanatory connections and the correlated probabilities come in degrees.³⁰

In sum, this gives us the following explanationist understanding of epistemic support.

Explanationism: Some evidence E is *sufficient to support* the belief that p iff: (i) E is part of the best explanation for p , or (ii) p is a logical consequence of some sufficiently good explanation for p and (iii) E *increases* the *likelihood* of p above the believability threshold t .

Before, I link this explanationist understanding to my general understanding of epistemic standards [§5.5.4], let me make some clarifications about clauses (i) – (iii).

First, Explanationism gives us a theory of when some evidence supports a proposition and when this makes the proposition believable. That is in accordance with the existence of some threshold principle and Lipton’s slogan that “Inference to the Best Explanation might be more accurately if less memorably called ‘Inference to the Best Explanation if the Best is Sufficiently Good’” (2004: 154).

Second, clause (ii) is added by McCain (2014) as a reply to persuasive arguments by Lehrer (1974) and Goldman (2011) that logical relations are not explanatory relations. However, and this connects to the third point below, in contrast to McCain (2014), I do not believe that p needs to be available to S as a logical consequence for E to support p . E supports p evidentially, regardless of whether S acknowledges or sees that p is a logical consequence. While, there is some inferential connection to be established, in my view this connection is established via a proper basing requirement [§5.6] as well as a higher-order requirement [ch.7] [ch.8].

Third, and most importantly for our investigations, Explanationism does not by itself require that S draws any inference to the best explanation. Rather Explanationism merely says that what

³⁰ As argued by Conee and Feldman (2008: 98) this is also true for non-hybrid explanationist understandings. For a critique see Berghofer (2022: 78-80).

matters for evidential support is that the required explanatory relations hold between the evidence E and the target proposition p .³¹

5.5.4 Subjectivism vs. Objectivism

A common worry is that Explanationism, and with it the proposed hybrid view, make the epistemic support relation purely subjective (Berghofer 2022: 79; Fratantonio forthcoming: §2.3.2; Lutz 2020). After all, many have argued that there is no objectively correct way to weigh and balance off theoretical virtues against each other (Kuhn 1977: 320-339; Longino 1996). So, different weightings and different catalogues of theoretical virtues give us different understandings of better or worse explanations. If then any catalogue and possible ranking of virtues is permitted, we end up with a radical subjectivist notion of epistemic support.

Given that in the proposed hybrid view explanatory constraints ground probabilistic constraints, this would lead to a position that looks like extreme Bayesian subjectivism. One can believe anything as long as one's credences form a probability function and one updates by conditionalising. That is, there are no non-probabilistic restrictions on the priors. However, what fuels the subjectivism, on the proposed analysis of epistemic support is subjectivism with respect to explanatory virtues, while the standards one subscribes to fix the priors (Hedden 2015a: 718).³² In a slogan: there is no freely picking the priors without picking the explanatory standards first.

On the other hand, if one believes that there are some or complete restrictions on the explanatory standards, we are permitted to have, one ends up with a position that looks like moderate subjectivism or even objectivism about evidential support. Following Hedden (2015a; 2015b) I do not think that the project of objectivism about explanatory standards is hopeless. That is, there is hope that there are some *a priori* considerations that might lead us closer towards those standards.³³

³¹ This is pointed out by Fratantonio (forthcoming: §2.3.2) and Hedden (2015: 720-721). Interestingly, however, McCain explicitly states that “seeming that p is part of the best explanation available to S at t for why S has the evidence she does” (McCain 2014: 78). This is problematic since it makes Explanationism vulnerable to charges of overintellectualisation and potentially triggers a regress (Appley & Stoutenburg 2017; Lutz 2020: 2635-2637; Berghofer 2022: 78). I will discuss those worries extensively in the context of higher order requirements of justification [ch.8].

³² Here is one way to look at this: in fixing the priors $P(E)$, $P(p)$ and $P(p \& E)$ explanatory standards also indirectly fix the conditional probability $P(p|E)$ and with it the evidential support relation.

³³ Hedden (2015a) points out that only *a priori* considerations could in principle help us settle what the objectively correct epistemic standards are since following considerations about immodesty (Lewis 1971: 56; Titelbaum 2010), empirical considerations can never settle the question of which explanatory standards to use.

Instead of subscribing either to subjectivism or objectivism at this point, let me say how this connects to my general view of epistemic standards. In chapter [ch.1], I took the observation that there are allegedly permissive cases as a primitive. Afterwards, I explained their structure and dynamics by highlighting that they have various features of choices under incommensurability. This supported the hypothesis that in these cases we face epistemic standard conflicts, rather than conflicts in which one set of standards equally supports multiple doxastic attitudes. Now, I remained neutral on the permissibility of these conflicts. Equally, I can remain neutral on Explanationism giving us an objective or subjective epistemic support relation. As with epistemic conflicts, it all boils down to metaphysical questions about the nature of incommensurability which I have not taken a stance on [§1.5]. If incommensurability is explained by a genuine fourth comparative relation, epistemic rationality is permissive and we get subjectivism about epistemic support and if incommensurability is explained by indeterminacy, we get weak permissibility (i.e., impermissivism with blameless divergence within the indeterminate range) and subsequently objectivism.³⁴

5.6 Epistemic Basing

Now that we have some understanding of evidence [§5.3], defined the evidential base $\langle E_{\text{BASE}} \rangle$ [§5.4] and established what it means for a piece of evidence to be sufficient to support a doxastic attitude [§5.5], we can take a closer look at the *epistemic basing relation* (1b).

- (1) S possesses some evidence E, which is (a) *sufficient to support* Bp, and (b) the belief that p is *properly based* on E.

The epistemic basing requirement is often expressed as follows (Turri 2010: 313; Silva 2015a: 946):

Basis: S's belief that p is *doxastically* justified iff (i) the belief is *propositionally* justified for S in virtue of S's having evidence E for p, and (ii) S believes p *on the basis of* E.

As noted above the basing relation enables us to distinguish between propositional and doxastic justification [§5.2]. Following Basis, doxastic justification is explained in terms of propositional

³⁴ On the rather subtle relationship of objectivism and subjectivism, on the one hand, and impermissivism and impermissivism on the other, see Hedden (2015a: 716-719).

justification, which in turn should be explained by a relationship between the respective doxastic attitude and the evidential base. If S believes that p and S's overall evidence $\langle E_{\text{BASE}} \rangle$ is sufficient to support p, S's belief that p is propositionally justified. On this view, some epistemic agents can fail to have a doxastically justified belief that p if they either fail to have sufficient evidence for p or fail to base their belief that p on this evidence. While my view on the relationship between doxastic and propositional justification differs from this traditional analysis (see [§6.7]) I think we can use this traditional analysis as a helpful heuristic to highlight the importance of the epistemic basing relation.

One prominent theory of basing understands the epistemic basing relation as a causal relation between the relevant evidence and the doxastic attitudes. According to such causal theories of basing, beliefs that are based on a particular piece of evidence are caused by it in an appropriate way.³⁵ The other prominent account of epistemic basing is doxastic, stating that having an appropriate meta-belief about the evidence is what it means for the belief to be based on that evidence, i.e., for an epistemic agent S to base the belief p on the evidence E, S needs to believe (or have the disposition to believe) that E provides evidence for p.

There are persistent problems for both understandings of basing. The most prominent objection to a causal understanding of basing is the problem of so-called *deviant* causal chains (Davidson 1963). For example, if my belief that I am late for the seminar causes me to hurry, which causes me to stumble and hurt my already injured knee, which in turn causes me to form the belief that I should go to the orthopaedist. In cases like these, we might say that my belief that I am late for the seminar caused my belief that I should go to the orthopaedist, but not that the latter is based on the former (see also Plantinga 1993: 69; or Brown 2022a: 13). While there are various promising answers to the deviancy problem, there is no agreed-upon solution.³⁶ The

³⁵ While being indecisive about the role of the causal nature of epistemic basing, in their evidentialist framework, Conee and Feldman seem at least inclined to think that causality is a necessary requirement of basing, when claiming that “it is unclear whether one can appreciate the evidence without being caused to have the belief by the evidence” (2004: 93). On another occasion they consider being caused by the relevant evidence at least as one sufficient way of basing (they speak of well-founding) a belief. For them, “it might be held that a basis for a belief is a mental cause of the belief. [...] We need not single out one mental cause of a belief as the relevant one. The idea used [...] is that one has a well-founded belief if it is believed on some basis or other that is epistemically okay. Presumably the basis is a cause of the belief” (Conee & Feldman 2004: 163).

³⁶ A promising set of solutions is built on the notion of cognitive dispositions. For example, Turri (2011) argues that some belief non-deviantly causes another iff the causing manifests a cognitive disposition of the agent. Others argue that for one belief to cause another, we need to have the disposition to revise the belief if we lose the former belief (or the former belief loses its normative status) (Lord & Sylvan 2019).

following investigations, therefore, rest on the assumption that we can give a suitable account of what it means to non-deviantly cause a belief.

There are multiple reasons why I will rely on a causal rather than a doxastic understanding of basing. First, I consider some of the arguments against higher-order doxastic requirements, which I discuss in [ch.8], to render doxastic basing requirements highly problematic.³⁷ Most importantly, such requirements do not only overintellectualise epistemic justification but also possibly induce an infinite regress. Furthermore, following Brown (2022a), I will demonstrate in [§5.7] that the causal understanding of basing is better equipped to make sense of basing in group reasoning. Groups often have various mechanisms in place that establish a causal relationship between their attitudes and the group's evidence. Then the causal relationship can be realised in multiple ways, such as through explicit deliberation of the evidence among group members or through a division of epistemic labour (Brown 2022a: 14).³⁸ In these cases it makes sense to say that the group bases their belief that *p* on some or multiple pieces of evidence possessed by different members without the group having any higher-order belief that those pieces of evidence support the formed belief.

In sum, for an epistemic agent to be doxastically justified to believe that *p* based on *E* the belief must be based on, which is to say non-deviantly caused by, *E* (1b), and *E* needs to be sufficient to support *p* (1a). If we acknowledge that the belief that *p*, as well as the mental state *M* that picks out the evidence *E* which serves as the basis for *B_p*, are mental states, the basing relation is a causal relation between the mental states of the epistemic agent. This explains the causal aspect of the mentalist supervenience thesis I subscribe to:

Mentalist Doxastic Supervenience: The justification of any epistemic agent's doxastic attitude toward any proposition at any time strongly supervenes on the mental states that the person has at the time as well as the causal relations that obtain between the person's mental states.

5.6.1 Proper Basing

With this causal understanding of basing at hand, we may now investigate what it means for a belief to be (1b) *properly* based on some respective evidence *E*. It is generally acknowledged that basing, as described in Basis is a necessary, but not sufficient condition for a belief to be

³⁷ For an up-to-date discussion see Korcz (2021) or Neta (2019).

³⁸ For a discussion of distributed cognition cases see Hutchins (1995a), Bird (2010; 2014), or Habgood-Coote (2019).

doxastically justified (putting issues about the relationship between propositional and doxastic justification aside). First, for E to be a proper basis for Bp, E in total needs to be *proper* evidence. That is, given our understanding of evidence [§5.3], E solely needs to consist of non-factive experiential mental states (ultimate evidence) or beliefs that are justified by being based on that ultimate evidence.

To demonstrate this we can consider the following example inspired by Turri (2010b: 315-316):³⁹

INFECTED BASE: A scientific researcher S possesses evidence E that is, taken together, sufficient to support the belief that p. While S recognises that the combined evidence E (relative to S's epistemic standards) makes it overwhelmingly likely that p S still fail to form the belief that p. One day a tarot card reader visits S and tells her that, based on one of her recent tarot card readings, she has determined that accepting E makes it overwhelmingly likely that p. As a consequence, S forms the belief that p.

This example is designed to illustrate that beliefs, while being supported by sufficient evidence E can be unjustified when they are reached by relying partly on an improper piece of evidence. While S should have inferred p from the evidence E but did not, adding an improper piece of evidence E* (the tarot card reading) to E made S believe that p based on E* and E.⁴⁰ Therefore, INFECTED BASE demonstrates that we need to exclude any improper evidence from the operational part of the evidence E, even if E would be sufficient to support p.

We could fix this problem by adding some kind of proper evidence requirement, so that sufficient to support means not only that E fulfils the explanatory requirements, but also that E *solely* consists of proper evidence for p; whereby proper evidence means ultimate evidence or beliefs that are justified by that ultimate evidence. So, for the belief that p to be doxastically justified, only proper pieces of evidence are allowed to cause it.

However, merely excluding improper pieces of evidence from E is not enough either. To illustrate this, consider the following case of improper basing which does not include any improper evidence, again inspired by Turri (2010b: 317):

ILL-FORMED INFERENCE: Imagine juror S, considering the guilt of Mr Suspicious and Mrs Unsuspicious. S paid close attention throughout the trial and, as a result, possesses

³⁹ Similar examples are also found in Neta (2002: 671) and Pryor (2004: 365).

⁴⁰ Notice that this does not mean that E needs to be the minimal set of evidence being sufficient to support p. E can also contain many independently sufficient reasons for p, i.e., p can be epistemically overdetermined by E.

overwhelming evidence E that Mr Suspicious is guilty, and Mrs Unsuspicious is not. Additionally, S knows that –trivially– either one, none or both are guilty. Based on the evidence E, S correctly infers that Mr Suspicious is guilty. While S also has sufficient evidence that Mrs Unsuspicious is innocent, S reaches that conclusion via a widespread logical fallacy, called *affirming the disjunct*: for any A and B: $A \vee B$; if A, then $\neg B$.

While INFECTED BASE showed that there can be some piece of ‘infectious’ evidence within E that makes our otherwise properly based belief improper, ILL-FORMED INFERENCE shows that something can go wrong with the inference (or basing process more generally) itself that makes our otherwise properly based belief improper.

Note, that we cannot advocate the same argumentative strategy as in INFECTED BASE, which was resolved by excluding improper evidence from E. Doing so would treat beliefs about the inference rules used to be part of E. However, we shall not treat inference rules to be part of our evidence (even evidence about proper inference rules such as modus ponens). How problematic this kind of confusion is, was famously shown by Lewis Carroll (1895), who demonstrated that counting inference rules employed in an argument among the premises of the argument, leads to a vicious infinite regress.⁴¹

The general lesson to be learned from ILL-FORMED INFERENCE is that we cannot simply think of doxastic justification as something that is sufficiently met by having proper evidence and forming our beliefs based on that evidence. Instead, we also need to take the nature of the basing relationship into account; i.e., base something in *the right way*. A minimal modification of Basis which incorporates this insight is:⁴²

Proper Basis: S’s belief that p is doxastically justified iff (i) the belief is propositionally justified for S in virtue of S’s having evidence E *sufficient to support* p, and (ii) S’s belief that p is *properly based* on E.⁴³

⁴¹ For a discussion of this problem in relation to epistemic basing see Volpe (2017: 31-35). Note, furthermore that Carroll’s insight extends from inference rules to belief-formation mechanisms in general. If I think that perceptual beliefs are justified, we need to think that they are evidence for their own content without relying on, e.g., higher-order evidence about the reliability of perception. Doing so would again start an infinite regress. At least if we also demand that this higher-order evidence is supported by higher-order evidence about the evidence and so on. I will come back to this objection when discussing higher-order requirements of justification in [ch.8].

⁴² While many authors approve that there is something amiss with the orthodox understanding of doxastic justification there is currently no consensus on how to dismantle the problems raised by Turri and others. The given minimal modification Proper Basis is inspired by Silva (2015b: 954).

⁴³ Again, this differs from the traditional understanding of the relationship between propositional and doxastic justification [§5.5]. I will come back to this in [§6.7].

The two introduced cases of improper group reasoning, INFECTED BASE and ILL-FORMED INFERENCE showed that there are at least two distinct things that can go wrong in the orthodox understanding of epistemic basing: the belief while based on proper evidence could additionally be based on improper evidence, or the belief formation processes used when forming the belief based on the evidence could be improper. Accordingly, Continuous Evidentialism is formulated in a way that is sensible to exclude improperness from the basing relation between E and the supported belief (1b).

I cannot hope to settle the debate of what it means to properly base a belief here. Nonetheless, I think that for our current purpose, an intuitive understanding of proper basing is enough: *modus ponens* is proper and affirming the disjunct is not. This is a general normative fact governing the inferences and belief-formation processes of epistemic agents. Any complete theory of epistemic justification needs to account for that. While I will argue that the higher-order evidential requirement that I'll defend in [ch.8] brings us a step closer towards a full understanding of proper basing, I remain open to the possibility that a proper basing requirement can only be spelt out by appeal to the general epistemic norms.⁴⁴ For now, the reader may think of Proper Basis as a placeholder that can be occupied once an adequate theory of proper basing is available.

5.7 Continuous Evidentialism

Now that we have answered the four fundamental questions for evidentialism, we can see whether the proposed package of views is epistemically continuous in the desired way. It is probably best to start with an example:

ROCKET LAUNCH: Several days before launching a crewed spacecraft for a mission to the moon the executive committee of NASA sets up a teleconference with several teams of experts sharing their expertise and deliberating their evidence. There is a group of engineers presenting how the spacecraft is attached to the launch vehicle and how the heat shield protects the crew at atmospheric entry. A group of chemists explains how the fuel is burnt in a controlled manner to guarantee steady acceleration. A group of software developers show how the board computer calculates the trajectory of the vehicle and communicates with the flight control centre, and so on. Each group of experts, as well as every individual expert, has a good amount of evidence that their subtasks of the space mission will be successful. At the conference, this evidence is shared among the executive

⁴⁴ For some discussion on this see, for example, Broome (2013: 157-159).

committee, and based on that, every member, as well as NASA as a group forms the belief that it is safe to launch the mission as planned (=p).

To evaluate how NASA's belief that p can be justified by being based on evidence that is sufficient to support p we need to start by answering what evidence NASA possesses. So what is NASA's evidence for the belief that p? Is it all of the relevant evidence possessed by the experts? The evidence possessed by the subgroups of engineers, chemists and software developers respectively? Or the evidence possessed by NASA as a group directly? To answer these questions we need a theory of group evidence.

According to Hedden (2019), we can roughly distinguish two classes of theories about group evidence. On the one hand, a group's evidence may be conceived as the function of the members' bodies of evidence. For example, we may think that the group's evidence is all the evidence possessed by at least one of the members or only the evidence that is possessed by the majority of the group's members. Following Brown (2022b), we can call this class of views *evidence summativism*. Hedden distinguishes three versions of evidence summativism (2019: 587):

Pooled Evidence: If at least one member's evidence includes E, then the group's evidence includes E.

Majority Evidence: If a majority of members have evidence E, then the group's evidence includes E.⁴⁵

Intersecting Evidence: If each member's evidence includes E, then the group's evidence includes E.⁴⁶

Based on different views of what the individual's evidence includes we could specify each of these summative views further. For example, given Statism and Pooled Evidence we could say that the group's evidence E is exactly all the propositions that are picked out by the sum of all the relevant mental states of all members.

On the other hand, we might think that the group possesses an independent body of evidence over and above the evidence possessed by the members. That is, groups like individual epistemic agents directly possess a body of evidence that does not depend on any of their

⁴⁵ This understanding of group evidence is defended by Buchak and Pettit (2015).

⁴⁶ Hedden (2019) calls this understanding of group evidence common evidence. I changed the label to avoid any association with the often-discussed common belief/knowledge condition which plays an important part in joint commitment accounts of group attitudes (Gilbert 1987).

member's evidence. We can capture the general idea of these so-called *non-summative views* as follows:

Group-Level Evidence: E is part of the evidence of a group G iff it is evidence directly possessed by the group.

Of course, there are many possible non-summative understandings of group evidence, depending on the exact nature of group doxastic attitudes and the underlying understanding of evidence. Hedden, for example, considers a functional understanding of group attitudes, which serve as the evidential base of the group. In other words, the evidence possessed by the group “depends in part on what the group believes [...] which in turn depends not only on its members' attitudes, but also on how it is structured” (2019: 586).⁴⁷ Given Statism and Group-Level Evidence this would give us a version of non-summativism that treats the group's evidence as all the propositions picked out by the relevant mental states of the group, that is, all (justified) group-level beliefs as well as all other relevant non-factive mental states of the group (if there are any).

From the point of view of Continuous Evidentialism, there are good reasons to prefer a non-summative view of group evidence. First, summativist views about evidence as other kinds of summativism are discontinuous since they explain evidence possession in the collective case in a way that cannot be used to explain evidence possession in the individual case [S4.7]. Furthermore, given Statism and our commitment to non-summative understandings of group attitudes (as defended in [ch.3]), suggest that one group-level belief can be based on another group-level belief without there being any corresponding member-level belief that we could consider supporting evidence.⁴⁸

Simultaneously, however, there are also good reasons to think that non-summativism cannot be the entire story. If we restrict the group's evidential base solely to group-level attitudes we thereby restrict the group evidence to anything that the group is concerned with and has beliefs (or other mental states) about. However, arguably the evidence possessed by some of the group's members can often play a crucial role when determining the epistemic status of the group attitude, even if the group doesn't form the respective attitude. Consider the following example:

⁴⁷ Brown (2022b) after raising various problems for both summative understandings (some of which I'll mention below) follows Hedden (2019) in defending a non-summative view of group evidence.

⁴⁸ Similarly, Brown (forthcoming: 2-18) has recently argued for various linkage claims between non-summative and understandings of group attitudes and group evidence and summative understandings of group attitudes and evidence.

HIRING COMMITTEE: Three of the four members of a committee appointed by the philosophy department to hire an assistant professor have a justified belief that the vending machine in the hallway is broken. All three of them failed to get some coffee from the vending machine before they had their meeting.

Does the committee have evidence that the vending machine is broken? It is often argued that it does not since the committee is a chartered group designed to form beliefs relevant to the hiring process and not about other things such as the vending machine.⁴⁹ Intuitively, however, it seems that the committee has evidence that the vending machine is broken that could support a respective belief if it would be formed, even though it appears to be irrelevant to the group's chartered purpose.

First, any proposition can in principle be relevant to deliberation about any other proposition. In other words, there are counterfactual scenarios in which the evidence about the broken vending machine would be relevant to the group's chartered purpose. For example, if the group knows that non-caffeinated department members are more likely to mess up in simple procedural tasks such as assigning the right CV to the anonymised writing sample, the member's evidence that the coffee machine is broken affects the epistemic state of all kinds of beliefs about the hiring process.

Second, groups arguably can form beliefs about propositions irrelevant to their chartered purpose. Even if the committee's purpose is hiring an assistant professor, they could deliberate seemingly irrelevant propositions and form robust attitudes about them. While this might be considered a professional or procedural defect of the group, this malfunctioning does not change the fact that these group attitudes can be epistemically justified.⁵⁰

I will argue for a hybrid understanding of (group) evidence possession that combines summative and non-summative elements:⁵¹

Evidential Base $\langle E_{\text{BASE}} \rangle$ [continuous version]: The evidential base, which is the total evidence possessed by an epistemic agent S, is exactly the sum of the propositions

⁴⁹ See for example Habgood-Coote (2019).

⁵⁰ The same is true for individuals and other epistemic agents. We may say that a human brain was not designed to have beliefs about e.g., the stock market, video games, or the completeness of Peano Arithmetic, but we still do have (justified) beliefs about those things. Or think of an AI (capable of having doxastic attitudes) designed to solve a particular problem that as a result of malfunctioning forms beliefs about propositions irrelevant to its intended purpose. I want to thank Andrew Peet for pointing this out.

⁵¹ Another important difference is that Hedden's understanding of evidence is factive (2019: 584-585).

described by the relevant non-factive mental states of S and any members and subgroups of S which are epistemic agents themselves.

This view of the evidential base $\langle E_{\text{BASE}} \rangle$ combines Pooled Evidence with Group-Level Evidence since $\langle E_{\text{BASE}} \rangle$ includes the evidence of each group member, as well as the subgroups, and the group itself. Importantly, this understanding of the evidential base is epistemically continuous since in the individual case, we can think of S as a single epistemic agent with no parts which are themselves epistemic agents.⁵² In the group case, we also need to consider the evidence possessed by all subgroups as well as the group itself, which are organised and structured in a way that makes them (independent) epistemic agents.⁵³

In ROCKET LAUNCH, for example, $\langle E_{\text{BASE}} \rangle$ is composed of the group's justified beliefs that there was consensus among the expert teams that p, as well as the expert teams having justified beliefs that their relevant subtasks are safe. Furthermore, the individual scientists have a plethora of relevant non-factive mental states and justified beliefs that bear on p. As such, $\langle E_{\text{BASE}} \rangle$ is intended to be very broad including all the implicit and explicit justified beliefs and other relevant mental states held by NASA as a whole, the executive committee, the expert panels, as well as any other member of the group not being at the teleconference (if there are any).

Furthermore, since there are no perspectival constraints concerning the doxastic relationship or accessibility $\langle E_{\text{BASE}} \rangle$ includes propositions described by mental states that are immediately available, those which could only be recalled given some effort, and such that are completely inaccessible to any of the group members or subgroups. This is crucial, if we think of groups as independent epistemic agents, capable of having (justified) beliefs, then presumably, some of their group-level beliefs can be based on other group or subgroup-level beliefs.

Before moving on, let me deflect some initial worries which might be raised at this point. In particular, I will shortly mention three alleged problems for summative accounts of group evidence, discussed by Brown (2022b: 166-172).

⁵² However, we do not need to model the individual as one unified epistemic agent. One advantage of thinking of the evidential base in the proposed way is that we can give a theory of justification that is compatible with a fragmented understanding of individuals (Kindermann & Onofri, 2021). In other words, the proposed understanding of the evidential base is epistemically continuous in a way that makes it flexible enough to account for individual justification even when the overall belief states of individuals are fragmented.

⁵³ This is (in the limit) the powerset of all possible combinations of the members. Given a group consists of the members A, B, C, and is organised in a way that each subgroup possesses relevant mental states, $\langle E_{\text{BASE}} \rangle$ would include the evidence possessed by the subgroups {A}, {B}, {C}, {AB}, {AC}, {BC}, and {ABC}.

First, Brown points out that pooling or sharing evidence easily leads to a group having an inconsistent set of evidence as its base, as long as we have a non-factive view of evidence. Brown illustrates this in the context of a pooled account of evidence on which it is sufficient for a proposition p to be part of one's evidence if one justifiably believes p . On this view, we might easily end up with a case in which p and $\neg p$ are part of the group's evidence because the former is justifiably believed by one member m_1 and the latter by another member m_2 . These cases while possible on the proposed understanding simply pose no problem for Continuous Evidentialism. On mentalist or abstractionist views of evidence, it is simply a psychological fact that most epistemic agents have an inconsistent evidential base and there is no reason to think that groups are different in this respect. As we will see, when evaluating the justification of doxastic attitudes, consistency only plays a role in the subpart of the evidential base which is *in fact* used to support the considered doxastic attitude [§5.6] as well as in the higher-order requirements that are necessary to account for potential defeaters [ch.6] [ch.7], but consistency of the entire evidential base is not required.

Second, Brown points out that any summative understanding of the evidential base massively inflates the evidence possessed by groups (a worry which is even worse for the proposed hybrid understanding of group evidence). However, there is nothing problematic *per se* about inflating the evidential base. It becomes only problematic when we think about accessibility and defeaters, both of which I will address in later parts of the thesis. On the contrary, I think that having a too strict understanding of what is part of the evidential base worsens defeater problems and issues raised with the intentional manipulation of the evidential base.

Finally, Brown is worried about problems for the doxastic requirements posed by aggregating evidence to the group level, since propositions that are (justifiably) believed or known by group members are *not* necessarily known by the group itself (given non-summativism about those attitudes). This again, is no problem for Continuous Evidentialism since it is something that we would expect given the proposed statist understanding of evidence, on which evidence of an epistemic agent is mostly determined by mental states the epistemic agent bears no specific doxastic relationship with [§5.3]. Furthermore, I will argue that non-perspectivalist understandings fare better with respect to the causal understanding of basing defended below since beliefs are often based on non-doxastic mental states or beliefs no longer had.

However, while I do deflect Brown's worries against summative views of group evidence, my way of responding to these worries highlights a more fundamental disagreement, that is often implicit in discussions on group evidence (and evidence more generally). Namely, there are two

ways to view cases such as HIRING COMMITTEE that have different consequences for our view of group evidence. On the one hand, one might treat counterfactual scenarios in which the functioning of the vending machine becomes relevant as cases in which the group *asks* their members about evidence regarding the vending machine. That is, regard the members of the group as “evidence channels” by which the group receives evidence (Buchak & Pettit 2015: 213; see also Brown 2022a: 7-9). Doing so likely leads to judgements that G is in a position to acquire evidence by surveying their members (adding to G’s evidential base $\langle E_{\text{BASE}} \rangle$) *rather* than G already being in a position to justifiably believe that p via reflection on already possessed evidence. Going one way or the other has downstream effects on various other issues, such as the distinction between propositional and doxastic justification [§6.7], normative and mental state defeaters [§6.3], as well as issues having to do with the intentional manipulation of the evidential base [§9.3.4]. Accordingly, the best way to measure the plausibility of my hybrid understanding of group evidence is by evaluating it with respect to how successful it is in making sense of these issues.

Let me sum up by demonstrating how we can use the provided understanding of group evidence together with the causal understanding of basing [§5.6] to show how NASA’s belief can be justified in ROCKET LAUNCH. NASA’s belief that it is safe to launch the mission as planned is directly caused by E, which we have previously identified as the propositional content of the expert’s and subgroup’s respective mental states, such as the engineers justified beliefs that their subtasks are safe. During the conference when the experts disclosed their evidence and summarised their reasoning, it was their evidence that caused them to form certain beliefs and behave in a certain way which then caused NASA to form the justified belief that p based on sufficient evidence E.

5.8 Conclusion

Let us recapitulate. Evidentialism is a promising candidate for a continuous theory of justification because not only individuals but also collective epistemic agents are said to possess evidence and base their beliefs on that evidence. Any complete evidentialist theory of epistemic justification needs to answer at least four questions:

Nature of Evidence: What is evidence?

Evidence Possession: What does it mean to possess evidence?

Evidential Support: What does it take for evidence to support a proposition?

Evidential Basing: What does it mean for one to base her belief on one's evidence?

I have taken some first steps towards answering those questions. In particular, driven by some internalist assumptions I committed myself to the following package of views. First, I defended a statist understanding of evidence that takes evidence to be the propositional content of non-factive representational mental states [§5.3]. Second, the entire evidence possessed by an epistemic agent, called $\langle E_{\text{BASE}} \rangle$, is the sum of all relevant mental states, whereas, in the case of groups, this means the entire evidence possessed by the group itself, all subgroups that are epistemic agents themselves, and the group members [§5.4] [§5.7]. Third, I have summarised two promising takes on the epistemic support relation, probabilism [§5.5.1] and Explanationism [§5.5.2], and speculated about a hybrid approach [§5.5.3]. This hybrid approach roughly understands E is sufficient to support the belief that p if it makes p sufficiently likely by providing a sufficiently good explanation for why S has E (see Explanationism). Fourth, I have defended a causal understanding of epistemic basing [§5.6]. On this understanding for a belief that p to be based on some evidence E, the belief needs to be non-deviantly caused by E. Furthermore, we need to rule out instances of so-called improper basing where a belief is based on proper evidence but in an improper way (see Proper Basis).

These observations can be used to fill out the italicised bits of clause (1) of Continuous Evidentialism:

- (1) S possesses some evidence E, which is (a) *sufficient to support* Bp, and (b) the Bp is *properly based* on E.

We can now demonstrate that Continuous Evidentialism diverges from classical evidentialist theories in important respects. As mentioned above [§5.2], two of the central theses of Conee and Feldman's (2000) evidentialism are:

Evidentialist Justification: A doxastic attitude D towards proposition p is epistemically justified for S at t if and only if having D towards p fits the evidence S has at t.

Evidential Supervenience: The epistemic justification of anyone's doxastic attitude toward any proposition at any time strongly supervenes on the evidence that the person has at the time.

Furthermore, given their mentalist understanding of evidence they subscribe to Mentalist Supervenience.

Mentalist Supervenience: The epistemic justification of anyone's doxastic attitude toward any proposition at any time strongly supervenes on the mental states that the person has at the time.

Continuous Evidentialism conflicts with Evidentialist Justification since it does not require the entire evidence possessed by S, that is $\langle E_{\text{BASE}} \rangle$ to fit or support p. Clause (1) only requires D to have some piece of evidence E sufficient to support p. I will provide some arguments why I think that we do not want as a general requirement for epistemic justification that $\langle E_{\text{BASE}} \rangle$ fits or supports p in the next chapter [ch.6].

Second, Continuous Evidentialism conflicts with Evidentialist Supervenience since the epistemic justification of anyone's doxastic attitude does not only depend on the evidence one has but also on the way the doxastic attitude is formed, i.e., whether it is properly based or not. That is there are cases in which two epistemic agents are evidentially alike but causal factors about their belief formation make it the case that some of their beliefs have a diverging epistemic status. (This is also the case even if we understand Evidential Supervenience as a requirement only for propositional rather than doxastic justification. After all, as mentioned above, (1) does not require the entire evidential base $\langle E_{\text{BASE}} \rangle$ to support p.)

Third, Evidential Supervenience is violated by the 3-place relational understanding of epistemic support. On this view, the evidence alone cannot support a doxastic attitude *simpliciter* but only with respect to some set of epistemic standards. So there can be two epistemic agents that are exactly alike evidentially but differ in what they are justified in believing about some issues given they have different epistemic standards. Note, however, that if we understand epistemic standards to be mental facts of the epistemic agents, as I do [§5.5], we can still retain Mentalist Supervenience.

As a consequence, Continuous Evidentialism only allows us to subscribe to the following supervenience thesis, which I introduced in [§5.2]:

Mentalist Doxastic Supervenience: The justification of any epistemic agent's doxastic attitude toward any proposition at any time strongly supervenes on the mental states that the person has at the time as well as the causal relations that obtain between the person's mental states.

So, if two epistemic agents are mentally and causally alike they are alike with respect to the evidence and epistemic standards they possess as well as the ways in which they utilise their evidence and base their doxastic attitudes on them.

With that being said, clause (1) of Continuous Evidentialism cannot be a complete account of epistemic justification, since it gives us the conditions for what we might call *prima facie* justification (though see also [S7.4]). That means we are still in need of a theory of *ultima facie* or *undefeated* justification since even if some evidence E is sufficient to support p when considered in isolation, there might be some additional evidence φ (inside or maybe outside of $\langle E_{\text{BASE}} \rangle$), such that S cannot be justified in believing p given both pieces of evidence. In other words, beliefs that fulfil the conditions spelt out in (1) are still *defeasible*. This is what clause (2) of Continuous Evidentialism is meant to deal with:

- (2) (a) S possesses some higher-order evidence E_H , which is *sufficient to support* q , the proposition that the total evidence E_{TOTAL} on balance supports p and (b) S 's belief that p is *properly based* on E_H .

However, before spelling out how (2) deals with defeater cases in [ch.7], I will spend the next chapter [ch.6] discussing various alternative no-defeater clauses and testing them against different kinds of defeater cases.

6

The No-Defeater Clause:

Evidentialism, Responsibilism and Defeaters

Rational or epistemically justified beliefs are often said to be *defeasible*.¹ For example, my *prima facie* justified belief that my PhD thesis is ready for submission seems to lose its justificatory status when I learn that φ , there is a logical error in one of my central arguments. Following Lehrer and Paxson (1969), we can call this fact or piece of evidence φ a *defeater*.² This defeasibility of *prima facie* justification is a widely acknowledged feature of both internalist accounts of justification, such as the defended version of evidentialism, as well as externalist accounts of justification, such as process reliabilism (Graham & Lyons 2021: 39).³ Accordingly, within an analysis of justification, there is a need for a so-called no-defeater clause, that specifies the conditions under which *prima facie* justified beliefs are *ultima facie* justified.

In this chapter, I will survey various possible evidentialist as well as responsibilist no-defeater clauses and develop a general taxonomy of defeater cases these clauses can be tested against.⁴ Afterwards, in [ch.7] I will propose an evidentialist clause, that can be used to supplement the account of justification defended in the previous chapter Continuous Evidentialism [ch.5]. Despite influential arguments that evidentialist understandings of justification are ill-equipped to handle the full spectrum of defeater cases, I will demonstrate that evidentialism has the right tools to make sense of all kinds of defeaters, including propositional and normative defeaters.⁵ However, it will do so in an unconventional way, that pushes us to rethink the notion of defeasibility altogether [§7.5].

¹ In the literature, the term defeater is also used in discussions of defeasibility responses to the Gettier problem. Usually, in these discussions, defeaters are understood to be unknown propositions rather than something within the perspective of the epistemic agent (Graham & Lyons 2021: 40; Sudduth 2008: §2). Despite my main concern being epistemic justification rather than knowledge, my discussions cover both kinds of defeaters [§6.3].

² Strictly speaking, Lehrer and Paxson (1969), while talking of epistemic defeasibility, do not use the term ‘defeater’.

³ For externalist understandings of defeat see Goldman (1979), Alston (1988), Plantinga (1993: 40-42) or Graham and Lyons (2021). However, not all externalists accept defeasibility. For example, strong knowledge first accounts (i.e. J=K, combined with the idea that knowledge is indefeasible) will entail that justified beliefs are not defeasible (Lasonen-Aarano 2010). For a general criticism about the notion of defeasibility see Baker-Hyatt and Benton (2015), Lasonen-Aarano (2014) or Fraser (forthcoming).

⁴ The term ‘responsibilist’ is taken from Cloos (2015).

⁵ For discussions of normative defeaters see, e.g., Kornblith (1986), Baehr (2011), Cloos (2015), Goldberg (2016, 2017, 2018), or Graham and Lyons (2021).

6.1. Introduction

The idea that epistemic justification is defeasible is motivated by the intuitively plausible thought that the justificatory status of some beliefs is conditional. Take the case of perception. Many epistemologists argue that immediate perceptual beliefs provide us with justification in the absence of reasons to doubt them. For example, Pollock and Cruz state “if something looks red to you and you have no reason to think that it is not red then you are permitted to believe it is red” (1999: 157). Wedgwood thinks that you are permitted to believe that *p* when “you have an experience or apparent perception as of *p*’s being the case, and have no special reason to think that your experiences are unreliable in the circumstances” (2002: 276). This gives rise to the notion of *prima facie* justification. Beliefs that have some initial justification-conferring property, such as being based on sufficient evidence, are *prima facie* justified. Let us call the doctrine that we can have *prima facie* justified beliefs which are defeasible, Defeatism:⁶

Defeatism: Doxastic attitudes can have the status of being *prima facie* justified. That is, some doxastic attitude *D* can be *prima facie* justified by having some justification-conferring property while being defeated and, therefore, lacking *ultima facie* justification.

Accepting Defeatism pushes us to add some kind of no-defeater clause to our theory of justification that specifies when some doxastic attitude counts as defeated. This procedure usually has two steps. First, we nail down what defeaters are, and second we specify the conditions under which defeaters exert their defeating force. Doing so one way or another provides us with different no-defeater clauses which can be tested against various kinds of cases. In this chapter, I will survey several no-defeater clauses and develop a general taxonomy of defeater cases these clauses can be tested against.

In the spirit of Continuous Epistemology [ch.4], my analysis of different defeater clauses is built on intuitive judgements concerning cases involving individual and collective epistemic agents. This aligns with recent developments in collective epistemology, where understanding the role of defeat (Schmitt 1994; Carter 2015; Lackey 2016, 2021; Silva 2019) and its relation to group evidence (Buchak & Pettit 2015; Hedden 2019; Brown 2022b, forthcoming) are seen as crucial in developing a theory of collective justification. Consequently, no-defeater clauses can

⁶ The term ‘Defeatism’ is taken from Baker-Hytech and Benton (2015), who understand defeatism more broadly as encompassing any view that puts a no-defeater condition on knowledge; i.e. they define defeatism as the doctrine, that “in addition to a belief’s being non-accidentally true and justified there is a no-defeater condition (perhaps built into the justification condition) which must be fulfilled for that belief to count as knowledge” (2015: 40).

not only be utilised to supplement accounts of individual justification but also collective and continuous accounts, such as Continuous Evidentialism.

Here is the outline. I will start with some conceptual remarks about evidence [§6.2] and defeat [§6.3]. These preliminaries enable me to illustrate how conventional evidentialist strategies fail to give us the right verdict concerning the full range of defeater cases [§6.4]. Analysing the ways in which these conventional evidentialist strategies fail, pushes us towards a responsibilist understanding of defeat [§6.5], and helps me to develop a general taxonomy of defeater cases [§6.6]. While this proposed taxonomy suggests a responsibilist solution, I will show in the next chapter [ch.7] that this solution only works given certain unconventional understandings of epistemic responsibility. Accordingly, I will propose a novel understanding of epistemic responsibility that, following Goldberg (2018), understands responsibility to be grounded in the socio-epistemic expectations we are entitled to have towards epistemic agents [§7.1]. However, in contrast to Goldberg, I think these expectations are best understood as expectations about the higher-order evidence possessed and utilised by epistemic agents [§7.2]. This reduces responsibilist requirements to higher-order evidentialist requirements, and, therefore, gives us an evidentialist understanding of defeater cases.

