



A University of Sussex PhD thesis

Available online via Sussex Research Online:

<http://sro.sussex.ac.uk/>

This thesis is protected by copyright which belongs to the author.

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

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

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

Please visit Sussex Research Online for more information and further details

The Uniqueness Thesis: A Hybrid Approach

Tamaz Tokhadze

Submitted for the degree of PhD

University of Sussex

January 2022

Declaration

I hereby declare that this thesis has not been and will not be, submitted in whole or in part to another University for the award of any other degree.

Signature:

Tamaz Tokhadze

University of Sussex

Submitted for the degree of PhD

The Uniqueness Thesis:
A Hybrid Approach

Abstract

Uniqueness is the view that a body of evidence justifies a unique doxastic attitude toward any given proposition. Contemporary defences and criticisms of Uniqueness are generally indifferent to whether we formulate the view in terms of the coarse-grained attitude of belief or the fine-grained attitude of credence. This dissertation proposes and defends a hybrid view I call *Hybrid Impermissivism*, which combines the following two theses: *Moderate Uniqueness* and *Credal Permissivism*. Moderate Uniqueness says that no evidence could justify both believing a proposition and its negation. However, on Moderate Uniqueness, evidence could justify both believing and suspending judgement on a proposition (hence the adjective “Moderate”). And Credal Permissivism says that more than one credal attitude could be justified on the evidence. Hybrid Impermissivism is developed into a precise theory by using normative bridge principles on how rational belief and credence ought to interact.

Hybrid Impermissivism is an attractive position in several respects: as I argue, it captures plausible motivating ideas behind permissive and impermissive epistemologies and avoids some important problems associated with them. Still, I show that Hybrid

Impermissivism faces a special problem, the *diachronic coordination problem*, which has to do with coordinating an agent's beliefs and credences over time. A significant part of this dissertation is dedicated to solving this new problem. I state a formal result I call the *coordination theorem*, which identifies plausible constraints on rational belief and credence under which Moderate Uniqueness and Credal Permissivism remain diachronically consistent.

My overall conclusion is that Hybrid Impermissivism is a coherent, plausible, and conciliatory position and provides a correct characterisation of the requirements of evidence on doxastic attitudes.

Contents

List of Figures	IX
Preface	X
Acknowledgements	XIV
1 Introduction	1
1.1 Setting the Stage: Key Distinctions and Concepts	5
1.1.1 Varieties of Uniqueness	5
1.1.2 Rationality	11
1.1.3 Evidence	14
1.2 Uniqueness: The State of the Debate	16
1.2.1 The Arbitrariness Argument	20
1.2.2 The Evidential Support Argument	22
1.3 What Lies Ahead	24
1.4 Appendix: Does Uniqueness Require Brilliant Insights?	29
1.4.1 Introduction	29
1.4.2 The Evidence-Change Response	31
1.4.3 To What Extent does Uniqueness Require Rational Insights?	36
2 Motivating Hybrid Impermissivism	38
2.1 The Intuitiveness of Hybrid Impermissivism	39
2.2 An Argument for Moderate Uniqueness	45

2.2.1	White’s Argument for Moderate Uniqueness	46
2.2.2	A New Argument for Moderate Uniqueness	49
2.2.3	Relational Objectivity	57
2.2.4	Conclusion	64
2.3	Summary	65
3	Belief, Credence, and their Relationship	66
3.1	The Setting	66
3.1.1	The Possible Worlds Account of Propositions	68
3.2	Belief	72
3.3	Credence	80
3.4	Worries with the Bayesian Model of Degree of Belief	87
3.4.1	Probabilism and Doxastic Precision	87
3.4.2	Conditionalisation and Conceptual Innovations	90
3.5	The Relationship between Belief and Credence	94
3.6	Summary	98
4	Hybrid Theories and the Coordination Problem	99
4.1	The Lockean Hybrid Impermissivism	100
4.2	The Lockean Hybrid Impermissivism and the Uniqueness Debate	103
4.2.1	The Epistemic Standards Objection Reconsidered	103
4.2.2	The Arbitrariness Argument Reconsidered	106
4.3	The Coordination Problem for Hybrid Impermissivism	110

4.3.1	The Coordination Problem in a Formal Setting	115
4.4	Two Initial Responses to the Coordination Problem	118
4.5	Summary	122
5	Hybrid Impermissivism and the Stability Theory	124
5.1	Introduction	125
5.2	The Stability Theory	129
5.2.1	Three Consequences of the Stability Theory	133
5.2.2	From Monotonicity to the Stability Theory	140
5.2.3	The Problem of Context-Sensitivity	142
5.3	The Humean Solution to the Coordination Problem	151
5.3.1	The Coordination Theorem	154
5.3.2	Order Uniqueness, Credal Uniqueness, and Ordinal Uniqueness	160
5.3.3	The Coordination Problem and Jeffrey Conditionalisation	165
5.4	Summary	171
5.5	Appendix to Chapter 5	172
5.5.1	The Coordination Problem for Strong Hybrid Impermissivism	172
5.5.2	A Coordination Theorem within AGM Belief Revision Theory	177
5.5.3	Order Uniqueness and Relational Objectivity	185
6	Hybrid Impermissivism without the Stability Theory	191
6.1	Is the Stability Theory Overly Demanding?	192
6.1.1	A Counterexample	195

6.2	An Alternative to Monotonicity	199
6.3	An Alternative to the Stability Theory: The Dominant Core Theory	203
6.4	The Dominant Core Theory vs The Stability Theory	210
6.5	Concluding Discussion	213
7	Belief Within Incomplete Doxastic Rankings	217
7.1	Introduction: Bayesianism and Likelihoodism	220
7.2	Two Tenets of Likelihoodism	224
7.3	The Problem of Belief Guidance	230
7.3.1	An Argument Against Likelihood-based Belief Guidance	235
7.4	The Case for Likelihood-based Belief Guidance	239
7.5	Summary	247
8	Conclusion	249
	Bibliography	253

List of Figures

Figure 3. 1: Logical relationships between three basic propositions	72
Figure 3. 2: Possible worlds over $W = \{w_1, w_2, w_3, w_4\}$	76
Figure 3. 3: $B_W = \{w_1, w_3\}$ for $W = \{w_1, w_2, w_3, w_4\}$	77
Figure 3. 4: Geometric representation of all coherent credence functions	83
Figure 5. 1: Nested spheres representing the nested stable propositions	131
Figure 5. 2: Piggybacking of belief revision on credal revision	152
Figure 5. 3: Geometric representation of order equivalent probability distributions	163

Preface

This dissertation is a study of how strongly a rational agent's beliefs are constrained by her evidence. Its main question is: does a rational agent's evidence fully determine what she ought to believe?

A positive answer to this question is called *Uniqueness*: the view that there is always one, unique rational response to any body of evidence. The dissertation is about Uniqueness and whether it provides a correct characterisation of the evidential constraints on a rational individual's beliefs.

The view put forward in the dissertation is a *hybrid view*, where Uniqueness, on the whole, is true if we think about belief in a coarse-grained, qualitative manner; but Uniqueness is false if we think about belief in a fine-grained manner, in terms of numerical degree of belief or credence.

More fully, I defend the view I call *Hybrid Impermissivism* that combines the following two theses: *Moderate Uniqueness* and *Credal Permissivism*. Moderate Uniqueness says that no evidence could justify both believing a proposition and its negation. However, on Moderate Uniqueness, evidence could justify both believing and suspending judgement on a proposition (hence the adjective "Moderate"). And Credal Permissivism says that more than one credal attitude could be justified on the evidence.

The guiding idea behind Hybrid Impermissivism is that the appeal and plausibility of Uniqueness and Permissivism are sensitive to which doxastic attitude-type we are focusing on, belief or credence. Permissivism seems almost obviously correct when we think about belief in an extremely fine-grained manner, in terms of (numerical) credence. For instance, it seems wildly implausible that your current evidence determines a unique, hyper-precise credence in any proposition you may consider. However, if, instead of credence, we focus on

(qualitative) belief, which includes just three attitudes: belief, disbelief, and suspension of judgment, then Uniqueness, especially Moderate Uniqueness, becomes significantly more plausible. So, a hybrid approach to the Uniqueness question aims to reconcile plausible aspects of Uniqueness and Permissivism into a single, precise, and coherent view.

To make Hybrid Impermissivism a precise position, we need to specify what rational beliefs and credences are and, most importantly, how they ought to interact. I will use the label (impermissivist) *hybrid theory* to refer to a position that endorses Hybrid Impermissivism together with the norms of how rational belief and credence ought to interact. My overall goal is to provide a plausible, formally precise, systematic hybrid account of doxastic rationality that reconciles important intuitions, norms and approaches to doxastic rationality defended by the opposing parties to the Uniqueness debate.

By a hybrid account of doxastic rationality, I only mean a hybrid account of the requirements of evidence on belief. As discussed in Section 1.1.3, defending a detailed, nuanced position about Uniqueness does not require us to answer all questions about doxastic rationality. To use Jon Williamson's (2010, 5) phrase, Uniqueness is "... but one piece in the jigsaw puzzle of normativity."

In this dissertation, I'll propose and motivate not one but two different hybrid theories, which differ solely in their conceptions of how rational belief and credence ought to interact. Each of these hybrid theories has its strengths and achieves the main goal of this dissertation by reconciling Moderate Uniqueness with Credal Permissivism in a precise, plausible way. Or so I will argue. The reason for discussing these two hybrid theories is to show that the proposed hybrid approach to Uniqueness is *flexible* (as it can be developed in more than one way) and does not depend on a singular, contentious way of understanding the relationship

between belief and credence. I will also show that, the choice between these two hybrid theories will not be consequential for this dissertation.

Besides the hybrid approach to the Uniqueness debate, another defining feature of this dissertation concerns its methodology. Almost all major arguments I develop use formal-mathematical methods and results, mostly involving basic set theory and the probability theory. All the formal tools that I use are fully explained when introduced. No advanced prior knowledge is needed to understand and appreciate their usefulness (and scope). Formal tools are used for two main purposes: first, to provide precise, *idealised models* or *explications* of vague and imprecise concepts such as: proposition, confirmation, belief, degree of belief and evidence; and second, to prove theorems about explicated concepts.

Why use formal methods for this project? Many topics relevant to the Uniqueness debate are technical, involving probability theory or require technical tools for an appropriate analysis. For instance, the topics of this dissertation require us to analyse an agent's entire *system* or *set* of beliefs instead of individual beliefs. As the question of how belief and credence ought to interact is not about individual beliefs and credences but sets of beliefs and credences. And formal methods are particularly suited for talking about a system or set of beliefs.