6.2. Types of Evidence

Before talking about defeaters, it is important to start with some preliminaries on evidence and its role in epistemic justification. While I have discussed most of these issues in the previous chapter [ch.5], the commitments underlying the following discussions are less specific and, therefore, it is worth explicating them separately. Moreover, doing so will enable me to define several restrictions we could make on the total evidence relevant to an epistemic situation which can be used to characterise different kinds of defeaters [§6.3].

First, I will remain neutral with respect to the *nature* of evidence. In other words, I will not take sides in the debate on whether pieces of evidence are mental states or propositions picked out by those mental states. While evidentialist understandings of justification usually fall into the mentalist camp, in the following discussions I will not take a stance on the nature of evidence.⁷ Relatedly, some argue that false mental states or propositions can be part of one's evidence, while those who defend a *factive* understanding of evidence claim that only true propositions can

⁷ For a discussion of externalist versions of evidentialism see Bergman (2018) or Berghofer (2022: 74-76).

constitute evidence. Following Williamson (2000), factive understandings of evidence are increasingly popular. However, there is still a considerable number of epistemologists who understand evidence to be non-factive [§5.3].⁸ The following discussions are organised as if I rely on a non-factive understanding of evidence. Since if we can provide an evidentialist no-defeater clause relying on a non-factive understanding of evidence, we have done more than required for the evidentialists who have a factive understanding of evidence.⁹

With these preliminaries about the nature of evidence in mind, we can now draw different distinctions between different types of evidence, such as *possessed* and *unpossessed*, *accessible* and *inaccessible*, or *available* and *unavailable* evidence. While drawing such distinctions will be helpful in defining various kinds of defeaters [§6.3], I will not commit myself to any particular understanding of these demarcations. (Partly because I will argue that we do not need these distinctions for the sake of formulating an evidentialist no-defeater clause.)

Let us start with the evidential base, that is, the entire evidence *possessed* by an epistemic agent.

The evidential base $\langle E_{\text{BASE}} \rangle$: The total evidence *possessed* by an epistemic agent S.¹⁰

There are different ways to understand evidence possession. For example, mentalists might define $\langle E_{\text{BASE}} \rangle$ to be the entirety of some relevant mental states, such as experiences; whereas perspectivalists hold that the evidence possessed by an epistemic agent is the evidence that they have a certain doxastic relationship with [§5.4].¹¹ This could, for example, be all of one's justified beliefs (if evidence is non-propositional) or all of the propositions one knows (if evidence is propositional).

⁸ In [§5.3] I have called the thesis evidence solely consist of true propositions Factivity. For objections to Factivity, see Comesana and Kantin (2010), Comesana and McGrath (2014) or Schroeder (2008).

⁹ As pointed out by Hedden (2019: 583-584), having a factive understanding of evidence implies that we can never have cases in which my total evidence supports p and $\neg p$ (given one set of epistemic standards [§1.2] [§5.4]). That means, that presupposing a factive understanding of evidence would rule out mental state defeater cases [§6.3] altogether. Furthermore, Brown (2022a: 166-167) points out that a factive understanding fares better with respect to summative understandings of group evidence, since it rules out cases in which a group possesses inconsistent evidence.

¹⁰ In the wider context of this thesis the evidential base shall be understood as defined in [§5.7], as roughly all relevant mental states had by S and all subgroups of S that are epistemic agents themselves. In the context of this chapter, an intuitive understanding of the evidential base is enough.

¹¹ For a discussion of perspectivalist constraints see, e.g., Alston (1986), Greco (1990). I have discussed such constraints in [§5.4].

Depending on our understanding of evidence and evidence possession, we can make some further distinctions here. For example, we may distinguish between the *accessible* possessed evidence and the *inaccessible* possessed evidence; where accessibility roughly means that S could become aware of the evidence upon reflection.

The accessible base: The total evidence *possessed* by and *accessible* to S.

Based on this notion of accessibility we can further distinguish between accessible evidence that *has been accessed* (that is, evidence that S was or is aware of) and evidence that is accessible but *has not been accessed* (that is, evidence that S could become aware of).

The accessed base: The total evidence which is *accessible* to and *has been accessed* by S.

Furthermore, we could consider some normative restrictions on the accessible evidence to differentiate between the part of the accessible base which we *should have accessed* from the part which we could blamelessly ignore. Let us call this the required base:

The required base: The total *unaccessed* evidence that S *should have accessed*.

So far, we have introduced the notion of an evidential base and distinguished various subparts of it. However, beyond the evidential base is a wider set of evidence pertinent to epistemic considerations. It includes evidence that is not possessed by S but is epistemically relevant for S in a looser sense. It includes all the propositions that S should know and/or could know given some effort, or equivalently, all the evidence that S should come to possess and/or should come to possess. I will call this the *available evidence*.

The available evidence: The total *unpossessed* evidence that S could *come to possess*.

The available evidence will be important in the following discussions of propositional and normative defeaters [§6.3]. One of the questions will be whether *prima facie* justified beliefs can be defeated only by unpossessed evidence that we *should have possessed*, or also by unpossessed evidence that we are not epistemically required to have but which is nonetheless available to S.

That is, we could further differentiate between the unpossessed but required and unpossessed and not required evidence:

The required evidence: The total *unpossessed* and *available* evidence that S *should have had*.

Finally, if we take the union of the evidence possessed by S and the evidence unpossessed by S we get what I call the total evidence:

The total evidence $\langle E_{TOTAL} \rangle$: The total evidence that is *relevant* to S's epistemic situation.

We can understand the total evidence as the entire evidence that bears on the proposition in question. As such, it includes the entire evidence possessed by S and also the entire evidence not possessed by S.¹²

In sum, we can distinguish between different subsets of the total evidence via the following descriptive relations *possession*, *availability*, and *accessibility*, as well as the two normative relations, *should have accessed* and *should have possessed*. This gives us the following tree diagram [Fig. 1].

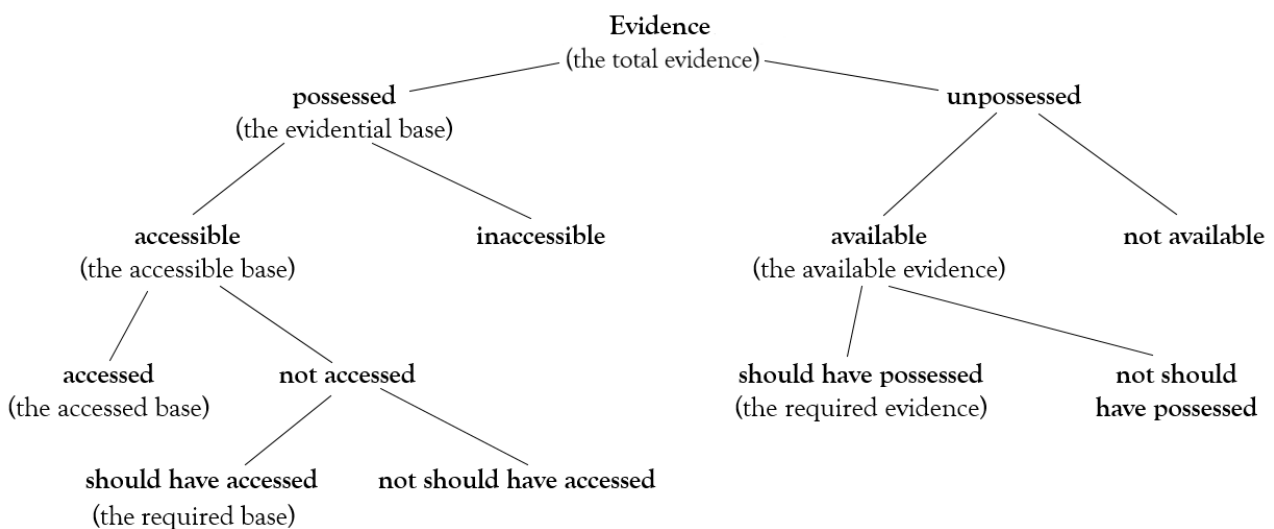


Fig. 1: Evidence

Much of the following discussions will be centred around the questions of whether we can or should make all of these distinctions and whether some of these distinctions are extensionally

¹² I will leave open whether we should restrict the unpossessed evidence relevant to the epistemic situation to the available evidence, that is, the evidence that could be acquired by S.

equivalent. Regardless of how we answer these questions, having an intuitive understanding of these distinctions will turn out to be useful when characterising different kinds of defeaters [§6.3] and discussing extant no-defeater clauses [§6.4] [§6.5].

6.3. Defeaters

Defeaters come in various types, and the two central kinds are *mental state* defeaters and *propositional* defeaters.¹³ In other words, defeaters can be true propositions external to the perspective of the epistemic agent (propositional defeaters), or conditions internal to and/or within the perspective of the epistemic agent (mental state defeaters) (Sudduth 2008; §1).¹⁴

Remaining neutral on the nature of evidence enables us to give a general evidential characterisation of defeaters that encompasses mental state defeaters, as well as propositional defeaters. That is, we can understand defeaters as evidence relevant to the epistemic situation of S that has some defeating force concerning some otherwise justified belief of S. Having this evidential understanding of defeat permits us to use the above-introduced relations (possessed, accessible, available...) to neatly distinguish various kinds of propositional and mental state defeaters.

First, *propositional defeaters* are outside of S's evidential base. That is, S's belief that p is factually defeated by some unpossessed evidence φ if acquiring φ would result in a loss of justification of S's belief that p (Sudduth 2008: §2; Bergman 2005: 154). For example, my otherwise justified belief that there is a barn in front of me might be defeated by the true proposition φ that I am in Fake Barn County (Goldman 1979). Accordingly, we can define propositional defeaters as follows:

Propositional Defeater: Some evidence φ is a *propositional defeater* iff, (i) φ has enough defeating force to render S's *prima facie* justified belief that p unjustified and (ii), φ is *unpossessed* (i.e. outside $\langle E_{\text{BASE}} \rangle$).

¹³ For general discussions of the distinction between propositional and mental state defeaters see, e.g., Bonjour (1980), Goldman (1986), Bergman (2006), Lackey (2008) or Sudduth (2008).

¹⁴ One question I have omitted here is whether all propositional defeaters need to be true, that is, whether all propositional defeaters are factual defeaters. While some have argued that false propositions can be defeaters, I will follow Baker-Hytch and Benton (2015) in arguing that in these cases not the false proposition but the fact that there is such a proposition is the relevant defeater. Or in Baker-Hytch and Benton's words, "for any case in which one's evidence supports a falsehood which ought to be believed, one might defer to the fact that one's evidence supports this as the relevant defeater: in this way, false normative defeaters reduce to factual defeaters after all" (2015: 61).

These propositional defeaters are traditionally understood as defeaters for knowledge rather than justified beliefs. However, more recently, many epistemologists have argued that some specific kind of propositional defeaters, called *normative defeaters*, can also defeat or weaken the epistemic status of an otherwise justified belief. If I should have known that I am in Fake Barn County, that is, if my ignorance with respect to φ is epistemically irresponsible (in a sense to be specified [§6.5] [ch.7]), φ is a normative defeater. Accordingly, we can characterise normative defeaters as “potential defeater[s] that the subject does not actually possess but should” (Graham & Lyons 2021: 45) or “evidence she does not possess but should have possessed” (Nottelmann 2021: 1183; see also Goldberg 2016; 2018: ch.6).

Accepting that there are normative defeaters raises further questions about the relationship between the evidence we should have possessed and its availability. In particular, we might ask ourselves whether ‘ought have possessed’ implies ‘could have possessed’ and/or whether ‘could have possessed’ implies ‘ought have possessed’. While some have suggested that all normative defeaters are defeaters that are available (Harman 1980: 164; Goldberg 2018: 191; see also Nottelmann 2021: 1186), the literature is less clear on whether all available defeaters are normative defeaters. In acknowledgement of this, I will assume that all normative defeaters are available defeaters but not *vice versa*. An unopened letter containing some important information concerning the whereabouts of my friend might be an available, an unavailable, or an available and normative defeater for my belief that she is in town. For example, if the letter is on my desk it might be available and normatively required, while it would only be available but not normatively required if it is at my colleague’s desk and neither if it were still at the post office.¹⁵

Accordingly, we can define two types of propositional defeaters, available defeaters and normative defeaters, where the latter is a subtype of the former:

Available Defeaters: Some evidence φ is an *available propositional defeater* iff, (i) φ has enough defeating force to render S ’s *prima facie* justified belief that p unjustified and (ii), φ is *unpossessed, available but not required* (i.e. outside $\langle E_{\text{BASE}} \rangle$ but inside the required evidence).

Normative Defeater: Some evidence φ is a *normative propositional defeater* iff, (i) φ has enough defeating force to render S ’s *prima facie* justified belief that p unjustified and (ii), φ is *unpossessed, available and required* (i.e. within the required evidence).

¹⁵ This example is based on a case from Harman (1980).

Second, *mental state defeaters* are within the evidential base $\langle E_{\text{BASE}} \rangle$ of S. That is, a mental state defeater is some evidence possessed by S that has some putative defeating force with respect to some *prima facie* justified belief of S. Internalists characteristically deny that there are propositional or normative defeaters but both internalists and externalists typically agree that mental state defeaters can defeat otherwise justified beliefs.

Mental State Defeater: Some evidence φ is a *mental state defeater* iff, (i) φ has enough defeating force to render S's *prima facie* justified belief that p unjustified and (ii), φ is *possessed* (i.e. φ is within $\langle E_{\text{BASE}} \rangle$).

Here, we could further distinguish between different kinds of mental state defeaters, including doxastic, reflective and inaccessible mental state defeaters.¹⁶ A *doxastic defeater* is some piece of defeating evidence φ of which one is aware (φ is within the *accessed base*). A *reflective defeater* is some piece of defeating evidence φ of which one is not aware but of which one could become aware upon reflection (φ is within the *accessible base* but outside the *accessed base*). An *inaccessible defeater* is some piece of defeating evidence φ which is possessed by S but which S cannot become aware of upon reflection (φ is outside the *accessible base*).

Doxastic Defeater: Some evidence φ is a *doxastic defeater* iff, (i) φ has enough defeating force to render S's *prima facie* justified belief that p unjustified and (ii), φ is *possessed* and *accessed* (i.e. φ is within the *accessed base*).

Reflective Defeater: Some evidence φ is a *reflective defeater* iff, (i) φ has enough defeating force to render S's *prima facie* justified belief that p unjustified and (ii), φ is *possessed* and *accessible* but *not accessed* (i.e. φ is within the *accessible base* and outside the *accessed base*).

Inaccessible Defeater: Some evidence φ is an *inaccessible defeater* iff, (i) φ has enough defeating force to render S's *prima facie* justified belief that p unjustified and (ii), φ is *possessed* and *inaccessible* (i.e. φ is within $\langle E_{\text{BASE}} \rangle$ and outside the *accessible base*).

Whether these distinctions between doxastic, reflective and inaccessible defeaters are important, or whether they even make sense depends on various background assumptions about the nature of evidence and evidence possession [§6.2], as well as assumptions about how defeaters exert their defeating force (see below). If we accept this distinction, one interesting question becomes

¹⁶ Similar distinctions are discussed by Bergmann (1997: 116-121) and Sudduth (2008: §5).

whether there are defeaters which we possess and are not aware of but should be aware of. In other words, whether there are *normative mental state defeaters* or not (Lackey 1999).

Normative Mental State Defeater: Some evidence φ is a *normative mental state* defeater iff, (i) φ has enough defeating force to render S's *prima facie* justified belief that p unjustified and (ii), φ is *possessed, accessible and not accessed but should have been accessed*.¹⁷

In sum, this gives us two general types of defeaters, propositional and mental state defeaters, as well as various subtypes which inhabit different subparts of the entire evidence relevant to the epistemic situation $\langle E_{\text{TOTAL}} \rangle$. Similar to how we characterised different subtypes of the relevant evidence above [Fig. 1] we can plug in different kinds of defeaters into a tree diagram [Fig 2]:

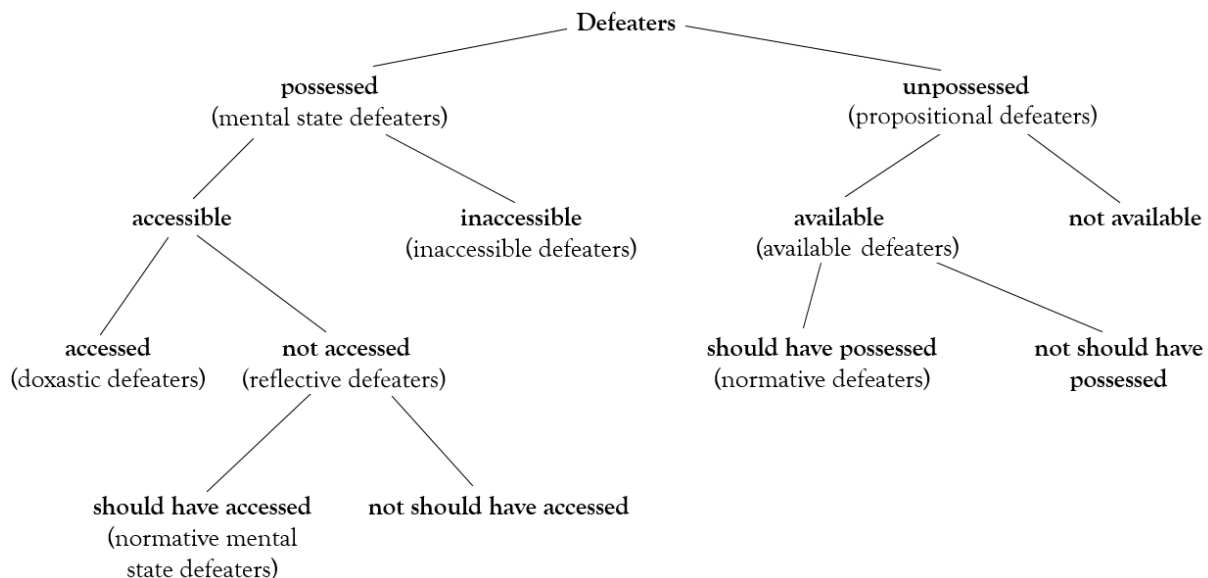


Fig. 2: Defeaters

¹⁷ Lackey distinguishes between two kinds of mental state defeaters which she respectively calls normative and doxastic defeaters (1999; 2006; 2016; 2021). In her framework a doxastic defeater for the belief that p is “a doubt or belief that is had by S that indicates that S’s belief that p is either false or unreliably formed or sustained” (Lackey 2006: 438); whereas normative defeaters are beliefs that S should have, that would constitute a doxastic defeater and that are supported by evidence possessed by S. As such, Lackey’s classification relies on various controversial assumptions about the way in which we possess defeaters and the way in which they exert their defeating force. For example, Lackey argues that doxastic defeaters can be mere beliefs that do not need to have any positive epistemic standing. While I will discuss some cases which are structurally similar to Lackey-style normative defeater cases, I will not rely on her wider understanding of normative mental state defeaters. For an extensive critique of Lackey’s framework see Graham and Lyons (2021).

Having categorised different types of defeaters we can now investigate how defeaters exert their defeating force.

First, we need to consider the question of whether defeaters need to have some positive epistemic standing. Does φ , to have a defeating force towards B_p , need to be sufficient to support $B_{\neg p}$ or at least to repudiate the support that S had for p in the first place? Here the traditional answer is that a defeater is a reason or piece of evidence φ such that given the initial evidence E and φ together are not sufficient to support p (while E alone would be) (Pollock & Cruz 1999: 195; Graham & Lyons 2021: 40-41). Others have argued that mere beliefs without any positive epistemic standing can defeat otherwise justified beliefs (Plantinga 2000: 364-365; Lackey 1999).¹⁸ Here I will follow the traditional conception and assume that defeaters need to have some positive epistemic standing.¹⁹

There are different ways defeaters can exert their defeating force. Both mental state and factive defeaters can be either *rebutting*, that is, provide evidence that the belief that p is false, or *undercutting*, that is, provide evidence that the belief that p is unreliably formed or sustained (Pollock 1986). In Pollock's words, a rebutting defeater attacks the conclusion while an undercutting defeater "attacks the connection between the evidence and the conclusion, rather than attacking the conclusion itself" (1986: 38). For example, reliable testimonial evidence that my colleague is currently in France is a rebutting defeater to my belief that I passed by her on campus today. In contrast, evidence about the unreliability of my facial recognition abilities, such as evidence that I have prosopagnosia or evidence about the unreliability of my short-term memory, is an undercutting defeater.

Finally, defeaters can exert their defeating force on a *higher-order* level, which arguably differs from these traditional means of defeat discussed by Pollock. For example, my friend telling me that she has drugged my coffee with a strong hallucinogenic substance may be a higher-order defeater for my belief that I passed by my colleague on campus today (Elga 2008; Christensen 2010). In contrast to rebutting or undercutting defeaters, such higher-order defeaters have a

¹⁸ This conception of defeat is especially pertinent in discussions of higher-order defeaters which are often argued to have some kind of defeating force even when they leave the first-order support relation intact.

¹⁹ On the Pollockian theory of defeat, nothing can defeat a belief that cannot also provide justification. So, defeaters for justified beliefs are themselves (potential grounds for) justified beliefs. Again, much here depends on various assumptions we make about the nature of evidence, evidential support and *prima facie* justification. While for the sake of developing a no-defeater clause, I will remain neutral on these issues, in the larger context of this thesis I am inclined to understand the defeating force of defeaters in the Pollockian way.

“retrospective aspect, providing a subject with evidence that her belief was never rational, reasonable or justified to start out with” (Lasonen-Aarnio 2014: 317).

Moreover, a defeater may itself be defeated (Lehrer & Paxson 1969: 228-229). In such circumstances, we may speak of a defeated defeater who is defeated by a *defeater-defeater*. For example, the allegedly reliable testimonial evidence that defeated my belief that my colleague is currently in France might itself be defeated by hearing from multiple independent sources that the person whose testimony I am relying on is a notorious liar. Differentiating between defeated and undefeated defeaters is important since it is often argued that while justified beliefs are incompatible with defeaters, we only need to worry about defeaters who have not been defeated themselves, that is, *undefeated defeaters*. (I will question this assumption below [§6.4]).

In sum, there are various kinds of defeaters including propositional (normative and non-normative) and mental state (including doxastic, reflective, inaccessible and normative) defeaters, which can defeat an otherwise justified belief in various ways (rebutting, undercutting, higher-order defeat), and which can itself be defeated. I will spend the next sections discussing conventional evidentialist [§6.4] and responsibilist [§6.5] strategies to account for the phenomenon of epistemic defeat. Based on these discussions, I will introduce a general taxonomy of defeater cases in [§6.6] and illustrate how extant evidentialist understandings fail to cover the entire range of defeater cases while responsibilist notions of defeat seem to give us the right verdict. However, they do so only given unconventional understandings of epistemic responsibility [ch.7]. This observation encourages me to introduce a novel higher-order evidentialist no-defeater clause [§7.3] that reduces the notion of epistemic responsibility to evidential considerations.

6.4 Evidentialism and Defeat

In this section, I will summarise conventional evidentialist ways of handling defeat and list some well-known, as well as some novel shortcomings of these approaches. Let me start with some reminders about evidentialism as understood by Conee and Feldman [§5.2].

The central evidentialist commitment in Conee and Feldman’s (2004) framework is the following thesis:

Evidentialist Justification: A doxastic attitude *D* towards proposition *p* is epistemically justified for *S* at *t* if and only if having *D* towards *p* fits the evidence *S* has at *t*.

Moreover, Conee and Feldman combine this evidentialist notion of justification with a mentalist understanding of evidence which gives rise to the following supervenience claim:

Mentalist Supervenience: The epistemic justification of anyone's doxastic attitude toward any proposition at any time strongly supervenes on the mental states that the person has at the time.

In combination, Evidentialist Justification and Mentalist Supervenience imply that two epistemic agents possessing exactly the same mental states would be exactly alike evidentially and, therefore, concerning what they are propositionally justified in believing about any given issue (Kelly 2016: §1).

So, how does this Conee and Feldman-style evidentialism do with respect to some simple mental state defeater cases?

DEFEATER: S's belief that p is *prima facie* justified. However, S is not aware that she possesses a *reflective defeater* φ for her belief.

DEFEATER-DEFEATER: S's belief that p is *prima facie* justified. However, S is not aware she possesses a *reflective defeater* φ for her belief. Luckily, φ is itself defeated by another *reflective defeater-defeater* λ .

The Conee and Feldman-style evidentialist seems to have a convenient way to account for mental state defeater cases such as DEFEATER and DEFEATER-DEFEATER. They can appeal to propositional justification. If the entire evidence possessed by an epistemic agent needs to support or fit the respective belief for the belief to be propositionally justified the belief cannot be defeated by a mental state defeater.²⁰ Since propositional justification is usually understood to be a necessary requirement for doxastic justification, S's belief cannot be justified. So strictly speaking, for Conee and Feldman there are no mental state defeaters since there is no (propositional) *prima facie* justification if the belief is not supported by the entire evidential base of S.

²⁰ More on Conee and Feldman's treatment of the relationship between doxastic and propositional justification below [§6.7].

Since we are looking for a no-defeater clause that can be added to all kinds of understandings of *prima facie* justification, we need to transform this insight into an independent no-defeater clause:

Evidentialist No-Defeater Clause <Evidentialism>: S's *prima facie* justified belief that p is undefeated iff S's evidential base *on balance supports* p.

According to <Evidentialism> in DEFEATER-DEFEATER, S is justified in believing that p while in DEFEATER S's belief that p lacks justification. Despite this being the seemingly desired result, I will demonstrate that handling defeater cases, via appeal to propositional justification is mistaken. In particular, I will discuss two kinds of propositional defeater cases [§6.4.1] and two kinds of mental state defeater cases [§6.4.2] on which <Evidentialism> fails to deliver the right verdict.

6.4.1 Evidentialism and Propositional Defeaters

My first thesis is that, since <Evidentialism> is only concerned with defeaters within the evidential base, it cannot handle cases of propositional or normative defeat.

PROPOSITIONAL DEFEATER: S's belief that p is *prima facie* justified. However, S is not aware that there is an unpossessed but *available defeater* φ for her belief that p.

NORMATIVE DEFEATER: S's belief that p is *prima facie* justified. However, S is not aware that there is an unpossessed but *available defeater* φ for her belief that p that S *should have possessed*.

Depending on whether S should have possessed φ or not we have a template to construct normative propositional or merely propositional defeater cases. I will, for now, focus on normative defeaters. However, I will revisit this distinction at the end of the section. (Furthermore, the account I'll offer in [§7.3] can make sense of normative and non-normative propositional defeater cases).

The unique threat normative defeaters have for evidentialism has been pointed out by many, including Kornblith (1983), Baehr (2009), Cloos (2015), as well as Goldberg (2017: 2891-2893; 2018: ch.6) among others. If we are merely concerned with the overall evidence an epistemic agent possesses at a time, we can never consider the evidence (defeating or not) that is not

possessed by S, but should be possessed by S. One illustrative example that has the structure of NORMATIVE DEFEATER is found in Kornblith (1983):

PERSISTENT PETE: Pete a young and stubborn physicist presents his newest pet hypothesis at a conference eager to hear the praise of his colleagues. Due to a personality disorder, Pete pays no attention to reasonable critics and strategically ignores important counterevidence. As a result, even devastating criticism fails to impact his beliefs not because he fails to take it into account but because he has not even heard it.

For Kornblith this example illustrates that Pete's "belief is unjustified, after his colleague presents his objection, and it is unjustified because of his culpable ignorance" (Kornblith 1983: 36). As such, Kornblith offered the case as a direct challenge to evidentialism, since if his interpretation is correct there are non-evidential factors which determine whether Pete is justified or not.²¹ While evidentialists have tried to dismiss the intuitive judgement we have towards cases such as PERSISTENT PETE, none of their strategies seems particularly convincing. For example, Conee and Feldmann (2004: ch.7) have argued that cases of normative defeat do not undermine epistemic justification but usually demark other shortcomings such as moral or professional failures. However, many disagree with this judgement (Goldberg 2016: 450; Lackey 2016: 374-375, Graham & Lyons 2021). After all, having these prudential or professional failures has devastating epistemic effects on Pete, not only concerning the reliability and truth-conduciveness of his belief-forming mechanisms but also concerning the evidence he possesses. There are good reasons to think that these allegedly professional failures are also epistemic failures (more on that below).

At this point, a plausible suggestion is to expand the evidence we are concerned with to include unpossessed evidence. That is, we might require that the total evidence, possessed and unpossessed, on balance supports p.

Extended Evidentialist No-Defeater Clause <Extended Evidentialism>: S's *prima facie* justified belief that p is undefeated iff the *total evidence on balance supports p*.

While <Extended Evidentialism> accounts for cases such as PERSISTENT PETE and other cases of normative defeat, it already demarks a significant departure from Conee and Feldman's version

²¹ For a discussion of this case see, e.g., Goldberg (2016; 2017).

of internalist evidentialism, since it denies the central mentalist supervenience claim.²² Furthermore, it cannot account for the following type of normative defeater cases:

LUCKY NORMATIVE DEFEATER-DEFEATER: S's belief that p is *prima facie* justified. However, S is not aware that there is an unpossessed but *available defeater* φ for her belief that p that S *should have possessed*. Luckily for S, there is another *normative defeater* λ which defeats φ .

To pump some intuitions, it might help to put some flesh on the bones:

SORROWLESS SARAH: Sarah forms a *prima facie* justified belief that p. However, the belief is based on some information E drawn from the daily tabloid, unknowingly to Sarah a very unreliable source. This fact constitutes a normative defeater φ for the belief that p, since Sarah could have easily and should have learned about φ . Instead of checking the quality of her sources Sarah unreflectively formed her belief based on E. Yet, while generally unreliable the daily tabloid is reliable in this particular instance since the one columnist Sarah based her belief on is exceptionally reliable. This fact serves as a defeater-defeater λ for φ .

Sarah's *prima facie* justified belief appears to be defective because she was just lucky that the defeater that she should have possessed is itself defeated. She could have easily learned about φ which, if rational, would have led her to abandon the belief. In not possessing φ Sarah behaved epistemically irresponsible just as Pete behaved epistemically irresponsibly when ignoring the counterevidence presented by his colleagues.

If that's the right verdict, we can easily see that extending <Evidentialism> to <Extended Evidentialism> does not solve the problem of normative defeat. Instead, normative defeater cases seem to suggest that we require epistemic agents to be responsible in the right kind of way. This is also illustrated by the following type of propositional defeaters cases:

UNLUCKY PROPOSITIONAL DEFEATERS: S's belief that p is *prima facie* justified. However, while being generally a very responsible investigator S is not aware that there is an *unpossessed* and *available* defeater φ for the belief that p.

²² As pointed out in the previous chapter [ch.5], while evidentialism is usually understood as a form of internalism, many have pointed out that we can consider externalist versions of evidentialism (Bergman 2018; Berghofer 2022: 74-76). Furthermore, the initial strategy to account for propositional defeaters advocated by Lehrer and Paxson (1969) is similar <Extended Evidentialism>. Note, however, that these traditional discussions were concerned with knowledge rather than justified beliefs.

Here is such a case:

METICULOUS MICHAEL: Michael's belief that the tip of his tongue is mainly responsible for tasting sweet is *prima facie* justified. Michael has good evidence that this is true. He has once learned it in school and even remembers (apparently) confirming it himself as a kid. Furthermore, he even double-checks his belief by looking it up in his old biology book. However, while being generally a very responsible investigator Michael could have easily found out (e.g., via a quick Google search) that the tongue-tasting map is a common scientific misconception that has repeatedly been proven wrong. This fact serves as a propositional defeater φ for Michael's belief.

Can Michael's belief that p be justified despite there being a propositional defeater φ ? It seems like it makes a difference if Michael acts meticulously or not. After all, he double-checked his belief and in general has good evidence for it being true. At least, there seems to be a difference to cases such as **PERSISTENT PETE**. Pete willfully ignored counterevidence and formed his belief partly based on the desire to prove his pet hypothesis.

If we judge Pete's and Michael's beliefs to have a different epistemic status, there are at least two ways to explain the different results. First, we might think that the way Pete and Michael form their beliefs make those beliefs epistemically responsible or irresponsible respectively and hence justified or unjustified. On the other hand, we might think that the different judgements suggest that **PERSISTENT PETE** is a normative defeater case, where Pete should have possessed φ , while **METICULOUS MICHAEL** is merely a propositional defeater case, and hence Michael had no epistemic obligation to possess φ . Either way, the analysis goes through the notion of epistemic responsibility, which suggests a responsibilist treatment of defeat.

However, before I explore this responsibilist treatment of defeaters further in [§6.5], let me discuss another set of defeater cases evidentialist approaches have trouble dealing with [§6.4.2].

6.4.2 Evidentialism and Mental State Defeaters

While <Evidentialism> could handle regular mental state defeater cases, such as **DEFEATER** or **DEFEATER-DEFEATER**, normative defeater cases illustrated that <Evidentialism> is too narrow. In trying to overcome this problem by extending <Evidentialism> to <Extended Evidentialism>, we made it too general, since it declared sorrowless Sarah's belief to be justified despite her epistemically irresponsible behaviour and meticulous Michael's belief to be

unjustified despite his epistemically responsible behaviour. We can find a similar pattern when we take a closer look at mental state defeater cases.

In particular, if we use propositional justification to declare all *prima facie* justified beliefs in all mental state defeater cases to be unjustified and in all defeater-defeater cases to be justified we get the wrong verdict with respect to two cases: *inaccessible defeater* cases, in which the epistemic agent is *responsible* [§6.4.2.1] and *accessible defeater-defeater* cases, in which a defeated defeater is *irresponsibly ignored* [§6.4.2.2]. Let me discuss those cases subsequently.

6.4.2.1 Inaccessible Defeaters. If we solely care about propositional justification, we rule out justification in *any* case in which we stipulate that there is a defeater φ within the $\langle E_{\text{BASE}} \rangle$. This is the case, irrespective of how strong the initial justification for Bp is (as long as φ has enough defeating force), regardless of whether the defeating evidence is doxastic, reflectively accessible or entirely inaccessible, and independently of how responsible the epistemic agent is. Here is a template for such cases:

UNLUCKY MENTAL DEFEATERS: S's belief that p is *prima facie* justified. However, while being generally a very responsible investigator S is not aware that she *possesses* a non-doxastic *inaccessible* or *hardly accessible* defeater φ for the belief that p .

We can utilise this template to get the following case which bears some similarities to METICULOUS MICHAEL [§6.4.1]:

CAREFUL CAROLINE: Caroline's belief that the tip of her tongue is mainly responsible for tasting sweet is *prima facie* justified. Caroline has good evidence that this is true. She has once learned it in school and even remembers (apparently) confirming it herself as a kid. Furthermore, she even double-checks her belief by looking it up in her old biology book. However, while being generally a very responsible investigator Caroline forgot that her university professor once told her that the tongue-tasting map is a common scientific misconception that has repeatedly been proven wrong. This currently inaccessible memory serves as a non-doxastic defeater φ for Caroline's belief.

CAREFUL CAROLINE aims to illustrate that it depends on the exact details of mental state defeater cases if we should consider defeated beliefs to be justified or not. It seems overly demanding to require that epistemic agents are always completely sensitive to the entire body of evidence they possess, independently of how accessible the defeater is. After all, as justification is

fallible concerning the truth of the formed beliefs it is also plausibly fallible with respect to the overall evidential support.

So, depending on the exact nature of the evidential base, <Evidentialism> may lead to an absurdly restrictive notion of justification. Take, for example, the so-called inclusive view of evidence discussed by Conee and Feldman (2004: 228) (see also [§5.4]). On this view, the evidence possessed by an epistemic agent is just the sum of all their non-factive mental states. That means the evidential base includes wrong beliefs and unconscious or in-principle inaccessible mental states. All of these mental states could serve as defeaters. In other words, even if S is overly responsible, attentive and aware of most of the evidence they possess the *prima facie* justified belief that p could ultimately be unjustified because they happen to possess an inaccessible or hardly accessible defeater.

Furthermore, the problem with inaccessible defeaters becomes even worse in the context of collective justification. When it comes to collective entities, the evidential base arguably includes the evidence possessed by many epistemic agents, as well as the collective G itself [§5.7]. So, in the worst case, the mere fact that some member of G has an entirely inaccessible mental state defeater would be enough to undercut or even rebut an otherwise justified belief. Consider this illustrative example:

CAUTIOUS PI: Paolo, the principal investigator of the research group G behaves in a way that constitutes the group to have a *prima facie* justified belief that p. However, while Paolo and all relevant members are excessively responsible, G possesses (via one of its members) an inaccessible defeater φ for the belief that p.

Faced with cases such as CAUTIOUS PI, many contemporary theorists of collective justification make restrictions on the evidence that can defeat the group's beliefs (either directly or by restricting what's part of the collective evidential base). Accordingly, some have proposed no-defeater clauses that require that the total evidence possessed by *enough* operative members does not include defeaters (Silva 2019), or that sufficient disclosure of *some* members' evidence accompanied by sufficient deliberation would not undermine the group attitude (Lackey 2016; 2021).²³

Confronted with these cases the evidentialist might be tempted to restrict the notion of the evidential base in a way that rules out inaccessible or hardly accessible evidence. However, then

²³ More on collective defeater cases in [ch.9].

the problem posed by normative defeaters becomes all the more important. In restricting the evidential base, we would push the defeater out of it into the unpossessed but available or required evidence and thereby transform cases such as CAREFUL CAROLINE and CATIOUS PI into normative defeater cases such as PERSISTENT PETE.

6.4.2.2 Normative Mental State Defeaters: Furthermore, both <Evidentialism> and <Extended Evidentialism> also give us the wrong verdict for some normative mental state defeater-defeater cases. To see this, let us start with this case inspired by Lackey (2008: 63):

IGNORANT ALICE: Alice is told by her optometrist that her vision is nearly completely unreliable, yet she refuses to accept his diagnosis, without having any rational basis for doing so. Afterwards, as she is walking out of the doctor's office, and sees a car accident. Based on that Alice forms the corresponding true belief that there was such an accident. Intuitively, however, her belief is defeated by the doctor's diagnosis even if she refuses to accept it.

Let us compare this case to some cases already discussed. On the face of it, IGNORANT ALICE looks like a regular defeater case rather than a normative defeater case such as PERSISTENT PETE. That is, ignorant Alice, in contrast to persistent Pete, possesses but ignores a defeater φ . Following Lackey, we could, nonetheless, call these kinds of defeater cases *normative* since Alice *should* form the corresponding belief (and Lackey thinks that this gives φ its defeating force).²⁴ In other words, we seem to have special normative obligations towards not only some evidence we should possess but also some evidence we already possess but *should access*.

Independently, of whether that is the right interpretation of the case, <Evidentialism> can make perfect sense of it, since Alice's belief is not propositionally justified. However, we can use Lackey's understanding of normative obligations to access some bits of our evidential base to construct the following type of defeater-defeater case which poses a problem for <Evidentialism>:

LUCKY MENTAL DEFEATER-DEFEATER: S's belief that p is *prima facie* justified. However, due to being ignorant and negligent, S is not aware that she possesses an *accessible defeater* φ which she *should have accessed* for the belief that p. Luckily, S also possesses an *inaccessible* or *hardly accessible defeater-defeater* λ for φ .

We can again add some flesh to the bones:

²⁴ For details of Lackey's framework see footnote 17.

CARELESS CARL: Carl's belief that he saw his colleague Hao on campus today is *prima facie* justified. Carl has a vivid memory of her passing by the café while he was drinking his after-lunch coffee. Carl when asked if Hao is in town responds that she is without further reflection. However, another colleague told Carl last week that Hao will fly out on holiday tomorrow, something that Carl could recall easily if he would just reflect on his belief. This would make Carl doubt his memory and likely lead him to judge that he must have mistaken Hao for her twin sister. Accordingly, this memory counts as a defeater φ for Carl's belief that Hao is in town. Luckily for Carl, he also possesses a *defeater-defeater* λ for φ ; namely that Hao told him that she did not plan any holiday this year and that she will definitely stay in town over the summer to work on her newest book. However, in contrast to φ , Carl cannot easily recall λ and hence would only remember it after a long period of reflection.

CARELESS CARL illustrates that not only in defeater cases but also in defeater-defeater cases the omitted details matter. Carl's belief, while undefeated and *prima facie* justified, seems epistemically defective for various reasons. First, Carl's belief is unstable. He could easily be aware of φ which, if rational, would lead him to dismiss the belief. So, it seems wrong to judge Carl's belief to be justified because of the luck involved. Carl, while having *prima facie* justified beliefs clearly forms his beliefs improperly. It is a mere coincidence that Carl's otherwise defeated belief ends up being undefeated.²⁵ So it seems like we have all the right to blame Carl for having this belief and generally forming beliefs in that way. And if this blameworthiness goes hand in hand with our epistemic evaluations we need to say that Carl's belief is unjustified. (More on the relationship between epistemic blameworthiness and justification in the next chapter [§7.3]).