I conclude this preface by sketching the contents of each of the chapters and commenting about the overall structure of this dissertation to help orient the reader. A detailed summary of the contents of the chapters and their function in this dissertation is given in section 1.3.

Chapter 1 introduces the Uniqueness debate and discusses its central topics and arguments. Chapter 2 motivates each of the two theses of Hybrid Impermissivism: Moderate Uniqueness and Credal Permissivism, but does not discuss in any detail how these theses may

hang together. Chapter 3 prepares the ground for articulating a precise hybrid theory of doxastic rationality. I give a detailed discussion of the standard logical account of rational belief and the standard Bayesian account of degrees of belief. I also provide a preliminary discussion of how these accounts may interact with each other. Chapters 4, 5 and 6 comprise the core of the dissertation. In Chapter 4, I'll state the first precise hybrid theory of doxastic rationality and present a general problem that all hybrid theories face: *the diachronic coordination problem*. Roughly, the problem is in coordinating an agent's beliefs and credences over time in a way that preserves both Moderate Uniqueness and Credal Permissivism. The subsequent two chapters, Chapters 5 and 6, put forward two distinct hybrid theories that avoid the diachronic coordination problem. In Chapter 7, I discuss rational belief within a setting where the available evidence does not justify a complete probability distribution over a set of propositions. I provide an impermissivist account of how one should think about rational belief and evidential support within such a setting. Chapter 8 concludes by summarising the main claims of the dissertation.

Acknowledgements

I consider myself incredibly fortunate to have been able to research the topics of this dissertation for the past three years, free from most material hurdles. Being a PhD researcher at Sussex and the luxury of living in the UK (for the first half of my PhD, before the global outbreak of the COVID-19 pandemic) provided me with many opportunities and enabled me to attend various conferences, workshops and a summer school that were instrumental in teaching myself some advanced philosophy, formal tools, and probability theory I used in this dissertation. I would have written a much lesser dissertation without this freedom and opportunities.

Many people have helped and influenced this project in one way or another. Alex Davies introduced me to the Uniqueness debate in his introductory epistemology class during my MA studies at the University of Tartu. Alex kindly agreed to be my MA supervisor – even though he was already quite busy with teaching, supervising, and researching – and gave me more time and instructions than I could have asked for. I owe him a huge debt of gratitude for all his help during and after my studies at Tartu. Thank you so much, Alex. From my time at Tartu, I also want to thank Francesco Orsi, Uku Tooming, and Eve Kitsik for their teaching and fostering philosophical rigour in me.

At Sussex, first, I want to thank Tony Booth, one of my two supervisors, who saw promise in my research proposal and played a crucial role in my becoming a researcher at Sussex. Tony's supervision, written feedbacks, and conversations have been an important influence from the beginning of this project. Without Tony's suggestions, I would not have worked on some of the topics in this dissertation at all. Specifically, the last chapter of the dissertation – which relates the central topics of this dissertation to a debate within the

philosophy of statistics – would not have been written if not for his reading suggestions and help.

I owe lots of gratitude to Corine Besson, my primary supervisor, for her many thoughtful, detailed feedback, conversations and encouragement. Corine has read and extensively commented on almost every part of this dissertation and several other papers not included here. She made this dissertation much more focused, rigorous and articulate than it would have otherwise been. Thanks, Corine, for being a great supervisor.

I am extremely indebted to the permanent participants of the reading group on Hannes Leitgeb’s “The Stability of Belief”: Simon Mcgregor, Sophio Machavariani, Richard Lohse. This reading group – organised during the height of COVID restrictions within Europe – was one of the highlights of my PhD and a great source of inspiration for my research. I owe them a great deal – Thanks, guys.

I want to thank Julien Murzi, who, very kindly, helped me to do a one-semester exchange at the University of Salzburg. Even though the exchange took place when the University of Salzburg was closed due to the pandemic, it was still very useful to me. Living in an outskirt of Salzburg during a very strict lockdown in Europe was as comfortable living in these times as one could hope for.

I’m also very grateful to every person who read or discussed various parts or topics of this dissertation, including Prasanta S. Bandyopadhyay, Brett Topey, Tamar Tskhadadze, Elisabeth Jackson, and Kevin Kelly. Thanks to all the audiences of various conferences and other talks for helpful comments. Some of the sections of the dissertation have improved considerably from the comments and criticisms of various anonymous referees from my

journal articles.¹ Many thanks also to Beka Jalagania, a fellow PhD candidate, for friendship and many talks regarding being a PhD student and an academic.

I also owe a lot to people who made free online resources on philosophical topics which are not taught in many universities. I'm very grateful to Stephan Hartmann and Hannes Leitgeb, who made a free Coursera course, "Introduction to Mathematical Philosophy". This course was an invaluable source of information for me and introduced a whole other domain of philosophy that I was completely unaware of. Thanks to Michael Strevens, whose "Notes on Bayesian Confirmation Theory" taught me the basic Bayesian philosophy of science. Thanks to Mike Titelbaum for making available his textbook-in-progress "Fundamentals of Bayesian Epistemology" that also taught me a lot. And thanks to Branden Fitelson for making this and many other resources available through his website.

I want to thank my family and especially my mother, Tamuna. At first, my choice to switch from studying economics to philosophy was not taken positively by my family (and understandably so), except by my mother. Without her unconditional support, I would not have been able to study at Tartu, and, hence, will not be finishing this dissertation at Sussex.

Finally, my principal gratitude is to Nina Abesadze, my best friend and partner. She was an attentive reader of this dissertation and a keen conversationalist. Our discussions helped shape many aspects of this dissertation. I dedicate this dissertation to her.

¹ The articles that are adapted for this dissertation are as follows: Chapters 4 and 5 are adapted and extended from "Hybrid Impermissivism and the Diachronic Coordination Problem" in *Philosophical Topics* (forthcoming b). Section 2.2 is adapted from "Extreme Permissivism Revisited" in *The European Journal of Analytic Philosophy* (forthcoming c). Chapter 7 is adapted and extended from "Likelihoodism and Guidance for Belief" in *Journal for General Philosophy of Science* (forthcoming d).

1 Introduction

Sometimes, our evidence unequivocally determines what we ought to believe. For instance, there is overwhelming evidence that smoking increases the risk of lung cancer. Hence, an informed individual is rationally *required* to believe that smoking increases the risk of lung cancer. But could a body of evidence rationally *permit* a belief without making this belief rationally required?

The question is especially pertinent in cases where it is unclear how to best interpret the available evidence. For instance, scientists have put forward various hypotheses on what caused Neanderthal extinction, and there is little agreement about which of these hypotheses, if any, is likely to be true. Could two scientists rationally believe incompatible theories about Neanderthal extinction on the very same evidence? Or could they rationally invest different levels of confidence in the competing hypotheses?

According to the Uniqueness Thesis (Uniqueness, for short), the answer to these questions is “No”. The view is that there is always one, unique rational or justified response to any body of evidence.

What does it mean to call a belief rational or justified?² Within this dissertation, calling a belief “rational” means that the belief is not ruled out by any norms or principles of rationality (more on this in Section 1.1.1).³ Given this characterisation of rational belief, Uniqueness can also be stated as follows:

² The normative terms “rational” and “justified” are used interchangeably throughout the dissertation. While in other contexts, the two notions are sometimes distinguished, such a distinction would serve no useful purpose in this dissertation.

³ This view of what it means to call a belief rational follows Gibbard (1990).

For any evidence E and propositions H , the norms of rationality do not permit more than one doxastic attitude towards H .

Many philosophers have defended norms of rationality that are inconsistent with Uniqueness. For instance, consider the following version of the pragmatist norm called *doxastic inertia* (Levi 1998) or *doxastic conservatism* (Harman 1999):

Any belief of an agent is rational by default until the agent acquires strong evidence for thinking that her belief is false.

If this norm of doxastic inertia is true, then Uniqueness is false (see Brueckner and Bundy 2012, Section 1). To illustrate this, suppose that an agent's who believes H finds out that her total evidence about H is *balanced*: that is, her evidence equally supports H and $\neg H$. Now, the balanced evidence about H does not seem to constitute strong evidence against H (or against $\neg H$). So, on doxastic inertia, the agent is rationally permitted to continue believing H . And, in general, if two equally informed agents have opposing doxastic attitudes towards a proposition, these attitudes could be equally rational if their evidence is balanced, given the norm of doxastic inertia.

By contrast, some have defended norms that favour Uniqueness (or a principle very similar to it). For instance, consider the following highly plausible *Pyrrhonian* norm (Feldman 2007):

If an agent's evidence E equally supports a proposition H and its negation $\neg H$, then the agent should suspend judgment on H .

This Pyrrhonian norm seems to be a necessary condition for the truth of Uniqueness. This is because, in the situations where an agent's evidence is balanced with respect to H , the Pyrrhonian norm prohibits the agent from adopting any other doxastic attitude towards H

except the suspension of judgement. Hence, epistemologies that accept Uniqueness should also accept this Pyrrhonian norm.

Epistemologies that accept norms implying some substantive version of Uniqueness are called *impermissive*. And epistemologies that accept norms incompatible with Uniqueness are called *permissive*. Hence, the negation of Uniqueness is called *Permissivism*: the view that more than one attitude towards a proposition can be rational on a body of evidence.

This dissertation is about Uniqueness; and hence, about Permissivism. The main question is:

Is Uniqueness the right way to think about the strength of evidential requirements on doxastic states?

I develop a particular approach to this question, where the answers that I shall offer are sensitive to how we specify the term “doxastic state”. If we focus on qualitative (categorical) belief, then my answer is “Yes, to a larger extent”. More precisely, I will defend a version of Uniqueness that I call *Moderate Uniqueness*: the view that no evidence could justify both believing a proposition and its negation. However, I will reject the stronger thesis that I call *Extreme Uniqueness*: the view that any evidence justifies a unique belief-attitude towards a proposition – either belief, disbelief, or suspension of judgement.⁴

⁴ Why reject Extreme Uniqueness? As we will see, there are two overall reasons for doing so. First, as we discuss in Section 4.2.1, often, what seems to make the difference between believing a proposition and suspending judgement is not the agent’s evidence, but her epistemic goals and interests, such as how much value she attaches to believing what is true. And second, there is a more technical reason for rejecting Extreme Uniqueness, as it does not cohere with Credal Permissivism; I discuss this in one of the appendices to Chapter 5 (Section 5.5.1).

By contrast, if instead of belief, we focus on degree of belief or credence, my answer to the dissertation's main question is "No". More precisely, I'll endorse the view called *Credal Permissivism* that a body of evidence could equally justify more than one credal attitude towards a proposition.

Hence, this dissertation defends a *hybrid view* where Moderate Uniqueness is combined with Credal Permissivism. I call this proposed view *Hybrid Impermissivism*.

Why consider Hybrid Impermissivism as an impermissivist view? As we shall see, even without endorsing Extreme Uniqueness, the proposed hybrid theory is closer to the impermissivist end of the spectrum rather than to the permissivist end, because it constrains significantly what a rational agent ought to believe and to what degree, given her evidence.

Now, Hybrid Impermissivism has an important initial appeal. Uniqueness about numerical credence, unlike Uniqueness about ordinary belief, seems to be "an extremely strong and unobvious claim" (Kelly 2010, 121). At least at first blush, it is unlikely that two equally informed people are always rationally required to adopt the exact same credence, say a credence of 0.62394, towards a given proposition; but they may still be rationally required to adopt the same belief attitude towards the proposition. For instance, as there is overwhelming evidence supporting the anthropogenic global warming hypothesis, the only rational option is to believe this hypothesis. However, there seems to be no unique credence that the evidence justifies towards the anthropogenic global warming hypothesis.

Hence, the initial appeal of Uniqueness is sensitive to which attitude type we are focusing on – belief or credence. While some have explicitly recognised this (e.g., Kelly 2014; Jackson 2019b), so far, no one has provided an in-depth discussion or defence of any hybrid view within the Uniqueness debate.

The dissertation aims to fill this gap and provides the case for Hybrid Impermissivism. To make Hybrid Impermissivism a precise position, we need to specify what rational beliefs and credences are and, most importantly, how they ought to interact. I will use the label (impermissivist) *hybrid theory* to refer to a position that endorses Hybrid Impermissivism together with the norms of how rational belief and credence ought to interact.

The goal of this dissertation is to articulate a coherent, systematic hybrid theory of doxastic rationality. Motivating and building such a theory is a complex task. The major steps in which I pursue this task are explained at the end of this introductory chapter, Section 1.3, where I discuss the subject matter and results of each of the subsequent chapters in detail.

The rest of this introduction will prepare the ground by discussing the main concepts and distinctions within the Uniqueness debate (Section 1.1) and the state of the debate (1.2).

1.1 Setting the Stage: Key Distinctions and Concepts

1.1.1 Varieties of Uniqueness

Over the years, Uniqueness has been defined in a variety of ways. Some differences in these definitions are relatively minor, and the arguments of the dissertation won't be sensitive to them.⁵ For the purposes of this dissertation, all significant varieties of Uniqueness correspond

⁵ From the outset, it would be useful to set aside one type of variation among the definitions of Uniqueness. Some (e.g., Feldman 2007, Matheson 2011) define Uniqueness as the thesis that evidence justifies *at most* one unique attitude towards a proposition. So, on this definition, all evidence is such that it either justifies a unique standard attitude towards a proposition, or *no attitude* at all. Some (e.g., Matheson *ibid.*) think that this "no attitude" option makes Uniqueness more plausible. However, in this dissertation, we do not have any use for this fourth type of qualitative doxastic attitude.

Of course, I grant that if an agent has never considered proposition H and H is logically and semantically independent of any other proposition that she has ever entertained, then we can say that the agent

Footnote continued on the next page.

to just two general distinctions. This section discusses these distinctions and explains which versions of Uniqueness we will be focusing on.

The first important distinction is between *Interpersonal* and *Intrapersonal* Uniqueness. Intrapersonal Uniqueness is the thesis that, if we pick any rational agent from the set of all rational agents, her evidence will fix a unique doxastic attitude towards any given proposition. More precisely:

Intrapersonal Uniqueness: For all rational agents, *S*, evidence, *E*, and propositions, *H*, if *E* is *S*'s total body of evidence, then there is a unique doxastic attitude that *S* can rationally adopt towards *H*.

So, Intrapersonal Uniqueness excludes the possibility of an agent's evidence permitting her to adopt more than one rational doxastic attitude towards a proposition. But Intrapersonal Uniqueness leaves open the possibility where the same evidence justifies one doxastic attitude for an agent and some other doxastic attitude for another agent. To exclude such a possibility, we need to endorse Interpersonal Uniqueness:

has no doxastic attitude towards *H*. For instance, most people do not *have* any doxastic attitude towards the proposition that "quark-gluon plasma can be created by colliding gold ions at nearly the speed of light". But *having* no attitude towards a proposition you never considered is different from *adopting* no attitude towards a proposition you consider. The former is about a simple lack of belief, while the latter is about adopting a certain postulated attitude, namely, "no attitude" towards a proposition. The phenomenon of lack of belief is real but irrelevant to the Uniqueness debate. The reality or theoretical usefulness of the "no attitude" attitude is controversial, and I won't make use of this notion in this dissertation (I discuss a closely related topic in Section 3.2). Either way, whether we define Uniqueness as "at most one attitude" view or simply "one attitude" view won't have any significance for the main arguments of this dissertation.

Interpersonal Uniqueness: For all evidence, E , propositions, H , there is a unique doxastic attitude such that, for all rational agents, S , if S 's total body of evidence is E , then there is a unique doxastic attitude that S can rationally adopt towards H .

As we can see, the difference between Intrapersonal and Interpersonal Uniqueness is the position of the quantifier phrase "there is a unique doxastic attitude". In Intrapersonal Uniqueness, this quantifier phrase has a narrow scope, and in Interpersonal Uniqueness – a wide scope. So, for Interpersonal Uniqueness, the existential quantifier comes first, and hence the universal quantifier "for all agents" falls under the scope of the existential quantifier. And for Intrapersonal Uniqueness, the quantifier positions are reversed.

Some of the earlier arguments for Uniqueness were ambiguous between the interpersonal and intrapersonal readings of the thesis (we will look at such an argument in section 1.2.1). So, it is important to keep the distinction between Intrapersonal and Interpersonal Uniqueness in mind when articulating and evaluating arguments for or against Uniqueness.

This dissertation is solely concerned with Interpersonal Uniqueness, the logically stronger thesis. Whenever I endorse a version of Uniqueness, I always endorse some interpersonal version of Uniqueness. Hence all the subsequent definitions of Uniqueness will be variations of Interpersonal Uniqueness. Because we are only interested in interpersonal versions of Uniqueness, we will mostly omit the adjectives "interpersonal" and "intrapersonal" when talking about Uniqueness.

The parallel distinction should also be made between Intrapersonal and Interpersonal Permissivism. As we will see (Section 1.2.1), most permissivists endorse Interpersonal

Permissivism but explicitly reject the stronger thesis, Intrapersonal Permissivism.⁶ I will mostly keep this distinction implicit. Although, it is worth noting here that all hybrid theories I develop are consistent with a moderate version of Intrapersonal Permissivism which permits the same agent to either believe a proposition or suspend judgement (more on this in Section 4.2.1).

Another central distinction that will preoccupy us to a far greater extent is the distinction between different versions of Uniqueness for different types of doxastic attitudes. It is useful to categorise doxastic attitudes into two types:

- (1) Categorical or nongraded (coarse-grained) doxastic attitudes.
- (2) Graded (fine-grained) doxastic attitudes.⁷

It is commonly assumed that there is only one categorical doxastic attitude referred to as “full belief” or simply “belief”. Belief is a categorical attitude of holding a proposition to be true; it is a categorical attitude because one either holds the proposition to be true or does not. By contrast, graded doxastic attitudes come in different strengths, where one can have different degrees of confidence in the truth of a proposition. I will use the term “belief” to denote the categorical, coarse-grained doxastic attitude. However, sometimes, I will use the same term,

⁶ While Jackson (2019a), like everybody else, thinks that Intrapersonal Permissivism is a more demanding view than Interpersonal Permissivism, she is unsure (see *ibid.*, Footnote 8) whether the former is logically stronger than the latter. For our purposes, it is both convenient and inconsequential to assume that Intrapersonal Permissivism is a logically stronger thesis.

⁷ One may argue that there is only one fundamental doxastic attitude. For instance, one may argue that nongraded doxastic attitudes are reducible to graded attitudes (or vice versa) or that nongraded attitudes do not exist. We shall discuss this issue in detail in Section 3.5. At this point, if the reader is sceptical about the existence of different doxastic attitude types, she can think of this distinction (between nongraded and graded doxastic attitudes) as the distinction between different *concepts* of doxastic attitudes (or beliefs).

“belief”, as an umbrella term to talk about doxastic attitudes in general (both graded and nongraded). In most cases, context will disambiguate in which sense the term is used. When appropriate, I’ll use the modifiers like “qualitative”, “categorical”, or “all-or-nothing” to emphasise that we are speaking about belief in the traditional, categorical sense.

Much of traditional epistemology is concerned with categorical belief, which comes in just three types: one can either believe a proposition, disbelieve it, or suspend judgment. So, with respect to traditional belief-attitude, we have the following version of Uniqueness that I will provisionally call *Belief Uniqueness*:

Belief Uniqueness: Given any body of evidence, E , and proposition, H , there is a unique belief-attitude (either belief, disbelief, or suspension) that any agent should adopt towards H .⁸

By contrast, some of the most significant work in contemporary epistemology is centred on graded belief or degree of belief. The best-known formal model of degree of belief is the Bayesian model, according to which rational degrees of belief are numerically graded and have a structure of mathematical probabilities. I will call the Bayesian conception of degree of belief – credence. The Bayesian model recognises infinitely many credal attitudes towards a proposition, where each credal attitude is represented by a real number in the unit interval. The credence of 0 represents the minimal confidence, while 1 represents the maximal confidence. So, for instance, I’m almost certain that the next US presidential election will be won by either a Democrat or by a Republican candidate; hence, my credence

⁸ As I’ve already explained, I distinguish two versions of Uniqueness about categorical belief: Extreme Uniqueness (equivalent to what I’ve just called Belief Uniqueness); and Moderate Uniqueness, a weaker thesis that does not exclude the possibility where the evidence permits both believing a proposition and suspending judgement. For the sake of simplicity, we do not draw the distinction between the two theses here.

in that proposition is nearly 1. By contrast, my credence in the proposition that a third-party candidate will win is almost 0. And, in general, an agent's credence can range anywhere from 0 to 1.

So, with respect to Bayesian credal attitudes, we have the following version of Uniqueness that I call *Credal Uniqueness*:

Credal Uniqueness: Given any body of evidence, E , and proposition, H , there is a unique credence that any agent should have toward H .