Similar to CAREFUL CAROLINE formulating a group version of the case might help to pump some intuitions. After all, if the evidential base is vastly bigger for groups, there is a larger pool of evidence that could luckily defeat any irresponsibly ignored and easily accessible defeaters. I will discuss such cases in greater detail in [§9.3.2].

In sum, conventional evidentialist strategies fail to give us the right verdict in at least two different kinds of defeater cases. First, they are insensitive to propositional and normative defeaters, and second, they are too hasty in declaring all mental state defeater cases to be unjustified and all mental state defeater-defeater cases to be justified. While extending the pool

²⁵ Interestingly, these cases have a similar structure to propositional defeater cases which are often framed as Gettier-style counterexamples to the traditional analysis of knowledge.

of considered evidence (such as in <Extended Evidentialism>) helps with some propositional and normative defeater cases it does not help with the latter types of cases discussed.

6.5. Responsibilism: Two Desiderata

As demonstrated above, conventional evidentialism has problems with a plethora of defeater cases; namely, normative defeater cases (PERSISTENT PETE), normative defeater-defeater cases (SORROWLESS SARAH), lucky propositional defeater cases (METICULOUS MICHAEL) inaccessible mental state defeater cases in which the epistemic agent is otherwise very responsible (CAREFUL CAROLINE and CAUTIOUS PI) and defeater-defeater cases in which the agent behaves in an epistemically defective way (CARELESS CARL). On the face of it, these misjudgements of evidentialist strategies stem from a common source: they do not take into account whether the epistemic agents behave the *epistemically responsible* or not.

Accordingly, one way to account for these cases is to introduce a no-defeater clause that obliges epistemic agents to be epistemically responsible in the right kind of way. On similar grounds, many have argued that we have epistemic normative obligations (Lackey 2016; 2021; Silva 2019), requirements to be sensitive to the evidence we should have had (Goldberg 2016; 2018: ch.6; Cloos 2015) or that we are required to know everything we should have known due to professional or prudential duties (Goldberg 2017). These strategies can be broadly categorised as *responsibilist* treatments of defeat.

Responsibilist No-Defeater Clause <Responsibilism>: S's *prima facie* justified belief that p is undefeated iff S is *epistemically responsible* in believing that p.

<Responsibilism> looks promising, and indeed, if understood correctly, it gives us the right verdict with respect to all of the cases discussed so far [§6.6]. However, as I will demonstrate in the next chapter [ch.7], there are more and less plausible ways to understand epistemic responsibility. For now, I will just, based on the discussions above, identify two desiderata any account of epistemic responsibility needs to fulfil:

Desideratum 1: Any responsible epistemic agent forms and maintains doxastic attitudes in a way that is *sensitive* to the evidence they *accessed* and *should have accessed*.

Desideratum 2: Any responsible epistemic agent forms and maintains doxastic attitudes in a way that is *sensitive* to the evidence they *should have possessed*.

These desiderata, while falling short of giving us a full-blown understanding of epistemic responsibility will be helpful when comparing different kinds of no-defeater clauses by checking their verdicts concerning the full taxonomy of defeater cases [§6.6].²⁶

Before doing so, however, let me briefly introduce a hybrid no-defeater clause that is neither exclusively evidentialist nor exclusively responsibilist.

Responsibilist-Evidentialist No-Defeater Clause <Responsible Evidentialism>: *S's prima facie* justified belief that *p* is undefeated iff *S's* evidential base *on balance supports p* and *S* is epistemically *responsible* in believing that *p*.

Both Baehr (2009) and Cloos (2015) point out that there is no conflict, in principle, with combining responsibilist notions of defeat with evidentialist notions of justification. Likewise, Silva (2019) adds a responsibilist no-defeater clause to his otherwise evidentialist understanding of collective justification.²⁷ However, these accounts are full-blown accounts of justification rather than merely no-defeater clauses that are used to supplement understandings of *prima facie* justification. As such, <Responsible Evidentialism> can be added to all kinds of understandings of *prima facie* justification including evidentialist, responsibilist or reliabilist accounts.

In sum, in the last two sections [§6.3] [§6.4], I have identified four potential responsibilist and/or evidentialist no-defeater clauses and discussed various cases. The next section uses these insights to define three independent parameters to give us a general taxonomy of defeater cases against which any no-defeater clause can be tested.

²⁶ I will leave it open for now, how exactly we should understand these desiderata. Especially different understandings of how we should interpret the sensitivity clause will turn out to give us different verdicts about epistemic responsibility [ch.7].

²⁷ Similarly, Lackey's *Group Epistemic Agent Account* of collective justification has a built-in no-defeater clause which combines responsibilist and evidentialist ideas (Lackey 2016: 381). However, Lackey's understanding of defeat, as described above, does not understand normative defeaters to be external to the perspective of the epistemic agents. I will discuss Lackey's understanding of collective justification in more detail in [§9.2].

6.6. Responsibility, Balance and Location

Most of the proposed no-defeater clauses make different judgements about the justificatory status of beliefs depending on the location of the defeater. For example, for <Evidentialism> and <Responsible Evidentialism> whether the belief that p is justified (in part) depends on whether the defeater is possessed or unpossessed (within or outside of $\langle E_{\text{BASE}} \rangle$). <Extended Evidentialism> in contrast declares the belief to be defeated as long as there is a defeater within the total evidence E_{TOTAL} . Let us call the parameter that demarcates whether the defeater is possessed or unpossessed the *location* of the defeater.²⁸

In addition to the location, most evidentialist accounts care about the propositional justification, that is, the overall balance of the evidence. If there is a defeater φ (within the evidential base) the belief is unjustified; however, if there is a defeater-defeater λ that undermines φ the belief's justification is restored again. That is, the overall evidence can be in favour of p because there is no defeater, or because there are only defeated defeaters. Let us call this parameter the evidential *balance*.

<Responsibilism>, declares beliefs to be un/justified regardless of the overall balance and location of the defeater(s), while <Responsible Evidentialism> declares Careless Carl's belief to be unjustified and Careful Caroline's to be only justified if the defeater is *not* part of her evidential base. In contrast, <Evidentialism> and <Extended Evidentialism> do not take the careless or careful nature of the epistemic agent into account at all. Let us call this parameter the *responsibility* of the epistemic agent.

Having identified the parameters, *responsibility*, *balance*, and *location* presents us with $2^3 = 8$ possible cases, which, in turn, gives us a general taxonomy of defeater cases. Here is a template, we can use to change the parameters accordingly to see which account gives us which result:

Defeater Template: S's belief that p is *prima facie* justified. There is a [*balance*: defeated/undefeated] and [*location*: possessed/unpossessed] defeater φ for S's belief that p . S is [*responsibility*: careless/careful] in forming and sustaining the belief that p .

In [Fig. 3] I plugged in the judgements made by <Evidentialism>, <Responsible Evidentialism>, <Extended Evidentialism>, and <Responsibilism> with respect to the different possibilities

²⁸ Of course, we could define the location more thoroughly and distinguish between the different kinds of mental state defeaters and propositional defeaters defined above [§6.3]. However, since none of the discussed no-defeater clauses is sensitive to the exact location of the defeater, I will restrict the following discussion to the possessed/unpossessed binary.

provided by Defeater Template, whereby ‘N’ stands for defeated or unjustified and ‘Y’ stands for justified or undefeated:

Taxonomy of cases:				<Evidentialism>	<Responsible Evidentialism>	<Extended Evidentialism>	<Responsibilism>
responsibility	balance	location	verdict				
careless	defeated	possessed	N	Y	N	Y	N
careless	undefeated	possessed	N	N	N	N	N
careless	defeated	unpossessed	N	Y	N	Y	N
careless	undefeated	unpossessed	N	Y	N	N	N
careful	defeated	possessed	Y	Y	Y	Y	Y
careful	undefeated	possessed	Y	N	N	N	Y
careful	defeated	unpossessed	Y	Y	Y	Y	Y
careful	undefeated	unpossessed	Y	Y	Y	N	Y

Fig.3: Responsibility, Balance, and Location

By plugging in the respective parameters into the Defeater Template we can see that only <Responsibilism> matches our intuitive judgements: careless Carl’s belief (careless; defeated; possessed) is not justified, such as persistent Pete’s and sorrowless Sarah’s (careless; defeated; unpossessed) while Careful Caroline’s belief (careful; defeated; possessed) remains justified [Fig. 3].

While this sounds like a full-blown responsibilist conclusion, I will spend the next chapter [ch.7] arguing that the best way to understand epistemic responsibility is as a higher-order evidentialist requirement. However, before plunging ahead, I want to show how my suggested treatment of defeater cases casts doubt on the conventional understanding of the relationship between propositional and doxastic justification below [§6.7].

6.7 Propositional and Doxastic Justification: Revisited

As pointed out in the last chapter [§5.6], Continuous Evidentialism rests on an unconventional understanding of the relationship between propositional and doxastic justification. While this was partly motivated by general reflections on how we want to restrict the evidential base $\langle E_{\text{BASE}} \rangle$ of a (collective) epistemic agent [§5.3] [§5.7], some of the lessons we have learned about defeaters above can be used to further motivate this view. I will use this opportunity to walk through different views about the relationship between propositional and doxastic justification that have

been proposed in the literature, to then identify where we should locate Continuous Evidentialism on the space of possible options.

Let me start with the view which treats propositional justification as a necessary requirement for doxastic justification. In [§5.6] I have characterised this view as follows:

Basis: S's belief that p is *doxastically* justified iff (i) the belief is *propositionally* justified for S in virtue of S's having evidence E for p, and (ii) S believes p *on the basis of* E.

Following Turri (2010b), we can call this the *orthodox view*. While this orthodox view is defended in numerous places (Pollock & Cruz 1999: 35-36; Swain 1979: 25; Feldman 2002), theorists who defend it are often ambiguous between two different readings of Basis.²⁹ Sometimes they require the entire body of evidence possessed by S, what I call $\langle E_{\text{BASE}} \rangle$, to support p for S to be propositionally justified. On the other hand, sometimes they *only* require the operational part E to be sufficient to support p, for S to be propositionally justified. That is, we can distinguish between a stronger, wide-scope version of the orthodox view as well as a weaker, narrow-scope version:

Orthodox view [narrow scope]: S's belief that p is *doxastically* justified iff (i) the belief is *propositionally* justified for S in virtue of S's having evidence E *sufficient to support* p, and (ii) S believes p *on the basis of* E.

Orthodox view [wide scope]: S's belief that p is *doxastically* justified iff (i) the belief is *propositionally* justified for S in virtue of S's evidential base $\langle E_{\text{BASE}} \rangle$ being *sufficient to support* p, and (ii) S believes p *on the basis of* E (some part of $\langle E_{\text{BASE}} \rangle$) that is *sufficient to support* p.

These two versions of the orthodox view differ in their understanding of propositional justification; that is, they give us different understandings *in virtue of what* S's belief that p is propositionally justified (i). Usually, these accounts are supplemented by a theory of epistemic basing that specifies condition (ii). However, some theorists, including Conee and Feldman propose an account that makes additional restrictions on condition (ii) which are not (solely)

²⁹ A more extensive list is found in Turri (2010b: 313-314). Note, however, that Turri does not differentiate between what I call the wide scope and the narrow scope versions of the orthodox view (see below). See also Korcz (2021) or Sylvan (2016) for a general discussion of epistemic basing.

captured by a theory of epistemic basing. In this context, they talk of well-founded attitudes, where they define the well-foundedness relation as follows (Conee & Feldman 2004: 93):

- S's doxastic attitude D at t toward proposition p is well-founded if and only if
- (i) having D toward p is justified for S at t; and
 - (ii) S has D toward p on the basis of some evidence E, such that
 - (a) S has e as evidence at t;
 - (b) Having D toward p fits E; and
 - (c) There is no more inclusive body of evidence E' had any S at t such that having D toward p does not fit E'.

This proposal is interesting in various respects. First, note that on Conee and Feldman's understanding of what it means to have a justified belief, condition (i) implies the orthodox wide-scope requirement (for them for D to be justified D needs to fit S's entire evidence). Second, Conee and Feldman require that the belief is based on the most inclusive part of $\langle E_{\text{BASE}} \rangle$ that supports p (ii) (c), for S to be doxastically justified.

Accordingly, Conee and Feldman's proposal could give us the right verdict with respect to defeater-defeater cases in which we use some bit of evidence E and have total evidential support by $\langle E_{\text{BASE}} \rangle$ for p but irresponsibly ignore some (easily accessible) defeater (that is itself defeated). One such case was CARELESS CARL [§6.4.2].³⁰ However, Conee and Feldman's resolution of such cases is very demanding since it requires epistemic agents to use large parts of their evidence when forming beliefs. More precisely, it requires epistemic agents to use at least as much of the relevant evidence possessed so that they are able rule out all of the possessed counterevidence when forming the belief that p. This is very demanding since it is often not possible to take all of the possessed evidence relevant to a given issue into account. This should be especially clear in the case of groups, where the $\langle E_{\text{BASE}} \rangle$ arguably includes evidence possessed by the group itself, subgroups as well as group members [§5.7]. (This is one reason why I find the above specified responsibilist understanding of defeater cases attractive [§6.5]).

Conversely, others have approached questions about the relationship between propositional and doxastic justification the other way around. Following Goldman (1979), some have started with doxastic justification and then, subsequently, defined propositional justification through

³⁰ This is strikingly similar to how Lackey understands normative defeat, as being sensitive to the evidence you *should have accessed*. See [§6.3], especially footnote 17.

doxastic justification, while they do not understand propositional justification to be a necessary requirement for doxastic justification. Goldman writes (1979: 124; *emphasis in the original*):

“Person S is *ex ante* [propositionally] justified in believing p at t if and only if there is a reliable belief-forming operation available to S which is such that if S applied that operation to his total cognitive state at t, S would believe p at t-plus-delta (for a suitably small delta) and that belief would be *ex post* [doxastically] justified.”

If we keep in mind that Goldman operates in a reliabilist framework, it makes sense that he treats doxastic justification to be more fundamental than propositional justification. After all, being justified by forming one’s belief via a reliable process can be independent of one having reasons or evidence sufficient to support p. So, for the reliabilist who wants to retain some concept of propositional justification, it makes sense to back-engineer it into their theory via their understanding of doxastic justification.

This is similar to a more recent approach found in Turri which we might dub the *unorthodox view* (2010b: 319):

Unorthodox View: S’s belief that p (at t) is *propositionally* justified if S at t possesses at least one means of coming to believe that p such that, were S to believe that p in one of those ways, S’s belief would thereby be *doxastically* justified.

Turri’s defence of the Unorthodox View is inspired by examples of improper basing cases which I have discussed elsewhere in this thesis [§5.6]. One thing to notice about the Unorthodox View is that explaining the presence of propositional justification in terms of an agent having the (current) ability to acquire doxastic justification, it follows that *the absence* of such an ability would imply the absence of propositional justification. This raises a cluster of questions about what it means to have or lack such an ability.³¹

Having spelt out these different orthodox and unorthodox understandings of the relationship between doxastic and propositional justification we can now turn to Continuous Evidentialism. What does Continuous Evidentialism imply about the relationship between propositional and doxastic justification? On the face of it, my analysis of justification is only peripherally concerned with propositional justification. There is no propositional justification

³¹ For a similar worry see Silva (2015: 947).

condition embedded in any of the clauses of Continuous Evidentialism. Nonetheless, there are two ways in which we could understand propositional justification in the context of Continuous Evidentialism.

First, we might argue that the concept of propositional justification becomes relevant when we think about epistemic support, where (1a) E needs to be sufficient to support p (and (2a) E_H needs to be sufficient to support q , that E_{TOTAL} supports q [§7.3]). So, another way to spell this out would be to say that E needs to *propositionally* justify p (and E_H needs to *propositionally* justify q). This looks like the narrow-scope orthodox view. However, I do not see any reason why we want to have this notion of propositional justification as part of our normative theory. Say S believes that p for no reason whatsoever; a lightning strike or a bump on the head made S believe that p , while coincidentally S 's belief is supported by some part of S 's evidence. Why should that belief be any better than any belief that is not supported by some part of S 's evidence, if S is completely insensitive to that evidence? Justified beliefs are beliefs that fulfil the complex conditions spelt out in Continuous Evidentialism. These beliefs enjoy a special normative status. Beliefs, however, that are coincidentally supported by some bit of the agent's evidence do not enjoy any special normative status.

Second, we could pick the unorthodox route and back-engineer propositional justification into our theory by saying that S is propositionally justified if S could become doxastically justified. (Given we can define this 'could' in some non-ambiguous way). However, this is not useful either because answering the question of whether S would be doxastically justified to believe that p requires us to answer whether the evidence E used as a basis for Bp is sufficient to support p . This can be understood as answering the question of whether S is propositionally justified in p , given the narrow-scope orthodox understanding. Yet, there is something peculiar about disentangling matters this way. Imagine again that S believes that p for no reason whatsoever. Now in cases like these S cannot be doxastically justified in believing that p . But how should we then answer the question of whether S is propositionally justified to believe that p ? We can't because there is no operational part E that enables us to figure out whether p is supported by E . And picking just any part of S 's evidence is surely not sensible. After all, S might possess evidence that is sufficient to support p while S might also possess evidence which is sufficient to support $\neg p$.

Where does this leave us? I take this to demonstrate, that there is no use for the concept of propositional justification within the context of these investigations.

6.8 Conclusion

In this chapter, we have reached a preliminary conclusion in our search for a no-defeater clause. I categorised different kinds of defeaters [§6.3] and discussed four different defeater clauses: <Evidentialism>, <Extended Evidentialism> [§6.4], <Responsibilism> and <Responsible Evidentialism> [§6.5]. I demonstrated that only <Responsibilism> as understood via the two identified desiderata gives us the right verdict with respect to the entire taxonomy of defeater cases [§6.6]. Accepting that responsibilist requirements can be sufficient to give us the verdict 'justified' even in some cases where there is a defeater possessed by S cuts the connection between propositional and doxastic justification [§6.7]. However, the presented understanding of epistemic responsibility underlying <Responsibilism> was admittedly vague. This urges us to take a closer look at the concept of epistemic responsibility. That is what I will do in the next chapter [ch.7].

7

Epistemic Responsibility:

Epistemic Responsibility and Higher-Order Evidence

In the previous chapter [ch.6] I demonstrated that we can use the notion of epistemic responsibility to formulate a no-defeater clause that gives us the intuitively correct verdict with respect to the full range of defeater cases. I called this clause <Responsibilism>. This was an important observation due to the widely accepted doctrine of Defeatism [§6.1], which states that doxastic attitudes can be *prima facie* justified by having some justification-conferring property while being defeated and, therefore, lacking *ultima facie* justification.

Therefore, one natural way to proceed would be to just take <Responsibilism> and add it as a second clause (2) to the first clause (1) of Continuous Evidentialism defended in [ch.5], where (1) gives us the conditions of *prima facie* justification, and (1) and (2) give us the conditions of *ultima facie* justification. However, doing so would be problematic in at least two respects. First, <Responsibilism> is quite schematic. It is *not* built on a general understanding of epistemic responsibility but just relies on the two desiderata introduced in [§6.5]. Second, adding <Responsibilism> as a no-defeater clause to Continuous Evidentialism would make Continuous Evidentialism a hybrid evidentialist-responsibilist account rather than a fully evidentialist theory of justification.¹ In this chapter, I will address these two issues. That is, I will provide a general framework for understanding epistemic responsibility and demonstrate that we can use this framework to reduce responsibilist requirements to higher-order evidential requirements.

7.1 Introduction

In this chapter, I will argue that epistemic responsibility is best understood via the general expectations we are entitled to have towards epistemic agents; and, that these expectations entitle us to expect epistemic agents to possess and utilise higher-order evidence that bears on their entire evidential situation. This results in a higher-order evidential requirement of the following form:

¹ See, for example, Cloos (2015) or Silva (2019) for such hybrid accounts.

S belief that p is *epistemically responsible* iff:

- (2) (a) S possesses some higher-order evidence E_H , which is *sufficient to support* q , the proposition that the total evidence E_{TOTAL} on balance supports p and (b) S's belief that p is *properly based* on E_H .

Here is the outline. I start by taking a closer look at the concept of epistemic responsibility. In particular, I will sketch one of the most developed frameworks of epistemic responsibility to date by Sanford Goldberg [§7.2]. Goldberg argues that epistemic responsibility is to be understood via the socio-epistemic expectations we are entitled to have of epistemic agents.² While I agree with Goldberg that epistemic responsibility is grounded in socio-epistemic expectations, I will argue that we shall *not* understand these expectations, as Goldberg does, as making us strictly liable for the evidence we should have had [§7.3]. In particular, I will discuss and dismiss two ceiling principles that ground Goldberg's understanding of epistemic responsibility as strict liability. One has to do with the relationship between first-order evidence and higher-order evidence [§7.3.1], while the other has to do with the relationship between the evidence we have and the evidence we should have had [§7.3.2].

In [§7.4], I will introduce and defend my alternative model, which gives us the above-stated understanding of epistemic responsibility encapsulated in clause (2). According to this model, we are entitled to expect from epistemic agents that they possess and utilise higher-order evidence when forming beliefs [§7.4.1], which can be used to motivate a clause of the form of (2) [§7.4.2]. One way to motivate this understanding can be derived from general observations about the normative force of unpossessed evidence [§7.4.3]. This enables us to reduce responsibility requirements of justification to higher-order evidentialist requirements.

I will conclude by using the insights gained about epistemic responsibility to revisit Defeatism, the general principle underlying our search for a no-defeater clause, and find it wanting [§7.5]. This leaves us with a positive higher-order requirement of justification that gives us the intuitive verdict with respect to a vast range of defeater cases while rejecting the doctrine of Defeatism. This is an important result since it makes Continuous Evidentialism immune to various influential arguments that have recently been put forward against Defeatism.³

² See Goldberg (2016; 2017; 2018).

³ While still prominent, many, including Bergmann (2005), Lasonen-Aarnio (2010, 2014, 2020) or Baker-Hyatt & Benton (2015), have recently argued against the doctrine of Defeatism.

7.2 Epistemic Responsibility: The Entitlement Model

One powerful framework to understand epistemic responsibility is found in Goldberg (2016; 2017; 2018). For Goldberg, epistemic responsibility is best understood via the socio-epistemic expectations we are entitled to have towards epistemic agents; where those expectations can be general expectations we have towards every epistemic agent, as well as the more specific expectations we are entitled to have of epistemic agents occupying certain social roles.⁴ As such, epistemic responsibility plays a key role in Goldberg's analysis of epistemic justification or epistemic propriety, as he calls it (2018: 13-47).

Goldberg proposes a theory of epistemic justification with the following hybrid structure, involving both *core criteria* and *general expectations*, both being grounded in social expectations. While core criteria can speak for or against the propriety, general expectations are implicit background criteria which can only disqualify someone from meeting the standard required for justification.

First, for Goldberg, a belief is *prima facie* justified when it fulfils the core criteria of epistemic propriety or justification. These criteria include reliabilist as well as coherentist conditions, where the coherence criteria capture the first dimension of the responsibility condition (2018: 202):

“avoiding incoherence is among the sorts of things we are entitled to expect of other epistemic subjects. Indeed, I argued that this is a ‘basic’ expectation we are entitled to have of other epistemic subjects: we are entitled to hold others to a standard on which they refrain from belief that p when their overall evidence clearly favors $\neg p$, or even when it provides no more support for the one than the other. So, since incoherence avoidance is properly expected of one, a failure to detect the incoherence in such cases does not enable a subject to evade responsibility.”

For Goldberg, we are entitled to expect epistemic agents to form beliefs “in a way that avoids bald incoherence with her background belief” (2018: 145). If we acknowledge that some of those background beliefs are evidence possessed and accessible by the epistemic agent, we can use this first dimension of epistemic responsibility to formulate a no-mental state defeater clause. In so doing, we can make sense of regular mental state defeater cases, as well as Lackey-style normative mental state defeaters (without overshooting the target as evidentialist requirements did). In other words, this makes sense of the first desiderata of epistemic responsibility identified in [§6.5]:

⁴ For a similar understanding of epistemic responsibility and normative defeat see Meeker (2014). For a concise summary of Goldberg's account see Habgood-Coote (2020).

Desideratum 1: Any responsible epistemic agent forms and maintains doxastic attitudes in a way that is *sensitive* to the evidence they *accessed* and *should have accessed*.

Second, we rightfully expect epistemic agents to fulfil what Goldberg calls general epistemic requirements. That is, we expect them to know certain things (2017) or possess certain pieces of evidence (2016; 2018: ch.6). Here we can identify two widely acknowledged sources of these expectations. On the one hand, Goldberg suggests that there are “basic moral and epistemic expectations that we have of any and all moral and epistemic subjects” (2017: 2875) and that these expectations entitle us to expect that epistemic agents possess certain pieces of evidence. On the other hand, we have distinct epistemic obligations that come with “professional or institutional role[s]”, or more generally emerge from “our interpersonal relations with one another” (2017: 2875).⁵ So, for example, we expect experts to know certain things that laypersons do not and rightly judge them to be epistemically irresponsible if they fail to live up to those expectations. For Goldberg, the best way to acknowledge the epistemic significance of these general expectations “is through the doctrine of normative defeat: if S should have known that p, then the proposition that p is a defeater for S” (2017: 2893; see also 2018: ch.6).

Accordingly, this second dimension of epistemic responsibility enables us to make sense of the second desideratum [§6.5]:

Desideratum 2: Any responsible epistemic agent forms and maintains doxastic attitudes in a way that is *sensitive* to the evidence they *should have possessed*.

In sum, we can use Goldberg’s framework to characterise epistemic responsibility as being sensitive to all of the possessed and unpossessed defeating evidence which others are entitled to expect us to be sensitive towards. In my jargon that gives us an understanding of epistemic responsibility that looks like this:

Epistemic Responsibility: S is epistemically responsible if S is *sensitive* to (i) the evidence S *possesses*, including those which S has *accessed* and *should access* according to basic normative expectations, and *sensitive* to (ii) the evidence S *should possess* according to the epistemic expectations we are entitled to have towards S.

⁵ For an account of collective justification that uses this understanding of epistemic responsibility see Silva’s (2019) *Responsibilist Evidentialism for Groups* which I will discuss in [ch.9]

This characterisation of epistemic responsibility is useful in our quest to formulate a no-defeater clause since it can be plugged into the responsibilist clause <Responsibilism> formulated above. Furthermore, it gives us a deeper explanation of epistemic responsibility by grounding it in epistemic expectations. This is promising since it provides us with a generic strategy of how to make and refine case judgements about defeater cases. For instance, we can judge whether someone's belief is (un)justified because of some defeater φ by asking whether we are entitled to expect them to be sensitive to φ .

However, a big question mark still hangs over the proposed understanding. While we have identified the source of epistemic responsibility in socio-epistemic expectations, and made sense of Desiderata 1 and Desiderata 2, we still need to say what exactly it means for an epistemic agent to be *sensitive* towards defeaters in the right way. Let me proceed by first sketching Goldberg's answer to this question and point towards some problems with it [§7.3], and then introduce my alternative model [§7.4].

7.3 Sensitivity and Epistemic Strict Liability

What does it mean to be *sensitive* towards the evidence relevant to the proposition in consideration in the formation of one's doxastic attitudes? In Goldberg's framework, we can understand being sensitive to defeaters via what he calls the doctrine of Epistemic Strict Liability, which roughly states that epistemic agents are answerable to the evidence they should have had.

Goldberg arrives at the doctrine of Epistemic Strict Liability by analysing the relationship between different types of evidence: the evidence one possesses E ,⁶ including possessed first-order evidence (in short: E_{FO}) and possessed higher-order evidence about unpossessed evidence one should have had (in short: E_H)⁷, and, on the other hand, unpossessed evidence you should have had (in short: E_{SHH}) (2018: 196-197). Goldberg observes that we come up with a general taxonomy according to variations along two dimensions characterising the epistemic bearing of E_{SHH} and E_H . Here, both E_H , as well as E_{SHH} can either support, disfavour or be neutral with respect to the proposition S believes based on E (while S may or may not possess any higher-order evidence E_H). This gives us a taxonomy of twelve different cases. In walking through these cases Goldberg arrives at two ceiling principles about the relationship between E , E_{SHH} and E_H :

⁶ It is not clear whether E refers to what I called the operational part of one's evidence E [§5.5], which is the evidence used in forming one's belief, or to what I have called the evidential base < E_{BASE} > [§5.4], which is the entire evidence possessed by an epistemic agent. More on that below [§7.3.1].

⁷ Note that Goldberg uses the abbreviation E_{HO} instead of E_H for higher-order evidence of this kind.

Ceiling Principle 1: S's evidence E [including E_{FO} and E_H] establishes an *epistemic ceiling*: in the final assessment the epistemic goodness of S's belief that p can be no greater than it is when it is formed on the basis of E.

Ceiling Principle 2: The epistemic goodness of S's belief that p can be no greater than it is when it is formed on the combination of $(E + E_{SHH})$.

These two ceiling principles then motivate what Goldberg calls the doctrine of Epistemic Strict Liability (2018: 215; *emphasis in the original*):

Epistemic Strict Liability: In a case in which you believe that p on evidence E, the epistemic goodness of your belief is answerable to the evidence you should have had E_{SHH} , in the sense that the goodness can be no greater than it would be if your belief were formed on the basis of $E + E_{SHH}$ — *whatever your higher-order evidence led you to anticipate regarding that further evidence E_{SHH} .*

While I will not discuss Goldberg's treatment of all of the cases, it is helpful to briefly summarise his reasoning in defence of Ceiling Principle 1 [§7.3.1] and Ceiling Principle 2 [§7.3.2] and point towards some problematic consequences of understanding sensitivity via those principles. Doing so enables us to identify some general reasons why we should not analyse epistemic responsibility via strict liability [§7.3.3], despite relying on Goldberg's general understanding of epistemic responsibility as grounded in social expectations [§7.2].

7.3.1 Epistemic Ceiling 1: First-Order and Higher-Order Evidence

Here is Ceiling Principle 1 again:

Ceiling Principle 1: S's evidence E [including E_{FO} and E_H] establishes an *epistemic ceiling*: in the final assessment the epistemic goodness of S's belief that p can be no greater than it is when it is formed on the basis of E.

Before we can investigate Ceiling Principle 1 in detail some general remarks about higher-order evidence are needed.

Higher-order evidence can be understood as evidence about the overall evidential situation an epistemic agent finds herself in. This can, for example, be evidence about the (im)properness of already possessed evidence, or the rational capacity of the epistemic agent to respond to evidence (Christensen 2010: 185-186). Consider the case of an expert, studying some phenomena

and acquiring a set of first-order evidence E taking it to support her conclusion C . Assuming that she is generally competent when it comes to evaluating evidence (she has the relevant expertise for) the belief that *she arrived at C* based on E can be considered as higher-order evidence E_H for the claim that E supports C . Subsequently, E_H not only supports E supports C , but also C itself.⁸

Goldberg arrives at Ceiling Principle 1 after discussing cases in which an epistemic agent S considers a proposition p while having first-order evidence E_{FO} supporting p and higher-order evidence E_H that there is evidence S should have possessed E_{SHH} which diminishes or defeats E_{FO} 's support for p (Goldberg 2018: 193-201). Here, Ceiling Principle 1 tells us that in all of these “unhappy” cases, the epistemic status of the belief is decreased to the exact degree that E_H indicates that E_{SHH} diminishes E_{FO} 's support (2018: 194). Or in Goldberg's words, “higher-order evidence bearing against the belief that p is, by itself, epistemically significant, in that when it is possessed by a subject it diminishes the epistemic goodness of her belief that p , and this effect is independent of the epistemic bearing of the evidence she should have had” (2018: 203).

At first glance, Ceiling Principle 1 seems plausible. First, one's belief that p cannot be more justified than it is supported by one's evidence E (including E_{FO} and E_H) for p . Second if we possess higher-order evidence E_H that there is further evidence relevant to p which bears on the justificatory status of my belief that p we need to take this evidence into account. So in the case where E_H is unhappy, this should have a negative effect. For instance, if I have good reasons to believe that I have ignored a defeater for my otherwise justified belief, the belief cannot be justified anymore. Since then, I have reason to believe that the belief is defeated which means that I already have a defeater for my belief. As Goldberg points out, this can be derived from the principle that evidence for evidence, that is higher-order evidence, is just evidence (Goldberg 2018: 201-202):⁹

“To begin, if S has higher-order evidence to think both that there is further evidence she should have, and that this evidence will bear against her belief that p , then to just that extent she ought to question the likelihood that the proposition that p is true. This would follow from a version of the principle that evidence of evidence is evidence. [...] In short,

⁸ Some have recently argued that it is more accurate to speak of higher-order evidential effects or higher-order evidential import, rather than higher-order evidence, since evidence can be relevant to a person's beliefs about different propositions, in a number of different ways (Hedden & Dorst 2022; Lasonen-Aarnio 2019). Readers sympathetic to this proposal can understand as some evidence that E_H primarily has higher-order import towards p .

⁹ Goldberg puts this in terms of the higher-order evidence indicating some unreliability in the entire belief-formation process, whereas “the subject ought to question the epistemic propriety of her belief that p to the degree that her higher-order evidence is unhappy—for it threatens reliability to precisely that degree” (2018: 201).

the subject ought to question the epistemic propriety of her belief that *p* to the degree that her higher-order evidence is unhappy –for it threatens reliability to precisely that degree.”

This result is further supported by the entitlements we intuitively seem to have towards epistemic agents. We seem entitled to expect epistemic agents to not form beliefs that contradict their higher-order evidence since this would violate the basic epistemic expectations we have towards them (Goldberg 2018: 202). In other words, Goldman uses his entitlement model of epistemic responsibility [§7.2] to justify Ceiling Principle 1.

However, I think that despite its initial plausibility we should reject Ceiling Principle 1. Ceiling Principle 1 prevents us from reaching the intuitively correct verdict with respect to unlucky mental state defeater cases, in which epistemic agents seem to behave correctly but are unlucky. In [§6.4], I have given the following template to construct such cases:

UNLUCKY MENTAL DEFEATERS: *S*'s belief that *p* is *prima facie* justified. However, while being generally a very responsible investigator *S* is not aware that she possesses a non-doxastic *inaccessible* or *hardly accessible* defeater φ for the belief that *p*.

One case of this form was CAREFUL CAROLINE, whose properly based and well-supported belief that the tip of her tongue is mainly responsible for tasting sweet was defeated by some currently inaccessible memory [§6.4]. According to Ceiling Principle 1, Careful Caroline's beliefs cannot be justified because her otherwise carefully formed belief is defeated by evidence she possesses but is not aware of due to no fault of her own.¹⁰ So, if the evidence one has, including first-order and higher-order evidence, establishes an epistemic ceiling there can be no unlucky cases.

However, since whether there are such unlucky cases depends on diverging case judgements, I do not want to rest too much on this observation. Instead, I will leave it at this dispute concerning intuitions until I discuss more general reasons to reject Ceiling Principle 1 and Ceiling Principle 2 with it [§7.3.3]. But before doing so let me take a closer look at Ceiling Principle 2 [§7.3.2].

¹⁰ Note that Goldberg's account *could* make sense of CAREFUL CAROLINE (depending on the omitted details) by pointing out that Caroline's belief does not satisfy some of the required reliability criteria at the level of the process type (and hence her belief does not meet the core-criteria [§7.2]). However, I am confident that we can construct cases with an analogous structure in which core criteria are met while we still want to judge Caroline's belief to be unjustified.

7.3.2 Epistemic Ceiling 2: Evidence You Should Have Had

We can now turn to Ceiling Principle 2:

Ceiling Principle 2: The epistemic goodness of S's belief that p can be no greater than it is when it is formed on the combination of (E + E_{SHH}).

Interestingly, Ceiling Principle 2 has some similarities to the extended evidentialist no-defeater requirement <Extended Evidentialism> I have discussed in [§6.4.1]. While the latter makes it a requirement of justification that <E_{TOTAL}> propositionally justifies the respective belief, the former sets a limit to how justified a belief that p can be depending on the overall support that is lent to p by <E_{TOTAL}>. As such, Ceiling Principle 2 gives us an understanding of what it means to be sensitive to normative defeaters and mental state defeaters because it covers the possessed and the unpossessed evidence. In short, we are liable to all of the evidence we should have possessed and all of the evidence we should have accessed in case our, otherwise justified, belief is defeated by it. That is, as Goldberg puts it, “instead of speaking of the expectation that others fulfill all of their social-epistemic responsibilities, we could speak instead of the expectation that others believe what they *would* believe if they *were* to fulfill all of those responsibilities” (2018: 223; *emphasis in the original*).

Unfortunately, this understanding of being sensitive to defeaters as being liable to them is problematic, since it gives us the wrong verdict with respect to unlucky propositional defeater cases, in which epistemic agents seem to behave correctly but are unlucky. These cases are similar to the cases that raise problems for Ceiling Principle 1 but instead of non-negligently missing a possessed defeater, the respective agent non-negligently misses an unpossessed defeater they could have possessed. Here is the template as stated in [§6.4] again:

UNLUCKY PROPOSITIONAL DEFEATERS: S's belief that p is *prima facie* justified. However, while being generally a very responsible investigator S is not aware that there is an unpossessed and available defeater φ for the belief that p.

One such case was METICULOUS MICHAEL, whose properly based and well-supported belief that the tip of her tongue is mainly responsible for tasting sweet could have been refuted by a quick Google search. According to Ceiling Principle 2, meticulous Michael's beliefs cannot be justified because his otherwise carefully formed belief is defeated by some evidence he is not aware of due to no fault of his own. In other words, if the evidence one has and the evidence one should

have possessed together establish an epistemic ceiling there can be no unlucky propositional defeater cases.

Again, whether there are such unlucky propositional defeater cases is controversial. If all propositional defeater cases are normative defeater cases this is questionable. We might think that one cannot be responsible while ignoring some evidence one should have possessed, at least in the case where this evidence has defeating force with respect to the proposition in consideration. After all, METICULOUS MICHAEL is a propositional rather than a normative defeater case. But more on that in the next section [§7.3.3].

7.3.3 Against Epistemic Strict Liability

The discussions above suggest that thinking about sensitivity via ceiling principles is mistaken because it does not allow us to make intuitively correct judgements about so-called unlucky cases. This shows that neither the evidence one possesses (including E_{FO} & E_H) nor the evidence one possesses and the evidence one should have possessed (E & E_{SHH}) necessarily establish an epistemic ceiling.

However, the defender of Epistemic Strict Liability will likely be unimpressed by that result. After all, we are only liable to the evidence we *should* be sensitive to, which is the evidence we have accessed, the evidence we *should* have accessed and the evidence we *should* have possessed.¹¹ If that's the case, there cannot be any cases in which one is responsible but unlucky since if one didn't take some evidence they should have accessed (UNLUCKY MENTAL DEFEATER) or some evidence they should have possessed (UNLUCKY PROPOSITIONAL DEFEATER) into account. Being responsible, according to this understanding, then simply is to take this evidence into account.

I think this result is mistaken. To see why in full detail, we need to wait until I have introduced my higher-order evidential understanding of epistemic responsibility [§7.4]. In short, I will argue that careful Caroline and meticolus Michael can be epistemically responsible despite missing defeating evidence if they have and utilise sufficient higher-order evidence that the entire evidence $\langle E_{TOTAL} \rangle$ supports their proposition (on top of first-order evidence sufficient to support p).

¹¹ Another way to think of this is that this kind of responsibility makes the distinction between possessed and unpossessed evidence redundant. The responsibility might say that the entire evidence relevant to an epistemic situation simply is the entire evidence that a perfectly responsible agent would pay attention to, whether it is possessed or not.

Before doing so, however, let me point towards another problem faced by an understanding of epistemic responsibility that rests on Ceiling Principle 1 and Ceiling Principle 2. Those principles together are not *sufficient* to give the right verdict with respect to all defeater cases since they allow for lucky propositional defeater cases, as well as lucky mental state defeater cases, in which epistemic agents seem to behave irresponsibly but are lucky.¹² Here are the templates for those cases again:

LUCKY NORMATIVE DEFEATER-DEFEATER: S's belief that p is *prima facie* justified. However, S is not aware that there is an unpossessed but *available defeater* φ for her belief that p that S *should have possessed*. Luckily for S, there is another *available defeater* λ which defeats φ .

LUCKY MENTAL DEFEATER-DEFEATER: S's belief that p is *prima facie* justified. However, due to being ignorant and negligent, S is not aware that he possesses an *accessible defeater* φ which she *should have accessed* for the belief that p . Luckily, S also possesses an *inaccessible* or *hardly accessible defeater-defeater* λ for φ .

In [ch.6] have discussed one such case respectively, namely, SORROWLESS SAHRA [§6.4.1] and CARELESS CARL [§6.4.2]. In these cases, the epistemic agents form their beliefs that p , while based on some evidence that is sufficient to support that p , ignorant of the entire epistemic situation they find themselves in. It is a mere coincidence that their otherwise defeated beliefs end up being undefeated. So, liability does not enable us to be attentive to the epistemic malfunctions associated with carelessness.

Importantly, lucky cases as well as unlucky cases are a dime in a dozen. The discussed examples are just representative illustrations. The basic intuition underlying my judgements in these cases is twofold. First, justification is fallible; that is, no justification-conferring defeater-detecting mechanism needs to be perfect.¹³ Second, we do not hold epistemic agents liable for the fallibility of their beliefs but for taking negligent risks when forming their beliefs in an

¹² In personal correspondence Goldberg pointed out that there is another way to account for cases such as CARELESS CARL which differs from the treatment proposed in Goldberg (2018). We might say that if there is any expansion of Carl's evidence so that it includes some of the further evidence he should have had, where doing so leads to normative defeat, then we get the result that Carl's belief is defeated. This seems puzzling, since if Carl acquires the defeater φ without acquiring the defeater-defeater λ , it seems irresponsible for Carl to consider his belief to be defeated since he should have also possessed the defeater-defeater λ .

¹³ The sceptical reader should keep in mind that we are looking for a no-defeater clause for justified beliefs, not knowledge.

otherwise justified way. Once we acknowledge that, we can easily see why we judge doxastic attitudes that are formed in a non-cautious way but which are luckily undefeated to be unjustified and attitudes that are formed cautiously but which are unluckily defeated to be justified.

At this point, some might worry that the proposed analysis confuses two dimensions of normative evaluations, epistemic blame and epistemic responsibility, where it is not clear whether those evaluations go hand in hand. This is also pointed out by Goldberg (2018: 175-184) who goes to great lengths to argue for an understanding of epistemic responsibility which is independent of epistemic praise or blame. In other words, Goldberg allows for cases in which subjects are blameless but irresponsible (or not responsible) or responsible but blameworthy.

In contrast, we do not have to engage in this kind of endeavour. Instead of introducing another layer of normative evaluation we can stick to our intuitive judgement and say that we judge careless Carl to be epistemically irresponsible and hence unjustified and careful Carloline to be epistemically responsible and hence justified.¹⁴ Introducing another layer of normative evaluation by pulling epistemic blame and epistemic responsibility apart introduces complications where none are needed. So, I take it that any theory that can avoid this complication is *ceteris paribus* a better theory of epistemic justification.¹⁵ But more importantly, as I will further argue below, making this distinction here conflicts with any account that wants to ground epistemic responsibility in the expectations we are entitled to have towards epistemic agents. As I will show in the next section, these expectations directly align with our habits of judging people to be responsible or irresponsible for their epistemic achievements.¹⁶

If we want to retain the intuitive verdicts concerning these cases we need to have a different understanding of epistemic responsibility that fits with the intuitive reading of *sensitivity* embedded in Desiderata 1 and Desiderata 2. I will propose such an understanding in the next

¹⁴ For an account that understands being epistemically blameworthy and being epistemically irresponsible to go hand in hand see Booth and Peels (2010).

¹⁵ Note that I am not committed to saying that every doxastic attitude for which we are responsible is epistemically blame- or praiseworthy. As in the moral case, there might be instances in which we are morally responsible for some action but are neither praise- nor blameworthy for it (McKenna 2012: 16-17). All I would need for the analogy between the moral and the epistemic domain to be successful is that there are no instances in which we are responsible for a bad outcome/attitude but not blameworthy or responsible for a good outcome/attitude but not praiseworthy.

¹⁶ Note that these diverging judgements might be caused by a fundamental difference in the kinds of epistemic evaluations we are concerned with. Externalists about epistemic justification already have a need for a distinction between blameworthiness and conformity to the epistemic standards right at the outset. So, for the externalist, one gains nothing by preserving a link between blamelessness and epistemic justification, since that link itself is not a desideratum. I want to thank Sandy Goldberg for pointing this out.

section [§7.4.1] and then utilise this understanding to defend higher-order evidence no-defeater clause [§7.4.2].

7.4 Reducing Responsibility: Higher-Order Evidence

To recapitulate: we want a no-defeater clause that makes epistemic agents responsible, that is, sensitive to evidence, possessed or unpossessed, they should take into account, in a way that makes room for lucky and unlucky cases. In so doing, we want epistemically responsible agents to be less likely to ignore defeaters but rule in cases in which they are responsible but miss some defeaters due to no fault of their own. Furthermore, we want to rule out cases in which epistemic agents are irresponsible in forming some beliefs but are lucky that those beliefs are *not* defeated. In this section, I will develop such an understanding of epistemic responsibility and utilise it to formulate a no-defeater clause. I will do so in three steps. First, I will motivate the general idea that to be epistemically responsible is to have and utilise higher-order evidence [§7.4.1]. Afterwards, I will introduce the higher order clause [§7.4.2]. In [§7.4.3] I will address a potential objection, and provide some additional reasons for why I think having and utilising higher-order evidence is a general requirement for justification.

7.4.1 Responsibility and Higher-Order Evidence

One promising way to retain our intuitive judgements about epistemic responsibility is to equate being epistemically responsible with utilising higher-order evidence about one's epistemic situation. That is, instead of taking epistemic agents to be liable for all of the evidence they have and should have had, we expect that epistemic agents have and utilise higher-order evidence about the entire evidence relevant to their epistemic situation $\langle E_{TOTAL} \rangle$.

However, there is another important element at the heart of the proposed understanding of epistemic responsibility. Instead of being entitled to expect people to not ignore their higher-order evidence about unpossessed and unaccessed evidence we are entitled to expect epistemic agents *to have such evidence in the first place*. This enables us to distinguish between lucky and irresponsible formed beliefs, such as careless Carl's beliefs, as well as unlucky and responsible formed beliefs, such as careful Caroline's. We judge Carl to be irresponsible because he does not have and utilise higher-order evidence while we judge Caroline responsible because she does.

Let me further illustrate this with another example. For instance, we trust a doctor who makes a diagnosis based on some blood test result because of the doctor's expertise in reading

and interpreting blood test results. This means that we are entitled to expect the doctor to have higher-order evidence supporting the general reliability of blood test results and his ability to interpret them. That is, we expect doctors to possess general information about blood tests via their medical education and specific information about their track record of making diagnoses based on blood tests. After all, if the doctor does not possess or use any such higher-order evidence in making their diagnosis we are entitled to blame them for that. So, for example, if the doctor has not had enough practice in making a diagnosis on blood test results, we expect them to indicate that. That is independent of whether the doctor's failure to use any such higher-order evidence leads him to misjudge the blood test results or not.

This indicates that we should *not* understand epistemic responsibility, while grounded in social expectations, via liability. The doctor is *not* liable for not having and utilising evidence he should have in the case this leads to a misjudgement of the situation but is generally required to have higher-order evidence about her situation. Let us make this more precise. Let us stipulate that Dr. S looked at the blood test results E and concluded that p . Furthermore, Dr. S's belief is also based on a bunch of background evidence E_H about the general reliability of blood tests and her competence with respect to interpreting these test results. That is, E , as well as E_H played a non-deviant causal role in the belief formation of S's belief and/or S would provide or has the disposition to provide E and E_H as reasons for p if asked (depending on the underlying understanding of epistemic basing [§5.6]).

Is there any sense in which Dr. S could be epistemically irresponsible? Let's say there is a professional norm of double-checking test results before making a diagnosis and that Dr. S did *not* double-check. Furthermore, let us stipulate that the patient has a very rare condition Y , rather than X , and that Y could in principle be diagnosed based on the test results (if S would only have double-checked!). If that's the case we might feel the temptation to judge S to be irresponsible.

But what would motivate such a judgement? The most plausible answer here is that in this case, the higher-order evidence E_H possessed by Dr. S is not sufficient to support that the total evidence $\langle E_{TOTAL} \rangle$ supports p . So, either there is a good evidential reason for the norm of double-checking to be in place or not. And if there is, Dr. S's belief that p is not justified, and if not, it is justified. Note, that Dr. S might be legally responsible for not double-checking or we might expect him, non-normatively speaking, to do so. However, epistemically speaking, double-checking needs to be part of the procedure to acquire strong enough higher-order evidence, for double-checking to be a requirement of epistemic responsibility.

Now some might be worried that this response is so flexible that it makes the proposed understanding of epistemic responsibility immune to counterexamples from extant responsibilist theories. After all, we have left open how much or how strong the higher-order evidence E_H needs to be. While I admit that this is a genuine worry, it is important to point out that we only face this worry as much as any other theory that relies on the notion of evidential support. The threshold for E_H to be strong enough is defined by whatever it means for E_H to be sufficient to support q . So, on the explanationist model of epistemic support defended in the previous chapter [§5.5], E_H just needs to be part of a sufficiently good explanation for q , and in so doing raises the probability of q above the required threshold.

In sum, this gives us the following picture of epistemic responsibility. First (i) running the risk of missing relevant evidence by not acquiring higher-order evidence about one's epistemic situation suffices for epistemic irresponsibility. Second, (ii) not running the risk but nonetheless missing evidence does *not* suffice for epistemic irresponsibility. Where responsibility is a necessary condition for having an undefeated and hence justified belief. This is the understanding of epistemic responsibility that underlies the evidentialist no-defeater clause I will specify below [§7.4.2].

7.4.2 A Higher-Order Evidence Clause

We can now turn to my preferred higher-order evidentialist understanding of epistemic responsibility:

S belief that p is *epistemically responsible* iff:

- (2) (a) S possesses some higher-order evidence E_H , which is *sufficient to support* q , the proposition that the total evidence E_{TOTAL} on balance supports p and (b) S's belief that p is *properly based* on E_H .

The attentive reader will have noticed, that this definition of epistemic responsibility is identical to clause (2) of Continuous Evidentialism. Accordingly, I will simply refer to it as (2).

Subclause (2a) states that epistemic agents need to have higher-order evidence E_H that is sufficient to support proposition q , the total evidence on balance supports p . That is, one way to understand clause (2a) is that epistemic agents need to have higher-order evidence that allows them to treat themselves as experts (regarding p) to be justified in believing that p ; whereby

‘expert’ is a placeholder for being generally competent in forming beliefs that are compatible with the total evidence $\langle E_{\text{TOTAL}} \rangle$ (more on that below [§7.4.3]).¹⁷

Let us now turn to subclause (2b). Merely stipulating that S possesses some second-order evidence E_H is not sufficient since this does not guarantee that E_H is involved in the belief formation of the belief that p. Without (2b), E_H could not bear on the epistemic standing of Bp. So, as with beliefs formed via first-order evidence, we need a proper basing requirement [§5.4].

To illustrate this, we might think about careless Carl and careful Caroline again. If careful Caroline is responsible by having and utilising E_H it does not matter whether there is a defeater (possessed or unpossessed) for her *prima facie* justified belief. She is responsible because she tries to ensure that there are no defeaters. And, she did so by utilising sufficient higher-order evidence in her reasoning. In contrast, Carl, who forms beliefs on a whim, is irresponsible, even if he does coincidentally possess some higher-order evidence E_H .

Note that none of this needs to be happening consciously, Caroline could just be acquainted with this procedure in a way that she naturally forms a belief because she has higher-order evidence that supports this belief-forming procedure. Subclause (2b) encapsulates this intuition. It ensures that E_H is not merely possessed by Caroline but actually used. That is, based on different understandings of epistemic basing, clause (2b) could be understood in multiple ways, it could, e.g., mean that the belief that p is (partly) caused by E_H and/or that Sarah when reflecting on her belief, would point towards E_H as part of their justificatory reasons.

With these clarifications about (2a) and (2b) it is easy to demonstrate that (2) gives us the right verdict with respect to the entire taxonomy of defeater cases [§6.6]. Defeated beliefs can be justified independent of the overall evidential balance and the location of the defeater if the belief is partly based on higher-order evidence that sufficiently bears on the overall evidential situation. So, careless Carl and sorrowless Sarah have unjustified yet ‘undefeated’ beliefs, while careful Caroline, meticulous Michael and cautious PI have justified, yet ‘defeated’ beliefs. This is the case because one lacks while the other possesses and uses higher-order evidence that the total relevant evidence $\langle E_{\text{TOTAL}} \rangle$ supports the formed doxastic attitude. That the higher-order evidence is actually (properly) used is important since it suggests that the agent in consideration is in fact responsible and not only possesses sufficient evidence.

¹⁷ So, according to evidentialist theories, such as Continuous Evidentialism E_H would be sufficient to make S’s belief that q *prima facie* justified if S would properly base it on E_H .

7.4.3 Why do we need Higher-Order Evidence?

One urgent question emerging from the proposed understanding of epistemic responsibility is the question of why we are entitled to expect epistemic agents to have higher-order evidence but not first-order evidence. We can put this in terms of an objection to (2). If we are entitled to expect epistemic agents to possess higher-order evidence about their evidential situation on which grounds can we deny Goldberg's initial proposal that we are entitled to expect epistemic agents to have first-order evidence they should have had?¹⁸

I have in part answered this question above by appealing to intuitions about case judgements. While we do *not* seem to be entitled to expect epistemic agents to have some unpossessed evidence they could have had if they have all the reasons to believe that there is no such evidence, we seem to be entitled to expect them to reflect on their epistemic situation. But this answer only pushes the question back one step. It raises the question, of why we seem to be entitled to expect epistemic agents to adequately reflect on their epistemic situation which requires higher-order evidence but not entitled to expect them to have first-order evidence (they 'should' have had)?

While I will not be able to provide a fully developed answer to this question, I think the answer lies in a categorical difference inherent in these expectations. We are entitled to expect from all epistemic agents that they have and utilise higher-order evidence that tells them something about the evidential situation they find themselves in because it is a *general feature* of evidential situations that there likely is relevant further unpossessed evidence.

One way –though probably not the only way– to arrive at the result that there is a general requirement for possessing and utilising higher-order evidence is by looking at what has been called the problem of unpossessed evidence (Ballantyne 2015; Milburn 2023). We are often in situations where we have to lower our confidence or dismiss some doxastic attitude based on the realisation that we might have missed some relevant evidence. Ballantyne uses the following illustrative example (2015: 325):

LIBRARY: You are wandering among rows of bookshelves at the library. These books concern matters about which you hold views. But you've read only a few of them. Let's imagine you think that free will and determinism are compatible, having studied a dozen journal articles and a couple of books years ago in graduate school. Scanning the shelves here, you appreciate that there are several dozen titles relevant to the question of whether compatibilism about free will is true. Some books contain arguments against your view. You hadn't considered this mass of work until now and you haven't yet looked at it.

¹⁸ I want to thank Sandy Goldberg for pushing a version of this objection.

Ballantyne argues that in cases like these, we should dismiss our belief in the compatibility of free will and causal determinism, via a so-called Meta-Defeater Argument.¹⁹ Ballantyne's reasoning goes roughly like this. Realising that there is a bookshelf full of books that bear on my belief that free will is compatible with determinism is evidence for the existence of a defeater for that belief.²⁰ Second, if we stipulate that we have learned from some colleagues or by reading the book titles that some of the books defend incompatibilism, we have evidence of the likely existence of a defeater for my belief. After all, if some books defend incompatibilism there are some books that likely contain defeaters for our belief in compatibilism. So, as long as we do not have a defeater-defeater, for the fact that there is unpossessed evidence that is expected to contain defeaters, we need to consider our belief defeated. How could such a defeater-defeater look like? Ballantyne (2015) lists several options, of which I will mention a few here.

First, we might have evidence that the sources containing the unpossessed evidence are biased and hence unreliable, such as when we are facing a report about the long-term health effects of smoking funded by a tobacco company (2015: 322). Second, there might be some beliefs we have epistemic authority over, such as introspective beliefs (2015: 326). Third, there are some matters we are experts about. Having written a PhD thesis on the correspondence theory of truth, being confident that you have almost read all of the relevant literature on the topic enables you to form justified beliefs on the subject matter, even if you find out that there are some books (arguing against the correspondence theory) you have not read. There you can reasonably accept that you possess most of the relevant evidence or at least that the evidence you possess is a representative sample of the entire evidence relevant to the proposition (2015: 329).

While this list of possible defeater-defeaters illustrates that we can sometimes avoid the sceptical threat posed by unpossessed evidence, it should be equally clear that we often cannot appeal to anything that would undermine the defeating force of unpossessed evidence. Or as Ballantyne puts it “we have reason to accept D1–D3 [the defeater-defeaters] only if we are in a rather strong evidential position, and we'll enjoy such a position only in unusual circumstances. Thus, often enough, [...] we will not have those defeater-defeaters, and the Meta-defeater Argument tells us why sticking with our opinions is not (fully) rational.” (2015: 324).

¹⁹ More precisely, Ballantyne discusses three arguments, which he calls the Meta-Defeater Argument, the Overlooked Defeater Argument and the Doubtful Fairness Argument. For a critical discussion of Ballantyne's arguments and an improved argument with a similarly sceptical conclusion see Milburn (2023).

²⁰ Here Ballantyne adheres to above-mentioned principle that evidence for evidence is evidence (Feldman 2007; 2014).

I think we can use Ballantyne's insight to motivate a general requirement to possess and utilise higher-order evidence about one's evidential situation when forming beliefs. If acknowledging that there is unpossessed defeating evidence undermines the justificatory status of our doxastic attitudes, we can reasonably expect epistemic agents to possess some kind of higher-order evidence that protects them from this acknowledgement. This aligns with the above-mentioned way to understand clause (2) as a requirement to have higher-order evidence that allows epistemic agents to treat themselves as experts (regarding *p*) to be justified in believing that *p*; where expert was understood as a placeholder for someone who has sufficient higher-order evidence about the subject matter.

This highlights the above-mentioned fundamental difference between Goldberg's and my proposal again. We are entitled to expect epistemic agents to have higher-order evidence about their situation because it is a *general feature* of every epistemic situation that there could be evidence that would have made a difference. And we can reasonably expect epistemic agents to acknowledge that, or at least behave in a way that takes this fact into account. In contrast, we can never tell (without higher-order evidence about the unpossessed evidence) what the evidence we do not possess supports, whether it is defeating, supporting or neutral with respect to our beliefs. Hence, epistemic agents can form their beliefs responsibly or irresponsibly in the former respect but not in the latter.

Importantly, this also makes the proposed understanding of responsibility compatible with an internalist understanding of justification (more on that in [§8.2]). In contrast, the kind and amount of evidence that we 'should have had' according to Goldberg's understanding differs from situation to situation and hence depends on coincidental features of the evidential situation the epistemic agent may or may not be aware of. This would make the justification of one's beliefs dependent on external factors in a way that is incompatible with the general internalist framework Continuous Evidentialism relies on.

We can now take a step back and reflect on the initial idea of the defeasibility of justified belief which motivated our search for a no-defeater clause more generally and our analysis of epistemic responsibility in specific.

7.5 Defeatism Defeated?

In the previous chapter, we turned to epistemic responsibility because it was a promising candidate in our search for a no-defeater clause. This search was motivated by the following doctrine:

Defeatism: Doxastic attitudes can have the status of being *prima facie* justified. That is, some doxastic attitude *D* can be *prima facie* justified by having some justification-conferring property while being defeated and, therefore, lacking *ultima facie* justification.

While prominent, Defeatism has recently been attacked on various fronts (Bergmann 2005; Lasonen-Aarnio 2010, 2014, 2020; Baker-Hytch & Benton 2015). While I won't discuss these arguments here let me illustrate how the responsibilist higher-order evidential requirement (2) bypasses the notion of *prima facie* justification and defeasibility.

Strictly speaking, in accepting (2) as a unified responsibilist solution to defeater cases we ironically do away with the notion of epistemic defeat. After all, there are no *prima facie* justified beliefs which are subsequently defeated. Instead, higher-order evidential requirements are a necessary condition for justification and once a belief fulfils whatever the first-order conditions are plus the specified higher-order requirements the belief is justified, independently of whether there is a defeater or not (within or without $\langle E_{\text{BASE}} \rangle$). This is in stark contrast to the other no-defeater clauses discussed, which all, one way or another, make justification dependent on the *de facto* existence of a defeater. In so doing, they rely on the notion of *prima facie* justification, the idea that beliefs that fulfil the first-order requirements of justification are justified but defeasible. This is an important result since it makes the proposed account sensible to the phenomenon of defeat by retaining the intuitive judgements we have towards alleged defeater cases while avoiding general problems put forward against Defeatism.

As a consequence, my treatment of defeater cases makes it the case that strictly speaking there is no evidence we *should* –epistemically speaking– have possessed or have accessed (except the higher-order evidence E_H that is required to be responsible).²¹ In that sense, most of the distinctions drawn in [§6.2] and [§6.3] are redundant. The only set of evidence we need to define to analyse the phenomenon of epistemic defeat is $\langle E_{\text{TOTAL}} \rangle$, the entire evidence that is relevant

²¹ Note, that we should have E_H is only a normative requirement to possess evidence in as much as all evidentialist theories require one to possess evidence that is sufficient to support our beliefs. So this might not be particularly worrisome to the evidentialist.

to an epistemic situation, to judge whether some evidence E_H is sufficient to support that, $\langle E_{TOTAL} \rangle$ is sufficient to support p .

Let me stress that, in so doing we deny that there are non-negligent forms of the phenomenon that people have called normative defeat. That is cases in which one has sufficient higher-order evidence about one's epistemic situation while this evidence is ignorant with respect to some relevant piece of evidence they could have possessed.²² Or to put it in yet another way, while the need for higher-order evidence arises from the problem of unpossessed evidence [§7.4.3], there is *no* specific problem having to do with unpossessed evidence more narrowly construed, that is, the unpossessed evidence we should have had.

The difference between the defended higher-order evidence clause (2) and conventional no-defeater clauses can be highlighted by distinguishing between negative and positive higher-order clauses or requirements:

Negative Higher-Order Requirement $\langle HOR-\rangle$: *S's prima facie* justified belief that p is *ultima facie* justified iff S lacks a *defeater* ϕ [that S should have had], that is [a belief/a justified belief/evidence...] *indicating* that the belief that p is *not prima facie* justified.

Positive Higher-Order Requirement $\langle HOR+\rangle$: If S 's belief that p is justified then S *necessarily* has a *defender* Φ , that is [a belief/a justified belief/ evidence/...] *indicating* that p is justified.

We can see that traditional no-defeater clauses take the form of negative higher-order requirements $\langle HOR-\rangle$. This includes the evidentialist clauses discussed in [ch.6], as well as some responsibility clauses such as Goldberg's understanding of epistemic responsibility as strict liability, discussed above [§7.2]. As such, $\langle HOR-\rangle$ directly rely on Defeatism, the idea that there are *prima facie* justified beliefs that lose their justificatory status in case certain defeating conditions are obtained.

In contrast, the higher-order evidentialist clause (2) takes the form of a positive higher-order requirement $\langle HOR+\rangle$ which neither relies on the notion of defeasibility, nor on *prima facie* justification. Instead, $\langle HOR+\rangle$ only requires doxastic attitudes to have an additional property, a *defender* Φ on top of the properties specified by the first-order requirement in order to be justified.

²² Such cases are, for example, discussed in Harman (1980).

While one kind of defender is the higher-order evidence E_H as specified in (2), we will encounter other possible defenders in the next chapter [ch.8].

7.6 Conclusion

In this chapter, I have shown that we can reduce epistemic responsibility to higher-order evidential requirements. In particular, I followed Goldberg's framework which understands epistemic responsibility as being grounded in socio-epistemic expectations [§7.2]. However, in contrast to Goldberg, I did not understand these expectations as making us strictly liable to the evidence we should have had [§7.3]. Instead, I argued that we are entitled to expect from epistemic agents that they possess and utilise higher-order evidence when forming beliefs. This enabled me to reduce epistemic responsibility to higher-order evidential requirements [§7.4]. Doing so, while retaining intuitive judgements about defeater cases, required us to do away with the notion of defeat as traditionally understood since no belief that is formed under the specified conditions can be defeated anymore [§7.5].

By reducing responsibility to evidential requirements, we can retain a wholly evidentialist understanding of justification. Accordingly, (2) if understood as a second-order clause of an evidentialist account of justification, such as Continuous Evidentialism gives us a purely evidentialist theory. This is important not only for those who have prior preferences for evidentialist views of justification but also for those who are worried about problems associated with hybrid views, combining responsibilist, evidentialist, and/or reliabilist elements (Cloos 2015; Goldberg 2018; Silva 2019; Lackey 2016).²³

Yet, cashing out epistemic responsibility in terms of higher-order evidential requirements (2) leaves us with a theory of justification that relies on positive higher-order requirement <HOR+>. These requirements face their own problems. I will address these problems in the next chapter [ch.8].

²³ In particular, avoiding a hybrid understanding of justification by appeal to a responsibilist understanding of *prima facie* justification is unpromising, due to the manifold problems associated with such responsibilist accounts (Baehr 2011; Sylvan forthcoming).

8

Higher-Order Requirements:

In Defense of a Two-tiered Theory of Justification

In the previous chapter [ch.7], I defended a higher-order evidentialist understanding of epistemic responsibility. This enabled me to motivate clause (2) of Continuous Evidentialism, which is necessary to make sense of the phenomenon of epistemic defeat. Yet, it does so in an unconventional way. Instead of adding a negative higher-order requirement in the form of a no-defeater clause to our theory of epistemic justification, (2) is a positive higher-order requirement [§7.5]. Solving the problem of defeat via such a positive higher-order requirement enabled us to avoid various problems that are usually associated with Defeatism, the doctrine that we can have *prima facie* justified but defeasible beliefs [§6.1].

In this chapter, I want to address another class of problems that are often put forward against positive higher-order requirements. I will do so by discussing different kinds of positive higher-order requirements, systematically identifying the problems they face and demonstrating that (2), and with it Continuous Evidentialism, is immune to them.

8.1 Introduction

As illustrated in [ch.6], rational or epistemically justified beliefs are typically understood to be defeasible. In the previous chapter, I followed in the footsteps of those who find that idea wanting [§7.5]. In particular, I have argued that there are good reasons to be sceptical that making these general restrictions allows us to account for the phenomenon of defeat. Accordingly, a more promising way to proceed, which also aligns with our intuitive judgments about the wide range of so-called defeater cases discussed in [ch.6], is to come up with a higher-order requirement. The aim of such requirements is not to rule out defeater cases in general but to ensure that our theory of epistemic justification makes having justified beliefs in the face of a defeater less likely. I have called such a condition a defender Φ . Taking this idea we can formulate positive higher-order requirements as follows:

Positive Higher-Order Requirement <HOR+>: If S's belief that *p* is justified then S necessarily has a *defender* Φ , that is [a belief/a justified belief/ evidence/...] indicating that *p* is justified.

Positive higher-order requirements are promising. They allow us to account for the phenomenon of defeat, without running into general problems faced by Defeatism [§7.5]. This, however, raises the question of why we want to have any kind of higher-order requirement built into our theory of justification. Why aren't first-order conditions, such as the ones spelt out in [ch.5] enough? After all, many who have argued against Defeatism have drawn the stronger conclusion that we should do away with higher-order requirements in general.¹

To address this worry, I will start by giving some additional motivations for thinking that we need to have some higher-order requirements built into our theory of epistemic justification [§8.2]. Next, I will give an overview of different types of higher-order requirements, that have been proposed in the literature, including doxastic, non-doxastic, and dispositional requirements [§8.3]. In so doing, I will identify three worries often associated with such requirements. Namely, the charge of overintellectualisation [§8.4.1], possible epistemic irrelevance [§8.4.2], as well as a lurking infinite regress [§8.4.3]. I will conclude that none of the doxastic, non-doxastic or dispositional requirements proposed in the literature provide a satisfactory reply to all of these worries. In [§8.4] I will demonstrate that the higher-order evidential requirement (2) defended in the last chapter can avoid these worries.

8.2 Defencism Defended

Before discussing any higher-order requirements in specific, we need to motivate the general need for why we want our theory of epistemic justification to have such a requirement. This is important for two reasons. First, if we are sceptical of Defeatism we might want to extend that scepticism to higher-order requirements in general (<HOR+> and <HOR->). And since I have raised general scepticism about Defeatism, I need to offer a reason why I think that making this extension is not warranted. Second, as we will see below [§8.4], many criticise positive higher-order requirements by questioning the epistemic benefit they provide. So, one of the central challenges to higher-order requirements is answering the question of how fulfilling the higher-

¹ See, for example, Lasonen-Aarnio (2014: 331-334)

order requirements makes a belief epistemically superior to ones that merely fulfil the first-order requirements.

Let us call the doctrine that any complete theory of epistemic justification needs to contain <HOR+> *Defencism*. In this section, I will motivate Defencism and with it the need for <HOR+>, by identifying two intuitions that any justified belief needs to fulfil higher-order requirements which I will then use to identify two desiderata [§8.2.1]. Afterwards, I will make some clarifications as to why <HOR+> cannot be replaced by a theory of epistemic basing [§8.2.2].

8.2.1 Higher-Order Requirements: Two Desiderata

Many think that a complete theory of epistemic justification needs to be able to distinguish properly and improperly formed *inferential* beliefs. For example, Fumerton (1995; 2004) argues that we need to differentiate cases in which someone infers a proposition *p* from a piece of evidence *E* while they are aware of the fact that *E* entails or supports *p* and cases in which someone is merely “caused to believe *p* as a result of *E* where *E* does in fact entail *p*, but where the entailment is far too complicated for *S* to understand” (2004: 154).² In other words, we need to distinguish properly inferred beliefs from the ones reached merely by “mental jogging”, as Broome (2013: 234) calls it.³ Tucker summarises Fumerton’s idea as follows (2012: 325):

“Fumerton’s argument [...] contrasts two cases where a subject accepts *E* therefore *p*, one in which the subject, call him ‘Seer’ sees that *E* supports *p*, and one where another subject, call him ‘Non-Seer’ doesn’t see that *E* supports *p*. Since *E* therefore *p* seems capable of justifying its conclusion only for the subject who sees that *E* supports *p*, we are supposed to infer the following explanation of the Seer and Non-Seer cases: *E* therefore *p* justifies its conclusion for only Seer because only Seer sees that *E* supports.”⁴

If we take this intuition seriously, we want to have an explanation for why we judge Seer to be justified, and Non-Seer to be unjustified, or at least have an explanation for why we have that intuition. One popular way to do this is via a higher-order requirement. Here we could distinguish between Seer and Non-Seer cases by referring to some higher-order awareness requirement for justified beliefs, which requires agents to have some kind of (non-doxastic)

² P changed to p.

³ For further discussion see also Wright (2014: 33).

⁴ P changed to p.

mental state indicating that ‘E supports p’ when successfully inferring p from E (Tucker 2012; Boghossian 2014; Chudnoff 2014; Valaris 2016).⁵

This gives us the first desideratum for our <HOR+>:

Desideratum 1: Higher-order requirements need to enable us to distinguish between *proper* and *improper inferential* beliefs; that is, between Seer and Non-Seer cases.

Yet, accepting the need for a higher-order requirement for justification in the case of inferential beliefs, we have not established that higher-order requirements are general requirements for justification.

Here, and this is the more controversial part, I think that there are good reasons to take higher-order requirements to be general. We can motivate this via the following thought experiment:

SENSORY SUBSTITUTION: Sandra, who is congenitally blind, grew up in a community of congenitally blind people. Sandra has never heard of eyesight before nor has she ever encountered a non-blind person. One day she walks into a room and finds a specific kind of sensory substitution device, a tactile-visual sensory substitution waistcoat (in short TVSS). After putting the TVSS on, the device converts the image from a camera into a tactile image on the skin of Sandra, pixel by pixel using an electrode array. After wearing it for a few minutes she has experiences phenomenologically similar to visual experiences. Based on these experiences Sandra forms the belief that there is a round object in front of her. She does not confirm or disconfirm this belief by using different sense modalities.

When using the TVSS for the first time Sandra has no evidence about the reliability of visual perception nor does she even know that such a sense modality exists. This raises the following question: Can Sandra, upon ‘looking’ at the external world for the first time form justified beliefs based on visual perception?⁶ Let us stipulate that she does form a perceptual belief that there is

⁵ Note that on some understandings of the epistemic support relation, such as some versions of explanationism [§5.5.2], epistemic support contains a higher-order requirement. Many explanationists are committed to the idea that for some evidence E to support a belief the subject must have awareness of some kind indicating that the explanatory relations that hold between the belief and the evidence (Conee & Feldman 2008; McCain 2014: ch.4). However, on my explanationist understanding of epistemic support no such awareness is required. Here, I followed Fratantonio (forthcoming), who indicated that no higher-order awareness of the explanatory relation seems required [§5.5.2]. For a discussion of McCain’s (2014) understanding of awareness in the context of explanationism see Lutz (2020: 2637). I will come back to this below [§8.5].

⁶ Note that giving an answer to this question relates to what has been called Molyneux’s problem (Denegar & Lokhorst 2021). The problem raised by William Molyneux in a letter to John Locke, in which Molyneux asks

a round object in the room similar to how non-blind persons do. After all, she has the cognitive capacity to do so. Furthermore, as experiments have shown, when congenitally blind persons use TVSS devices similar brain areas to those who are responsible for visual perception are activated. Moreover, in experiments congenitally blind persons were able to rapidly learn to interpret the tactile stimulations and use them to detect simple shapes and orient themselves.

So let's say there is a round table in the room, and further, that Sandra forms the belief that there is a round in front of her object based on a vague visual impression of that table. Intuitively, this belief cannot be justified as regular perceptual beliefs are justified. For it to be justified Sandra would have to confirm it by using some of her other sense modalities. Furthermore, she would likely need to perform several tests until she has some feel for the general reliability of this newly acquired way of perceiving the world. We can derive this judgment from a very central intuition underlying the internalist project of epistemic justification. I have spelt out this intuition in [§5.3] as follows:

Internalist Intuition: Epistemic agents can have justified beliefs independently of whether they are in a lucky or an unlucky case.

As mentioned in [§5.3], the Internalist Intuition can be supported by two kinds of cases. First, cases in which a suspect is unlucky but forms beliefs in a seemingly proper way, as in the new evil demon case (Cohen & Lehrer 1983). Second, cases in which an epistemic agent is lucky but seemingly reasons improperly, such as in Bonjour's (1980) reliable clairvoyant case or Lehrer's (1990) Mr. Truetemp case. Here is Bonjour's famous case (1980: 41):

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

whether a congenitally blind person who can distinguish geometrical shapes by touch, would be able to distinguish and name these shapes simply by sight if they would acquire visual perception.

The basic idea is this. If we judge Norman the reliable clairvoyant who forms beliefs via clairvoyance without having reasons to regard his clairvoyance as reliable as having *no* justified beliefs, we should judge seeing Sandra's beliefs to be unjustified as well.⁷ After all, for an agent like seeing Sandra, there is no difference between beliefs formed via visual perception and beliefs formed via clairvoyance. How could there be? Seeing Sandra does not possess any evidence about her newly acquired belief-forming mechanism. So, as I will argue below, for seeing Sandra to form beliefs via perception, Sandra needs to have and utilise some higher-order evidence about perception. But, regardless of whether that is the correct explanation or not, taking the Internalist Intuition seriously is to say that there needs to be some kind of explanation that enables us to judge regular beliefs formed via visual perception to be justified and seeing Sandra's beliefs to be unjustified.

Here is another way to put it. At some point, call it t_1 , seeing Sandra if she continues to wear the TVSS device can form justified beliefs via visual perception. In this respect, Sandra is similar to the reliable clairvoyant who can, at some point t_1 , form justified clairvoyance beliefs. Now the interesting question is what happens at t_1 or what distinguishes the beliefs formed after t_1 and before t_1 . Here different higher-order requirements will give us different answers. For example, we might say, following BonJour (1980), that after t_1 both Sandra and Norman can form justified beliefs because they have higher-order beliefs that support that their respective belief-formation method is reliable.⁸ Another possibility, which is the one which I defend, is that at t_1 Sandra and Norman have acquired higher-order evidence which they can then utilise to form justified beliefs via their respective method.

Independently, however, of the details of this answer, the only thing we need to accept the second desideratum is that our theory of justification is in need of giving us some explanation that explains these intuitive judgements:⁹

Desideratum 2: Higher-order requirements need to enable us to distinguish between *proper* and *improper non-inferential* beliefs.

⁷ Importantly, many externalists seem to have the clairvoyance intuition as well. See, e.g., Goldberg (2018: 34-44).

⁸ This is one of the lessons that BonJour (1980) draws from the case.

⁹ Note that accepting Desideratum 2 is to deny that there are any baseless but justified beliefs, which is controversial even among evidentialists. See, for example, Silva (2019: 274; fn 20).

Importantly, as I will discuss in [ch.9] accepting Desideratum 2 can also be motivated when thinking about certain cases in collective epistemology. Namely cases in which a newly assembled collective agent *G* starts to form beliefs via a novel belief-forming mechanism *B* whereby this agent does not possess any information about the reliability of *B*. In these cases, we should judge *G*'s beliefs to be unjustified.¹⁰

8.2.2 Epistemic Basing and Higher-Order Requirements

Before plunging ahead a clarification is needed. Let me specify the relationship between what I call positive higher-order requirements <HOR+> and epistemic basing. As pointed out in various places, many theorists employ the distinction between propositional and doxastic justification to distinguish between well-founded and non-well-founded beliefs (Conee & Feldman 2004), *ex-ante* justified and *ex-post* justified beliefs (Goldman 1979) or properly and improperly based beliefs (Turri 2010b). While I rejected various orthodox and unorthodox understandings of the relationship between propositional and doxastic justification [§6.7], I still endorse a basing requirement for doxastic justification. In particular, I opted for a causal understanding of basing for reasons specified in [§5.6]. Accordingly, one way to understand the above-presented arguments in favour of <HOR+> is as arguments that we need to go beyond a causal understanding of basing. That is, for a belief *B* to be based on some evidence *E* we need some additional non-causal requirement such as a meta-belief or some higher-order evidence to fix the basing relation between *B* and *E*. However, there are at least two complications which make this comparison unhelpful. First, those sympathetic to hybrid understandings of basing typically understand basing to be disjunctive; i.e. that either sufficient causal requirements or non-causal doxastic requirements are met (Korcz 2000). In contrast, the defended requirement would be conjunctive. Second, accepting a basing requirement on justification does not necessarily commit one to evidentialism or a reasons-based understanding of justification. That is, defending a hybrid understanding of the basing relation is in principle compatible with baseless but justified beliefs.¹¹ However, if we accept that Desideratum 2 motivates a general higher-order requirement we accept that there are no baseless but justified beliefs.

¹⁰ In other words, Desideratum 2 urges us to ask whether newly assembled collective epistemic agents bear similarities to Norman the Clairvoyant or seeing Sandra. I will discuss such a case in [§9.3.4].

¹¹ See, for example, Turri (2011).

8.2.3 Summary

In sum, the role of higher-order requirements is twofold. First, we need them to distinguish properly inferred from improperly inferred beliefs (Desideratum 1). Seer is justified because she sees that E supports p by having a *defender* Φ [belief/justified belief/ evidence/...] *indicating* that E supports p while Non-Seer does not. Second, we need them to distinguish proper and improper non-inferential beliefs (Desideratum 2). Non-blind persons are justified in forming beliefs via visual perception because they have a *defender* Φ [belief/justified belief/ evidence/...] *indicating* that the way they formed their belief is justified, while seeing Sandra does not. Accordingly, many epistemologists have defended higher-order requirements for inferential beliefs (Fumerton 1995, 2004; Tucker 2012; Boghossian 2014; Chudnoff 2014; Valaris 2016) or as general requirements for justification (BonJour 1980; Leite 2008; McCain 2014).

Despite their alleged usefulness in making sense of various normative judgements, many have pointed out that higher-order requirements are problematic for various reasons.¹² For example, many think that if we require epistemic agents to have higher-order cognitive states we are in danger of *overintellectualising* epistemic justification. Second, many higher-order requirements, especially doxastic ones, have been shown to lead to a *vicious regress*. Third, many have questioned the motivation behind higher-order requirements and the presumed *epistemic benefit* they provide. Based on these, and other to-be-discussed worries, many epistemologists deny that there are any higher-order requirements for justification or at least that higher-order requirements are *not* necessary for every kind of justified belief.¹³

8.3 Higher-Order Requirements: Overview

To discuss different kinds of positive higher-order requirements and the problems they face systematically we need some broad categorisations. Here is the general characterisation of <HOR+> again:

¹² See, for example, Bergmann (2005), Lasonen-Aarano (2014), McHugh and Way (2016), or Baker-Hyatt and Benton (2015).

¹³ Bergmann, for example, argues that while <HOR+> is not a necessary requirement for justification it “is often contingently the case that in order for beliefs to avoid having defeaters, they must satisfy a higher-level requirement” (2005: 420). Others explicitly defend a <HOR+> for inferential beliefs while remaining neutral or rejecting a general <HOR+> (Tucker 2012; Boghossian 2014).