All the currently discussed versions of Uniqueness are varieties of either Belief Uniqueness or Credal Uniqueness. The negation of Belief Uniqueness is called "Belief Permissivism", and the negation of Credal Uniqueness, "Credal Permissivism". Belief and Credal Permissivism make existential claims: according to these theses, some body of evidence is such that it justifies more than one belief/credence in a proposition.

Belief Uniqueness and Credal Uniqueness do not imply each other; the same is true about Belief Permissivism and Credal Permissivism.⁹ For instance, assume that Belief Uniqueness is true: that is, evidence E always justifies a unique belief-attitude towards a proposition, H . But the same evidence E can *permit* a range of different credences towards H . As Jackson (2019b, 2782) correctly points out: "The evidence could allow one to believe H and have a credence of 0.8 or to believe H and have a credence of 0.9, but not allow for withholding belief or belief that not- H ."¹⁰ Therefore, an argument for Belief Uniqueness, if successful, may not establish Credal Uniqueness.

⁹ See Jackson (2019b, Section 3) for another discussion on the relationship between Credal and Belief Uniqueness/Permissivism.

¹⁰ In all quoted passages, the notation is adapted for uniformity of reading.

The converse is also the case. Even if we have good reasons for thinking that Credal Uniqueness is true, these reasons may not establish Belief Uniqueness. For instance, suppose we have good reasons to think that an agent's evidence uniquely determines her credences. Still, we may think that whether an agent should outright believe a proposition depends on some non-evidential factors, such as the agent's epistemic goals and interests. Hence, it is possible to argue for Credal Uniqueness without also arguing for Belief Uniqueness.

To wrap up this section: both Uniqueness and Permissivism come in (i) interpersonal and intrapersonal versions and (ii) coarse-grained and fine-grained versions. Regarding (i): in this dissertation, we are solely concerned with logically stronger, interpersonal versions of Uniqueness. This means that whenever we discuss an argument for Uniqueness, this argument needs to establish some interpersonal version of Uniqueness. Regarding (ii): we will discuss both coarse- and fine-grained versions of Uniqueness and Permissivism in detail.

In the next section, we shall discuss two concepts at the heart of the Uniqueness debate: rationality and evidence.

1.1.2 Rationality

As is often the case within epistemology, the term "rational" is used to evaluate an agent or her belief(s).

By "rational agent", I mean that the agent does not violate any norms of rationality: i.e., logical, probabilistic, or evidential norms that we assume in a particular context of the discussion. Similarly, by "rational belief", I mean that the belief is not ruled out by norms of rationality. For instance, we generally assume the probabilistic norm called *Probabilism* according to which rational degrees of belief satisfy the standard axioms of probability (as discussed in Section 3.3). So, a rational agent won't have degrees of belief that violate any axioms of probability.

It is widely recognised now – after a series of seminal papers by Kahneman and Tversky (e.g., 1971, 1974) – that humans are bad probabilistic reasoners and very often violate the axioms of probability. Hence, on the assumed definition of rational agents, humans are not rational in general. However, humans are rational in a narrower sense, as we are *capable* of avoiding probabilistic and deductive mistakes in particular reasoning contexts, when thinking about concrete, well-delineated problems or topics (more on this in Section 3.2).

I will call the norms like Probabilism and the familiar logical norms about belief (like Deductive Cogency, discussed in Section 3.2) *coherence norms* (or, sometimes, coherence principles/requirements). Coherence norms concern how an agent's attitudes fit or cohere with each other. For instance, on Probabilism, an agent's credence in a proposition H and its negation $\neg H$ should sum to unity. Probabilism does not demand any particular credence in either H or $\neg H$; it only constrains how the agent's credences in H and $\neg H$ should fit together. I'll discuss the coherence norms on rational belief and credence in detail in Sections 3.2 and 3.3, respectively.

Besides coherence norms, we will also be concerned with some evidential or epistemic norms about how an agent's doxastic attitudes fit with the agent's evidence. For instance, the *Pyrrhonian* norm that we've discussed in Section 1 – If an agent's evidence E equally supports a proposition H and its negation $\neg H$, then the agent should suspend judgment on H – is an evidential/epistemic norm, because it concerns what an agent ought to believe given her evidence.

So, the norms of rationality will include both coherence and evidential norms on belief. And we will be concerned with the agents and beliefs that do not violate any of these norms of rationality.

This setup does not prejudge the issue either in favour or against Uniqueness. Here is why. Whether Uniqueness is true depends on the norms of rationality and not on whether the agents in question are ideally rational. For instance, it is well-known that if Probabilism is the only norm on rational credence, then two agents can have widely different credences in response to the same body of evidence (more on this in section 3.3). So, the relevant question concerning Uniqueness is whether Probabilism is the only norm on rational credence. And the answer to this question is by no way determined by focusing on ideal agents.

While rational agents satisfy the norms of rationality, I do not assume that rational agents have other special intellectual abilities. I do not assume that rational agents believe only truths. An agent can be perfectly rational but believe some falsehoods: for instance, the agent may form the belief that it won't rain in the next few hours based on strong evidence (say, the evidence of clear blue sky). But, as the strong evidence for a proposition is not a guarantee that the proposition is true, this perfectly rational agent may end up believing a falsehood. So, while the agent may not violate any norms of rationality, her beliefs may still be false.

I also do not presuppose that rational agents are required to articulate or discover new, brilliant hypotheses that best explain their evidence. For instance, I do not assume that rationality required Darwin and Wallace to discover the theory of natural selection. Instead, I'm concerned with what an agent is rationally permitted to believe, given all the relevant evidence she has, including which set of competing propositions or hypotheses are available to her. So, the norms of rationality, as I understand them, don't require an agent to add anything to the already available evidence. The norms of rationality only concern what the agent ought to believe, relative to her evidence and set of competing hypotheses. This does not mean that there is no sense in which rationality requires us to seek or enrich our evidence.

But the kind of rationality we are concerned with is solely about what an agent ought to believe on the available evidence; and not about how or when an agent should seek new evidence.

Douven (2009) and Jackson (2019a) have argued that the view that rationality does not require brilliant insights already contradicts Uniqueness. I disagree. I will consider and respond to their arguments in detail in the appendix to this chapter.

I will make one further qualification before turning to the notion of evidence. As I've already noted in the introduction, I use the terms "rational belief" and "justified belief" interchangeably. Normally, epistemologists distinguish two types of justification: *propositional* and *doxastic*. If a proposition, H , is justified on evidence, E , – in the sense of propositional justification – then there is a justification for believing H on E . By contrast, doxastic justification concerns whether an agent's belief is held appropriately; that is, whether the agent came to have the belief by being responsive to her evidence and not by a fluke.

In this dissertation, nothing will turn on this distinction between propositional and doxastic justification. If a proposition H is propositionally justified on evidence E , and if an ideal agent believes H on evidence E , then we will always assume that her belief is doxastically justified.

1.1.3 Evidence

Following Jon Williamson (2010), I assume that an agent's (total) evidence consists of those propositions that the agent takes for granted in her context of reasoning. This includes the agent's relevant background information, topic-specific beliefs, and a set of competing propositions that the agent considers in her context of reasoning (more on this in the appendix to this chapter).

What the agent takes for granted may change from reasoning context to reasoning context, depending on which question(s) or problem(s) the agent is concerned with. An example from Williamson (ibid., 5) nicely illustrates why this may be so:

A medical consultant, for example, is rational to take for granted much of her education and medical training as well as uncontroversial studies in the medical literature and observations of the patient's symptoms, given her purpose of treating a patient. On the other hand, the purposes of a philosophical sceptic preclude her from taking much for granted at all. Therefore, if the medical consultant by day studies philosophy by night, her evidence base may change radically.

Whether an agent is rational in taking certain propositions for granted is an important question within epistemology. But this question goes beyond the Uniqueness debate. As I understand it, the central question about Uniqueness is a *relational* question:

Given that an agent's evidence is E , does E fix one unique doxastic attitude towards H for the agent?

Hence, as far as Uniqueness is concerned, we are interested in the relationship between the set of propositions assumed to represent the agent's evidence and the agent's beliefs. Whether this set of propositions correctly identifies the agent's evidence or whether the agent is rational in taking these propositions for granted are important but separate questions from the Uniqueness debate.

It is in this sense that a theory about the requirements of evidence on belief contributes but a piece in the jigsaw puzzle of epistemic normativity. There are many other important questions within epistemology involving the concepts of evidence and rationality, the answers to which are not determined by whether Uniqueness is true.

1.2 Uniqueness: The State of the Debate

The term “Uniqueness” has been coined relatively recently by Richard Feldman (2007).¹¹ But, it would be a mistake to see the Uniqueness debate as a recent one. William James’s seminal “The Will to Believe” (1897) deals with many central issues surrounding Uniqueness. And, to this day, James’s essay contains one of the strongest criticisms of Uniqueness (a Jamesian permissivist view is discussed in Section 4.2.1).

The decades-old and still very active debate between the so-called subjective and objective Bayesians also relates to many central themes relevant to the Uniqueness debate. According to subjective Bayesianism (which is discussed in numerous sections of this dissertation, most fully in Section 3.3), the only norm on rational credence is Probabilism: the norm that rational credences ought to satisfy the axioms of probability. By contrast, on objective Bayesianism (which is most fully discussed in Chapter 7), besides Probabilism, there are important evidential norms (like *Calibration* and *Equivocation*; see Williamson 2010) that, in many cases (though not always), fully determine what credences the agent ought to have given her evidence. Trivially, whether subjective Bayesianism is true and Probabilism is the only norm on rational credence has significant implications about the Uniqueness debate (more on this in Sections 2.2.2 and 2.2.3).¹²

¹¹ Though, Feldman’s manuscript had been in circulation a couple of years before its publication. This is why (White 2005) is the first *published* paper that uses the term “Uniqueness”.

¹² There is also an important debate in the philosophy of science about the underdetermination of theory by evidence that closely relates to the Uniqueness debate. As Jackson and Turnbull (Forthcoming) discuss, Permissivism is a version of an underdetermination thesis: Given Permissivism, some bodies of evidence can underdetermine what doxastic attitude an agent should take towards a proposition. Though, as Jackson and Turnbull (*ibid.*, 7) have correctly pointed out, by endorsing Permissivism, one is not committed to the stronger view that evidence together with various *non-evidential factors* also underdetermine an agent’s doxastic

Footnote continued on the next page.

So, while Uniqueness connects to some important earlier debates within philosophy, it has been only 15 years or so since epistemologists have paid exclusive, in-depth attention to Uniqueness and various issues surrounding it. This recent interest in Uniqueness was triggered by the debate about the epistemic significance of disagreement between epistemic peers: i.e., disagreement between equally informed, intelligent, and competent individuals. Nowadays, most articles discuss Uniqueness independently from the epistemology of disagreement.¹³ This is also what we do in this dissertation. This said, I consider the application and implications of the developed hybrid views to the epistemology of disagreement and more general issues concerning higher-order evidence as work to be done.

attitudes. As I discuss in numerous sections, including in Sections 1.2.1 and 4.2.1, most permissivists think that evidence, together with agent-relative, subjective factors, such as the agent's epistemic standards, fully determine what the agent ought to believe. So, Permissivism should not be equated with a strong version of the underdetermination thesis.

¹³ Feldman (2007) has argued that Uniqueness entails a *conciliatory* view about peer disagreement, or Conciliationism: the view that the evidence of peer disagreement rationally requires an agent to revise her original doxastic attitude regarding the disputed proposition closer to that of her epistemic peer. For instance, according to a popular version of Conciliationism, called the *Equal Weight View*, if you believe a proposition H and you learn that your epistemic peer believes $\neg H$, then you are rationally required to suspend judgement on H . Because Feldman thought that Uniqueness was highly plausible, he argued that Conciliationism was also highly plausible. By contrast, Kelly (2010) turned Feldman's modus ponens into a modus tollens and argued against Conciliationism because he (ibid., 121) thought that Uniqueness is "an extremely strong and unobvious claim". Since Feldman's and Kelly's articles, many authors have called the connection between Uniqueness and Conciliationism into question, including Lee (2013), Levinstein (2017), Titelbaum and Kopec (2019), and Weisberg (2020). So, now, there is an overall agreement within the literature that Conciliationism and Uniqueness do not imply one another, at least in a straightforward sense that Feldman and Kelly had thought.

To understand the current state of the debate about Uniqueness, we will discuss and categorise important arguments in favour of Uniqueness and popular responses to them.

Overall, I think there are three major distinct types of arguments for Uniqueness that have been proposed so far. I list these arguments below (the references indicate the first detailed statement of these arguments within the literature):

1. The arbitrariness argument (White 2005)
2. The evidential support argument (White 2014)
3. The metaepistemic argument (Greco and Hedden 2016)

The arbitrariness argument is a *reductio ad absurdum* argument (explained in detail in Section 1.2.1). Briefly put, the argument aims to show that if Permissivism is true, then it can be rationally permissible to arbitrarily favour believing H over believing $\neg H$ (as both H and $\neg H$ may be rationally permissible to believe). And as forming beliefs arbitrarily is irrational, Permissivism cannot be true; or so the argument goes.

The evidential support argument (explained in detail in Section 1.2.2) is built around the plausible principle that the support relation between evidence and a proposition (or hypothesis) is absolute and does not change from agent to agent. And given some plausible assumptions, this principle seems to entail that rational individuals cannot respond differently to the same evidence.

The metaepistemic argument for Uniqueness comes in various forms.¹⁴ These arguments are called metaepistemic because they are concerned with “the roles that epistemically evaluative talk – attributions of rationality, justification, and the like – play in our lives” (Greco and Hedden 2016, 4). For instance, Greco and Hedden (*ibid.*), following

¹⁴ See Thorstad (2019) for a discussion and responses to this type of arguments.

Craig (1990), have argued that epistemic evaluations are closely tied to deference. They articulate this connection via the following principle:

Deference: If agent S_1 judges that agent S_2 's belief that H is rational on evidence E , then S_1 expresses her commitment to defer to S_2 's belief, unless S_1 has some other relevant evidence that S_2 lacks.

Greco and Hedden argue that Permissivism is in tension with Deference: if Permissivism is true and E rationally permits believing H and believing $\neg H$, then Deference “yields inconsistent deferential commitments” (ibid. 10); and this, according to them, cannot be right.¹⁵

This dissertation is not concerned with such metaepistemic arguments. It seems to me that to argue from a metaepistemic principle to Uniqueness is to do things in the wrong order; as one's attitude towards Uniqueness would determine one's views about which metaepistemic principles are correct. Hence, it is more productive to debate Uniqueness by appealing to more direct epistemic considerations (rather than metaepistemic considerations). This is what we shall do in this dissertation.

So, in the next two sections, I explain and discuss the arbitrariness and evidential support arguments, respectively.

¹⁵ As noted, permissivists have responded to such arguments (see Thorstad 2019). See also Schoenfield (2019) who presents a metaepistemic argument for Permissivism. She argues that it is more difficult for impermissivists than for permissivists to explain why being rational is valuable.

So, it is far from clear whether metaepistemological considerations favour Uniqueness over Permissivism.

1.2.1 The Arbitrariness Argument

The arbitrariness argument is the most widely discussed and criticised argument within the Uniqueness debate.¹⁶ It is a *reductio ad absurdum* argument: it assumes that Permissivism is true and tries to derive an apparently absurd conclusion from this assumption. Here is how it can be stated.

Assume that e is some permissive evidence: that is, e equally justifies the attitude of belief and disbelief towards some proposition, h . Now, if an agent's total body of evidence about h is e , then, the agent can rationally believe h based on e or believe $\neg h$ based on the same evidence e . But, if this agent knows that e is permissive about h , then she knows that her evidence won't provide guidance on whether she should believe h . So, instead of looking at her evidence, she can form her belief via some arbitrary method; say, by rolling a die and believing h if she rolls one and believing $\neg h$ otherwise. After all, if e is permissive, then she would form a justified belief either way. Therefore, if Permissivism is true, then it is rationally permitted to form beliefs by epistemically arbitrary methods; or so the argument goes.

This argument does not convince permissivists. Some (e.g., Smith 2020) simply deny that *acknowledged* permissive cases ever obtain: that is, permissive cases where an agent correctly recognises or acknowledges that her evidence is permissive. On this view, while a body of evidence e could be permissive, it is unknowable that e is permissive. So an agent can never be rational in thinking that there is more than one justified response to her evidence. However, most permissivists accept that acknowledged permissive cases are

¹⁶ Its first, detailed version is due to White (2005).

possible, and hence they don't appeal to this line of defence against the arbitrariness argument.

Instead, most permissivists respond that the arbitrariness argument is problematically ambiguous between the *intrapersonal* and *interpersonal* versions of Permissivism (Kelly 2014, Schoenfield 2014; and Kopec and Titelbaum 2019). Most permissivists accept the intrapersonal thesis that, for any single agent, her evidence justifies a unique doxastic attitude towards a proposition.¹⁷ But all permissivists deny the interpersonal thesis that equally informed individuals are always rationally required to adopt the same doxastic attitude towards a proposition. Permissivists stress the difference between intrapersonal and interpersonal principles because they think that whether an agent should believe a proposition depends not only on the agent's evidence but also on some "third", agent-relative factor(s). Over the years, permissivists have developed several different interpretations of the relevant agent-relative factor. The most common interpretation of this agent-relative factor is in terms of *epistemic standards*. Epistemic standards are the norms of evaluating and reasoning about evidence that are deemed to be reliable or *truth conducive* (Schoenfield 2014). Arguably, no unique set of epistemic standards are required for all agents; hence, two individuals can rationally respond to the same body of evidence differently if they endorse different epistemic standards.

For impermissivists, the appeal to epistemic standards raises the arbitrariness problem yet again. The worry is well-summarised by Simpson (2017, 529):

¹⁷ To my knowledge, the sole exception is Jackson (2019a) who explicitly defends both Intra- and Interpersonal Permissivism.

Where before our problem was arbitrarily favoring one among two rationally permissible doxastic attitudes, now our problem is arbitrarily favoring one among two rationally permissible sets of epistemic standards.

Now, permissivists have proposed different answers for why an agent may have a non-arbitrary reason for favouring one among many rationally permissible epistemic standards (e.g., Peels and Booth 2014; Schoenfield 2014; Simpson 2017; Kopec and Titelbaum 2019). For instance, Schoenfield thinks that our epistemic standards are (at least partly) justified by our epistemic standards themselves. As she explains:

... regardless of whether you are a permissivist, a justification for our standards of reasoning is not something we can provide independent justification for and the demand for such justification would result in widespread skepticism.

Schoenfield's answer to the arbitrariness argument has not convinced most, including permissivists like Simpson (2017) and Weisberg (2020). But the alternative permissivists proposals have not been any more popular.

I return to the arbitrariness argument in Section 4.2.2, after developing a hybrid theory of doxastic rationality. I propose a response to the argument that grants that the choice between different epistemic standards is arbitrary but argues that such arbitrariness is not problematic, given that we endorse an appropriate hybrid theory.

Next, I'll turn to the second prominent argument for Uniqueness that I've introduced in the previous section.

1.2.2 The Evidential Support Argument

The evidential support argument for Uniqueness is built around a plausible principle about evidential support which I call *the Objectivity of Evidential Support* (Objectivity, for short). According to Objectivity, the support relation between evidence and a proposition (or hypothesis) is absolute and does not change from agent to agent; so that if the evidence

supports a hypothesis, e.g., that there is anthropogenic climate change, then the evidence supports the hypothesis for all agents.¹⁸

Roger White (2014) has used Objectivity to argue for Uniqueness; more specifically for Moderate Uniqueness: the view that evidence could not justify both believing a proposition and its negation. The argument is quite straightforward. It appeals to the belief guiding role of evidence, where an agent is rational to believe a proposition H if and only if her total body of evidence supports H . And given that the evidential support relation is objective, all agents with the same body of evidence should adopt the same doxastic attitude towards a proposition.¹⁹ I call this the *evidential support argument* for Moderate Uniqueness.

Many permissivists have responded to White's argument that Objectivity presupposes "a superseded view of evidential support" (Douven 2009, 347). According to this standard permissivist response, it is a mistake to view the support relation as a two-place relation between evidence and a hypothesis (or a proposition); instead, evidential support can be sensitive to various third, agent-relative factors, such as epistemic standards, personal credence functions, epistemic goals or cognitive abilities.²⁰ For instance, according to this line of thought, equally informed jurors may come to different but equally justified

¹⁸ In the context of the Uniqueness debate, Objectivity has been explicitly expressed and endorsed by White (2013), Hedden (2015), and Weisberg (2020).

¹⁹ A similar argument can be found in Feldman (2007) and Matheson (2011). All the other criticisms of EP that I'm aware of are committed to Objectivity: e.g., Hedden (2015), Dogramaci and Horowitz (2016), Greco and Hedden (2016), Stapleford (2019).

²⁰ The view that evidential support, at least in some cases, is agent-relative (or requires some agent-relative factor) has been defended by Douven (2009), Titelbaum (2010), Decker (2012), Kelly (2014), Meacham (2014), Schoenfield (2014), Peels and Booth (2014), Kopec and Titelbaum (2016, 2019), Podgorski (2016), Simpson (2017), Jackson (2019), Jackson and Turnbull (forthcoming).

conclusions about whether a defendant is guilty because they have different epistemic standards on what counts as sufficient and relevant evidence for the defendant's guilt.