<HOR+>: If S's belief that p is justified then S *necessarily* has a *defender* Φ , that is a [belief/justified belief/ evidence/...] *indicating* that p is justified.

If we look at this characterisation we can see that <HOR+> can take various forms. In particular, we can distinguish four different versions by distinguishing four different kinds of *defenders* Φ :

<HOR+ doxastic>: If S's belief that p is justified then S *necessarily* has Φ , a *doxastic* mental state that *indicates* that p is justified.

<HOR+ mental>: If S's belief that p is justified then S *necessarily* has Φ , a *non-doxastic* mental state that *indicates* that p is justified.

<HOR+ dispositional>: If S's belief that p is justified then S *necessarily* has Φ , a *disposition* to form beliefs in a way that *indicates* that p is justified, and this disposition caused S to believe that p.

<HOR+ evidential>: If S's belief that p is justified then S *necessarily* has Φ , sufficient *evidence* that *indicates* that p is justified.

One detail that this categorisation leaves open is what is needed for a defender to *indicate* that p is justified. Here different versions of different <HOR+> will provide us with different answers to that question. I will work out these answers in the next section when I discuss potential worries faced by these requirements [§8.4]. Before doing so, however, let me make some clarification on how we should understand the different types of requirements identified above.

<HOR+ doxastic>: The basic idea behind <HOR+ doxastic> is simple. For S to be justified to believe that p, the belief does not only need to fulfil some first-order requirements, such as being based on sufficient evidence, but also S needs to have a defender Φ in the form of another (justified) belief that indicates that the belief that p is justified (or at least have the disposition to form such a belief upon reflection). For example, to be justified in believing that it is cloudy based on a quick glance through your office window you need to have an accompanying higher-order (justified) belief that indicates that visual perceptions under these conditions are reliable. One encounters such requirements often in the defence of a doxastic (Audi 1993: 233-273; Leite 2008) or causal-doxastic (Korcz 2000) requirement of epistemic basing.¹⁴ Leite, for example, states

¹⁴ For an early defence of <HOR+ doxastic> outside the context of epistemic basing see Bonjour (1980).

that “in order for one to have positive epistemic status δ in virtue of believing p on the basis of R , one must believe that R evidentially supports p , and one must have positive epistemic status δ in relation to that later belief as well” (2008: 422).

<HOR+ mental>: Instead of requiring an accompanying belief or doxastic state <HOR+ mental> only requires epistemic agents to have some accompanying mental state that falls short of being genuinely doxastic. In the context of inferential beliefs <HOR+ mental> requirements have recently seen a revival. In order to explain the difference between Seer and Non-Seer cases, many authors have argued that epistemic agents need to *take* the reasons for which a belief is formed (premises) to support the inferential belief (conclusion) (Tucker 2012; Boghossian 2014; Chudnoff 2014; Valaris 2016).¹⁵ Boghossian speaks of a taking condition for inference, whereby S successfully “inferring from p to q is for S to judge that q because S takes the (presumed truth of) p to provide support for q ” (2014: 4).¹⁶ However, while many defend a condition similar to Boghossian’s taking condition, there is significant disagreement about what this taking p to provide support for p requires. For example, Chudnoff (2014) understands taking to mean having an intuition or seeming that E supports p . Boghossian (2014), on the other hand, understands taking to be a *sui generis* mental state that is not only distinct from full-blown belief or a disposition to have such a belief but also from other mental states such as intuitions. In contrast, Tucker (2012), takes the relevant relation to be one of awareness, while awareness can take multiple forms.¹⁷ That is, for Tucker, I can be “aware of E ’s supporting P by believing (justifiably or unjustifiably) that E supports P , by its seeming to me that E supports P , or by being acquainted with the fact that E supports P ” (2012: 326-327).¹⁸

<HOR+ dispositional>: In contrast to the other higher-order requirements discussed <HOR+ dispositional> do not require the epistemic agent to have a defender Φ in the form of a

¹⁵ Although <HOR+ mental> requirements have been defended as general requirements for epistemic justification, they are, nowadays, most commonly defended when discussing inferential beliefs (Tucker 2012; Chudnoff 2014; Boghossian 2014, Valaris 2016). Accordingly, these accounts are what I will be focusing on in the following discussion. Despite this, I think that most of the dialectical structure can be generalised to more extensive mental higher-order requirements covering inferential and non-inferential beliefs.

¹⁶ In earlier works Boghossian defended hybrid dispositional-doxastic understanding of inference whereas inference involves that “ S is able to know by reflection alone that his premises provide him with a good reason for believing the conclusion” (2003: 268).

¹⁷ Another place where <HOR+ mental> show up is in explanationist understandings of the epistemic support relation. For example, McCain argues that disposition “to have a seeming that, for example, the truth of p and if p , then q ensure the truth of q .” (2014: 78).

¹⁸ Here acquaintance to an object is understood following Russell as “a direct cognitive relation to that object” which does not “constitute[s] judgment, but [...] which constitutes presentation” (1911: 108).

mental state but to have some disposition to form their attitudes in a particular way. The basic idea behind <HOR+ dispositional> is that something within a well-functioning cognitive architecture manifests a disposition to react in a certain way when epistemic agents successfully form beliefs or engage in reasoning. An early account of such a higher-order requirement is found in Greco, who suggests that justification requires countenance of the rules of reasoning (1993, 1999).¹⁹ For Greco “a belief *p* is subjectively justified for a person *S* [...] if and only if *S*'s believing *p* is grounded in the cognitive dispositions that *S* manifests when *S* is thinking conscientiously”; where “by ‘thinking conscientiously’ I intend the usual state that most people are in as a kind of default mode the state of trying to form one's beliefs accurately” (1999: 289).

<HOR+ evidential>: Depending on different understandings of the nature of evidence <HOR+ evidential> will take different forms. If we need to have some kind of doxastic relation to the evidence in order to possess it, <HOR+ evidential> will be similar to doxastic requirements. If evidence, however, is understood in mentalist or abstractionist terms, evidential requirements will be very close to mental requirements. However, while similar in contrast to <HOR+ mental>, we do not need any special kind of mental state that needs to indicate that *p* is justified (since many kinds of mental states could play the required evidential role). One example of an evidential requirement is clause (2) the higher-order evidentialist clause of Continuous Evidentialism.

8.4 Higher-Order Requirements: Worries

Having identified different kinds of higher-order requirements we can now discuss various worries raised against them and see how well the respective versions do concerning these worries. In particular, I will discuss three different worries: over-intellectualisation [§8.4.1], the charge of epistemic irrelevance [§8.4.2], as well as worries having to do with lurking vicious regresses [§8.4.3].

8.4.1 Overintellectualisation

One urgent and often-discussed problem with <HOR+> is the so-called *overintellectualisation* worry, or as I prefer to call it the *think-of-the-children* worry.

¹⁹ A more recent defence of a dispositional requirement is found in Sosa (2015: ch.1).

Worry 1: Overintellectualisation: We often judge less cognitively sophisticated epistemic agents (such as children, non-human animals, or certain types of collectives) to have justified beliefs. Higher-order requirements that conflict with such judgements implausibly overintellectualise our theory of epistemic justification.

Let us start with <HOR+ doxastic>. Having a doxastic requirement built into your account of epistemic justification makes it the case that less cognitively sophisticated epistemic agents such as small children or animals, who allegedly lack the capacity to form higher-order beliefs, cannot have any justified beliefs.²⁰ Likewise, some groups we intuitively judge to have justified beliefs seem to lack this requirement as well, since they do not have the right institutional structure to reflect on their belief formation processes.

Many think this result is problematic since we usually judge these epistemic agents to be able to acquire justified beliefs or knowledge via perception, simple inferences or testimony. Similarly, we think that in many instances also sophisticated epistemic agents have direct knowledge or justified beliefs about a vast number of propositions via, for example, immediate perception or rational insight while lacking any higher-order cognitive states about their doxastic attitudes.²¹

To address this second point, we might want to weaken the requirement such that we only require the epistemic agent to have a *disposition* to form the respective higher-order mental state when reflecting on one's belief-formation process.²² This is reasonable since it tracks the way epistemic agents usually reflect on their belief formation processes, when asked. So when asked why we believe that there is a bumblebee on the balcony we likely are disposed to give certain answers to that questions having to do with visual perception, distinctive characteristics of bumblebees and so on. However, it does not dismantle the worry with respect to the judgements we have towards children and other less cognitively sophisticated agents. If those do not have the

²⁰ For general discussions see, for example, Greco (1993; 1999), Mc Hugh and Way (2016), Bohgossian (2014: 6-7) or Berghofer (2022: 78). For parallel discussions about testimonial knowledge/justification see, for example, Lackey (2005) or Goldberg (2008).

²¹ For example, phenomenal conservativists (Huemer 2001), think that seeming provides *prima facie* justification without the need for an higher-order belief state such as the belief that seemings provide justification. As a result, phenomenal conservativists rely on a negative no-defeater clause, which I considered problematic [§7.5]. For further discussion see Hawthorne and Lasonen-Aarnio (2021).

²² Note that these dispositional accounts are different from dispositional accounts <HOR+ dispositional> which require an agent to exercise a cognitive disposition when forming their beliefs. I will discuss <HOR+ dispositional> below.

capacity to form such higher-order mental states they can neither have the disposition to do so.

23

This pushes us towards non-doxastic requirements. Non-doxastic requirements of the structure <HOR+ mental> do not face the overintellectualisation problem to the same extent. After all, the required higher-order mental state does not need to be as cognitively sophisticated as the higher-order doxastic attitude. And as we have seen above, there are many possible candidate mental states including seemings (Chudnoff 2014), higher-order awareness (Tucker 2012) or a *sui generis* taking state (Boghossian 2014).

Keeping the disagreements about the required mental state in mind, we need to ask ourselves what is needed for <HOR+ mental> to play its intended role. As McHugh and Way (2014), as well as Wright (2014), point out, any mental state proposed needs to be understood as an attitude with content which helps to explain how epistemic agents come to believe what they believe. That is, in the context of inferential beliefs, whatever the required mental state is, “it is not just a placeholder for whatever it is that distinguishes inference from mere causation. Nor is saying that you take your premises to support your conclusion just another way of saying that you reason from the former to the latter” (McHugh & Way 2014: 317). It is only when understood in these more substantive terms a <HOR+ mental> can place restrictions on justified beliefs. However, this brings us back to the overintellectualisation worry. If the accompanying mental state needs to have some content this requires the epistemic agent to have such a contentful higher-order mental state, which is a capacity that at least some of the less cognitively sophisticated agents we judge to have justified beliefs lack. So, neither <HOR+ doxastic> nor <HOR+ mental> avoid the overintellectualisation worry completely.

But what about <HOR+ dispositional>? Take for example Greco’s (1993, 1999) proposal that justification requires countenance of the rules of reasoning. Tucker portrays Greco’s dispositional requirement as follows (2012: 330):

Greco’s Requirement: S’s argument E therefore p can justify S’s belief in p only if S’s believing that p is the result of cognitive dispositions that S manifests when S thinks conscientiously.

²³ This is for example argued by Berghofer (2022: 78) in his critical investigation of McCain’s explanationist understanding of epistemic support. McCain argues that “to understand something as an answer to a why-question or the disposition to have a seeming that, for example, the truth of p and if p, then q ensure the truth of q” (Berghofer 2014: 78).

Greco's proposal is promising since it not only gives us an intuitive understanding of what distinguishes Seer from Non-Seer cases but it also seems to bypass the overintellectualisation worry completely. That is, as Greco (1993: 424) points out, countenance with the rules of logic as understood above is compatible with small children having justified beliefs. The idea here seems to be that these kinds of dispositions do not require us to have concepts such as 'argument', 'premises', or 'entailment' nor any kind of additional mental state that secures the connection between E and p. All that is required to have the disposition to reason from E to p in instances in which these inferences are proper, whereby this may involve explicit reasoning and/or conceptual understanding or not. In Greco's words (1993: 424-425):

“Typically only logicians have beliefs about the deductive rules which we use in our reasoning, and it is agreed on all sides that no one has successfully characterised the rules which govern our non-deductive reasoning. But if we typically do not have beliefs about the rules which govern correct reasoning, how are we to include sensitivity to such rules [...] I would suggest that although we do not typically have beliefs about such rules, we do countenance such rules in our reasoning.”

While <HOR+ dispositional> are immune to the overintellectualisation worry I will demonstrate below that these requirements fail to explain the epistemic relevance of higher-order requirements [§8.4.2]. However, before doing so, let me make some general remarks about the overintellectualisation worry.

As I understand it, many motivations behind the overintellectualisation worry rely on more fundamental intuitions regarding the general nature of epistemic justification. Every account of justification needs to make some cut-off when it comes to the question of how cognitively sophisticated systems need to be to count as epistemic agents, that is, potential bearers of epistemic states. And different theorists will draw this line differently. While I think that any theory that makes sense of our intuitive judgments with respect to children, non-human animals etc. is *ceteris paribus* superior to any theory that does not, we need to balance these judgments with other intuitions. Especially when these intuitive judgements conflict with the desiderata identified above [§8.2]. In other words, if we take the intuitions underlying Desiderata 1 and Desiderata 2 seriously we likely require epistemic agents to be more cognitively sophisticated to account for them.

In sum, both doxastic and non-doxastic mental requirements face the overintellectualisation worry to varying extents depending on the complexity of the required mental state. While dispositional requirements are immune to these worries. However, following considerations having to do with the trade-off between the intuitions underlying the desiderata worked out above [§8.2], I do not take any worries built on overintellectualisation to be especially conclusive. So, I think that any higher-order requirement needs to strike the right balance between overintellectualising justification and retaining our judgements with respect to Seer and Non-Seer as well as Clairvoyance and Sensory Substitution cases.

8.4.2 Epistemic Relevance

We can now turn to the second worry, which concerns the epistemic relevance of positive higher-order requirements. Many who are critical of higher-order requirements including Bergmann (2005: 431) and Mc Hugh & Way (2016: 319) worry that adding a higher-order requirement to some first-order requirements does not contribute to achieving the overall aim of justification.

Worry 2: Epistemic Relevance: Higher-order requirements need to contribute to achieving the overall aim of justification. It is not clear how higher-order requirements such as the requirement to have a higher-order mental state do that.

The epistemic relevance worry can take various forms, depending on what we think the aim of justification is. One way to raise the worry is to question whether the higher-order requirement increases the *truth-conduciveness* of the belief. Inspired by Tucker we might spell out this worry as follows (2012: 326-327):

P1: If one's *defender* Φ of Bp does not have a *positive epistemic status*, that is, if Φ is neither justified nor caused by a reliable or properly functioning faculty, then Φ cannot contribute to the *truth-conduciveness* of Bp.

P2: If Φ cannot contribute to the *truth-conduciveness* of Bp, then it cannot contribute to the Bp's justification.

C1: Hence, if Φ is to contribute to Bp's justification, then Φ must have a *positive epistemic status* as well.

Tucker's discussion of this worry is inspired by Bergmann (2005) who makes a plausible case for P1 and P2 in discussing Bonjour's (1980) doxastic higher-order requirement. Namely, having a higher-order mental state that bears on the epistemic status of one's first-order belief does not seem to be epistemically relevant to the justifiedness of the belief, unless the higher-order belief itself has some relevant epistemic properties. After all, we can imagine cognitively sophisticated but systematically faulty epistemic agents. That is, epistemic agents which always have accompanying higher-order beliefs or other mental states which do not increase their epistemic performance. Consequently, Bergmann doubts that any higher-order state, including "irrational, irresponsible or insane belief(s)" can "contribute to the justification of S's belief that p, any more than having no doxastic attitude at all [...] does" (Bergmann 2005: 431). In response to this Bergmann speculates that defenders of <HOR+> will likely suggest that the higher-order attitude needs to have some positive epistemic standing as well. Yet, as we will see this suggestion is problematic since it likely faces a regress that many consider vicious [§8.4].

How threatening the epistemic relevance worry is depends on whether we think that the overall and sole aim of justification is truth-conduciveness, which is something that internalists about justification would likely deny. After all, it is baked into the fundamental Internalist Intuition that epistemic agents can be justified despite their means of belief formation being fundamentally misguided. So, as Tucker (2012: 327) illustrates, all we need to dismantle P2 is to point out that doxastic justification concerns more than just the truth connection.²⁴

Tucker (2012) argues that the higher-order inferential requirement's main job is *not* to make the belief more truth-conducive but to secure some *mental connection*. In Tucker's own words (2012: 328; *emphasis in the original*):

"Apparently, then, basing one's belief in P on one's belief in E is not needed to secure the truth connection. Why, then, do we impose a basing requirement on inferential justification? Because, unless the subject satisfies this requirement, she will fail to secure the required *mental connection*. [...] Basing a belief in P on a belief in E contributes to the inferential justification of my belief in P, not because it makes the belief more likely to be true, but because it constitutes, at least in part, the required mental connection between my belief in E and my belief."

²⁴ For a discussion on the role of truth in epistemic justification and epistemology more generally see, for example, David (2013) and Kvanvig (2013).

While Tucker uses this response to argue for some <HOR+ mental> on inferential justification we can apply his strategy to justified beliefs more broadly. That is, we might say, higher-order requirements are exactly relevant because they enable us to explain the difference between Seer- and Non-Seer cases as well as Clairvoyance and regular perception cases by appeal to mental connections. These requirements contribute to the aim of justification by establishing some kind of mental connection between the formed belief and the belief-forming mechanism.

Unfortunately, this is at best a partial response to the epistemic relevance worry. As pointed out above [§8.4.1], any mental state proposed needs to be understood as an attitude with content, which helps to explain how epistemic agents come to believe or should come to believe their conclusions (McHugh & Way 2014; Wright 2014). In other words, the mental state accompanying the belief that *p* needs to have some content that has some kind of epistemic significance with respect to the performed inference. Merely taking a premise to support a conclusion does *not* establish any epistemic benefit. After all, we can easily imagine someone being acquainted with poor inference rules such as counter-inductivism or affirming the disjunct (while incidentally reaching the right conclusion). This is similar to observations about basing requirements of doxastic justification we have encountered in [§5.6.1]. We can only value forming beliefs for the right kinds of reasons in the right kind of way if the reasoning itself is proper. So even if my belief that *p* is supported by some evidence *E* and this evidence non-deviantly causes me to believe that *p*, and I have the respective higher-order mental state indicating that my belief is based on *E*, I could do so by following improper rules (while taking these rules to be proper). That is, if we want to have a properness constraint on the basing relation we also want to have a properness constraint on the higher-order attitude.

Let us now turn to <HOR+ dispositional>. These requirements seem to establish the required mental connection between the evidence and the belief. After all, the basic idea behind those requirements is that properly forming a belief is exerting a cognitive disposition. However, as Tucker (2012) points out these sorts of cognitive dispositions do *not* seem necessary to establish the right inferential connection. Suppose *S* is *not* disposed to believe *p* when he believes *E* (and not disposed to form the belief that *E* supports *p*) but, nonetheless, happens to recognise that *E* supports *p*. Then, it seems that in some particular instances, *S* can believe that *p* based on her justified belief in *E*, in part because she justifiably believes that *E* supports *p*, without having the disposition to believe that *p* because of *E*. Nonetheless, in such a case, *E* therefore *p* seems to make the conclusion justified. This is because when (i) *S* justifiably believes *E* and that *E* supports

p, (ii) E does in fact support p, and (iii) S bases her belief in p on her belief in E, S seems to have all (and perhaps more than) S needs for the belief in p to be justified.²⁵ So, at the very least, there are two ways to infer correctly, one which involves a disposition and another which involves a higher-order doxastic state.

This brings me to a final worry, put forward by Neta (2019) in his discussion of doxastic basing requirements, which affects all higher-order requirements equally. Neta (2019: 193-195) demonstrates that neither dispositions to form a relevant mental state, doxastic or not, justified or unjustified nor actually having that higher-order mental state is sufficient to establish the right kind of mental connection. To illustrate this Neta considers a case in which someone has a problematic disposition to form an inferential belief based on a CNN news report (stipulated to be an unreliable source) but a diverging but unproblematic representation of how the belief is based. For instance, I might have the disposition to believe that my belief is based on information received from Al-Jazeera (stipulated to be a reliable source), while it is actually based on the CNN news report (and this belief is formed by a cognitive disposition).

I take these examples to demonstrate two things. First, dispositional requirements miss the aim of higher-order requirements. If the purpose of these requirements is to establish some kind of mental connection between the evidence and some doxastic attitude it is irrelevant whether I have that connection as a result of a disposition or not. That is, cognitive dispositions are not only insufficient and unnecessary as Tucker (2012) argues but also not a necessary component of any proper way to form beliefs. In other words, even if there are multiple valid ways to secure the mental connection between the belief and the belief-forming mechanism none of them involves dispositions.

8.4.3 Vicious Regress

To avoid the epistemic relevance worry, some have suggested a doxastic higher-order requirement that demands that the higher-order beliefs are justified as well (BonJour 1980; Fumerton 1995, Leite 2008). If one, however, understands this higher-level requirement as general, so that the

²⁵ At this point Tucker points out that “some may fail to get the intuition that E therefore p justifies its conclusion in the above circumstances because they find it weird that Seer would accept E therefore p when he is not disposed to do so” (2012: 330); however, there is nothing puzzling about doing something we are not disposed to do. On the contrary, we develop and cultivate new dispositions by repeated action. As we develop our disposition to exercise regularly and crave healthy food by repeatedly engaging in these kinds of activities we develop our dispositions to reason correctly by doing logic or engaging in conceptual analysis.

required higher-level belief must itself be justified by having an accompanied justified higher-order belief, and so on, then the requirement gives rise to a regress viewed by many as vicious (Bergmann 2005; McHugh & Way 2016: 318; Boghossian 2003; Tucker 2012):

Worry 3: Regress: If the initial reason to posit that first-order requirements need to be supplemented by second-order requirements can be extended to higher-order requirements we end up with an infinite hierarchy of requirements.

Bergmann describes the regress problem as follows (2005: 431; *emphasis in the original*):

“the only way in which [a higher-order requirement] can sensibly be required for the justification of S’s belief that p is if we stipulate that it must be made true by S’s having a *justified* belief [...] But to require that [a higher-order requirement] must be satisfied by S’s having a justified belief [...] is to face the dilemma noted earlier in the paper: either this requirement is a general one applying to all beliefs, which leads to a vicious regress, or it applies to beliefs at the object-level and perhaps at some higher levels but not at all higher levels, an *ad hoc* restriction made for no reason other than to avoid the regress.”

So according to Bergmann, the proponent of higher-order requirements faces a trilemma: (i) either we do *not* require the higher-order belief to be justified as well, which makes it *epistemically irrelevant*, or (ii) we do require the higher-order belief to be justified, which triggers a *regress*, or (iii) we provide some *ad hoc* reason that at some higher-order level beliefs do not need to be justified to make lower level beliefs justified, to avoid the regress.

As we have seen above, defenders of <HOR+ mental> tend to grab the first horn of the trilemma. That is, they argue that higher-order beliefs or other mental states need not be justified to be epistemically relevant for the justificatory status of the first-order belief. Instead, they are epistemically relevant by establishing the right kind of mental connection between the evidence and the belief. That strategy is promising since it avoids the charge of epistemic irrelevance, as well as the regress worry in one swipe. However, even if we grant that this strategy is successful with regard to the epistemic relevance worry we might still be concerned that it cannot avoid the regress entirely (Boghossian 2003; Wright 2014; McHugh & Way 2014). Any kind of mental requirement adds one additional mental layer M_2 to the initial mental layer M_1 consisting of the belief B_p and the evidence E that supports it. If this additional layer M_2 is needed to govern the relationship between B_p and E then it plausibly needs to be accompanied by a third mental layer

M₃ that plays the same governing role between M₂ and Bp and E. As a consequence we either need a *non-ad hoc* reason why this does not generalise or we face a regress again.

This is, among others, pointed out by Wright in his discussion of Boghossian's taking condition (Wright 2014: 31; *emphasis in the original*):²⁶

“Now I confess to seeing no alternative interpretation of this ‘taking that’ than to say that it requires an *information-bearing state* [...] –we will have to add that this registration state somehow *controls* S’s movement in thought from judgement of the truth of his premises to judgement of the truth of the conclusion. [...] We confront a dilemma. Suppose the content of the registration state is *general*: for instance, that any transition of the appropriate kind is licensed when the system is in a state of acceptance of the relevant kind of premises. Then it seems that we will have to understand the control exerted by the registration state on the specific movement in question as mediated by an ‘appreciation’ that the latter comes with the ambit of the former –so as an instance, in effect, of the inference: transitions of such and such a kind are mandated; this is a transition of such and such a kind; so it is mandated [...]. And now, fatally, we have represented inference as involving another; regress ensues.”

Faced with this problem, many have drawn the comparison to Lewis Carroll's (1895) parable of the tortoise which I already discussed in [§5.6].²⁷ Carroll has shown that counting inference rules employed in an argument among the premises of the argument leads to a regress. It would require us to add further and further inference rules which tell us how to infer from the premises and the initial inference rules, which expands the body of premises *ad infinitum*. Similarly, if the content of the required higher-order mental state accompanying the inference needs to be accompanied by another higher-order mental state that tells us how the first-order and second-order mental states hang together we expand the set of required mental states to the infinite (as the tortoise expands the set of premises in Carroll's parable). The crucial premise here to get this regress going is (Chudnoff 2014):

P_{INFERENCE}: In order to make an inference from some premises to a conclusion because of some higher-order mental state that those premises support that conclusion, one must take the claim that those premises support that conclusion as a premise in an inference.

²⁶ Ironically Boghossian (2003) uses similar arguments to dismiss other mental requirements such as intuition-based accounts. See Chudnoff (2014: 18-20) for a discussion of Boghossian's arguments.

²⁷ See, for example, Boghossian (2008: 274-275), McHugh and Way (2014: 319) Chudnoff (2014: 8)

$P_{\text{INFERENCE}}$ looks plausible, and in contrast to Chudnoff (2014), I do not think that we can dispute it to defend a $\langle \text{HOR+ mental} \rangle$.²⁸ But that is not our concern right now. Instead, I am interested in the question of whether there is a more general version of $P_{\text{INFERENCE}}$ that encompasses inferential and non-inferential beliefs alike.

P_{GENERAL} : In order to form a belief based on some E because of some higher-order mental state D that E supports Bp, one must take the claim that D as part of your basis for Bp.

P_{GENERAL} looks plausible as well. Indeed it can be motivated by taking the need for a higher-order requirement seriously that accommodates desiderata identified in [§8.2].

Before I move on, let me mention another way to try to avoid the trilemma. This response grabs the third horn and argues that there is some *non-ad hoc* stopping point. This can be achieved by linking discussions of higher-order requirements to traditional discussions about the structure of justified belief systems; that is, the debate between coherentists and foundationalists.²⁹

On the one hand, foundationalists may hope to avoid the vicious regress by arguing that it bottoms out when it comes to fundamental or basic beliefs. So, foundationalists might say that only non-basic beliefs, such as inferential beliefs, require accompanying higher-order beliefs about more fundamental beliefs; whereas basic perceptual beliefs or rational insights do not. However, independently of the plausibility of foundationalism, admitting this is already to deny higher-order requirements to be general since they do not apply to basic beliefs then. That is, any higher-order requirement that is not general does not make sense of the intuitions we had towards Clairvoyance and Sensory Substitution cases.

Alternatively, one could opt for a coherentist understanding of justification which implies that beliefs are at least sometimes justified solely by virtue of their relationship to other beliefs. If that is the case a higher-order justified belief can be justified by virtue of being embedded in a sufficiently large and coherent web of mutually supportive beliefs.³⁰ This would avoid the regress

²⁸ One possible response to avoiding the regress is to think of the needed mental state as an enabling condition on inference (McHugh & Way 2014: 319). One way in which this suggestion might be pursued is to say that the mental state plays a role in the application of a rule of reasoning. But then the question of epistemic relevance reappears. If I am already following the rules of reasoning, why is this higher-order mental state required, that makes me appreciate that I am following the rules of reasoning. Why is following the rules of reasoning not enough?

²⁹ For a contemporary overview of the debate between foundationalism and coherentism see, for example, Elgin and Van Cleve (2013).

³⁰ Another suggestion, which I mentioned above is to say that $\langle \text{HOR+} \rangle$ are not meant to increase the justifiedness of the belief but to secure some mental connection (Tucker 2012: 327-329).

since higher-level beliefs could be justified by virtue of their coherence with lower-level beliefs and vice versa. In explanationist terms we could put it like this: higher-level beliefs can support lower-level beliefs and vice versa if they are part of the best explanation of S's overall epistemic situation [§5.5]. However, while promising, it is important to point out that coherentism can at most account for propositional justification but not doxastic justification. That is in my framework it can at most help us to understand the epistemic support relation. Accordingly, this coherentist explanation of a higher-order requirement would collapse the doxastic requirement into an evidential higher-order requirement. Thus, I will explore this strategy further when I discuss evidential requirements in the next section [§8.5].

8.4.4 Summary

In this section, I introduced three worries:

Worry#1: Over-intellectualisation: We often judge less cognitively sophisticated epistemic agents (such as children, animals, or certain types of collectives) to have justified beliefs. Higher-order requirements conflict with such judgements and, therefore, implausibly overintellectualise our theory of epistemic justification.

Worry#2: Epistemic Relevance: Higher-order requirements need to contribute to achieving the overall aim of justification. It is not clear how higher-order requirements such as the requirement to have a higher-order mental state do that.

Worry#3: Regress: Higher-order requirements induce a vicious regress. If we think that the initial reason to posit that first-order requirements need to be supplemented by second-order requirements can be extended to higher-order requirements we end up with an infinite hierarchy of requirements.

Afterwards I examined these worries as worries against the following three higher-order requirements <HOR+>:

<HOR+ doxastic>: If S's belief that p is justified then S *necessarily* has Φ , a *doxastic* mental state that *indicates* that p is justified.

<HOR+ dispositional>: If S's belief that p is justified then S *necessarily* has Φ , a *disposition* to form beliefs in a way that *indicates* that p is justified, and this disposition caused S to believe that p.

<HOR+ mental>: If S's belief that p is justified then S *necessarily* has Φ , a *non-doxastic* mental state that *indicates* that p is justified.

All of these requirements succumb to worries 1-3 to varying extents.

First, <HOR+ doxastic> leads to an overintellectualisation of justification since having higher-order doxastic states is no easy cognitive task. Furthermore, it is not clear how <HOR+ doxastic> can add anything of epistemic relevance without producing a regress. Mere beliefs are too weak since they lack any epistemic relevance while independently justified beliefs are too strong by leading to a vicious regress.

Second, <HOR+ dispositional> while avoiding the overintellectualisation worry, as well as the regress worry could not explain the epistemic relevance of higher-order requirements. Here out of the three requirements discussed <HOR+ dispositional> could make the least sense of the epistemic relevance of higher-order requirements and, therefore, could not live up to the desiderata identified in [§8.2]. This is a reason to take dispositional requirements to be a non-starter.

Third, <HOR+ mental> seem to strike the right balance between the complexity of <HOR+ doxastic> and the epistemic irrelevance of <HOR+ dispositional>. Those requirements promise to deliver a reasonably good answer to the epistemic relevance worry as well as the overintellectualisation worry. However, the replies given still left some questions unanswered. More importantly, however, <HOR+ mental> still could not avoid the regress. At least without making further (controversial) assumptions about the structure of justification.

We can finally turn to <HOR+ evidential>, the kind of higher-order requirements I want to defend in the context of my theory of epistemic justification.

8.5 Evidential Requirements

Here is the above-given characterisation of evidential requirements again:

<HOR+ evidential>: If S's belief that p is justified then S *necessarily* has sufficient *evidence* that *indicates* that p is justified.

Depending on different understandings of the nature of evidence <HOR+ evidential> will take different forms. If evidence possession is some kind of doxastic or perspectival relation, those requirements will be similar to doxastic requirements. If evidence and evidence possession, however, are understood in mentalist or abstractionist terms, evidential requirements will be very close to non-doxastic mental requirements. Hence, <HOR+ evidential> can be understood as a subtype of <HOR+ mental>.

Nonetheless, in the proposed evidentialist framework <HOR+ evidential> are superior to <HOR+ mental> for at least two reasons. First, in contrast to the above-discussed mental requirements evidential requirements do *not* need the mental state to be of any specific kind to play its intended role. Instead of requiring some specific seeming (Chudnoff 2014), taking (Boghossian 2014), or higher-order awareness (Tucker 2012), on my statist understanding of evidence every (set of) representational non-factive mental state(s) with the relevant content can establish the higher-order relation. Second, evidential requirements can make immediate sense of what it means for some mental state to *indicate* that *p* is justified; namely, by evidential supporting some higher-order proposition that bears on *p*.

While some other higher-order evidential requirements have been discussed in the literature, I will focus on condition (2) of Continuous Evidentialism, which I take to be the most promising candidate for an evidential higher-order requirement.³¹

S belief that *p* is *epistemically responsible* iff:

- (2) (a) *S* possesses some higher-order evidence E_H , which is *sufficient to support* *q*, the proposition that the total evidence E_{TOTAL} on balance supports *p* and (b) *S*'s belief that *p* is *properly based* on E_H .

In section [§8.2.1], when motivating a general need for positive higher-order requirements we have identified two desiderata that any higher-order requirement needs to have: the ability to distinguish between Seer and Non-Seer cases (Desideratum 1) as well as the ability to distinguish between Clairvoyance and Sensory-Substitution, on the one hand, and regular perception cases on the other (Desideratum 2). Afterwards, in [§8.3], I have identified and discussed three worries having to do with overintellectualisation, epistemic relevance and vicious regress. I will now demonstrate how (2) can respond to these worries and in so doing live up to Desiderata 1 & 2.

³¹ For some discussion of <HOR+ evidential> see, for example, Tucker (2012: 328-333).

This shall be seen not only as a specific argument for (2) but as a more general proof of principle about the versatile character of <HOR+ evidential>.

Let me start with the first worry, *overintellectualisation*. As stressed above, every account of justification needs to make some cut-off when it comes to the question of how cognitively sophisticated a system needs to be to count as an epistemic agent, that is, a potential bearer of epistemic states. So how demanding is the evidential higher-order requirement (2)? Not very demanding. All it requires is that an epistemic agent needs to be capable of having two mental states with the relevant propositional content E and E_H that are sufficient to support two different propositions and that non-deviantly cause S to believe that p . This explains why Norman and Sandra cannot form justified beliefs via clairvoyance or sensory substitution respectively (at least not right away) while toddlers and other less sophisticated agents who possess a plethora of evidence about the kinds of inferences they are using can.

Importantly, while I happen to understand epistemic support in explanationist-probabilist terms [§5.5], I do not think that for E to support p (for S) S needs to be aware of or even has the cognitive capacities to regard E to be part of the best explanation for p .³² In this respect, my account significantly differs from other explanationist accounts proposed in the literature (Conee & Feldman 2008; McCain 2014). For example, McCain tries to dismiss the charge of over-intellectualisation as follows (2014: 78):

“this does not require one to have well-developed concepts of ‘evidence’, ‘explanation’, ‘logical consequence’, or ‘entailment’. All that is required is the ability to understand something as an answer to a why-question or the disposition to have a seeming that, for example, the truth of p and if p , then q ensure the truth of q .”

Here, Berghofer, among others, points out that it is plausible to assume “that when a dog hears its favourite person approaching and then sees the person, the dog knows (or has justification to believe) that the person is there” despite lacking “the cognitive capabilities ‘to understand something as an answer to a why-question’” as required by McCain (Berghofer 2022: 78).³³ I agree with Berghofer that McCain’s understanding of epistemic support makes justification too cognitively demanding. My account avoids this problem. All it means for E to support p is that p and E explanatorily cohere according to (some permissible) epistemic standards that S

³² As pointed out in [§5.5] this is inspired by Fratantonio (forthcoming).

³³ See also Lutz (2020: 2637).

subscribes to. Furthermore, S can subscribe to epistemic standards without being able to spell out what these standards are, or even without having a concept of epistemic standards. Consequently, on my account, the dog can have a justified belief that his favourite person is approaching them, without having any higher-order beliefs or the capacity to form such beliefs.

We can now turn to the second worry, the worry that higher-order requirements are *epistemically irrelevant*. I think the best way to address this worry is twofold. First, we need to demonstrate how (2) incorporates Desiderata 1 and Desiderata 2 [§8.2.1]. Second, we need to give an argument on why and how the respective explanation makes Seer's belief epistemically superior to Non-Seer's and regular perception cases epistemically superior to Clairvoyance and Sensory Substitution cases.

Here is how (2) enables us to distinguish Seer and Non-Seer cases. Seer, in contrast to Non-Seer, properly infers her belief that p because her belief is based on (that is, non-deviantly caused by) the premise E as well as some higher-order evidence E_H about how the premises support the conclusion. So, Seer's and Non-Seer's beliefs that p which are supported and non-deviantly caused by E differ in the following respect. Seer's belief is also non-deviantly caused by some higher-order evidence E_H which is sufficient to support that E_{TOTAL} supports p . This supports the judgement had by many about the nature of inference on which inference is taken to be more than a causal process.

This is the same intuition that motivated internalist dismissals of clairvoyance (BonJour 1980) and sensory substitution cases. We judge the clairvoyant who (reliably) forms clairvoyance beliefs without having reasons to regard her clairvoyance as reliable as having unjustified beliefs. Now for an agent like Sandra, there is no difference between beliefs formed via visual perception or clairvoyance. How could there be? After all, Sandra like Norman does not possess any evidence about the reliability of her newly acquired belief-forming mechanism. So, for Sandra to form beliefs on perception such as 'there sits a pigeon on the tree outside my office', she needs to have some evidence E_H sufficient to support that for all I know perception under these circumstances is a valid source of justification. Compare this with a regular toddler who is familiar with the causal structure of the mesoscopic environment and forms beliefs about immediate physical happenings. The toddler, in contrast to Sandra, can justifiably form beliefs based on perception, without having the cognitive capacities to reflect on these belief formations.

But this is only one step towards the answer to the epistemic relevance worry. How do these higher-order evidential requirements on inferential and non-inferential beliefs increase epistemic

performance? Here the answer provided by evidential requirements is straightforward (in contrast to the other above-discussed higher-order requirements). Having sufficient evidence for a belief contributes to the overall epistemic quality of the belief, especially given antecedent evidentialist commitments. Hence having higher-order evidence about the entire evidential situation does so as well, at least indirectly. That is if E supports p and E_H supports E_{TOTAL} supports p we have increased the overall explanatory coherence between S 's evidence and p .

Note that this is not to say that evidential requirements such as (2) cannot also appeal to other resolutions of the relevance worry such as the mental connection provided by the higher-order evidence. E_H , given my statist understanding of evidence, is after all the propositional content of a mental state. So whatever connection can be achieved by other non-doxastic mental states such as seemings (Chudnoff 2014) or a taking state (Boghossian 2014) can reasonably be realised by E_H as well.

We can now turn to the *regress worry*. To see how evidential requirements can avoid the kind of regress that doxastic and other mental requirements face, we need to take another look at what is causing the regress. In the case of doxastic requirements, which require an accompanying justified belief the regress is triggered because the higher-order belief to be justified needs to have another accompanying justified belief, which requires to have another justified belief and so on. Now, with respect to other mental requirements or doxastic requirements which do not need the higher-order belief to be justified, the situation is different. Here the regress is triggered because of the supposed content of the mental state t connecting E to Bp . Here t supposedly has the content that E mandates that the belief that p and t is used in getting us from E to Bp . But if t needs to be used, to mandate us to get us from E to Bp , why is there no higher-order taking state t_2 that needs to be used that mandates us to use t in getting us from E to Bp .

Now as I have demonstrated above, following Chudnoff (2014), whether we can trigger a regress in this way depends on whether we accept the following premise:

$P_{GENERAL}$: In order to form a belief based on some E because of some higher-order mental state D that E supports Bp , one must take the claim that D as part of your basis for Bp .

We can immediately see that (2) works without accepting $P_{GENERAL}$. First, while E_H is required to be part of the basis alongside E on which S believes that p , E_H does not concern the relationship between E and Bp directly. Instead, it is only concerned about the relationship between E_{TOTAL} and p . Second, in contrast to the <HOR+ mental> discussed above, we do not need to *take* E_H

to support E_{TOTAL} in the sense that we have some higher-order mental state. Instead E_H is just evidence we have all along that bears on our overall evidential situation with respect to p .

So if there is a lurking regress it must rest on a different assumption than $P_{GENERAL}$. Here we might, for example, appeal to a recent argument put forward by Appely and Stoutenburg (2017) against explanationist understandings of evidential support. Appely and Stoutenburg point out that if the explanationist relies on a mental understanding of evidence and requires the epistemic agent to be aware that p is part of the best explanation for why S has E (in order to support p) we trigger a regress. This is because the required awareness is itself a mental state and, therefore, a piece of evidence. Thus, for every awareness, there is an emerging new piece of evidence such that a new awareness mental state is required, and so on *ad infinitum* (Appely & Stoutenburg 2017: 3077).

Two things to note about this. First, this worry is a general worry about explanationist understandings of evidential support. So it only indirectly concerns the defended higher-order requirement since it relies on such an understanding. But in principle, we could rest (2) on a different understanding of evidential support. Second, the argument, while plausibly successful against other explanationist frameworks such as McCain's (2014) or Conee and Feldman's (2008), does not concern my preferred version of explanationism. In my view, S does not need to be aware or have a seeming (or a disposition to) that governs the support relation between (1a) E and p and (2a) E_H and q . All that is needed is that p is part of the best explanation of why S has E and q is part of the best explanation of why S has E_H , according to the relevant epistemic standards.

Before I conclude, let me mention another way to avoid the regress, which grabs the third horn of Bergmann's trilemma [§8.4.3]. That is by acknowledging that (2) as a responsibilist clause is grounded in the socio-epistemic expectations we have towards epistemic agents. These expectations plausibly do not require epistemic agents to have an infinite hierarchy of higher-order evidence. So, in response to Bergmann's (2005) formulation of the worry, we can say that there is some *non-ad hoc* stopping point at which the epistemic agent is not irresponsible in not possessing higher-order evidence about their epistemic situation. Here I argued that this stopping point is a universal feature arising from the problem of unpossessed evidence [§7.4.3]. We expect epistemic agents to have higher-order evidence about their evidential situation when forming beliefs. But do we expect them to have higher-order evidence about their higher-order evidence? Maybe. But plausibly we do not expect them to have higher-order evidence all the way up.

In sum (2), and with it possibly other <HOR+ evidential> live up to Desiderata 1 and 2 and, therefore, provide us with a straightforward explanation of the case judgements about Seer and Non-Seer as well as Clairvoyance and Sensory Substitution Cases. Second, in contrast to other requirements discussed, (2) provides us with a satisfactory answer to prominent worries often raised against higher-order requirements.

8.6 Conclusion

There are various reasons why we might think that epistemic justification is subject to higher-order requirements. Namely, the reasons that motivate us in our case judgements with respect to Seer and Non-Seer, as well as Clairvoyance and Sensory Substitution cases. However, as we have seen, many such requirements face a plethora of well-known objections. In walking through these objections and discussing several types of higher-order requirements I have given a systematic overview of the pros and cons of various higher-order requirements. I concluded that the evidentialist higher-order requirement (2) is the only requirement that allows us to give a satisfactory response to the discussed objections. Having defended (2) against possible objections, makes my defence of Continuous Evidentialism complete. This means that we are now ready to see Continuous Evidentialism in action [\[ch.9\]](#).

9

Continuous Evidentialism[†]

In the second half of the thesis [ch.5] – [ch.8], I developed a continuous evidentialist account of epistemic justification which I called Continuous Evidentialism. In this final chapter, I want to put all of the conceptual pieces together to recapitulate and demonstrate the applicability of Continuous Evidentialism. In particular, I will compare the defended account to other influential frameworks and apply it to a large range of cases discussed in collective epistemology. This enables me not only to demonstrate the versatile nature of the proposed account but also allows me to highlight some interesting features of it.

9.1 Introduction

Let me start by restating Continuous Evidentialism and then summarising how I have understood the conceptual pieces embedded in the respective clauses.

Continuous Evidentialism:

S justifiably believes that p *iff*:

- (0) S believes that p (call this proposition “Bp”).
- (1) (a) S possesses some evidence E, which is *sufficient to support* Bp, and (b) Bp is *properly based* on E, and
- (2) (a) S possesses some higher-order evidence E_H, which is *sufficient to support* q, the proposition that the total evidence E_{TOTAL} on balance supports p and (b) S’s belief that p is *properly based* on E_H.

In [ch.5] I introduced the general evidentialist framework Continuous Evidentialism rests upon, and discussed the central terms embedded in clauses (1a) and (1b). I defined and critically discussed the evidentialist core tenets as defended by Conee and Feldman (2004), and explored the nature of evidence, evidence possession, evidential support and epistemic basing. In so doing, I defended a specific version of abstractionism, called Statism, which takes evidence to be the

[†] Some parts of this chapter are based on the following published work:

Graf, S. (2021). Review: The Epistemology of Groups by Jennifer Lackey. *Perspectives: International Postgraduate Journal of Philosophy*, 9 (1), 380-387.

propositional content of representational non-factive mental states. Analogously, I understood the entire evidence possessed by an epistemic agent S , called $\langle E_{\text{BASE}} \rangle$, to be the propositional content of all of the relevant mental states of S and all of its subgroups which are epistemic agents themselves. Furthermore, I defended a probabilist-explanationist understanding of epistemic support and a causal understanding of epistemic basing.

In [ch.6], I discussed various no-defeater clauses as potential supplements for accounts of epistemic justification and developed a general taxonomy of defeater cases. Out of the clauses considered only the responsibilist clause could give us the right verdict with respect to the full range of cases. This led me to discuss Sanford Goldberg's social entitlement model of epistemic responsibility in [ch.7]. While I agreed with Goldberg that epistemic responsibility is grounded in socio-epistemic expectations, I argued that we shall not understand these expectations as making us strictly liable for the evidence we should have. Instead, we are entitled to expect epistemic agents to possess and utilise higher-order evidence when forming beliefs. This resulted in a responsibilist condition that takes the shape of the higher-order evidentialist clause (2). I argued that this clause, while enabling us to make sense of defeater cases, shall not be understood as a no-defeater clause in the traditional sense, due to various problems linked to Defeatism. This led to further challenges which I defended (2) against in [ch.8].

It might help to depict the structure of Continuous Evidentialism as follows:

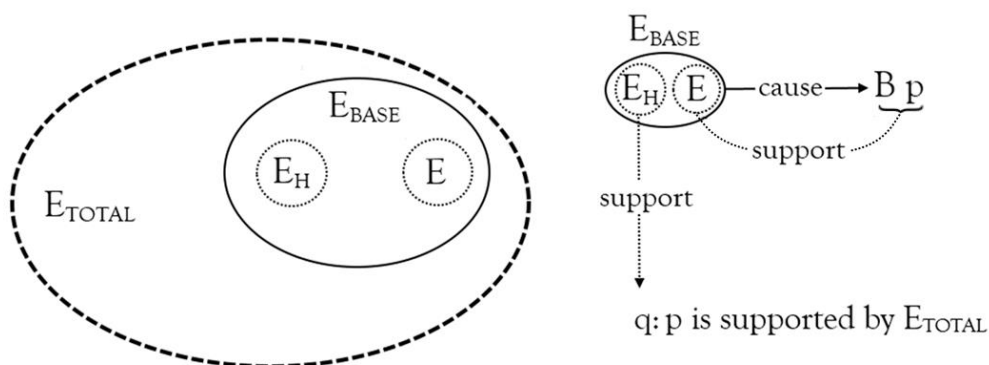


Fig. 1: Continuous Evidentialism

Merely formulating Continuous Evidentialism and defending it against various objections is not enough to demonstrate that it should be accepted. Another important aspect of whether a particular epistemic theory should be accepted hinges on how well the theory handles the cases.

Of course, I have discussed a wide range of cases in previous chapters already when developing my understanding of the various above-mentioned conceptual building blocks. In this chapter, I will prolong this journey. In particular, I will apply Continuous Evidentialism to a wide range of problematic cases that have been discussed in the literature on collective justification. I will do so by directly comparing Continuous Evidentialism to prominent extant accounts of collective justification [§9.2]. Afterwards, I will walk through different types of cases discussed by those accounts. Namely distributed evidence cases [§9.3.1], collective defeater cases [§9.3.2], collective basing cases [§9.3.3] and evidential manipulation cases [§9.3.4]. In the final section [§9.4], I will highlight some consequences of Continuous Evidentialism which became visible when elaborating on how we shall understand some of the discussed cases. Specifically, I will revisit the notion of epistemic divergence [§9.4.1] and discuss where we should locate Continuous Evidentialism on the spectrum between distributed and non-distributed accounts of collective justification [§9.4.2].

9.2 Continuous Evidentialism: Comparison

There are multiple epistemological frameworks that served as a major influence in the development of Continuous Evidentialism. Some of which I have discussed extensively, such as Earl Conee and Richard Feldman's evidentialist framework or Sandy Goldberg's social entitlement model of epistemic responsibility. But there are others, which I discussed less extensively but which are of similar importance, such as Jessica Brown's recent work on group evidence and collective epistemic basing, or Brian Hedden's probabilist-explanationist understanding of evidential support.¹ In this section, I want to draw attention to two accounts of collective justification that served as a primary source of inspiration when thinking about Continuous Evidentialism in the context of collective epistemology. Namely, Jennifer Lackey's (2016, 2021) *Group Epistemic Agent Account* <GEAA>, as well as Paul Silva's (2019) *Evidentialist Responsibility for Groups* <ERG>. Let me, for the sake of reference, state these accounts in full here.

¹ As discussed in Brown (2022a, 2022b, 2023a, 2023b, forthcoming) and Hedden (2015a, 2015b, 2019), respectively.

Evidentialist Responsibilism for Groups <ERG>:

A group, G, justifiably believes that p on the basis of evidence E iff:

- (1) E is a sufficient reason to believe p, and the total evidence possessed by enough of the operative members of G does not include further evidence, E*, such that E and E* together are not a sufficient reason to believe p, and
- (2) G is epistemically responsible in believing p on the basis of E.

Group Epistemic Agent Account <GEAA>:

A group G justifiably believes that p iff:

- (1) A significant percentage of the operative members of G (a) justifiably believe that p, and (b) are such that adding together the bases of their justified beliefs that P yields a belief set that is coherent, and
- (2) (a) Full disclosure of the evidence relevant to the proposition that p, (b) accompanied by rational deliberation about that evidence among the members of G (c) in accordance with their individual and group epistemic normative requirements, would not (d) result in further evidence that when added to the bases of G's members' beliefs that p, yields a total belief set that fails to make [it] sufficiently probable that p.

Immediately, we can see some similarities between Continuous Evidentialism, <GEAA> and <ERG>. For example, all three accounts embrace a view of collective justification that contains evidentialist elements. One way or another for a group to hold a justified belief it needs to be sensitive to the evidence possessed by the group and/or the group's members. Second, all three accounts rely on some notion of epistemic responsibility; while <ERG> speaks of epistemically responsible believing, <GEAA> speaks of epistemic normative requirements.² Furthermore, all three accounts have a higher-order condition that is supposed to make sense of defeater cases (together with the responsibilist condition).

Despite these similarities, there are significant differences that will come to the forefront when we analyse how these accounts give us different verdicts with respect to different cases [§9.3]. However, apart from these case-specific differences, there are some general contrasting features that I want to highlight here.

First, <ERG>'s scope is limited by its left-hand side since it only concerns beliefs that are based on evidence. That is, Silva “leaves open the possibility of a more general account of group justification that permits some non-evidence-based justified group beliefs” (2019: 274). As I have argued in [ch.8], such restrictions are misguided since following evidentialist commitments and

² Here we have to keep in mind that I understand clause (2) of Continuous Evidentialism to be a responsibilist condition that rests on the assumption that we can reduce responsibility to having and utilising higher-order evidence in the right kind of way [§7.4.1].

the internalist intuitions they rely on, I do not think that epistemic agents can form justified beliefs without basing them on evidence that is sufficient to support their beliefs.

Second, both <ERG>, as well as <GEAA>, make our epistemic theorising discontinuous. Not solely because they are theories about collective justification (rather than justification in general) but because for them the justification of a group directly depends on various idiosyncratic features of collective agents. For example, both accounts rely on the distinction between non-operative and operative members who play an important role in analysing the epistemic status of collective attitudes.³ This completely divorces our understanding of individual justification and collective justification (since, after all, individuals do not have any operative members).

Relatedly, <GEAA>, and to a lesser extent <ERG>, retain summativist elements, concerning their understanding of group belief, justification, evidence and/or epistemic basing. <GEAA>, in particular, is directly built on a summativist understanding of group belief (Lackey 2021: 48-49).⁴ This conflicts with the divergence argument I have discussed in [ch.3] (and which I will revisit below [§9.3.4]).⁵

Having pointed out these general differences we can now proceed to discuss a selection of cases and compare the verdicts delivered by Continuous Evidentialism, <GEAA> and <ERG> against each other.

9.3 Continuous Evidentialism: Application

In this section, I want to give an overview of various types of cases that have been considered in the literature and illustrate how Continuous Evidentialism can make sense of them. In so doing, I will mainly focus on cases that have been discussed by Lackey (2016, 2021) and Silva (2019) in their works on collective justification. However, I will also introduce some (novel) cases those accounts have trouble dealing with. Despite the huge variety of cases discussed in the literature, we can group them together as follows: distributed evidence cases [§9.3.1], collective defeater cases [§9.3.2], collective basing cases [§9.3.3] and evidential manipulation cases [§9.3.4].

³ For the concept of an operative member see also Tuomela (2004: 113).

⁴ <GEAA> is directly built on the so-called Group Agent Account of group belief. Following this account “a group, G, believes that p if and only if: (1) there is a significant percentage of G’s operative members who believe that p, and (2) are such that adding together the bases of their beliefs that p yields a belief set that is not substantively incoherent” (Lackey 2021: 48-49).

⁵ As I have demonstrated in [§4.7], summativist accounts are subtypes of parasitist accounts, which neither accept <Unity> nor <non-Primacy>.

9.3.1 Distributed Evidence

Many theorists of collective knowledge and justification, including Bird (2010, 2014), Goldman (2014), Lackey (2016, 2021), Silva (2019), Habgood-Coote (2019) and Brown (forthcoming), have pointed out that groups seem to be able to base their beliefs on evidence that is distributed among its different members.⁶ To illustrate this we can use an often-discussed case from Goldman (2014: 16):⁷

DIFFERENT BASES: *G* is a group of 60 British Museum guards (M_1 - M_{60}). There are three subgroups consisting of 20 guards each, M_1 - M_{20} , M_{21} - M_{40} , and M_{41} - M_{60} . Each of those subgroups has a justified belief about some specific guard planning an inside theft of a famous painting. M_1 - M_{20} believe that Albert is planning an inside theft of a famous painting based on some evidence E_1 . M_{21} - M_{40} believe that Bernard is planning an inside theft of a famous painting based on some evidence E_2 . M_{41} - M_{60} believe that Cecil is planning an inside theft of a famous painting based on some evidence E_3 . By deduction each of the subgroups (M_1 - M_{20} , M_{21} - M_{40} , M_{41} - M_{60}) infers p , that there is a guard who is planning such a theft.

The consensual judgement by Goldman (2014: 16), Lackey (2016: 359) or Silva (2019: 263-264) about DIFFERENT BASES is that *G* is justified in believing the proposition that p . (Of course, there is significant disagreement among those authors about why *G* is justified in believing that p .)⁸

Using Continuous Evidentialism we can reach the same verdict. *G*'s belief that p is justified as long as E_1 - E_3 together are sufficient to support p and *G*'s belief is based on p (and the higher-order conditions (2) are met). Here the causal understanding of epistemic basing underlying Continuous Evidentialism becomes important. All it takes for E_1 - E_3 to be able to justify *G*'s belief is for it to non-deviantly cause the group to believe that p . For that to be the case E_1 - E_3 must cause *G* to believe that p . (And this is left open by the description and may or may not be the case.) We

⁶ One interesting question of course is whether this really is a genuine difference between groups and individuals. On some understandings of testimony, for example, we might think that individuals can also base their beliefs on evidence, which is possessed by the testifier, or that the evidence or the justificatory status of the attitude is transmitted from the testifier. For a convincing argument against such a transmission view see, for example, Peet and Pitcovski (2017).

⁷ For the sake of narrative ease, I simplified the case by reducing the number of members from 100 to 60 and the number of subgroups from five to three. For further discussion of this case see also Brown (2023b).

⁸ For example, Lackey highlights how her judgement about DIFFERENT BASES differs from Goldman's initial treatment by considering two variations of the case (Lackey 2016: 361-369) which she takes to pose a problem for Goldman's understanding of collective justification. I will discuss these variations in the next section [§9.3.2].

neither have settled on an understanding of group belief nor are the relevant details about how the group belief is formed specified in DIFFERENT BASES. The case leaves it open whether G believes that p and even if we assume that G does believe that p, the case also leaves open if E_1 - E_3 , some subset of E_1 - E_3 , or some other evidence E^* causes G to believe that p (or whether the higher-order conditions specified in (2) are met).

This treatment of DIFFERENT BASES differs substantially from how Lackey's <GEAA> reaches the verdict that G's belief is justified. For <GEAA> it is important that a significant percentage of the operative members of G justifiably believe that p and that their bases can be coherently combined. In contrast to <GEAA>, for Continuous Evidentialism group-level justification does not run through member-level justification, so it does not matter how many of G's members, if any, justifiably believe that p.⁹

Consequently, Continuous Evidentialism can also make sense of cases in which none of the members possess sufficient evidence for a complex inferential belief, but the group as a whole does. One representative case of this sort is found in Bird (2010: 34-36; 2014: 57-58):¹⁰

DISTRIBUTED COGNITION: Dr X., a physicist, and Dr Y., a mathematician, collaborate on a project to demonstrate the truth of the conjecture that p. Their project can be broken down into three parts. Part one is a problem in physics, the problem of showing that q, which will be the work of Dr X. alone. Part two is a problem in pure mathematics, that of proving that if q then p, for which Dr Y. takes sole responsibility. Part three is an application of modus ponens to the results of parts one and two. They arrange for an assistant to publish a paper [on behalf of their collaboration G] if and only if the assistant receives from X the demonstration that p is true and from Y the proof of $q \rightarrow p$ (the brief final part with the application of modus ponens has been pre-written). We can imagine that X and Y have no other communication with each other or with the assistant and so do not know at the time of publication that p has been proven.

For Bird, cases such as DISTRIBUTED COGNITION illustrate that groups, or social epistemic subjects, often function by distributing cognitive labour among its members and or other parts

⁹ Note, that <GEAA> lacks a belief requirement similar to that embedded in Continuous Evidentialism (clause (0)). Here the summative understanding of group belief (Lackey 2021: 48-49) is embedded in <GEAA>'s first clause (1); i.e. in any case in which clauses (1) and of <GEAA> are satisfied, for Lackey the group believes that p.

¹⁰ A good overview of different distributed cognition cases is found in Habgood-Coote (2019), who discusses them in the context of group know-how.

of its system.¹¹ That is, distributed cognition “identifies a system for producing knowledge, and studies how the various components, usually performing distinct sub-tasks, contribute to that overall task” (2014: 45); where these cases illustrate that “what the social epistemic agent knows does not depend on what the individuals know” (2014: 58).¹²

For the sake of the argument, we can presuppose that *G* believes that *p* in DISTRIBUTED COGNITION. If so, Continuous Evidentialism presents us with a simple way to explain why this belief can be justified (and, therefore, potentially amount to the social knowledge Bird is concerned with). *G*’s belief that *p* is caused by both the evidence E_1 sufficient to support *q* (possessed by Dr. X), as well as the evidence E_2 sufficient to support $q \rightarrow p$ (possessed by Dr. Y) and by transmission of E_1 and E_2 to the assistant, *G* was caused to believe that *p*. Now, whether we should judge *G* to justifiably believe that *p* also depends on whether condition (2) of Continuous Evidentialism is met. Here condition (2) can be realised in multiple ways. For example, the assistant might (2a) have sufficient higher-order evidence E_H about the overall evidential situation and (2b) utilise E_H when putting together the proof and handing in the paper for publication. Or it could be met by *G* itself having a (justified) group-level belief that supports that (2a) $\langle E_{TOTAL} \rangle$ is sufficient to support *p* which (2b) plays the relevant causal role in *G*’s belief formation.

This draws our attention to a third type of case having to do with distributed evidence. Namely, cases in which *G*’s believes that *p* is (partly) based on evidence directly possessed by *G* rather than one of its members. Take a variation of DISTRIBUTED COGNITION, in which *G* is just a subgroup of a larger collaboration G^+ whereby *G*’s belief that *p* (based on Dr. X’s demonstration that *q* and Dr. Y’s proof that $q \rightarrow p$) is just part of another inference performed at the group level. Here it is plausible to say that *G*’s justified belief that *p* is part of G^+ ’s evidence and hence could be used in support of a group-level of G^+ . For instance, G^+ might infer that *t*, and as a consequence form the respective belief, based on a proof by some of its members that $p \rightarrow t$. This belief then is partly justified by evidence that is not possessed by any of G^+ ’s members, but by G^+ ’s subgroup *G*; in particular, on *G*’s justified belief that *p*.

Importantly, both $\langle GEAA \rangle$ and $\langle ERG \rangle$ have problems with distributed cognition cases. This should be obvious for $\langle GEAA \rangle$ which explicitly relies on a summative understanding of

¹¹ Bird emphasises, following Hutchins (1995a, 1995b), that not only members, but also other component parts of collective agents such as infrastructure or machinery often plays a determining role in the belief formation of those agents. More on that below.

¹² For an earlier account of distributed cognition see also Hutchins (1995a, 1995b).

group beliefs and justification.¹³ But also <ERG> has problems with the second kind of distributed cognition case due to its summative understanding of collective basing (which I will discuss in more detail below [§9.3.3]). For the members of G+ neither need to believe that p, justifiably believe that p or even jointly accept that p for p to play its justifying role in G+'s belief formation.

With that being said, much of the plausibility of judging groups to be able to acquire justified beliefs via distributed cognition hangs on the underlying account of group belief. While I haven't settled on an account of group belief, the most compelling candidate accounts, namely functionalist and interpretationist accounts allow us to ascribe beliefs to groups in at least some distributed cognition cases.¹⁴ If these accounts are correct Continuous Evidentialism has an advantage over <ERG> and <GEAA> since it provides us with conditions under which a belief formed by a group via distributed cognition can be justified. Furthermore, Continuous Evidentialism can make sense of another feature of groups that distributed cognition accounts have often emphasised; namely, that "almost always in such systems there are non-human vehicles for representing various pieces of information; these must be taken fully into account in articulating how the system achieves its cognitive goals" (Bird 2014: 45). If the information stored in these non-human components influences the beliefs on the group level, groups can have beliefs that can serve as an evidential base of other group level beliefs that cannot possibly be shared by the members of the group agents.

In sum, given the understanding of evidence possession and epistemic basing underlying Continuous Evidentialism, we can make sense of cases in which some group G justifiably believes that p based on *different* bodies of evidence, one *distributed* body of evidence, or evidence possessed by the group *directly*.

9.3.2 Collective Defeat

In Lackey's discussion of DIFFERENT BASES, she considers two variations which pose a problem to Goldman's initial understanding of the case (2016: 361-369).¹⁵ In the first variation, called CONFLICTING BASES the respective subgroups (M₁-M₂₀, M₂₁-M₄₀, ...) have conflicting bases that

¹³ Note that Lackey (2021:111-127) extensively discusses and dismisses distributed cognition accounts. I have mentioned some of her worries with those accounts in [§3.6]. Due to considerations of space, I decided not to discuss these worries here. However, I will revisit some of them below [§9.4.2].

¹⁴ See, for example, Tollefson (2002b, 2015), List and Pettit (2011), Brouwer, Ferrario and Porello (2021) or Brown (2023a).

¹⁵ For an early account of collective defeat see Carter (2015).

cannot be coherently combined. In other words, the evidence E_1 possessed by M_1 - M_{20} supports that *only* Albert is planning a theft, while E_2 supports that *only* Bernard is planning a theft and E_3 supports that *only* Cecil is planning a theft. In the second variation, called NON-CONFLICTING BASES, the bases of the subgroups (M_1 - M_{20} , M_{21} - M_{40} , ...) can be coherently combined but each of the subgroups has a defeater for the other subgroup(s) beliefs.¹⁶

Depending on our exact understanding of group evidence and defeaters we might consider these variations to be defeater cases. As such, both Lackey and Silva judge G in both variations to be unjustified. In CONFLICTING BASES, the body of evidence G relies on (what I called the operational part E) cannot be coherently combined and is hence defeated. While in NON-CONFLICTING BASES some evidence within the overall evidence possessed by G (what I called $\langle E_{\text{BASE}} \rangle$) defeats G 's belief. Reaching this verdict is possible, as pointed out by Brown (2022a: 7-9), because of the summative understanding of group evidence $\langle GEAA \rangle$ and $\langle ERG \rangle$ rely on.

So how does Continuous Evidentialism treat those cases? First, condition (1a) enables us to rule G 's belief in CONFLICTING BASES to be unjustified because the evidence G relies on E (E_1 - E_3) needs to be sufficient to support p . This condition cannot be met since E_1 - E_3 defeat each other's support for the respective proposition. In contrast, reaching the desired verdict in NON-CONFLICTING BASES is not as straightforward. That is a consequence of my general treatment of defeater cases [ch.6] [ch.7]. The mere existence of a defeater possessed by some member of G (and, therefore, within $\langle E_{\text{BASE}} \rangle$) does not undermine the justificatory status of G 's belief. What's important is whether there is and whether G utilises higher-order evidence about the entire evidential situation of G in forming their belief (as specified in condition (2)). Since this is not the case as NON-CONFLICTING BASES is spelt out, G 's belief cannot be justified according to Continuous Evidentialism.

This, however, is not to say that there are no cases similar to NON-CONFLICTING BASES in which G justifiably believes that p despite possessing some defeater. Acknowledging this was what motivated the verdict that G is justified in CAUTIOUS PI. Here is an extended version of the case discussed in [§6.4.2], in which I have filled in the relevant details.

CAUTIOUS PI [EXTENDED]: Paolo is the principal investigator of research group G . G possesses evidence E_1 - E_n for p which is distributed among the group's members. Simultaneously, Paolo has sufficient higher-order evidence E_H that E_1 - E_n supports p and that there is no further evidence within $\langle E_{\text{TOTAL}} \rangle$ that would nullify E_1 - E_n support for p .

¹⁶ For related discussions of the group defeater problems see Silva (2019: 274) or Brown (2022a: 9).

Based on this Paolo puts out a statement on behalf of the group that constitutes the group forming a belief that p . However, unbeknownst to Paolo, G possesses (via one of its members) a defeater φ for the belief that p .

This case is structurally analogous to the NON-CONFLICTING BASES case, in any way that is relevant for $\langle \text{GEAA} \rangle$ and $\langle \text{ERG} \rangle$. However, the further specifications added to CAUTIOUS PI suggest the opposite verdict. Here the higher-order mechanism that governs G 's belief formation enables us to judge that G 's belief that p is justified. This judgement is driven by the proposed understanding of epistemic responsibility [ch.7]. The way G has arrived at their belief that p is responsible independently of whether there is some defeater within (or outside) of G 's evidential base $\langle E_{\text{BASE}} \rangle$. This is something that $\langle \text{GEAA} \rangle$ and $\langle \text{ERG} \rangle$ are not sensitive to.

Analogously, Continuous Evidentialism in contrast to $\langle \text{ERG} \rangle$ and $\langle \text{GEAA} \rangle$ enables us to make sense of the following kinds of defeater-defeater cases:

CARELESS PI: Peter is the PI of research group G . G possesses evidence E_1 - E_3 for p which is distributed among the group's members. Each of the individual bodies of evidence E_1 , E_2 and E_3 are sufficient to support p . When deliberating E_1 - E_3 at the weekly meeting G reaches the result that the possessed evidence sufficiently supports p . Based on this discussion, Peter puts out a report that displays the evidence for p and states that p is the case, which constitutes the group forming a belief that p . However, unbeknownst to the other members of G , Peter possesses a defeater φ that undercuts the support E_1 - E_3 provides for p . Simultaneously, unbeknownst to Peter some member possesses a defeater-defeater λ that would have been revealed if Peter would have disclosed φ during the meeting.

CARELESS PI has a similar structure to CARELESS CARL, the defeater-defeater case discussed in [§6.4.2]. Both $\langle \text{ERG} \rangle$ and $\langle \text{GEAA} \rangle$ have problems giving us the right verdict concerning the case. First, $\langle \text{GEAA} \rangle$ would judge G 's belief in CARELESS PI to be justified because luckily and unbeknownst to the PI the defeater would have been defeated (condition (2) of $\langle \text{GEAA} \rangle$). Second, $\langle \text{ERG} \rangle$ would judge G to be justified because of the summative understanding of group evidence since neither λ nor φ is possessed "by enough of the operative members of G " (condition (1) of $\langle \text{ERG} \rangle$). This is, as I have argued, the wrong verdict because the PI and with it the way the group's belief is formed is irresponsible; that is, G acts epistemically irresponsible by not having and utilising higher-order evidence about their respective situation [ch.7].

At this point, the proponent of either $\langle \text{ERG} \rangle$ or $\langle \text{GEAA} \rangle$ might resort to the responsibilist conditions embedded in the respective accounts. However, both of these conditions fail to give

us the right verdict with respect to CARELESS PI. First, <GEAA> includes the responsibility condition (2c) as a subclause of the no-defeater condition (2); where, as formulated, (2) can only have a negative impact on the justificatory status of G's belief if full disclosure of the entire evidence relevant to the proposition would have a negative impact on the belief's status. This is not the case in defeater-defeater cases such as CARELESS PI. In these cases, the entire evidence is sufficient to support the considered proposition.

Second, clause (2) of <ERG> states that for G to justifiably believe that p, G is epistemically responsible in believing p on the basis of E; while for Silva, a "group, G, is epistemically responsible in believing P on the basis of E iff (a) enough of the operative members of G satisfy their G-relevant epistemic duties, and (b) G properly bases its belief on E" (2019: 276). Given this summativist understanding of collective epistemic responsibility, G is epistemically responsible in CARELESS PI because most members behave epistemically responsible.

This brings us to normative defeater cases, the final kind of defeater cases I want to consider. Lackey in motivating <GEAA>'s responsibility condition (2c) considers the following example (2016: 372-373):

GROUP NORMATIVE OBLIGATIONS: G is a group whose members consist of three nurses employed at a nursing home, N_1 - N_3 , each of whom justifiedly believes that the patient is not at risk of dying. N_1 is aware that she forgot to give O'Brien his first medication, but she also justifiedly believes that this act of negligence alone is not sufficient to put him in danger of death. N_2 is aware that she forgot to give O'Brien his second medication, but she also justifiedly believes that this act of negligence alone is not sufficient to put him at risk of serious health problems. And N_3 is aware that she forgot to give O'Brien his third medication, but she also justifiedly believes that this act of negligence alone is not sufficient to put his life in jeopardy. At the same time, however, N_1 - N_3 all justifiedly believe that O'Brien's missing all three of his medications would put him at serious risk of dying. [...] Thus, [...] N_1 - N_3 lack crucial evidence that they should have had and that would reveal the epistemic deficiency of their beliefs that O'Brien is not at risk of dying.

Lackey argues that in GROUP NORMATIVE OBLIGATIONS the nursing unit is epistemically irresponsible since the nurses should have shared their evidence ϕ about O'Brien's medication. Lackey insists that these normative obligations are "epistemic rather than merely professional or prudential" since the duty to consult with other members of the nursing unit "concerns evidence that [...] is highly relevant to their beliefs about O'Brien's health that directly affects whether they

know he is not at risk of dying” (2016: 374). In other words, Lackey (2016) understands the case to be an example of normative defeat.

Yet, it is not clear whether this means that G possesses and did not but *should have accessed* the defeating evidence or whether G did not but *should have possessed* it. This has to do with the interplay between Lackey’s understanding of normative defeat and her summative understanding of group evidence.¹⁷ As pointed out in [§6.3], in Lackey’s framework of normative defeat, as developed in the context of individual justification, normative defeaters are evidence that one possesses but irresponsibly ignores. However, if <GEAA> gives us the right verdict in GROUP NORMATIVE OBLIGATIONS only because the nursing unit already possessed the relevant evidence but did not access it <GEAA> runs into problems with other normative defeater cases, where some unpossessed evidence seems to have defeating force with respect to the proposition in consideration.¹⁸

Continuous Evidentialism gives us the right verdict, independently of whether we say that G possessed and should have accessed ϕ or whether G did not possess but should have possessed ϕ . The nursing unit is not justified in believing that O’Brien is *not* at serious risk of dying because condition (2) is not met. There is no higher-order evidence E_H sufficient to support q , that < E_{TOTAL} > is sufficient to support p (2a); and, even if there is such evidence this evidence was not used in the nursing unit forming the respective belief (2b). Note again, that realising condition (2) could take various forms. For example, beliefs had by N_1 - N_3 about records of the medication could be higher-order evidence of the relevant kind. However, if these records would have actively figured in G’s belief formation, as for example, in some deliberation among the nurses about O’Brien’s medication, conditions (2b) would as well have been satisfied. Yet, since there are no records and the nurses did not disclose any evidence about O’Brien’s medication, and since there is no other basis for q , the group cannot be justified in believing that p .

¹⁷ While Lackey (2016, 2021) never explicitly takes a stance on the nature of group evidence, Brown (2022a, 2023b) convincingly argues that the best way to understand <GEAA> is, as relying on a summative understanding of group evidence.

¹⁸ This interpretation is suggested by Lackey’s idiosyncratic understanding of normative defeat [§6.3]. Furthermore, Lackey appeals to normative obligations determining a sufficient amount of deliberation that would lead to further evidence being discovered by the group. This suggests that Lackey thinks that the evidence is already possessed by G. Interestingly, however, <GEAA> is often interpreted as relying on a traditional understanding of normative defeat; that is, as stating that G *should have possessed* the relevant defeating evidence (Silva 2019: 266-267; Brown 2022b: 174-176). For further discussion of Lackey’s understanding of normative defeat see, for example, Graham and Lyons (2021).

How does <ERG> handle this case? Here, Lackey rightly argues that Silva is not well-equipped to offer a general solution to cases such as GROUP NORMATIVE OBLIGATIONS. Mainly because of <ERG>'s reliance on G-relevant epistemic duties (2). This reliance is problematic since it allows groups to acquire justified beliefs based on duties that are considered epistemically defective from the outside. Or as Lackey puts it, G-relevant epistemic duties could “deviate in important ways from general epistemic duties [...] for instance, might exclude the relevance of scientific testimony from non-approved sources, where the criteria for approval is explicitly bound up with financial motivations. Thus, the group satisfies its G-relevant epistemic duties, but such duties are objectively deeply epistemically problematic” (2021: 84).

I want to wrap up by making a general point about how Continuous Evidentialism handles (collective) defeater cases. Often when looking at cases such as CONFLICTING BASES, NON-CONFLICTING BASES or GROUP EPISTEMIC OBLIGATIONS we are confronted with the question of whether the relevant defeaters are part of the group's evidential base or not. As Brown (2022a, 2022b, 2023b) points out, different summative and non-summative understandings of group evidence give us different verdicts about those cases. From my point of view, we do not need to engage in debates about group evidence to resolve these cases. As in the individual case, the important point is not whether there is a defeater that S possesses (within <E_{BASE}>) or that S should possess (within <E_{TOTAL}>) but rather whether S is epistemically responsible, that is, whether S has and utilises higher-order evidence about their epistemic situation. As such, Continuous Evidentialism makes sense of normative and non-normative defeater cases in the same manner.

The defended hybrid understanding of <E_{BASE}>, which includes group-level and member-level evidence is only important in showing what kind of evidence could be part of E, the operational part that could be used to base the group-level belief. And here it should be clear that given a causal understanding of basing any collective epistemic agent can base their beliefs on group-level evidence or member-level evidence as long as there can be a non-deviant causal chain from the evidence to the group-level belief.

Furthermore, neither Silva (2019) nor Lackey (2016, 2021) give us a detailed analysis of epistemic responsibility. Instead, they just specify summativist conditions which spell out under what circumstances G's member's epistemic responsibility amounts to G being epistemically responsible. In contrast, the understanding of epistemic responsibility Continuous Evidentialism is reliant on is not summative. As a result, a group G could be responsible by having and utilising higher-order evidence in the right kind of way without any of its members being responsible.

Finally, Lackey's and Silva's treatment of normative defeater cases via epistemic responsibility seems to rely on a no-defeater condition. In [§7.5], I have pointed out that many take such conditions to be problematic, because of some general problems faced by the doctrine of Defeatism. That is, <GEAA>'s and <ERG>'s treatment of defeater cases falls prone to general problems faced by Defeatism, while Continuous Evidentialism avoids those problems by handling those cases via a positive higher-order defender clause [ch.8].

9.3.3 Collective Basing

The way I have formulated the proper basing requirement in [§5.6] makes Continuous Evidentialism sensitive towards the distinction between properly and improperly based beliefs. This is important since these cases have traditionally spelt trouble for various accounts of (collective) justification.¹⁹ For example, <GEAA>, as pointed out by Silva (2019), fails to give us the right judgement with respect to cases in which a group is (doxastically) justified in believing that *p* because the belief is properly based on the evidence while none of the members properly bases their beliefs on the evidence. Take the following example, based on a case discussed by Silva (2019: 268-270):

IMPROPER JURY: Every juror of a jury *G* believes that *p*, the suspect *S* is guilty as charged, whereby their respective beliefs are improperly based on some evidence *E* sufficient to support *p*. Instead of believing that *p* because of *E* all of the individual jurors believe that *p* because the tea leaves say that *E* makes *p* overwhelmingly likely. However, every juror respects the fact that only court-presented evidence should be used in the context of legal deliberations. So, when the jury members deliberate, they never once include their private views about the tea leaves in their discussions nor do they in any other way take the tea leaves into account in arriving at a collective position on the *S*'s guilt. Rather, when they deliberate as a group and reach the view that *S* is guilty, they reach it in a proper way by directly relying on *E*. Thus, when *G* reports on its belief that *S* is guilty the jury explains how they arrived at that belief in the following way: It is the view of this jury that *S* is guilty. We hold this view because *E* makes it overwhelmingly likely that *S* is guilty, and *E* is true.

¹⁹ The most extensive discussion of epistemic basing in the context of collective epistemology to date is found in Brown (2022a; forthcoming: ch.3).

IMPROPER JURY illustrates two things. First, every account of justification needs to be sensitive to improper basing cases and, therefore, include some kind of proper basing clause.²⁰ This is something that <GEAA> misses but <ERG> is explicitly designed to handle. Second, the justificatory status of group-level beliefs can diverge when there is a divergence in the way the belief is based on sufficient evidence E , at the member level and the group level.

However, while <ERG> can make sense of IMPROPER JURY, it relies on a problematic summative, or more precisely, conjunctive understanding of collective basing (Brown 2022a: 7-9). For Silva (2019) if different members of a group believe that p for different reasons, the group's belief that p is based on the conjunction $E_1 \& E_2 \& \dots \& E_n$ if "enough operative members believe [...] p on the basis of some subset of E_1 - E_n , and E_1 - E_n are each part of the basis of enough of the operative member's belief in [...] p " (2019: 275). To illustrate why such a conjunctive understanding of basing is problematic, we need to imagine a case in which E_1 causes the group to believe that p while most of the members believe that p on other evidential grounds E_2 - E_n . So if those respective members receive some defeating evidence with respect to E_1 they would no longer believe that p . However, assuming non-summativism about group beliefs, this does not mean that the group would also stop believing that p . That is, any summative understanding of basing has the problem that if G believes p based on evidence E_1 - E_n distributed among its members, it is not sensitive to the evidence E_1 - E_n in the way we expect epistemic agents to be sensitive if they base their beliefs on E_1 - E_n . Something similar is also pointed out by Brown, who, in defence of a causal understanding of collective basing, imagines the following modification of Goldman's DIFFERENT BASES case [§9.3.1] (Brown 2022a: 9):

DIFFERENT METHODS: Suppose that a team of 100 scientists is investigating an issue while the different scientists employ different methodologies. Of the 100 members, 60 believe that p for different but compatible justifying evidence: M_1 - M_{20} believe that p for E_1 ; M_{21} - M_{40} for E_2 ; and M_{41} - M_{60} for E_3 (furthermore, we will suppose that the belief that p is not undermined by evidence other members of the group have or should have had). However, while M_{41-60} base their beliefs on E_3 , the other two subgroups, M_1 - M_{20} and M_{21} - M_{40} do not even grasp E_3 .

In discussing DIFFERENT METHODS Brown observes that due to the structural similarities to the DIFFERENT BASES case, summativists about basing, such as Silva (2019) would hold that G

²⁰ Although as Brown points out there is a cluster of different summative understandings of group basing compatible with <GEAA> (2022a: 6-7).

believes that p based on the conjunction of E_1 - E_3 . This is problematic, however, given that neither M_1 - M_{20} nor M_{21} - M_{40} nor M_{61} - M_{100} grasp E_3 . Therefore, most of the group members cannot take the conjunction of E_1 - E_3 as evidence for their belief that p .

Another problematic consequence of this is that it makes the group's belief counterfactually insensitive to possible defeating evidence. If, for instance, G receives evidence that defeats E_1 and E_2 but not E_3 "then M_1 - M_{40} would no longer believe that p , and so there would no longer be a majority of operative members who believe that p . Thus, on the summative approach to group belief, the group would not believe that p . That's not what we'd expect if the group believes that p on the basis of the conjunction" (Brown 2022a: 9).