I'll return to White's argument in Section 2.2. There I put forward a novel argument for Moderate Uniqueness, which improves on White's original argument; or so I argue. This novel argument substitutes White's Objectivity with a less demanding, logically weaker thesis and is immune to the popular permissivist objection that I've outlined above.

1.3 What Lies Ahead

The dissertation proceeds as follows. In the appendix to Chapter 1, I discuss and evaluate a preliminary objection against Uniqueness: the objection that Uniqueness requires brilliant insights. The appendix is connected to Sections 1.1.2 and 1.1.3, where I have discussed the notions of rationality and evidence. As I have explained there, I do not assume that rationality requires an agent to discover brilliant hypotheses that best explain their evidence. Some (Douven 2009; Jackson 2019a) have argued that this assumption contradicts Uniqueness. In the relevant appendix, I explain and respond to this objection. My response is in line with the earlier sections on rationality and evidence.

In Chapter 2, I motivate Hybrid Impermissivism by providing reasons for accepting each of its two theses: Moderate Uniqueness and Credal Permissivism. I do not discuss how these two theses may hang together in detail. Instead, I aim to show that both Moderate Uniqueness and Credal Permissivism, considered independently, are plausible.

I provide both intuitive and theoretical considerations in favour of Hybrid Impermissivism. First, I argue that both Moderate Uniqueness and Credal Permissivism are intuitive and moderate positions that capture the most plausible aspects of Uniqueness and Permissivism. Next, I put forward a novel argument for a controversial part of Hybrid Impermissivism, Moderate Uniqueness. The argument improves on a similar style of

argument due to Roger White (2005, 2014) and is built on the plausible principle that the evidential support relation is objective. As I'll show, the proposed principle about evidential support is wholly consistent with the view that subjective, agent-relative factors have a rational influence on what an agent ought to believe and to what degree. For this reason, the presented argument avoids the standard permissivist criticism levelled against White's argument.

The overall function of Chapter 2 within this dissertation is to motivate the project of developing a hybrid theory of doxastic rationality. As I've already explained, a hybrid theory not only endorses Hybrid Impermissivism but also gives an account of how these two theses hang together (via an account of belief-credence interaction). Hence, by arguing that each of the two theses of Hybrid Impermissivism is individually plausible, I'm offering a justification for pursuing the project of developing such a theory.

Chapter 3 provides the building blocks for articulating a hybrid theory of doxastic rationality. I will discuss popular norms of rational belief and credence and provide a preliminary discussion on how they might interact. Both beliefs and credences will be discussed within the framework of possible worlds account of propositions, where a proposition is represented by a set of possibilities or possible worlds. On the side of belief, I will introduce and discuss the logical model of belief according to which rational belief is coherent and closed under logical entailment. While this logical model will be assumed for most of the dissertation, in Chapter 6 (Section 6.5), I explain that the assumption of logical closure is not essential for this dissertation. And on the side of credence, I will provide a detailed discussion of the standard Bayesian view of degrees of belief, according to which rational degrees of belief are numerically graded, have a structure of mathematical probabilities, and are updated via the principle of Conditionalisation. While I provide reasons

to call the standard Bayesian account into question, I conclude that its shortcomings won't affect the main argument of this dissertation.²¹

Chapters 4, 5, and 6 are the central chapters of the dissertation, as I start to build, evaluate and modify precise hybrid theories of doxastic rationality in these chapters. In Chapter 4, I put forward the first formally precise, simple hybrid theory about the evidential constraints on belief. As I will show, this hybrid theory runs into a diachronic problem that I call the *diachronic coordination problem*. The problem is about how to coordinate an agent's beliefs and credences over time in a way that preserves the required combination of Uniqueness and Permissivism. As we shall see, two equally informed and rational agents who have the same relevant beliefs but different credences may adopt different beliefs upon learning the same new information. The chapter's overall conclusion is that the success of any (impermissivist) hybrid theory is conditional on solving the diachronic coordination problem.

Chapter 5 is the longest and most demanding chapter of the dissertation. Its formal vocabulary and results are more abstract (and challenging) than what the reader encounters in the previous chapters. In this chapter, I put forward a more sophisticated hybrid theory (compared to the hybrid theory from the previous chapter) that enables us to solve the diachronic coordination problem. It utilises Leitgeb's theory (2014, 2017) about how belief and credence ought to interact. Leitgeb's theory – the *stability theory* – is the only theory that satisfies (i) the standard logical and probabilistic constraints on belief and credence, and (ii)

²¹ A brief comment on the placement of Chapter 3: this chapter does not follow the introduction because it is connected with the task of building a hybrid theory of doxastic rationality (as pursued in Chapters 4, 5, and 6). So, for the sake of continuity, it is more appropriate to place Chapter 3 right behind the three core chapters of the dissertation.

the highly plausible *Monotonicity principle*: the view that any proposition that an agent believes should be more probable for the agent than any proposition that she does not believe. As I show, to solve the coordination problem within the framework of the stability theory, we only need to endorse a relatively undemanding principle, *Order Uniqueness*, according to which for any evidence and proposition, the evidence justifies the unique plausibility order of relevant possibilities (or possible worlds) associated with this proposition.

Chapter 5 has an appendix that consists of three independent sections. In Section 5.5.1, I argue the coordination problem is fatal for a stronger hybrid theory that substitutes Moderate Uniqueness with Extreme Uniqueness. In Section 5.5.2, I provide a solution to the coordination problem for Hybrid Impermissivism by using independent diachronic norms on belief; namely, the AGM theory of belief revision (the best-known formal model of belief revision). And finally, I discuss the relationship between Order Uniqueness and Relational Objectivity (the principle about evidential support from Chapter 2).

Chapter 6 proposes another hybrid theory that also solves the diachronic coordination problem without satisfying the stability theory. This hybrid theory endorses a bridge principle that is less demanding than the stability theory. I call this bridge principle the *dominant core theory*. Even though the dominant core theory violates the Monotonicity principle, I argue that it is a highly plausible alternative to the stability theory. The chapter's main goal is to illustrate that a hybrid approach to Uniqueness is *flexible*: it does not depend on a very specific understanding of how belief and credence should interact. This is an important result because, as we shall see, there are legitimate worries about the stability theory.

Most of the dissertation assumes a setting where we have *complete* probability distributions over a set of propositions (or possibilities/possible worlds). In chapter 7, I switch to the setting where this assumption no longer holds: i.e., where the available evidence

does not justify complete probability distributions over a set of propositions. I provide an impermissivist account of how we can think about rational belief and evidential support in such a setting, using a broadly *likelihoodist* framework from the philosophy of statistics.

I conclude in Chapter 8 by summarising the main claims of the dissertation, reflecting on their main unintuitive/problematic aspect, and briefly touching on an important undiscussed topic – the epistemic significance of disagreement.

1.4 Appendix: Does Uniqueness Require Brilliant Insights?

As I've explained in Section 1.1.2, I don't assume that rational agents have genius-like abilities to discover original hypotheses that best fit their evidence. In short, I don't assume that rationality requires *brilliant insights*. Some (Douven 2009; Jackson 2019a) have argued that this assumption already contradicts Uniqueness. Briefly, their reasoning is as follows: if rationality does not require brilliant insights, you can rationally believe a hypothesis H , but then discover another hypothesis Q that explains your evidence better than H does. After your discovery, you rationally believe Q on the same evidence on which you rationally believed H . And this contradicts Uniqueness. I respond by arguing that considering a new, hitherto unavailable hypothesis is a matter of evidence change. So, in line with how I characterise the notions of rationality and evidence in Sections 1.1.2 and 1.1.3, I argue that a set of relevant hypotheses available to an agent is part of the agent's total evidence.

The appendix is structured as follows. In Section 1.4.1, I'll discuss in more detail the objection that Uniqueness requires brilliant insights. I respond to this objection in Section 1.4.2 and conclude in Section 1.4.3.

1.4.1 Introduction

Rationality does not seem to require brilliant insights. You may be wholly rational in believing a hypothesis, H , because out of the set of all available relevant hypotheses H provides the best explanation of your evidence E . But, at some later time, you may have an "Aha!" moment and realise that there is another hypothesis Q that you hitherto have not considered, and Q provides a better explanation of E than H does. Hence, by the inference to the best explanation, you start believing Q instead of H . Now consider the following question:

Question: Prior to considering Q , were you rational in believing that H ?

It seems that the answer to the question is “Yes”, as rationality does not seem to require you to have brilliant insights. It also seems plausible that considering a new hypothesis does not always change your evidence *E*. Hence, (i) the claim that rationality does not require brilliant insights and (ii) the assumption that considering a new hypothesis does not always change the relevant evidence, taken together, entail that Uniqueness is false.

Such an argument against Uniqueness, which I call the *IBE argument* (IBE for the Inference to the Best Explanation), has been first put forward by Douven (2009) and further elaborated by Jackson (2019a).

One of the most straightforward responses to the IBE argument is that considering a new hypothesis is a matter of evidence change. This *evidence-change* response has an important intuitive appeal. For instance, consider two hypothetical biologists, where the sole epistemic difference between them is that only one has conceptualised the theory of natural selection (and the other has never entertained the thought of natural selection). It seems highly plausible that, due to the explanatory power of the idea of natural selection, these hypothetical biologists possess different total bodies of evidence (solely because they have access to a different set of competing hypotheses).

But proponents of the IBE argument do not need to claim that considering a new hypothesis never changes evidence. They only need to argue that, in *some cases*, considering a new hypothesis is not a matter of evidence change.

In this appendix, I defend a version of the evidence-change response to the IBE argument. As I argue, the evidence-change response fits well with the main idea behind Uniqueness: whether an agent is rational in her belief is independent of subjective, agent-relative factors such as which set of epistemic standards and goals the agent endorses and which bits of her evidence she deems relevant. As the notion of evidence is notoriously

slippery, I think it is unproductive to debate the IBE argument by appealing to a criterion of evidence that is likely to be controversial for many. Instead, I propose to consider the following, more specific questions:

- (1) Does the evidence change response fit with the main guiding idea behind Uniqueness?
- (2) Does the evidence-change response make Uniqueness an uninteresting or trivial thesis?

I defend the positive answer to the first question and the negative answer to the second question. And if my answers are correct, I submit that the evidence-change response is well-motivated.