9.3.4 Evidential Manipulation

Thinking about epistemic justification in the context of collective epistemic agents highlights an unappreciated aspect of epistemic agents. Groups seem to have a degree of intentional control over their own constitution and design that surpasses any degree of intentional design present in human agents. Acknowledging that this control can then be put to better and worse use opens up a new dimension of normative evaluation.

Lackey (2016) demonstrates that many theories of collective justification allow groups to gain or sustain beliefs by manipulating their evidence. While the problem most obviously applies to joint acceptance accounts of collective justification (see below), many other accounts including <GEAA> (as Lackey admits) and <ERG> face versions of this objection. Lackey calls this the Illegitimate Manipulation of Evidence Problem (2016: 353).

Lackey introduces the problem in her discussion of Gilbert-style (1989) joint acceptance accounts of epistemic justification (Schmitt 1994; Hakli 2011; Schwengerer 2021). One feature of these accounts is that group justification depends on the evidence possessed by the group and that a group possesses some evidence only when it is mutually accepted by the group members as evidence. Lackey demonstrates that this leads to problematic consequences in the following case (2016: 351):

IGNORING EVIDENCE: Philip Morris is one of the largest tobacco companies in the world, and each of its operative members is individually aware of the massive amounts of scientific evidence revealing not only the addictiveness of smoking but also the links it has with lung cancer and heart disease. Moreover, each individual member believes that the dangers of smoking give the company a reason to believe that warning labels should be placed on cigarette boxes. However, because of what is at stake financially and legally,

none of these members would properly express a willingness to accept that the dangers of smoking give Philip Morris a reason to believe that it should put warning labels on cigarette boxes.

Lackey discusses two versions of this case, where in the above-described version Philip Morris decides to *ignore* some evidence revealing the addictiveness of smoking and in the latter, the company actively *manufactures* evidence that is meant to undermine the evidence revealing the addictiveness of smoking (2016: 351).

Lackey uses these examples to argue against the joint acceptance views defended by Schmitt (1994) and Hakli (2011).²¹ Importantly, however, there are other ways in which collective epistemic agents can manipulate their evidential base intentionally. For example, if we have some kind of summative understanding of group evidence (such as intersecting or majoritarian understandings [§5.7]) we can actively manipulate the evidential base by adding or removing group members. To illustrate this version of the evidence manipulation problem, Lackey considers an example in which a group is “achieving epistemic justifiedness simply by extending membership to more and more people, with no regard whatsoever to the grounds of the individual beliefs beyond independence” (Lackey 2016: 376). Accordingly, this is a problem for any understanding of collective justification that is directly dependent on the evidence possessed by the group’s members or the justificatory status of their attitudes.

What does Continuous Evidentialism tell us about such manipulation problems? First, manipulation cases, while possible, work differently in the context of Continuous Evidentialism. While adding or removing members intentionally can change G ’s evidential base, for any manipulation of $\langle E_{\text{BASE}} \rangle$ to affect the status of G ’s beliefs, this manipulation needs to affect the operational basis (E or E_H) of G ’s belief. This is because according to the proposed treatment of defeater cases, adding or removing defeating evidence from $\langle E_{\text{BASE}} \rangle$ alone does not necessarily change the epistemic status of G ’s belief. After all, in opting for a positive higher-order requirement to explain defeater cases I have dealt away with the notion of defeat.²² Relatedly, Continuous Evidentialism also relies on an unorthodox understanding of doxastic justification that does *not* require that for a belief that p to be justified, p needs to be propositionally justified by $\langle E_{\text{BASE}} \rangle$ [§6.7].

²¹ For a defence of joint-acceptance accounts of collective justification against these charges see Schwengerer (2021).

Nonetheless, here is one class of manipulation cases that poses a challenge to Continuous Evidentialism, which I will call *assemblage manipulation* cases. Here is one such example:

BIASED PI: A biased PI gathers scientists to form a scientific research group G based on the PI's knowledge of the positions defended and evidence possessed by the respective scientists concerning proposition p . The biased PI does so to create a group that possesses a body of evidence E sufficient to support p , and an evidential base $\langle E_{\text{BASE}} \rangle$ containing no defeaters that undercut or rebut E 's support for p . After formation G deliberates their evidence and forms the belief that p .

Is G 's belief that p justified? That depends on the omitted details. First, if the PI plays a causal role in the belief formation of the group, the group has no sufficiently strong higher-order evidence E_H that bears on the overall evidential situation. This is because the PI's awareness of the manipulation undercuts the support that any higher-order evidence E_H could provide for the proposition that $\langle E_{\text{TOTAL}} \rangle$ supports p . So, for G to be justified the PI would need to leave the group after assembly or not be aware of their bias. (Alternatively, we can imagine a scenario in which we randomly select a group of scientists out of the pool of scientists resulting, against all odds, in a group justifiably denying an otherwise widely accepted proposition).

So, if we add those details, should we still be worried about G being epistemically justified? I say no. These cases are not problematic since we cannot hold G accountable for their faulty beliefs. This bears on the fundamental internalist intuitions that motivated Continuous Evidentialism. So, cases such as BIASED PI are similar to cases in which the evidential manipulation is done from the outside. If we do not judge an individual epistemic agent who is systematically deceived, as in a global sceptical scenario, irresponsible for being fed misleading evidence from the outside we need to say the same about G . And since I argued that epistemic justification and responsibility go hand in hand, we need to think of G 's belief as being epistemically justified. This means if all other conditions of Continuous Evidentialism are met G 's belief that p is justified.

Moreover, this treatment of assemblage manipulation cases helps me to highlight a general lesson; namely, that newly assembled groups need higher-order evidence about their belief-forming mechanism that these mechanisms can produce justified beliefs. In other words, a just assembled group that did not inherit any evidence about the reliability or properness of the way it forms beliefs is in the same position as Norman the clairvoyant or Seeing Sandra [§8.2]. Here is a case that illustrates this:

ASSEMBLY: Two groups of probands in a sociology experiment G_1 and G_2 are independently guessing the number of marbles in a large urn. Upon deliberation both the members of G_1 and G_2 agree to all make independent guesses and then take the weighted average as the guess of the group (in a way which constitutes the group forming a belief). Based on this procedure G_1 believes that there are 679 marbles in the urn and G_2 believes that there are 701 marbles in the urn. G_1 reached this verdict by taking evidence about Condorcet's results and the reliability of such a procedure into account which was brought to bear by one of the probands during the deliberation phase. In contrast, none of G_2 's members was aware of the reliability of this procedure. Instead, they agreed on this method as a matter of fairness.

It should be clear from the set-up of the case that G_1 's belief can be justified while G_2 's cannot. After all, G_1 had good evidence to pick the Condorcet procedure as a way of forming a belief about the number of marbles in the urn while G_2 's reasons were entirely non-epistemic.

9.4 The Big Picture

We have seen how the different conceptual components from the earlier chapters come together to create Continuous Evidentialism. We have also seen how Continuous Evidentialism fits with the intuitively correct verdicts of many often-discussed cases; including some cases that have caused troubles for other influential accounts of collective justification such as Lackey's *Group Epistemic Agent Account* (2016, 2021) or Silva's *Evidentialist Responsibility for Groups* (2019). It is now time to turn to the big picture. What are the general lessons we have learned about collective justification from the proposed treatment of those cases? First, we can revisit the epistemic divergence arguments discussed in [ch.3] and draw some general lessons about the divergence between individual and collective justification [§9.4.1]. Second, we can locate Continuous Evidentialism on the spectrum of distributed and non-distributed accounts of collective justification [§9.4.2]. Third, we can revisit the criteria underlying Continuous Epistemology and demonstrate that Continuous Evidentialism is epistemically continuous in the desired sense [§9.4.3].

9.4.1 Epistemic Divergence: Revisited

In [ch.3] I have discussed so-called epistemic divergence cases, which are cases in which the epistemic status of a group-level attitude diverges from the epistemic status of the member-level attitudes. As such, these cases are meant to refute minimal epistemic summativism <MESC>.

Minimal Epistemic Summativist Commitment <MESC>: A group's doxastic attitude *D* towards *p* has the epistemic status *S* only if there is at least one member of *G* who has the doxastic attitude *D* towards *p* and for whom *D* has the epistemic status *S*.

I have captured the general structure of epistemic divergence arguments as follows [§3.2]:

- (1) If <MESC>, then necessarily, *G*'s doxastic attitude *D* towards *p* has the epistemic status *S* only if at least one member *m* of *G* has the doxastic attitude *D* towards *p* and *m*'s attitude has the epistemic status *S*.
- (2) It is possible that *G* has the doxastic attitude *D* towards *p* with the epistemic status *S* while no member *m*₁-*m*_{*n*} of *G* has the doxastic attitude *D* towards *p* with the epistemic status *S*.
- (C) <MESC> is false.

By looking at the argumentative structure of epistemic divergence arguments we can see that these arguments can work in two different ways. First, they can show that a group can have an attitude *D* with the epistemic status *S* while no member has the relevant attitude *D*. Second, the divergence can arise because the epistemic status of the relevant member-level attitude diverges from the epistemic status of the group-level attitude. Now, in [ch.3] I have argued that we can utilise epistemic permissivism, or more generally the Conflict View and the underlying understanding of epistemic standards, to generate epistemic divergence cases of the second kind.

We are now able to see that these latter kinds of divergence cases can be constructed in a variety of ways, where I have discussed distributed evidence cases [§9.3.1] collective defeater cases [§9.3.2] and collective basing cases [§9.3.3]. Here, each of these cases inspired the denial of some summativist assumptions (built into <GEAA> and <ERG>) about group belief, group evidence or collective basing. Let me add a final case to this list, inspired by Brown's recent work on collective justification (forthcoming: 93-99). This case is a divergence case of the second kind; that is, a case in which a group justifiedly believes that *p* while no member believes that *p*.

Lackey (2008), in her work on testimony, has presented us with a convincing example that demonstrates that one can be a good testifier without believing the proposition one testifies about

and so does not know or justifiably believe that proposition. To illustrate this, Lackey imagines the case of a creationist teacher who does not believe in Darwinian evolution but who eloquently teaches students about the scientific evidence underlying certain evolutionary principles. Following Brown (forthcoming: 98-99), we can use these testimony-without-belief cases to construct another divergence case:

CREATIONIST COMMITTEE: A group *G* of creationist biology teachers meet to put out a statement about teaching Darwinian evolution in school. All of them, due to their professional training in biology possess lots of evidence about the scientific standing of Darwinian evolution. However, none of them, due to their religious conviction, believe in Darwinian evolution. Nonetheless, they keep secret about it. After deliberating lots of evidence about Darwinian evolution they put out a statement in which they recommend *p*, that Darwinian evolution should be integrated into the regular syllabus. This constitutes *G* forming a belief that *p*.

Plausibly, *G* justifiably believes that *p*, since *G* possesses sufficient evidence supporting *p* as well as higher-order evidence gained via the deliberation about the entire evidential situation with respect to *p*. This is possible despite none of *G*'s members believing that *p*, which means that justified group-level belief does not require member-level (justified) belief.

9.4.2 The Best of Both Worlds

Confronted with the plethora of different divergence cases discussed above, some might be worried that Continuous Evidentialism disconnects member-level and group-level justification in an unhealthy way. More precisely, we might worry that the proposed treatment of collective justification is vulnerable to worries raised against other accounts of collective justification that disconnect collective and individual justification and knowledge such as Bird's (2010, 2014, 2019) distributed cognition model.

On the one hand, these models are attractive since they make sense of cases in which evidence is spread among the group in a compartmentalized way such as DISTRIBUTED COGNITION [§9.3.1]. In so doing, these models do justice to the fact that groups can be structured and process information in manifold ways.

On the other hand, Lackey (2021) raises worries about Bird's (2010, 2014, 2019) distributed cognition account. First, she argues that Bird divorces the vital connection between group knowledge, group action, and collective responsibility (2021: 115-123). Second, she argues that

allowing groups to know or justifiably believe things via distributed cognition makes them insensitive to defeater cases (2021: 123-127). Without spelling out the details of these worries, we can see that Continuous Evidentialism is immune to them. First, since Continuous Evidentialism is not reliant on a specific theory of group belief, we can outsource this worry to such a theory. Here, as mentioned, promising candidates able to supplement Continuous Evidentialism are functionalist or interpretationist accounts of belief because they are epistemically continuous. So if a group can justifiably believe something according to theory X without being able to act on this belief a continuous theory would also allow for cases in which this is the case for individuals, which might be considered a vice of theory X and hence a general reason to reject it. Furthermore, if you are an interpretationist about beliefs you will read off beliefs from the behaviour of the respective system. This means that beliefs and knowledge cannot be disconnected from actions as feared by Lackey.

Second, Continuous Evidentialism can make better sense of defeater cases than distributed cognition accounts can, because of the higher-order requirement (2). This requirement makes groups, and epistemic agents in general, sensitive to possessed and unpossessed defeaters by requiring them to have higher-order evidence bearing on their epistemic situation. Bird's understanding of distributed cognition lacks such a requirement. This is problematic because it makes group beliefs that are justified based on distributed evidence insensitive to defeaters spread among the group members. Furthermore, distributed cognition models give us the wrong verdict with respect to cases such as ASSEMBLY [§9.3.4]. There it judges G_1 's and G_2 's respective beliefs to be justified independently of how they reached the agreement to adopt Condorcet belief-formation procedure. This treatment of defeater cases was motivated by a general understanding of epistemic responsibility that underlies our intuitions about those cases. So, if this treatment is correct it generalises to the collective realm, and with it to distributed and non-distributed cognition cases.

In sum, Continuous Evidentialism is the best of both worlds. It is not restricted to small-scale non-distributed socio-epistemic entities such as <GEAA> or <ERG>. It makes room for the possibility of large-scale distributed social cognition and the epistemic evaluation of those socio-epistemic systems. As such, it enables us to reach the intuitively correct verdict concerning distributed cognition cases. However, it does so without facing well-known problems faced by other accounts which try to capture that phenomenon. This is mostly due to Continuous Evidentialism's neutrality with respect to the nature of group beliefs, as well as its higher-order evidentialist treatment of defeater cases (2).

9.4.3 The Possibility of Continuous Epistemology

I want to end by highlighting that in developing Continuous Evidentialism, defending it against various objections and applying it to a wide range of cases from within individual and collective epistemology, I have demonstrated the possibility of a continuous theory of epistemic justification. That is, I have developed a theory that relies on the following three principles spelt out in [ch.4]:

Conceptual Unity <Unity>: Epistemological concepts are general. That is, whenever we evaluate some type of epistemic state of different types of epistemic agents, we evaluate them under one unified concept.

Methodological non-Primacy <non-Primacy>: Epistemological theorising has no preferred starting point. That is, all case judgements should *ceteris paribus* be weighted in the same way.

Epistemic Agency <Agency>: There is a general way to characterise and identify epistemic agents. That is, there is one universal set of characteristics that enables us to identify potential bearers of epistemic states.

In [§4.6], I have argued that any continuous theory that relies on these principles is *ceteris paribus* superior to its non-continuous competitors. This has been motivated by general reflections on the nature of epistemic explanations, where continuous explanations have more explanatory power in unifying different phenomena in a way that maximises explanatory depth and explanatory range. However, I also noted there, that increased explanatory power is only a good if *ceteris* is indeed *paribus*. But that's what I showed throughout the second half of the thesis. Indeed, I understand the results from this chapter to go even a step further. Instead of showing that Continuous Evidentialism is at least as promising as other accounts in the literature, I have demonstrated that it fares better than other influential accounts of collective justification on various metrics.

10

General Conclusion

10.1 Recapitulation

In this thesis, I investigated the nature of epistemic justification and epistemic rationality. I started with some loosely correlated improvisations on various themes arising from issues having to do with the (im)permissibility of epistemic rationality. In particular, I developed a novel view of allegedly permissive cases [ch.1], defended it against an objection from self-fulfilling beliefs [ch.2] and demonstrated that the proposed view has some important implications for collective epistemic theorising [ch.3].

In contrast, the second half of the thesis [ch.5] – [ch.9] had more of a monograph-like structure, in which I stepwise, chapter by chapter, developed a theory of epistemic justification that is meant to be used to investigate the epistemic status of collective and individual epistemic agents. While I took an unconventional approach, I ended up defending a conventional-looking theory, evidentialism. However, as we have seen, while built on a well-explored epistemological framework, to make my theory work we had to make some unusual moves, open up some new avenues, bite some bullets and dwell on some strange consequences. It is now time to recapitulate this journey in detail.

Over the course of this thesis, I have defended the following propositions:

1. The Conflict View: Apparently permissive cases are epistemic standard conflicts. That is, cases in which one body of evidence supports different doxastic attitudes relative to different incommensurable sets of (permissive) epistemic standards.

1.1. More precisely, allegedly permissive cases are not cases in which we are confronted with multiple equally rational attitudes but with multiple incommensurable attitudes that are supported by different sets of epistemic standards.

1.2. Like value conflicts, epistemic standard conflicts have various idiosyncratic properties of choices under incommensurability, namely: sweetening insensitivity, bindingness, and angst.

1.2.1. Sweetening insensitivity: improving either of the alternatives does not necessarily make the sweetened alternative rationally obligatory and the other rationally impermissible.

1.2.2. Bindingness: ending up with one alternative binds one to appeal to similar resolutions in analogous future conflicts.

1.2.3. Angst: resolving the conflict often leaves a persistent uncertainty about whether one has responded rightly.

1.3. Treating allegedly permissive cases as epistemic standard conflicts gives us a novel understanding of the so-called arbitrariness objection to permissivism as well as the bindingness of epistemic conflict resolutions.

1.4. This allows us to shift the debate between permissivists and impermissivists by reducing normative questions about epistemic rationality to more fundamental questions about the nature of incommensurability.

1.4.1. Different extant understandings of incommensurability provide us with different answers to the question of whether epistemic rationality is genuinely permissive.

2. We cannot utilise cases involving self-fulfilling beliefs to cook up counterexamples to Evidential Uniqueness, the thesis that for any body of evidence *E* and any proposition *p* there is at most one doxastic attitude that any epistemic agent *A* could rationally take.

2.1. If self-fulfilling belief cases are epistemically permissive they would pose a threat not only to Evidential Uniqueness but also to the Conflict View.

2.2. My treatment of self-fulfilling belief cases makes them synchronically and diachronically impermissible. In other words, for any self-fulfilling situation, there is not only a unique rational doxastic attitude to be in but also only a unique rational doxastic attitude to transition into.

2.2.1. There are two types of self-fulfilling belief cases, cases in which we prefer one of the self-fulfilling attitudes and cases in which we are genuinely indifferent towards the different self-fulfilling attitudes.

2.2.1.1. In the former, practical considerations about prospective behaviour together with general assumptions about the degree to which belief formation is voluntary generate additional evidence that tips the evidential balance towards the preferred self-fulfilling doxastic attitude.

2.2.1.2. In the latter, epistemic agents are permitted to pick between the different actions of trying to induce one belief or another. In so doing, they generate evidence that singles out a uniquely rational doxastic attitude.

2.2.1.2.1. This means that genuinely indifferent self-fulfilling belief cases are practically but not epistemically permissive.

3. Epistemic conflicts, permissible or not, provide us with a template to generate epistemic divergence cases; that is, cases in which the epistemic status of group-level attitudes and member-level attitudes diverges.

3.1. Epistemic divergence cases can be used to support epistemic divergence arguments.

3.2. Epistemic divergence arguments can be used to argue against summativist understandings of epistemic justification and knowledge. Hence they motivate an inflationary understanding of collective epistemology.

3.3. There is a general template for constructing divergence cases based on epistemic permissivism.

3.3.1. Permissive Divergence Template: There are cases in which a group G is rationally required to have the doxastic attitude D towards p in relation to a body of evidence E and its permissible epistemic standards S . While all of G 's members m_1 - m_n possess the same body of evidence E , they have different permissible epistemic standards S_1 - S_i which recommend different doxastic attitudes D_1 - D_n towards p .

3.4. The presented divergence argument is superior to other divergence arguments that have been proposed in the literature because it makes divergence possible even under very stringent conditions, namely: (a) full disclosure of the evidence within the group, (b) mutual awareness of the disclosure and the modus operandi of the group, even if, (c) the group G itself, as well as all members m_1 - m_n are fully rational (or at least not determinately irrational).

4. Continuous Epistemology: Doing epistemology is doing epistemology *per se*. That is, when we engage in epistemological conceptual analysis we analyse general domain-overarching concepts that enable us to capture and uniformly evaluate the epistemic states of (every kind of) epistemic agent.

4.1. There are two often-distinguished epistemological endeavours, individual and collective epistemology. While individual epistemology is concerned with individual epistemic agents and their epistemic states, collective epistemology is concerned with collective epistemic agents and their epistemic states.

4.2. Both collective and individual epistemology are done by doing conceptual analysis following the method of cases.

4.3. The method of cases rests on three principles: Conceptual Unity, Methodological non-Primacy and Epistemic Agency.

4.3.1. Conceptual Unity: Epistemological concepts are general. That is, whenever we evaluate some type of epistemic state of different types of epistemic agents, we evaluate them under one unified concept.

4.3.2. Methodological non-Primacy: Epistemological theorising has no preferred starting point. That is, all case judgements should *ceteris paribus* be weighted in the same way.

4.3.3. Epistemic Agency: There is a general way to characterise and identify epistemic agents. That is, there is one universal set of characteristics that enables us to identify potential bearers of epistemic states.

4.4. If we apply the principles (Conceptual Unity, Methodological non-Primacy, Epistemic Agency) consequently in our epistemic theorising we are doing epistemology the continuous way.

4.5. Analysing whether these principles are accepted or not enables us to categorise different extant views in collective epistemology and highlights the novelty of a continuous approach.

5. Continuous Evidentialism: S justifiably believes that p iff: (0) S believes that p. (1)(a) S possesses some evidence E, which is sufficient to support Bp, and (b) Bp is properly based on E, and (2)(a) S possesses some higher-order evidence E_H, which is sufficient to support q, the proposition that the total evidence <E_{TOTAL}> on balance supports p and (b) S's belief that p is properly based on E_H.

5.1. E is an ultimate piece of evidence only if it is the propositional content of an experiential non-factive mental state M.

5.2. The evidential base <E_{BASE}> of an epistemic agent S is exactly the sum of the propositions picked out by the relevant mental states of S and any members and subgroups of S which are epistemic agents themselves.

5.3. Explanationism: Some evidence E is sufficient to support the belief that p iff: (i) E is part of the best explanation for p, or (ii) p is a logical consequence of some sufficiently good explanation for p and (iii) E increases the likelihood of p above the believability threshold t.

5.4. A belief B is properly based on some evidence E iff it is non-deviantly caused by E and fulfils the properness requirements.

6. Many have defended the doctrine of Defeatism, the idea that doxastic attitudes can be *prima facie* justified by having some justification-conferring property while being defeated and, therefore, lacking *ultima facie* justification.

6.1. To specify when a *prima facie* justified belief is *ultima facie* justified we need to have a no-defeater clause.

6.1.1. There are several kinds of defeaters depending on their location within the total body of evidence relevant to the epistemic situation. The main types are normative and mental state defeaters.

6.1.2. There are two main types of no-defeater clauses, evidentialist and responsibilist clauses.

6.2. Evidentialist clauses fail to give us the right verdict in at least three different kinds of defeater cases: normative defeater cases, lucky defeater-defeater cases and unlucky defeater cases.

6.3. In contrast a responsibilist understanding of defeat can give us the right verdict with respect to the entire taxonomy of defeater cases.

6.3.1. This responsibilist strategy requires epistemic agents to be sensitive towards the evidence they should have possessed and the evidence they should have accessed.

6.4. The proposed treatment of defeater cases suggests an unorthodox understanding of the relationship between propositional and doxastic justification on which propositional justification is not a requirement for doxastic justification.

7. We can understand epistemic responsibility via the socio-epistemic expectations we are entitled to have of epistemic agents.

7.1. These expectations entitle us to expect from epistemic agents that they possess and utilise higher-order evidence when forming beliefs.

7.1.1. This enables us to reduce epistemic responsibility to a higher-order evidential requirement.

7.2. S is epistemically responsible iff: S possesses some higher-order evidence E_H , which is sufficient to support q, the proposition that the total evidence $\langle E_{TOTAL} \rangle$ on balance supports p and (b) S's belief that p is properly based on E_H .

7.3. This understanding of responsibility takes the form of a positive higher-order requirement rather than a negative higher-order requirement.

7.3.1. Positive Higher-Order Requirement: If S's belief that p is justified then S necessarily has a defender Φ , that is a [belief/justified belief/ evidence/...] indicating that p is justified.

7.3.2. Negative Higher-Order Requirement: S's *prima facie* justified belief that p is *ultima facie* justified iff S lacks a defeater ϕ [that S should have had], that is [a

belief/justified belief/evidence...] indicating that the belief that p is not *prima facie* justified.

7.4. Positive higher-order requirements can make sense of defeater cases while rejecting the doctrine of Defeatism.

8. Positive higher-order requirements can be motivated by general internalist intuitions.

8.1. There are four different types of positive higher-order requirements:

8.1.1. Doxastic: If S 's belief that p is justified then S necessarily has Φ , a doxastic mental state that indicates that p is justified.

8.1.2. Mental: If S 's belief that p is justified then S necessarily has Φ , a non-doxastic mental state that indicates that p is justified.

8.1.3. Dispositional: If S 's belief that p is justified then S necessarily has Φ , a disposition to form beliefs in a way that indicates that p is justified, and this disposition caused S to believe that p .

8.1.4. Evidential: If S 's belief that p is justified then S necessarily has Φ , sufficient evidence that indicates that p is justified.

8.2. Positive higher-order requirements face the following three objections.

8.2.1 Over-intellectualisation: We often judge less cognitively sophisticated epistemic agents (such as children, animals, or certain types of collectives) to have justified beliefs. Higher-order requirements conflict with such judgements and, therefore, implausibly overintellectualise our theory of epistemic justification.

8.2.2 Epistemic Relevance: Higher-order requirements need to contribute to achieving the overall aim of justification. It is not clear how higher-order requirements such as the requirement to have a higher-order mental state do that.

8.2.3 Regress: Higher-order requirements induce a vicious regress. If we think that the initial reason to posit that first-order requirements need to be supplemented by second-order requirements can be extended to higher-order requirements we end up with an infinite hierarchy of requirements.

8.3. Only the evidential higher-order requirements can withstand all three objections.

9. Continuous Evidentialism gives us the right verdict with respect to a wide range of cases often discussed in works on collective justification.

9.1 In particular, Continuous Evidentialism can make sense of distributed evidence cases, collective defeater cases, collective basing cases as well as evidential manipulation cases.

9.2 This demonstrates that Continuous Evidentialism fares better concerning a wide range of cases than some of the most promising extant accounts of collective justification, including Jennifer Lackey's Group Epistemic Agent Account, as well as Paul Silva's Evidentialist Responsibility for Groups.

10.2 Outlook

Having recapitulated the general argumentative structure, as well as the key contributions of the thesis enables us to hypothesise about possible applications and further developments. Let me list a few of them here.

First, the Conflict View [ch.1] opened up possible new avenues to explore different questions related to the debate surrounding epistemic permissivism, such as synchronism vs diachronism or epistemic relativism. Take first the debate between synchronism and diachronism about epistemic rationality. While diachronists believe that epistemic rationality is partly governed by diachronic norms such as conditionalisation [§5.5.1], synchronists believe that “the relationship between two time-slices of the same person is not importantly different, for purposes of rational evaluation, from the relationship between time-slices of distinct persons” and that “the locus of rationality, so to speak, is the time-slice rather than the temporally extended agent” (Hedden 2015b: 424). Different understandings of bindingness [§1.6] will support diachronic norms of rationality to different extents. Here, my preferred understanding of bindingness, weak bindingness, might be compatible with synchronism as defended by Moss (2015) or Hedden (2015b).

Second, in shifting the debate between permissivists and impermissivists by reducing normative questions about epistemic rationality to more fundamental questions about the nature of incommensurability, we have raised new questions about epistemic relativism, the view that rationality is framework-relative and/or that there is no uniquely superior epistemic system, that is, an interconnected set of epistemic standards (Kusch 2017; Baghramian & Carter 2022: §4.4). While the former understanding of incommensurability aligns well with full-blown relativism, the latter is compatible with anti-relativist approaches to epistemology.

Third, in [ch.2] I presented an argument against the permissibility of self-fulfilling belief cases. I think we can extend this argument to argue against a cluster of views that appeal to the doctrine of epistemic freedom (Velleman 1989; Antill 2020), which we might call epistemic liberalism.

Epistemic liberalism, as defended by Velleman (1989) combines permissivism with non-evidentialism as well as doxastic voluntarism, the view that sometimes epistemic agents can choose to acquire a belief directly via an act of will. The basic idea behind Velleman's proposal is that in self-fulfilling belief cases, we can acknowledge that forming an attitude one way or another generates evidence for that formed attitude, which is to acknowledge that we have generated practical reasons to form either of the self-fulfilling attitudes, which then enables us to voluntarily pick one of those attitudes. While many have argued for or against various aspects of epistemic liberalism, by using the model discussed in [ch.2] we can give a unified argument against all three allegedly liberal aspects of self-fulfilling belief cases. Interestingly, however, this argument illustrates that while the overall narrative of the liberalist treatment of SFB-cases is correct it cannot be used to motivate any of the theses underlying epistemic liberalism. In other words, I think that if we carefully fill out various omitted details in the deliberation process proposed by the liberalist, we will find out that this process is neither epistemically permissive, directly employs practical reasons, nor supports doxastic voluntarism.

Fourth, while I have not defended any theory of group belief, Continuous Evidentialism motivates the search for a theory of group belief that is epistemically continuous and hence compatible with Continuous Evidentialism. Plausible candidates are functionalist as well as interpretationist theories. My proposed understanding of epistemic divergence and group-level epistemic standards raises new questions for interpretationist views of collective attitudes (Tollefsen 2002b, 2015; List & Pettit 2011; Brouwer, Ferrario & Porello 2021). Most of those views fix the doxastic attitudes of an epistemic agent by interpreting their behaviour in the most rationalising manner (Lewis 1974; Dennett 1981; Davidson 1984; Williams 2020). In so doing, we need to appeal to some epistemic standards according to which their behaviour is maximally rational. Here, group-level epistemic standards, in case they exist, are the most plausible candidates to get a fix on the most rational interpretation of groups.

Bibliography:

- Adam, J. & Sosa, E. (2022). Metaepistemology. *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/spr2022/entries/metaepistemology/>.
- Alston, W. (1988) An Internalist Externalism. *Synthese*, 74, 265-83. <https://doi.org/10.1007/bf00869630>.
- Amico, R. (1988). Reply to Chisholm on the Problem of the Criterion. *Philosophical Papers*, 17 (3), 235-236. <https://doi.org/10.1080/05568648809506303>.
- Amico, R. (1993) *The Problem of the Criterion*. Rowman & Littlefield.
- Antill, G. (2019). Evidence and Self-Fulfilling Belief. *American Philosophical Quarterly*, 56 (4), 319-330. <https://doi.org/10.2307/48563046>.
- Antill, G. (2020). Epistemic freedom revisited. *Synthese*, 197 (2), 793-815. <https://doi.org/10.1007/s11229-018-1735-6>.
- Appley, B. & Stoutenburg, G. (2017). Two new objections to explanationism. *Synthese* 194 (8), 3069-3084. <https://doi.org/10.1007/s11229-016-1093-1>.
- Audi, R. (1993). *The structure of justification*. Cambridge University Press.
- Backes, M. (2021). Can Groups be Genuine Believers? The Argument from Interpretationism. *Synthese* 199, (3-4), 10311-10329. <https://doi.org/10.1007/s11229-021-03246-w>.
- Baehr, J. (2009). Evidentialism, vice, and virtue. *Philosophy and Phenomenological Research*, 78 (3), 545-567. <https://doi.org/10.1111/j.1933-1592.2009.00255.x>.
- Baehr, J. (2011). *The Inquiring Mind*. Oxford University Press.
- Baghrarian, M. & Carter A. (2022). Relativism. *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/spr2022/entries/relativism/>.
- Baker-Hytech, M. & Benton, M.A. (2015). Defeatism Defeated. *Philosophical Perspectives*, 29, 40-66. <https://doi.org/10.1111/phpe.12056>
- Ballantyne, N. (2015). The significance of unpossessed evidence. *Philosophical Quarterly*, 65 (260), 315-335. <https://doi.org/10.1093/pq/pqu096>.
- Beddor, B. (2015). Evidentialism, circularity, and grounding. *Philosophical Studies*, 172 (7), 1847-1868. <https://doi.org/10.1007/s11098-014-0375-z>.
- Berghofer, P. (2022). *The Justificatory Force of Experiences: From a Phenomenological Epistemology to the Foundations of Mathematics and Physics*. Springer (Synthese Library).
- Bergmann, M. (2005). Defeaters and higher-level requirements. *Philosophical Quarterly* 55 (220), 419-436. [10.1111/j.0031-8094.2005.00408.x](https://doi.org/10.1111/j.0031-8094.2005.00408.x).
- Bergmann, M. (2018). Externalist Versions of Evidentialism. In McCain, K. (Ed.) *Believing in Accordance with the Evidence*. Springer (Synthese Library). https://doi.org/10.1007/978-3-319-95993-1_8.
- Biebel, N. (2023). Pragmatic encroachment and Justified Group Belief. *Synthese* 202 (2), 1-20. <https://doi.org/10.1007/s11229-023-04271-7>

- Bird, A. (2007). Justified judging. *Philosophy and Phenomenological Research* 74 (1), 81–110. <https://doi.org/10.1111/j.1933-1592.2007.00004.x>
- Bird, A. (2010). Social Knowing: The Social Sense of ‘Scientific Knowledge’. *Philosophical Perspectives*, 24, 23–56. <https://doi.org/10.1111/j.1520-8583.2010.00184.x>.
- Bird, A. (2014). ‘When Is There a Group that Knows?’: Distributed Cognition, Scientific Knowledge, and the Social Epistemic Subject. In Lackey Jennifer (Ed.), *Essays in Collective Epistemology* (pp. 42–63). Oxford University Press.
- Bird, A. (2019). Group Belief and Knowledge. In Fricker M. et al. (pp. 274-283) *The Routledge Handbook of Social Epistemology*. Routledge.
- Blackburn, S. (1996). Dilemmas: Dithering, Plumping, and Grief. In H. E. Mason (Ed.), *Moral Dilemmas and Moral Theory* (pp. 127-139). Oxford University Press.
- Block, N. (1978). Troubles with functionalism. *Minnesota Studies in the Philosophy of Science*, 9, 261-325. <https://hdl.handle.net/11299/185298>.
- Boghossian, P. (2003). Blind Reasoning. *Proceedings of the Aristotelian Society*, 77 (1), 225–248. <https://doi.org/10.1111/1467-8349.00110>.
- Boghossian, P. (2014). What is Inference? *Philosophical Studies* 169 (1), 1-18. <https://doi.org/10.1007/s11098-012-9903-x>.
- Bonjour, L. (1980). Externalist Theories of Empirical Knowledge. *Midwest Studies in Philosophy*, 5 (1), 53-73. <https://doi.org/10.1111/j.1475-4975.1980.tb00396.x>
- Booth, A. R. & Peels, R. (2010). Why responsible belief is blameless belief. *Journal of Philosophy*, 107 (5), 257-265. <https://doi.org/10.5840/jphil2010107519>.
- Bradley, D. (2015). *A Critical Introduction to Formal epistemology*. Bloomsbury.
- Bradley, D. (2019). Are There Indefeasible Epistemic Rules? *Philosophers’ Imprint* 19. <http://hdl.handle.net/2027/spo.3521354.0019.003>.
- Bratman, M. (1999). *Intention, Plans, and Practical Reason*. The David Hume Series of Philosophy and Cognitive Science Reissues. CSLI Publications.
- Bright, L. & Dang, H. & Heesen, R. (2018). A Role for Judgment Aggregation in Coauthoring Scientific Papers. *Erkenntnis*, 83 (2), 231-252. <https://doi.org/10.1007/s10670-017-9887-1>.
- Broome, J. (2001). Are Intentions Reasons? And How Should We Cope With Incommensurable Values? In Morris C. & Ripstein A. (Eds.) *Practical Rationality and Preference: Essays for David Gauthier* (pp. 98–120). Cambridge University Press.
- Broome, J. (2013). *Rationality Through Reasoning*. Wiley-Blackwell.
- Brouwer, T. & Ferrario, R. & Porello, D. (2021). Hybrid collective intentionality. *Synthese* 199 (2), 3367-3403. <https://doi.org/10.1007/s11229-020-02938-z>.
- Brown, J. (2018). What is Epistemic Blame? *Noûs*, 54 (2), 389-407. <https://doi.org/10.1111/nous.12270>.
- Brown, J. (2022a). Group Belief for a Reason. *Aristotelian Society Supplementary*, 96 (1), 1-22. <https://doi.org/10.1093/arisup/akac004>.

- Brown, J. (2022b). Group Evidence. *Philosophical Issues*, 32 (1), 164-179.
<https://doi.org/10.1111/phis.12220>.
- Brown, J. (2023a). Group belief and direction of fit. *Philosophical Studies*, 180, 3161-3178.
<https://doi.org/10.1007/s11098-023-02037-y>
- Brown, J. (2023b). Lackey on group justified belief and evidence. *Asian Journal of Philosophy*, 2 (2), 1-7.
<https://doi.org/10.1007/s44204-023-00091-1>.
- Brown, J. (forthcoming). *Groups as Epistemic and Moral Agents*. Oxford University Press.
- Brun, G. (2016). Explication as a Method of Conceptual Re-engineering. *Erkenntnis*, 81 (6), 1211-1241.
- Buchak, L., & Pettit, P. (2015). Reasons and rationality: the case of group agents. In Hirose I. & Reisner A. (Eds.), *Weighing and Reasoning: Themes from the Philosophy of John Broome* (pp. 207–231). Oxford University Press.
- Cappelen, H. (2020). Conceptual Engineering: The Master Argument. In Burgess A., Cappelen H., & Plunkett D. (Eds.), *Conceptual Engineering and Conceptual Ethics* (pp. 132-151). Oxford University Press.
<https://doi.org/10.1093/oso/9780198801856.003.0007>.
- Carnap, R. (1962 [1950]). *Logical Foundations of Probability*. 2nd edition. University of Chicago Press, Routledge & Kegan Paul.
- Carroll, L. (1895). What the Tortoise Said to Achilles. *Mind*, 4, 278–280.
<https://doi.org/10.1093/mind/IV.14.278>.
- Carter, A. & Kallestrup, J. (2019). Varieties of Cognitive Integration. *Noûs*, 54 (4), 867-890.
<https://doi.org/10.1111/nous.12288>.
- Carter, J. A. (2015). Group Knowledge and Epistemic Defeat. *Ergo: An Open Access Journal of Philosophy*, 2. <https://doi.org/10.3998/ergo.12405314.0002.028>.
- Chakravartty, A. (2018). Realism, Antirealism, Epistemic Stances, and Voluntarism. In Juha Saatsi (ed.), *The Routledge Handbook of Scientific Realism* (pp. 225-236). Routledge.
- Chalmers, D. (2011). Verbal Disputes. *The Philosophical Review*, 120 (4), 515–566.
<https://doi.org/10.1215/00318108-1334478>.
- Chang, R. (1997). Introduction. In R. Chang (Ed.) *Incommensurability, Incomparability, and Practical Reason* (pp. 1-34). Harvard University Press.
- Chang, R. (2002). The Possibility of Parity. *Ethics*, 112 (4), 659-688.
<https://doi.org/10.1086/339673>.
- Chang, R. (2005). Parity, Interval Value, and Choice. *Ethics* 115 (2), 331-350.
<https://doi.org/10.1086/426307>.
- Chisholm, R. (1973). *The Problem of the Criterion*. Marquette University Press.
- Chisholm, R. (1977). *Theory of Knowledge*. Prentice Hall.
- Chisholm, R. (1982). *The Foundations of Knowing*. University of Minnesota Press.
- Christensen, D. (2007). Epistemology of Disagreement: The Good News. *Philosophical Review*, 116, 187–217. <https://doi.org/10.1215/00318108-2006-035>.