In the next section, after a brief discussion of how a set of competing hypotheses influences what an agent ought to believe, I will defend the positive answer to question (1). As I'll argue, there is an important difference between paradigmatically non-evidential influences on belief due to *epistemic standards* and influences due to a set of rival hypotheses. Next, I'll defend the negative answer to question (2) by responding to criticism due to Jackson (2019a). I conclude in Section 1.4.3 by outlining an impermissivist view about the extent to which Uniqueness requires rational insights.

1.4.2 The Evidence-Change Response

The evidence-change response that I develop appeals to the standard idea behind IBE that whether a hypothesis, H , provides the best explanation of an agent's evidence E depends on a set of rival or competing hypotheses available to the agent.²² More fully:

²² See, for instance, Bird's (2017, 109) description of how IBE is "typically conceived".

For any agent S and evidence E , a hypothesis, H , provides the best explanation of E for S if and only if H provides a better explanation of E than its available rivals (given that S has been assiduous in considering the alternatives to H).

So, on IBE, whether H provides the best explanation of S 's evidence depends not only on E and H but also on the rivals of H that are available to S .²³ Let R be a set of rival hypotheses available to an agent; then, assuming that both E and R may be represented as sets of propositions, we can simply define an agent's total (relevant) body of evidence to be the union of E and R : $E \cup R$. The idea that an agent's total body of evidence includes a set of rival hypotheses is consistent with Feldman's (2007, 205) original definition of Uniqueness. As he writes (emphasis added):

... [Uniqueness] is the idea that a body of evidence justifies at most one proposition out of a competing set of propositions (e.g., one theory out of a bunch of exclusive alternatives).

Following Feldman's definition, if we specify Uniqueness as the view about what evidence E together with a set of rival hypotheses, R , supports, we have a straightforward response to the IBE argument. Because considering a new hypothesis changes an agent's total body of evidence, Uniqueness may permit an agent to rationally believe a hypothesis H , even if another hypothesis Q , which is unavailable to the agent, provides a better explanation of her evidence. Hence, on this response, Uniqueness does not require brilliant insights.

As I've explained in the introduction, I'll motivate the evidence-change response by arguing that it fits well with the main guiding idea behind Uniqueness. That is, even if an

²³ The idea that the evidential support relation depends on a set of competing hypotheses is not exclusive to IBE and is explicitly endorsed by many theorists in statistics and philosophy (e.g., by Royall 1997; Sober 2008; Bandyopadhyay et al. 2016).

agent's beliefs are sensitive to a set of available rival hypotheses, this does not imply that the agent's beliefs are "subjective" in the permissivist sense.

To show this, first, consider an agent who is fully aware that there are two competing sets of epistemic standards: ES1 and ES2. For the sake of argument, suppose that the agent can rationally believe H only if she endorses ES1 and $\neg H$ if she endorses ES2. So, whether she endorses ES1 over ES2 changes what she is rational in believing. By contrast, whether an agent (deliberately) endorses or favours one set of competing hypotheses R over the other set R^* does not influence which available hypothesis, if any, best explains her evidence. On the supposition that IBE is a reliable inference rule (in a given epistemic situation), an agent is rationally required to believe an available hypothesis that best explains her evidence, even if she knowingly excludes this hypothesis from her set or competing hypotheses. After all, one cannot rationally believe a hypothesis, H , by ignoring an alternative to H that provides a superior explanation of one's evidence. So, whether an agent (subjectively) favours one set of competing hypotheses over the other has no influence on which available hypothesis best explains her evidence. For instance, a religious apologist cannot be rational in her belief that the God hypothesis provides the best explanation of the apparent design in animals if she (deliberately) refuses to include the theory of natural selection into her set of competing hypotheses.

This illustrates an important asymmetry between the doxastic influences due to epistemic standards and due to competing sets of rival hypotheses. From a permissivist perspective, the fact that an agent favours epistemic standard ES1 over ES2 may rationally influence what she ought to believe. By contrast, even if R is an agent's favoured set of competing hypotheses and $Q \notin R$, the agent may still be rationally required to believe Q (via IBE) if she is aware of Q but does not include Q into R . In other words, on the supposition

that an agent is fully aware of Q , the fact that she favours a set of rival hypotheses R over another set R^* where $R^* = R \cup \{Q\}$, does not influence whether Q provides the best available explanation of the evidence.

One may respond to my argument that two equally rational individuals may disagree on whether to include Q into a set of relevant rival hypotheses, because they adopt different epistemic standards. For instance, one agent may assign a very low, negligible prior (probability) to Q , while the other agent may assign a higher prior to Q and treat Q as a realistic candidate for being true. I agree that the above-described situation is a possibility from a permissivist perspective. But if such permissive cases are possible, they are possible due to differing epistemic standards. So, from a permissivist perspective, two agents may agree that a hypothesis provides a better explanation of their evidence than its alternative, but they may still disagree about whether to take this hypothesis seriously. Such cases provide typical (alleged) counterexamples against Uniqueness from epistemic standards and do not support the distinctive IBE argument against Uniqueness.

In summary then: there is an important difference between non-evidential influences on belief, like influences due to epistemic standards and influences due to a set of competing hypotheses available to an agent. From a permissivist perspective, if an agent favours one set of competing standards over the other, this may have a rational influence on what the agent ought to believe. By contrast, on the supposition that IBE is a reliable inference rule, which hypotheses from a set of available hypotheses best explains an agent's evidence is independent of which hypotheses the agent deems relevant. Therefore, the positive answer to our question (1) is well-motivated.

In what follows, I'll consider two objections against the evidence-change response that have been articulated by Jackson (2019a). In my view, the most important criticism out

of the two is that the evidence-change response makes Uniqueness an uninteresting, trivial thesis (the same point has been made by Titelbaum and Kopec 2016; 2019). As she explains, if we claim that all types of changes in an agent's epistemic situation amount to evidence change, then we make Uniqueness an uninteresting or even a trivial thesis. But Uniqueness is not a trivial thesis. Hence, we should not think that considering a new hypothesis is a matter of evidence change.

First off, let me say that I agree with the premise of this response: proponents of Uniqueness should not claim that all types of epistemic changes are a matter of evidence change. I grant that epistemic standards are not part of an agent's evidence. For instance, consider a Jamesian pragmatist view (Kelly 2014). On the supposition that such a view is correct, an agent may believe H or fail to believe H , depending on how much weight she gives to two fundamental epistemic norms: "Believe Truth!" and "Avoid Error!". I fully agree that how much weight an agent gives to these norms does not change her evidence. And if such a Jamesian view is correct, then I concede that Uniqueness is false.

This being said, this criticism of the evidence-change response is wholly unwarranted. Even if considering a new hypothesis changes an agent's (total body of) evidence, it still may be true that evidence E supports more than one hypothesis relative to a set of rival hypotheses R . Again, take a Jamesian pragmatist view: it may be true that two agents who share the same evidence E and consider the same set of rival hypotheses R come to different but equally rational conclusions about which hypothesis from R is true, due to their differing epistemic goals. Such a view is not excluded by the claim that considering a new hypothesis is a matter of evidence change. Hence, the criticism that the evidence-change response makes Uniqueness an uninteresting or trivial thesis is completely unwarranted.

According to Jackson's other criticism, the conception of evidence used in the evidence-change response is "... contrary to how those in the inference to the best explanation literature think of evidence" (Jackson 2019a, 6). As she explains:

In this [IBE] literature, there is a clear demarcation between, on the one hand, one's evidence, and on the other, the hypotheses that are live for an agent. Considering a new hypothesis is not a matter of evidence change; live hypotheses are explained by one's evidence, not a part of one's evidence.

Now, I agree that considering a new hypothesis does not (always) change the evidence that an agent is interested in explaining. But the evidence that the agent is interested in explaining is not the only evidentially relevant factor in IBE. As discussed in the previous section, the support relation is sensitive to both the evidence being explained and a set of rival hypotheses, which comprise an agent's total body of evidence. Hence, my response does not undermine the demarcation between evidence an agent is interested in explaining and a set of rival hypotheses *R* available to the agent.

So my response to the IBE argument avoids both of these criticisms. In the next, concluding section, I will outline an impermissivist picture of the extent to which Uniqueness requires rational insights.

1.4.3 To What Extent does Uniqueness Require Rational Insights?

On the view that I've defended, Uniqueness does not require an agent to articulate a new, brilliant hypothesis that provides the best available explanation of the evidence. Rather Uniqueness is wholly concerned with how an agent evaluates already available evidence (relative to a set of rival hypotheses available to her). So, for instance, Uniqueness did not require Darwin and Wallace to discover natural selection. But, once the theory of natural selection has been conceptualised, Uniqueness required everyone familiar with Darwin's and Wallace's works to see natural selection as the best available explanation of the relevant evidence.

Certainly, coming up with a new, original hypothesis is considerably harder than recognising which way the evidence points (from the available competing hypotheses). Still, the second task may also require important intellectual insights. As Dennett (2001) likes to point out, using the phrase of one of Darwin's earliest critics, Robert MacKenzie, appreciating the idea of natural selection requires a strange inversion of reasoning; and especially in the dawn of Darwinism, such a strange inversion of reasoning required important intellectual virtues, such as open-mindedness, insightfulness, and impartiality.

In summary, while Uniqueness does not require brilliant insights in devising a new, hitherto unknown hypothesis, it may require fairly demanding insights in determining which way the evidence points.

2 Motivating Hybrid Impermissivism

Hybrid Impermissivism is the conjunction of two theses:

Moderate Uniqueness: For any hypothesis H and evidence E , it is not the case that E rationally permits belief that H and belief that $\neg H$.

Credal Permissivism: For some evidence, E , and proposition, H , E rationally permits more than one credence towards H .

This chapter motivates Hybrid Impermissivism by providing reasons for accepting each of these theses. I won't discuss in much detail how Moderate Uniqueness and Credal Permissivism may hang together. My aim is only to argue that these two theses, considered independently, are highly plausible. It is a task of a *hybrid theory* to give a full, precise account of how Moderate Uniqueness and Credal Permissivism cohere with each other (via the account of belief-credence interaction). The hybrid theories are the topic of Chapters 4, 5, and 6. By contrast, this chapter aims to justify the project of developing such a theory of doxastic rationality by showing that its two key theses, Moderate Uniqueness and Credal Permissivism, are individually plausible and well-motivated.

The reader should see the arguments of Chapter 2 as a part of the cumulative case for Hybrid Impermissivism instead of the sole justification for endorsing it. Many attractive features of Hybrid Impermissivism, its overall strengths and weaknesses, are only visible within a precise hybrid theory. For instance, we will see the details of how Hybrid Impermissivism can deal with the objections from epistemic standards and arbitrariness in Chapter 4 (Section 4.2) and show how stringent the requirements of evidence on credences are on Hybrid Impermissivism in Chapter 5 (and also in Chapter 6).