- Christensen, D. (2010). Higher Order Evidence. *Philosophy and Phenomenological Research*, 81 (1), 185-215. <https://doi.org/10.1111/j.1933-1592.2010.00366.x>.
- Christensen, D. (2022). Embracing Epistemic Dilemmas. In McCain, K., Stapleford, S. & Steup, M. (Eds.) *Epistemic Dilemmas: New Arguments, New Angles* (pp. 165-181). Routledge.
- Chudnoff, E. (2014). *The Rational Roles of Intuition*. In Anthony Booth & Darrell Rowbottom (eds.), *Intuitions* (pp. 9-35). Oxford University Press.
- Clark, A. & Chalmers, D. (1998). The Extended Mind. *Analysis*, 58 (1), 7-19. <https://doi.org/10.1093/analys/58.1.7>.
- Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.
- Cloos, C. (2015). Responsibilist Evidentialism. *Philosophical Studies*, 172 (11), 2999-3016. <https://doi.org/10.1007/s11098-015-0454-9>.
- Cohen, J. (1989). Belief and Acceptance. *Mind*, 98 (391), 367-389. <https://doi.org/10.1093/mind/XCVIII.391.367>
- Cohen, J. (1995). *An Essay on Belief and Acceptance*. Clarendon Press.
- Cohen, S. (2013). A Defense of the (Almost) Equal Weight View. In Christensen D. & Lackey J. (Eds.) *The Epistemology of Disagreement: New Essays* (pp. 98-117). Oxford University Press.
- Cohen, S., Lehrer, K. (1983). Justification, Truth, and Knowledge. *Synthese*, 55, 191-207. <https://doi.org/10.1007/BF00485068>.
- Comesana, J., & Kantin, H. (2010). Is evidence knowledge? *Philosophy and Phenomenological Research*, 80 (2), 447-454. <https://doi.org/10.1111/j.1933-1592.2010.00323.x>.
- Comesana, J., & McGrath, M. (2014). *Having false reasons*. In Littlejohn C., Turri J. (Eds.), *Epistemic Norms* (pp. 59-79). Oxford University Press.
- Conee, E. & Feldman, R. (2004). *Evidentialism: Essays in Epistemology*. Oxford University Press.
- Conee, E. (2020). Evidential support and best explanations. *Philosophical Issues*, 30 (1), 71-85. <https://doi.org/10.1111/phis.12172>.
- Conee, E., & Feldman, R. (2008). Evidence. In Q. Smith (Ed.), *Epistemology: New Essays* (pp. 83-104). Oxford University Press.
- Dahlback, M. (forthcoming). Infinitely Permissive. *Erkenntnis*.
- David, M. (2013). Truth as the Primary Epistemic Goal: A Working Hypothesis. In Matthias Steup, John Turri & Ernest Sosa (eds.), *Contemporary Debates in Epistemology* (pp. 363-377). Wiley-Blackwell..
- Davidson, D. (1963). Actions, Reasons, and Causes. *Journal of Philosophy*, 60 (23), 685-700. <https://doi.org/10.2307/2023177>.
- Davidson, D. (1984). *Inquiries Into Truth And Interpretation*. Oxford University Press.
- Decker, J. (2012). Disagreement, Evidence, and Agnosticism. *Synthese*, 187, 753-783. <https://doi.org/10.1007/s11229-011-0010-x>.

- Degenaar, M. and Lokhorst G. (2021) Molyneux's Problem. *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/win2021/entries/molyneux-problem>.
- Dennett, D. (1981). *The Intentional Stance*. MIT Press.
- Dennett, D. (1991). Real patterns. *Journal of Philosophy*, 88 (1), 27-51. <https://doi.org/10.2307/2027085>.
- Dougherty, T. (2011). In Defense of Propositionalism about Evidence. In Dougherty T. (ed.) *Evidentialism and Its Discontents* (pp. 226–235). Oxford University Press.
- Drake, J. (2017). Doxastic Permissiveness and the Promise of Truth. *Synthese* 194, 4897-4912. <https://doi.org/10.1007/s11229-016-1176-z>.
- Dunn, J. (2021). Reliable group belief. *Synthese* 198, 5653–5677. <https://doi.org/10.1007/s11229-018-02075-8>.
- Elga, A. (2007). Reflection and disagreement. *Noûs*, 41 (3), 478–502. <https://doi.org/10.1111/j.1468-0068.2007.00656.x>.
- Elga, A. (2008). *Lucky to be rational*. Unpublished manuscript. Online: <http://www.princeton.edu/~adame/papers/bellingham-lucky.pdf>.
- Elga, A. (2010). Subjective Probabilities Should be Sharp. *Philosophers' Imprint*, 10, 1-11. <http://hdl.handle.net/2027/spo.3521354.0010.005>.
- Elgin, C. & Van Cleve, J. (2013). Can Belief Be Justified Through Coherence Alone? In Matthias Steup, John Turri & Ernest Sosa (eds.), *Contemporary Debates in Epistemology* (pp. 244-273). Blackwell.
- Evans, G. (1982). *The Varieties of Reference*. Oxford University Press.
- Faria, D. (2021). Group Belief: Defending a Minimal Version of Summativism. *Epistemology and Philosophy of Science*, 58 (1), 82–93. <https://doi.org/10.5840/eps202158111>.
- Feldman, R. (2000). The Ethics of Belief. *Philosophy and Phenomenological Research*, 60 (3), 667–695. <https://doi.org/10.2307/2653823>.
- Feldman, R. (2002). *Epistemology*. PrenticeHall.
- Feldman, R. (2005). Respecting the Evidence. *Philosophical Perspectives* 19, 95–119. <https://doi.org/10.1111/j.1520-8583.2005.00055.x>.
- Feldman, R. (2007). Reasonable Religious Disagreements. In L. Antony (ed.) *Philosophers Without Gods* (pp. 194– 214). Oxford University Press.
- Fraser, R. (forthcoming). *The Will in Belief*. Oxford Studies in Epistemology.
- Fratantonio, G. (forthcoming). Evidentialism. In Kurt Sylvan (ed.) *The Blackwell Companion to Epistemology*, 3rd edition. Wiley-Blackwell.
- Fricker, M. (2010). Can There Be Institutional Virtues? In Szabó Gendler T. & Hawthorne J. (Eds.), *Oxford Studies in Epistemology* 3 (pp. 235-252). Oxford University Press.
- Friedman, J. (2013). Suspended Judgment. *Philosophical Studies*, 162, 165–181. <https://doi.org/10.1007/s11098-011-9753-y>.

- Friedman, M. (1974). Explanation and scientific understanding. *Journal of Philosophy*, 71 (1), 5-19. <https://doi.org/10.2307/2024924>.
- Fumerton, R. (1995). *Metaepistemology and Skepticism*. Rowman & Littlefield.
- Fumerton, R. (2004). Epistemic probability. *Philosophical Issues*, 14, 149-164. <https://doi.org/10.1111/j.1533-6077.2004.00025.x>.
- Fumerton, R. A. (2008). The Problem of the Criterion. In Greco J. (ed.) *Oxford Handbook of Skepticism* (pp. 34-52), Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780195183214.003.0003>.
- Gettier, Edmund L., (1963) Is Justified True Belief Knowledge?. *Analysis*, 23(6), 121-123. <https://doi.org/10.1093/analys/23.6.121>.
- Gibbard, A. & Harper, W. L. (1978). Counterfactuals and Two Kinds of Expected Utility. In Harper, W.L., Stalnaker, R., Pearce, G. (eds.) *The University of Western Ontario Series in Philosophy of Science* 15. https://doi.org/10.1007/978-94-009-9117-0_8.
- Gilbert, M. (1987). Modelling Collective Belief. *Synthese*, 73 (1), 185-204. <https://doi.org/10.1007/bf00485446>.
- Gilbert, M. (1989). *On Social Facts*. Routledge.
- Gilbert, M. (2004). Collective Epistemology. *Episteme* 1, 95-107. <https://doi.org/10.3366/epi.2004.1.2.95>.
- Gilbert, M., & Pilchman, D. (2014). Belief, Acceptance, and What Happens in Groups. In Lackey J. (ed.), *Essays in Collective Epistemology* (pp. 189-214). Oxford University Press.
- Goldberg, S. (2007). Norms of Trust, De Re Trust, and the Epistemology of Testimony. In Jäger C. & Löffler W. (eds.), *Epistemology: Contexts, Values, Disagreement. Papers of the 34th International Ludwig Wittgenstein-Symposium in Kirchberg, 2011* (pp. 229-256). The Austrian Ludwig Wittgenstein Society.
- Goldberg, S. (2008). Testimonial knowledge in early childhood, revisited. *Philosophy and Phenomenological Research*, 76 (1), 1-36. <https://doi.org/10.1111/j.1933-1592.2007.00113.x>.
- Goldberg, S. (2010). *Relying on others: An essay in Epistemology*. Oxford University Press.
- Goldberg, S. (2016). On the Epistemic Significance of Evidence You Should Have Had. *Episteme*, 13 (4), 449-470. <https://doi.org/10.1017/epi.2016.24>.
- Goldberg, S. (2017). Should have known. *Synthese*, 194 (8), 2863-2894. <https://doi.org/10.1007/s11229-015-0662-z>
- Goldberg, S. (2018). *To the Best of Our Knowledge: Social Expectations and Epistemic Normativity*. Oxford University Press.
- Goldman, A. (1979). What is Justified Belief? In G. S. Pappas (ed.) *Justification and Knowledge* (pp. 1-23). Springer.
- Goldman, A. (1980). The Internalist Conception of Justification. *Midwest Studies in Philosophy* 5 (1), 27-52. <https://doi.org/10.1111/j.1475-4975.1980.tb00395.x>.
- Goldman, A. (1986). *Epistemology and Cognition*. Harvard University Press.

- Goldman, A. (1999). *Knowledge in a Social World*. Oxford University Press.
- Goldman, A. (2001). Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, 63 (1), 85-110. <https://doi.org/10.1111/j.1933-1592.2001.tb00093.x>.
- Goldman, A. (2009). Williamson on Knowledge and Evidence. In Greenough P. & Pritchard D. (eds.), *Williamson on Knowledge* (pp. 73-91). Oxford University Press..
- Goldman, A. (2010). Why Social Epistemology is Real Epistemology. In Haddock A., Millar A. & Pritchard D. (eds.) *Social Epistemology* (pp. 1-29). Oxford University Press.
- Goldman, A. (2011). Toward a Synthesis of Reliabilism and Evidentialism? Or: Evidentialism's Troubles, Reliabilism's Rescue Package. In Dougherty T. (eds.) *Evidentialism and Its Discontents* (pp. 254-280). Oxford University Press.
- Goldman, A. (2014). 'Social Process Reliabilism': Solving Justification Problems in Collective Epistemology. In Lackey J. (ed.) *Essays in Collective Epistemology* (pp. 11-41). Oxford University Press.
- Goldman, A. and Beddor, B. (2021). Reliabilist Epistemology. *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/sum2021/entries/reliabilism>.
- Goodman, N. (1955). *Fact, Fiction, and Forecast*. Harvard University Press.
- Graham, P. & Lyons, J. (2021). The Structure of Defeat: Pollock's Evidentialism, Lackey's Framework, and Prospects for Reliabilism. In Brown J. & Simion M. (eds.), *Reasons, Justification, and Defeat*. Oxford University Press.
- Grasswick, H. (2018). Feminist Social Epistemology. *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/fall2018/entries/feminist-social-epistemology/>.
- Greco, D., Hedden, B. (2016). Uniqueness and Metaepistemology. *The Journal of Philosophy*, 113 (8), 365-395. <https://doi.org/10.5840/jphil2016113825>
- Greco, J. (1990). Internalism and epistemically responsible belief. *Synthese*, 85 (2), 245 - 277. <https://doi.org/10.1007/bf00484794>.
- Greco, J. (1999). Agent reliabilism. *Philosophical Perspectives*, 13, 273-296. <https://doi.org/10.1111/0029-4624.33.s13.13>.
- Greco, J. (2005a). Common sense: A contemporary Defense. *Notre dame Philosophical Reviews* (Review of Noah Lemos). <http://ndpr.nd.edu/review.cfm?id=3081>.
- Greco, J. (2005b). Justification is not internal. In Steup M. & Sosa E. (eds.), *Contemporary Debates in Epistemology* (pp. 257-269). Blackwell.
- Griffin, J. (1986). *Well-Being: Its Meaning, Measurement and Importance*. Clarendon Press.
- Habgood-Coote, J. (2019). Group Knowledge, Questions, and the Division of Epistemic Labour. *Ergo: An Open Access Journal of Philosophy*, 6, 925-966. <https://doi.org/10.3998/ergo.12405314.0006.033>.
- Habgood-Coote, J. (2020). To the Best of Our Knowledge: Social Expectations and Epistemic Normativity By Sanford Goldberg. *Analysis*, 80 (3), 608-610.
- Hakli, R. (2011). On Dialectical Justification of Group Beliefs. In Schmid B., Sirtes D. & Weber M. (eds.), *Collective Epistemology* (pp. 119-154). Ontos.

- Harman, G. (1980). Reasoning and Evidence One Does Not Possess. *Midwest Studies in Philosophy*, 5 (1), 163-182. <https://doi.org/10.1111/j.1475-4975.1980.tb00403.x>
- Harris, K. (2021). Collective Intellectual Humility and Arrogance. *Synthese*, 199, 6967–6979. <https://doi.org/10.1007/s11229-021-03101-y>
- Hawthorne, J. & Lasonen-Aarnio, M. (2021). Not So Phenomenal! *The Philosophical Review*, 130 (1), 1–43. <https://doi.org/10.1215/00318108-8699669>.
- Hedden, B. & Dorst, K. (2022). (Almost) all evidence is higher-order evidence. *Analysis*, 82 (3), 417-425. <https://doi.org/10.1093/analys/anab081>.
- Hedden, B. (2015a). A Defense of Objectivism about Evidential Support. *Canadian Journal of Philosophy*, 45 (5), 716-743. <https://doi.org/10.1080/00455091.2015.1123038>.
- Hedden, B. (2015b). Time-Slice Rationality. *Mind*, 124 (494), 449-491. <https://doi.org/10.1093/mind/fzu181>.
- Hedden, B. (2019). Reasons, Coherence, and Group Rationality. *Philosophical and Phenomenological Research*, 99, 581-604. <https://doi.org/10.1111/phpr.12486>.
- Hempel, C. (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. The Free Press.
- Hess, K. (2018). The peculiar unity of corporate agents. In Hess K., Ignieski V. & Isaacs T. (eds.), *Collectivity: Ontology, Ethics, and Social Justice*. Rowman & Littlefield International.
- Horvath, J. & Koch, S. (2020). Experimental philosophy and the method of cases. *Philosophy Compass*, 16 (1), 1-13. <https://doi.org/10.1111/phc3.12716>.
- Huemer, M. (2001). *Skepticism and the Veil of Perception*. Rowman & Littlefield.
- Hughes, N. (2019). Dilemmic Epistemology. *Synthese*, 196 (10), 4059-4090. <https://doi.org/10.1007/s11229-017-1639-x>
- Hughes, N. (forthcoming). Epistemic Dilemmas: A Guide. In N. Hughes (Ed.) *Essays on Epistemic Dilemmas*. Oxford University Press.
- Hutchins, E. (1995a). *Cognition in the Wild*. MIT Press.
- Hutchins, E. (1995b). How a Cockpit Remembers its Speeds. *Cognitive Science*, 19 (3), 265-288. https://doi.org/10.3998/10.1207/s15516709cog1903_1.
- Ichikawa, J. & Steup, M. (2018). The Analysis of Knowledge. *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/sum2018/entries/knowledge-analysis>.
- Jackson, E. & Turnbull, M. (forthcoming). Permissivism, Underdetermination, and Evidence. In Littlejohn C. & Lasonen-Aarnio M. (eds.) *The Routledge Handbook of the Philosophy of Evidence*. Routledge.
- Jackson, E. (2021). A Defense of Intrapersonal Belief Permissivism. *Episteme*, 18 (2), 313–327. <https://doi.org/10.1017/epi.2019.19>
- James, W. (1895). Is Life Worth Living? *International Journal of Ethics*, 6 (1), 1-24. <https://doi.org/10.1086/intejethi.6.1.2375619>.
- Jeffrey, R. C. (1965). *The Logic of Decision*. McGraw-Hill.

- Kallestrup, J. (2016). Group virtue epistemology. *Synthese*, 197 (12), 5233-5251. <https://doi.org/10.1007/s11229-016-1225-7>.
- Kallestrup, J. (2022). Nonreductive Group Knowledge Revisited. *Episteme*, 1-24. <https://doi.org/10.1017/epi.2022.19>.
- Kelly, T. (2008). Evidence: Fundamental Concepts and the Phenomenal Conception. *Philosophy Compass* 3, (5), 933-955. <https://doi.org/10.1111/j.1747-9991.2008.00160.x>.
- Kelly, T. (2010). Peer Disagreement and Higher Order Evidence. In Feldman R. and Warfield M. (eds.) *Disagreement* (pp. 111-174). Oxford: Oxford University Press.
- Kelly, T. (2013). Evidence Can Be Permissive. In Steup M., Turri J., and Sosa E. (eds.), *Contemporary Debates in Epistemology*, (2nd ed., pp. 298-311). Wiley-Blackwell.
- Kelly, T. (2016). Evidence. *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/win2016/entries/evidence/>.
- Kim, J. (1994). Explanatory knowledge and metaphysical dependence. *Philosophical Issues*, 5, 51-69. <https://doi.org/10.2307/1522873>.
- Kindermann, D., & Onofri, A. (2021). An Introduction. In Borgoni, C., Kindermann, D., & Onofri, A. (Eds.) *The Fragmented Mind*. Oxford University Press.
- Kitcher, P. (1981). Explanatory Unification. *Philosophy of Science*, 48 (4), 507-531. <https://doi.org/10.1086/289019>.
- Kopec, M. & Titelbaum, M. G. (2016). The Uniqueness Thesis. *Philosophy Compass*, 11 (4), 189-200. <https://doi.org/10.1111/phc3.12318>.
- Kopec, M. (2015). A Counterexample to the Uniqueness Thesis. *Philosophia*, 43 (2), 403-409. <https://doi.org/10.1007/s11406-014-9579-x>.
- Korcz, K. (2000). The Causal-Doxastic Theory of the Basing Relation. *Canadian Journal of Philosophy*, 30 (4), 525-550. <https://doi.org/10.1080/00455091.2000.10717542>.
- Korcz, K. (2021). The Epistemic Basing Relation. *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/spr2021/entries/basing-epistemic/>.
- Kornblith, H. (1983). Justified Belief and Epistemically Responsible Action. *Philosophical Review*, 92 (1), 33-48. <https://doi.org/10.2307/2184520>.
- Kuhn, T. (1977). *The Essential Tension: Selected Studies in Scientific Tradition and Change*. University of Chicago Press.
- Kusch, M. (2017). Epistemic relativism, scepticism, pluralism. *Synthese*, 194 (12), 4687-4703. <https://doi.org/10.1007/s11229-016-1041-0>.
- Kvanvig, J. L. (2013). Truth is Not the Primary Epistemic Goal. In Steup M. & Turri J. (eds.), *Contemporary Debates in Epistemology* (pp. 285-295). Blackwell.
- Lackey, J. (1999). Testimonial knowledge and transmission. *The Philosophical Quarterly*, 49, 471-90. <https://doi.org/10.1111/1467-9213.00154>.
- Lackey, J. (2005). Testimony and the Infant/Child Objection. *Philosophical Studies*, 126 (2), 163-190. <https://doi.org/10.1007/s11098-004-7798-x>.

- Lackey, J. (2006). Knowing from testimony. *Philosophy Compass*, 1 (5), 432–448.
<https://doi.org/10.1111/j.1747-9991.2006.00035.x>.
- Lackey, J. (2008). *Learning from Words: Testimony as a Source of Knowledge*. Oxford University Press.
- Lackey, J. (2010). What Should We Do When We Disagree? In Szabo Gendler T. and Hawthorne J. (eds.) *Oxford Studies in Epistemology* (pp. 274-293). Oxford University Press.
- Lackey, J. (2016). What is Justified Group Belief? *Philosophical Review*, 125 (3), 341–396.
<https://doi.org/10.1215/00318108-3516946>.
- Lackey, J. (2021). *The Epistemology of Groups*. Oxford University Press.
- Lamb, E. (2012). 5 Sigma What's that. *Scientific American*. 17 July, 2012
<https://www.scientificamerican.com/blog/observations/five-sigmawhats-that/>
- Lasonen-Aarnio, M. (2010). Unreasonable Knowledge. *Philosophical Perspectives*, 24 (1), 1-21.
<https://doi.org/10.1111/j.1520-8583.2010.00183.x>.
- Lasonen-Aarnio, M. (2014). Higher-Order Evidence and the Limits of Defeat. *Philosophy and Phenomenological Research*, 88 (2), 314-345. <https://doi.org/10.1111/phpr.12090>.
- Lasonen-Aarnio, M. (2019). Higher-Order Defeat and Evincibility. In M. Skipper, and Asbjørn Steglich-Petersen (eds.), *Higher-Order Evidence: New Essays* (pp. 144-172).
<https://doi.org/10.1093/oso/9780198829775.003.0007>.
- Lasonen-Aarnio, M. (2020). Enkrasia or Evidentialism? Learning to love mismatch. *Philosophical Studies* 177, 597–632. <https://doi.org/10.1007/s11098-018-1196-2>.
- Lehrer, K. & Paxson, T. (1969). Knowledge: Undefeated justified true belief. *Journal of Philosophy*, 66 (8), 225-237. <https://doi.org/10.2307/2024435>.
- Lehrer, K. (1974). *Knowledge*. Oxford University Press.
- Lehrer, K. (1990). *Theory of Knowledge*, Routledge.
- Leibniz, G. W. (1985 [1710]). *Theodicy: Essays on the Goodness of God, the Freedom of Man, and the Origin of Evil*. Open Court.
- Leite, A. (2008). Believing One's Reasons are Good. *Synthese*, 161, 419–441.
<https://doi.org/10.1007/s11229-006-9093-1>.
- Lemos, N. (2004). *Common sense: A contemporary Defense*. Cambridge University Press.
- Levi, I. (1974). On Indeterminate Probabilities. *Journal of Philosophy*, 71 (13), 391–418.
<https://doi.org/10.2307/2025161>.
- Levi, I. (1986). *Hard Choices: Decision Making Under Unresolved Conflict*. Cambridge University Press.
- Lewis, D. (1971). Immodest inductive methods. *Philosophy of Science*, 38 (1), 54-63.
<https://doi.org/10.1086/288339>.
- Lewis, D. (1974). Radical interpretation. *Synthese* 23, 331-344. <https://doi.org/10.1007/bf00484599>.
- Lewis, D. (1986). Causal explanation. In *Philosophical Papers Vol. Ii.* (pp. 214-240). Oxford University Press.

- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74 (4), 549 – 567. <https://doi.org/10.1080/00048409612347521>.
- Lewis, D. (1999). *Papers in Metaphysics and Epistemology*. Cambridge University Press.
- Lipton, P. (2004). *Inference to the best explanation*. Routledge/Taylor and Francis Group.
- List, C. & Pettit, P. (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press.
- Littlejohn, C. (2013). No Evidence is False. *Acta Analytica*, 28 (2), 145–59. <https://doi.org/10.1007/s12136-012-0167-z>.
- Longino, H. (1996). Cognitive and Non-Cognitive Values in Science: Rethinking the Dichotomy. In Lynn Hankinson N. & Nelson J. (eds.), *Feminism, Science, and the Philosophy of Science* (pp. 39-58). Kluwer Academic Publishers.
- Longino, H. (2001). *The Fate of Knowledge*. Princeton University Press.
- Longino, H. (2022). What's Social about Social Epistemology? *Journal of Philosophy*, 119, (4), 169-195. <https://doi.org/10.5840/jphil2022119413>.
- Lord, E. & Sylvan K. (2019). Prime-time (for the basing relation). In Carter A. and Bondy P. (eds.), *Well-Founded Belief: New Essays on the Basing Relation* (pp. 141-174), Oxford University Press,.
- Ludwig, K. (2014). Proxy Agency in Collective Action. *Noûs*, 48 (1), 75-105. <https://doi.org/10.1111/nous.12013>.
- Lutz, M. (2020). Background beliefs and plausibility thresholds: defending explanationist evidentialism. *Synthese*, 197 (6), 2631-2647. <https://doi.org/10.1007/s11229-018-1849-x>.
- Marxen, C. (2021a). Closing the Case on Self-Fulfilling Beliefs. *Australasian Journal of Philosophy*, 101 (1), 51-64. <https://doi.org/00048402.2021.1967416>.
- Marxen, C. (2021b). On the Epistemic Rationality and Significance of Self-Fulfilling Beliefs. *Synthese*, 199, 4243–4260. <https://doi.org/10.1007/s11229-020-02977-6>.
- Matheson, J. (2011). The Case for Rational Uniqueness. *Logos and Episteme*, 2, 359–373.
- Mathiesen, K. (2011). Can Groups Be Epistemic Agents?. In Schmid H. B., Sirtes D., and Weber M. (Eds.), *Collective Epistemology* (pp. 23–45). Frankfurt: Ontos Verlag.
- McCain, K. (2014). *Evidentialism and Epistemic Justification*. Routledge.
- McHugh, C. & Way, J. (2016). Against the Taking Condition. *Philosophical Issues*, 26 (1), 314-331. [10.1111/phils.12074](https://doi.org/10.1111/phils.12074).
- McKenna, M. (2012). *Conversation & Responsibility*. Oxford University Press.
- Meacham, C. (2013). Impermissive Bayesianism. *Erkenntnis*, 79 (6), 1185-1217. <https://doi.org/10.1007/s10670-013-9547-z>.
- Meeker, K. (2014) Justification and the Social Nature of Knowledge. *Philosophy and Phenomenological Research*, 69, 156–172. <https://doi.org/10.1111/j.1933-1592.2004.tb00388.x>.
- Miragoli, M. & Simion, M. (2020). Gender, Race & Group Disagreement. In Adam Carter & Fernando Broncano-Berrocal (eds.), *The Epistemology of Group Disagreement* (pp. 125-138). Routledge.

- Mitova, V. (2014). Truthy psychologism about Evidence. *Philosophical Studies*, 172 (4), 1105-1126
<https://doi.org/10.1007/s11098-014-0339-3>.
- Milburn, J. (2023). Unpossessed Evidence: What's the Problem? *Topoi*, 42, 107-120.
<https://doi.org/10.1007/s11245-022-09867-1>.
- Mokyr, J. (2002). *The Gifts of Athena*. Princeton University Press.
- Moss, S. (2014). Credal Dilemmas. *Nous*, 48 (3), 665-683. <https://doi.org/10.1111/nous.12073>.
- Moss, S. (2015). Time-Slice Epistemology and Action under Indeterminacy. *Oxford Studies in Epistemology* 5, 172-194.
- Muralidharan, A. (2021). Permissivism and Self-Fulfilling Propositions. *Ratio*, 34 (3), 217-226.
<https://doi.org/10.1111/rati.12301>.
- Nagel, T. (1987). The Fragmentation of Value. In Gowans C. (Ed.) *Moral Dilemmas* (pp. 74–87). Oxford University Press.
- Neta, R. (2002). S Knows That P. *Nous*, 36 (4), 663–681 . <https://doi.org/663-681>. 10.1111/1468-0068.00406.
- Neta, R. (2019). The Basing Relation. *Philosophical Review*, 128 (2), 179–217. <https://doi.org/10.1215/00318108-7374945>.
- Nguyen, C. T. (2020). Echo Chambers and Epistemic Bubbles. *Episteme*, 17 (2),141-161.
<https://doi.org/10.1017/epi.2018.32>.
- Nottelmann N. (2021). Against Normative Defeat. *Mind*, 130 (520), 1183–1204,
<https://doi.org/10.1093/mind/fzaa079>
- Palermos, O., & Tollefsen, D. (2018). Group Know-How. In Carter A., et. al (eds), *Socially Extended Epistemology* (pp. 112-131). <https://doi.org/10.1093/oso/9780198801764.003.0007>
- Parfit, D. (1987). *Reasons and Persons*. Oxford University Press.
- Peels, R. (2015). Believing at Will is Possible. *Australasian Journal of Philosophy*, 93 (3), 1-18.
<https://doi.org/10.1080/00048402.2014.974631>.
- Peet, A. & Pitcovski, E. (2017). Lost in transmission: Testimonial justification and practical reason. *Analysis*, 77 (2), 336-344.
- Plantinga, A. (1993). *Warrant and Proper Function*. Oxford University Press.
- Plantinga, A. (2000). *Warranted Christian Belief*. Oxford University Press.
- Podgorski, A. (2016). Dynamic permissivism. *Philosophical Studies*,173 (7), 1923-1939.
<https://doi.org/10.1007/s11098-015-0585-z>.
- Pollock, J. (1986). *Contemporary Theories of Knowledge*. Rowman & Littlefield.
- Pollock, J., Cruz J. (1999). *Contemporary Theories of Knowledge*. Rowman & Littlefield.
- Poston, T. (2014). *Reason and explanation*. Palgrave Macmillan.
- Priest, G. (2022). Rational Dilemmas and Their Place in the Bigger Picture. In McCain, K., Stapleford, S. & Steup, M. (eds.) *Epistemic Dilemmas: New Arguments, New Angles* (pp. 7-20). Routledge.

- Pryor, J. (2004). What's Wrong with Moore's Argument?. *Philosophical Issues*, 14 (1), 349–378. <https://doi.org/10.1111/j.1533-6077.2004.00034.x>.
- Quanbeck, Z. & Worsnip, Alex (forthcoming). A Permissivist Alternative to Encroachment. *Philosophers' Imprint*.
- Quinton, A. (1976). Social objects. *Proceedings of the Aristotelian Society*, 76 (1), 1-27. <https://doi.org/10.1093/aristotelian/76.1.1>.
- Raleigh T. (2017). Another Argument Against Uniqueness. *Philosophical Quarterly*, 67, 327-346. <https://doi.org/10.1093/pq/pqw058>.
- Raleigh, T. (2015). An Argument for Permissivism from Safespots. In van der Hoek W. et al. (eds.) *Logic, Rationality, and Interaction* (pp. 308–315). Springer-Verlag.
- Raleigh, T. (2023). The Argument from Small Improvement is a Red Herring. *Inquiry: An Interdisciplinary Journal of Philosophy*, 66, 1-16. <https://doi.org/10.1080/0020174X.2023.2205441>.
- Raz, J. (1986). *The Morality of Freedom*. Clarendon Press.
- Raz, J. (1997). Incommensurability and Agency. In Chang R. (ed.), *Incommensurability, Incomparability, and Practical Reason* (pp. 110–128). Harvard University Press.
- Reisner, A. (2013). Leaps of Knowledge. In Chan T. (ed.), *The Aim of Belief* (pp. 167-183). Oxford University Press.
- Rescher, N. (1967). Choice without Preference. A Study of the History and of the Logic of the Problem of 'Buridan's Ass'. *Journal of Symbolic Logic*, 32 (4), 545-546. <https://doi.org/10.2307/2270241>.
- Ross, R. (2021). Alleged Counterexamples to Uniqueness. *Logos and Episteme*, 12 (2), 203-213.
- Russell, B. (1911) Knowledge by Acquaintance and Knowledge by Description. *Proceedings of the Aristotelian Society*, 11, 108-128. <https://doi.org/10.1093/aristotelian/11.1.108>.
- Salmon, W. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton University Press.
- Sankey, H. (2010). Witchcraft, Relativism and the Problem of the Criterion. *Erkenntnis*, 72 (1), 1-16. <https://doi.org/10.1007/s10670-009-9193-7>.
- Sartre J. P. (1966 [1948]). *Existentialism and Humanism*. Translation by Mairet P.. Methuen & Co. LTD.
- Schmitt, F. (1994). The Justification of Group Beliefs. In Schmitt F. (ed.), *Socializing Epistemology: The Social Dimensions of Knowledge* (pp. 257–87). Rowman & Littlefield.
- Schoenfield, M. (2014a). Decision Making in the Face of Parity. *Philosophical Perspectives*, 28 (1), 263-277. <https://doi.org/10.1111/phpe.12044>.
- Schoenfield, M. (2014b). Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief. *Noûs*, 48 (1), 193–218. <https://doi.org/10.1111/nous.12006>.
- Schoenfield, M. (2023). Meditations on Beliefs Formed Arbitrarily. In Szabó Gendler T., Hawthorne J. & Chung J. (eds.), *Oxford Studies in Epistemology* 7, (pp. 278-305).
- Schroeder, M. (2008). Having reasons. *Philosophical Studies*, 139 (1), 57–71. <https://doi.org/10.1007/s11098-007-9102-3>.

- Schwengerer, L. (2021). Defending Joint Acceptance Accounts of Justification. *Episteme*, First View, 1-20. <https://doi.org/10.1017/epi.2020.55>.
- Sen, A. (2017). *Collective Choice and Social Welfare: An Expanded Edition*. Harvard University Press.
- Shafer-Landau, R. (2003). *Moral Realism*. Oxford University Press.
- Shope, R. (1983). *The Analysis of Knowing: A Decade of Research*. Princeton University Press.
- Silins, N. (2005). Deception and Evidence. *Philosophical Perspectives*, 19 (1), 375-404. <https://doi.org/10.1111/j.1520-8583.2005.00066.x>.
- Silva, P. (2015a). Does Doxastic Justification Have a Basing Requirement?. *Australasian Journal of Philosophy*, 93 (2), 371-387. <https://doi.org/10.1080/00048402.2014.962553>.
- Silva, P. (2015b). On Doxastic Justification and Properly Basing One's Beliefs. *Erkenntnis*, 80, 945-955. <https://doi.org/10.1007/s10670-014-9690-1>.
- Silva, P. (2019). Justified Group Belief is Evidentially Responsible Group Belief. *Episteme*, 16 (3), 262-281. <https://doi.org/10.1017/epi.2018.5>.
- Silva, P. (2023). Self-Fulfilling Beliefs: A Defence. *Australasian Journal of Philosophy*, 101 (4), 1012-1018. <https://doi.org/10.1080/00048402.2023.2172594>
- Simion, M. (2021). *Shifty Speech and Independent Thought: Epistemic Normativity in Context*. Oxford University Press.
- Simion, M. (2022). Skepticism about Epistemic Dilemmas. In McCain, K., Stapleford, S. & Steup, M. (Eds.) *Epistemic Dilemmas: New Arguments, New Angles* (pp.109-121). Routledge.
- Simpson, R. M. (2017) Permissivism and the Arbitrariness Objection. *Episteme*, 14 (4), 519-538. <https://doi.org/10.1017/epi.2016.35>.
- Smith, J. J. (2020). Unacknowledged Permissivism. *Pacific Philosophical Quarterly*, 101 (1), 158-183. <https://doi.org/10.1111/papq.12299>.
- Sosa, E. (2015). *Judgment & Agency*. Oxford University Press.
- Srinivasan, A. (2015). The Archimedean Urge. *Philosophical Perspectives*, 29 (1), 325-362. <https://doi.org/10.1111/phpe.12068>.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Clarendon Press.
- Stanley, J. and Williamson, T. (2001). Knowing How. *Journal of Philosophy*, 98 (8), 411-444. <https://doi.org/10.2307/2678403>.
- Steele, K., H. O. Stefánsson (2020). Decision Theory. *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/win2020/entries/decision-theory/>.
- Sudduth, M. (2008). Defeaters in Epistemology. *The Internet Encyclopedia of Philosophy*. <https://iep.utm.edu/defeaters-in-epistemology/>.
- Swain, M. (1979). Justification and the Basis of Belief. In Pappas G. (ed.) *Justification and Knowledge*. Philosophical Studies Series in Philosophy, 17. Springer.
- Sylvan, K. (2016). Epistemic Reasons II: Basing. *Philosophy Compass*, 11 (7), 377-389. <https://doi.org/10.1111/phc3.12329>.

- Sylvan, K. (forthcoming). Responsibility within Reason. In Kelp C. & Greco J. (eds.), *Virtue-Theoretic Epistemology: New Methods and Approaches*. Cambridge University Press.
- Talbert, M. (2023). Moral Responsibility. *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/fall2023/entries/moral-responsibility/>.
- Titelbaum, M. (2010). Not Enough There There: Evidence, Reasons, and Language Independence. *Philosophical Perspectives*, 24, 477–528. <https://doi.org/10.1111/j.1520-8583.2010.00201.x>.
- Tollefsen, D. (2002a). Challenging epistemic individualism. *Protosociology*, 16, 86–120. <https://doi.org/10.5840/protosociology20021622>.
- Tollefsen, D. (2002b). Organizations as True Believers. *Journal of Social Philosophy*, 33 (3), 395–410. <https://doi.org/10.1111/0047-2786.00149>.
- Tollefsen, D. (2007). Group testimony. *Social Epistemology*, 21 (3), 299–311. <https://doi.org/10.1080/02691720701674163>.
- Tollefsen, D. (2015). *Groups as Agents*. Polity Press.
- Tucker, C. (2012). Movin' on up: higher-level requirements and inferential justification. *Philosophical Studies*, 157 (3), 323-340. 10.1007/s11098-010-9650-9
- Tuomela, R. (2004). Group Knowledge Analyzed. *Episteme*, 1, 109–127. <https://doi.org/10.3366/epi.2004.1.2.109>.
- Tuomela, R. (2011). An Account of Group Knowledge. In Schmid H. B., Sirtes D., and Weber M. (Eds.), *Collective Epistemology* (pp. 75–117). Ontos Verlag.
- Turri, J. (2009). The Ontology of Epistemic Reasons. *Noûs*, 43 (3), 490-512. <https://doi.org/10.1111/j.1468-0068.2009.00715.x>.
- Turri, J. (2010a). Epistemic Supervenience. In Dancy, J., Sosa, E., and Steup M., (Eds.) *A Companion to Epistemology* (pp. 340–343). Blackwell.
- Turri, J. (2010b). On the Relationship between Propositional and Doxastic Justification. *Philosophy and Phenomenological Research*, 80 (2), 312-326. <https://doi.org/10.1111/j.1933-1592.2010.00331.x>.
- Turri, J. (2011). Believing For a Reason. *Erkenntnis*, 74, 383–97. <https://doi.org/10.1007/s10670-011-9271-5>.
- Ullmann-Margalit, E. & Morgenbesser, S. (1977). Picking and Choosing. *Social Research: An International Quarterly*, 44 (4), 757-785.
- Ullmann-Margalit, E. (2006). Big Decisions: Opting, Converting, Drifting. *Royal Institute of Philosophy Supplement* 58, 157-172. <https://doi.org/10.1017/s1358246106058085>.
- Valaris, M. (2017). What reasoning might be. *Synthese*, 194, 2007–2024. <https://doi.org/10.1007/s11229-016-1034-z>
- van Fraassen, B. (1990). Figures in a Probability Landscape. In Dunn J.M. & Gupta A. (eds.) *Truth or Consequences: Essays in Honor of Nuel Belnap*, (pp. 345–56). Kluwer & Springer.
- van Fraassen, B. (2002). *The Empirical Stance*. Yale University Press.
- Väyrynen, P. (2013). Grounding and Normative Explanation. *Aristotelian Society Supplementary*, 87 (1), 155-178. <https://doi.org/10.1111/j.1467-8349.2013.00224.x>.

- Velleman, D. (1989). *Practical Reflection*. Princeton University Press.
- Volpe, G. (2017). Propositional and Doxastic Justification: Their Relationship and a Questionable Supervenience Claim. In Brozek, B. et al. (eds.), *Supervenience and Normativity: Law and Philosophy Library*, 120 (pp. 25–47). https://doi.org/10.1007/978-3-319-61046-7_2.
- Webber, J. (2007). *The Existentialism of Jean-Paul Sartre*. Routledge.
- Wedgwood, R. (2002). The Aim of Belief. *Noûs*, 36, 267–297. <https://doi.org/10.1111/1468-0068.36.s16.10>.
- White, R. (2005). Epistemic Permissiveness. *Philosophical Perspectives*, 19 (1), 445–459. <https://doi.org/10.1111/j.1520-8583.2005.00069.x>.
- White, R. (2013). Evidence Cannot Be Permissive. In Steup M., Turri J., and E. Sosa (Eds.) *Contemporary Debates in Epistemology*, (2nd ed., pp. 312–323). Wiley-Blackwell.
- Wigmore, J. (1904). The History of the Hearsay Rule. *Harvard Law Review*, 17, 437–458. <https://doi.org/10.2307/1323425.#>.
- Willholt, T. (2009). Bias and Values in Scientific Research. *Studies in History and Philosophy of Science*, 40 (1), 92–101. <https://doi.org/10.1016/j.shpsa.2008.12.005>.
- Williams, R. (2014). Decision-Making Under Indeterminacy. *Philosophers' Imprint*, 14 (4), 1–34. <http://hdl.handle.net/2027/spo.3521354.0014.004>.
- Williams, R. (2016). Angst, Indeterminacy and Conflicting Values. *Ratio*, 29 (4), 412–433. <https://doi.org/10.1111/rati.12141>.
- Williams, R. (2020). *The Metaphysics of Representation*. Oxford University Press.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford University Press.
- Wright, C. (2014). Comment on Paul Boghossian ‘What is inference’. *Philosophical Studies*, 169, 27–37. <https://doi.org/10.1007/s11098-012-9892-9>
- Ye, R. (2019). The Arbitrariness Objection Against Permissivism. *Episteme*, 18 (4), 654–673. <https://doi.org/10.1017/epi.2019.42>.