The chapter proceeds as follows. In Section 2.1, I argue that Hybrid Impermissivism is an intuitive, moderate, and conciliatory position. In Section 2.2, I put forward a novel argument for a controversial part of Hybrid Impermissivism: Moderate Uniqueness. I conclude/summarise in Section 2.3 that these considerations provide a good justification for developing a hybrid theory of doxastic rationality.

2.1 The Intuitiveness of Hybrid Impermissivism

... the question of whether Uniqueness is true has many of the trappings of a classic philosophical puzzle. The thesis seems obviously false to many philosophers, and obviously true to many others.

Titelbaum and Kopec (2016,189)

A number of authors (e.g., Kelly 2010, 2014; Jackson 2019b; Stapleford 2019) have noted that the intuitive appeal of Uniqueness is sensitive to which attitude type we are focusing on: belief or credence. For instance, Kelly (2014, 300), who is a permissivist, notes that:

To my mind, uniqueness seems most plausible when we think about belief in a maximally coarse-grained way, so that there are only three options with respect to a given proposition that one has considered: belief, disbelief, or suspension of judgment. On the other hand, as we begin to think about belief in an increasingly fine-grained way, the more counterintuitive Uniqueness becomes. ... as one cuts up the psychology more and more finely, Uniqueness looks increasingly counterintuitive.

Kelly's thought is shared by Stapleford (2019), who is an impermissivist. As he (*ibid.*, 3) writes:

Uniqueness seems very intuitive to me – almost obviously right. So what am I missing? Why would anyone deny Uniqueness?

And as he answers the posed question:

[Uniqueness] loses some of its luster when you start thinking in terms of fine distinctions... So there's definitely something going for permissivism, especially the moderate form.

Kelly and Stapleford emphasise the same problematic aspect of Uniqueness: even when we think about doxastic attitudes in terms of the fine-grained attitude of credence, Uniqueness entails that any evidence fixes a unique doxastic attitude towards any given proposition. And this is highly implausible. It is hard to believe that in every evidential situation, there is always a unique credence, say, a credence of 0.623491, towards a proposition; and “any slight deviation ... [from this credence] counts as a deviation from perfect rationality” (Kelly 2014, 300). For instance, suppose that the only evidence you have about whether it will rain in the next few hours is the qualitative perceptual evidence that the sky above you is mostly blue, with a few clouds scattered here and there. Does such evidence of blue sky fix a unique credence in the proposition that it will rain? It is hard to believe that it does. Certainly, you may be quite confident – say, 90% confident – that it won’t rain, based on the evidence. But slightly more or slightly less confidence seems just as rational.

Even in many scientific settings, the available evidence can be too vague, meagre, or complicated to license a unique credence in a proposition. For instance, most climate scientists strongly believe the anthropomorphic global warming hypothesis. However, there seems to be no unique credence such that all equally informed climate scientists are rationally *required* to have this credence in the hypothesis.

Such considerations provide a good pro tanto reason for rejecting the credal version of Uniqueness, Credal Uniqueness, and endorse Credal Permissivism. And, as a matter of fact, Credal Uniqueness is widely rejected within contemporary epistemology. To quote Douven (2009, 348):

... to the best of my knowledge no one calling him - or herself a Bayesian thinks that we could reasonably impose additional constraints that would fix a unique degrees-of-belief function to be adopted by any rational person.

Douven is completely right. Even contemporary *objective Bayesians* like Williamson accept a *version* of Credal Permissivism. As Williamson (2010) has explicitly noted, his brand of objective Bayesianism in some cases allows subjective, agent-relative factors to influence how strongly an agent believes a proposition. So, it is a mistake to equate contemporary objective Bayesianism with Credal Uniqueness.

We may further note that, contemporary objective Bayesianism is primarily concerned with situations when an agent has mathematically well-described statistical data. And the defenders of objective Bayesianism are not concerned with situations where an agent draws conclusions from rough, fragmentary, qualitative evidence. So, even if objective Bayesianism is true with respect to certain evidential situations, it does not follow that there is one unique credence function that a rational agent should adopt in every evidential situation.

Some (e.g., Christensen 2007; Kelly 2014) think that Credal Uniqueness could be defended against the above-considered obvious objections by appealing to so-called *imprecise probabilities*. According to this response, while in many cases, our evidence does not require adopting a unique credence to a proposition, it may still require adopting a unique *range* or *interval* of credal attitudes. For instance, the evidence of blue sky may not justify a unique point valued credence but, instead, a unique range of credal attitudes with upper and lower bounds, say the range represented by the interval $[0.9, 0.95]$ (or a set of credence functions).²⁴

²⁴ For instance, Christensen (2007, 195, Footnote 8) writes that:

In situations in which the evidence bearing on some proposition P is relatively meager, it does not seem that one unique number could possibly be singled out as the uniquely rational degree of belief in P . But rejecting permissive conceptions of rationality need not commit one to representing the rational

Footnote continued on the next page.

This impermissivist response has been criticised in detail by Castro and Hart (2017). They've argued that it leads to a problem involving changing one's credences in light of new evidence.²⁵

However, as I will argue, there is a more straightforward reason why this strategy is unsuccessful. Even setting aside any potential worries with the imprecise probability framework, Credal Uniqueness does not become any more plausible within this framework. I do agree that, in some cases, an agent's degrees of belief can be modelled via ranges of credence distributions (more on this in Section 3.4). However, in some cases, evidence does not seem to justify a unique range-valued credence any more than a unique point-valued credence. For instance, what credal range should you adopt towards the proposition R "it will rain in the next few hours" if the only evidence you have is the qualitative evidence of blue sky above you? We can grant that you must be quite confident in R , so that your credence in R , $P(R)$, should be high. But is it more rational that $P(R) = [0.9, 0.99]$ rather

response to every evidential situation with a single probability function. . . . One can hold that the uniquely rational response to an evidential situation is representable by a particular set of probability assignments, and the uniquely rational attitude toward proposition P is represented by a particular range of values between 0 and 1.

²⁵ Briefly put, the problem is as follows: suppose that you have no evidence with respect to a proposition H and for that reason, rationality requires you to spread your credence over range $[0, 1]$. Now, it is easy to show that if $P(H) = [0,1]$ then for any evidence E , $P(H|E) = [0, 1]$. This is so, because via Bayes' theorem, $P(H|E) = P(H) * \frac{P(E|H)}{P(E)}$. And for any real number m in $[0, 1]$, there is some value of $P(H)$ such that $P(H) * \frac{P(E|H)}{P(E)} = m$. Hence, $P(H|E) = [0, 1]$. So, if you update your credences via the principle of conditionalisation, no new evidence will change your range-valued credence in H . And this is clearly absurd.

Castro and Hart (2017) show that all available strategies for addressing this and similar difficulties are problematic for imprecise impermissivists.

than $P(R) = (0.9, 0.999]$, or $P(R) = [0.87, 0.96]$? The problem, of course, is that such rough, qualitative evidence does not fix a unique range-valued credence any more than it fixes a point-valued credence. Or take the anthropomorphic climate change hypothesis. It seems equally implausible that instead of a point-valued credence there is some unique *range* with upper and lower bounds such that all equally informed people should adopt this credal range towards the anthropogenic global warming hypothesis. After all, if one is reluctant to accept that evidence always justifies a precise, point-valued credence towards a proposition, then why should one accept that evidence justifies a precise credal range?²⁶

So, contra to Christensen (2007) and Kelly (2014), Credal Uniqueness is not more plausible within the framework of *imprecise probabilities* than within the framework of *precise probabilities*. That being said, I'm happy to grant that the imprecise probability framework provides a better general model for representing rational doxastic states. However, Credal Uniqueness is not more plausible if, instead of point-valued credence, it requires that any given evidence justifies a unique range-valued credence.

Hence, independent of whether we work within the precise or imprecise probability framework, Credal Uniqueness still is "an extremely strong and unobvious thesis" (Kelly 2014).

²⁶ Essentially the same problem remains if we appeal to vague credal ranges, as it is still implausible that there is a unique vague credal range that any evidence justifies. Just ask yourself: what is this unique vague credal range that you should have towards R ("It will rain soon") based on the evidence of mostly blue sky? So, substituting point-valued credences with either precise or vague credal ranges does not turn Uniqueness more appealing.

For what it is worth, Uniqueness about point-valued credences seems more intuitive to me than Uniqueness about precise or vague credal ranges; as range probabilities (either precise or vague) seem more epistemically weird than point probabilities.

By contrast, as both Kelly (a permissivist) and Stapleford (an impermissivist) agree, Uniqueness about belief is not remotely as unintuitive or unpopular as Credal Uniqueness; on the contrary, some find Belief Uniqueness almost obviously right.

This intuitive appeal is much stronger with respect to a version of Belief Uniqueness that I defend in this dissertation – Moderate Uniqueness. Moderate Uniqueness, by definition, is compatible with weak permissivist view that evidence could permit both believing a proposition and suspending judgement; Call this version of Permissivism *Moderate Permissivism*:

Moderate Permissivism: There are some bodies of evidence E , such that E rationally permits two belief-attitudes towards a proposition, where suspension of judgment about the proposition is among the permitted attitudes.

We shall call the permissivist thesis incompatible with Moderate Uniqueness – *Extreme Permissivism*:

Extreme Permissivism: There are some bodies of evidence E , such that E rationally permits believing that H and believing that $\neg H$.

I will argue against Extreme Permissivism in the next section (Section 2.2).

In summary, both Moderate Uniqueness and Credal Permissivism are intuitive and moderate views. Hence, there is a good initial motivation for investigating the prospects of the hybrid view that combines Credal Permissivism with Moderate Uniqueness.

I've begun this section with a quote from Titelbaum and Kopec about the opposing intuitions that Uniqueness elicits. While sources of these intuitions may be varied, it is clear that one of the important sources are which attitude type we are focusing on when we think about Uniqueness. Hybrid Impermissivism offers a way of reconciling these *prima facie*

conflicting intuitions by endorsing Uniqueness about belief but rejecting Uniqueness about credence.

However, the fact that Hybrid Impermissivism is a compromising, conciliatory view does not imply that it is probably true. After all, some find even moderate versions of Uniqueness highly problematic and provide arguments against it (e.g., Schoenfield 2014; Jackson 2019a;). So, I do not expect that Hybrid Impermissivism will be accepted solely on the grounds of being a compromising view.

To motivate Hybrid Impermissivism further, I will put forward a novel argument for its controversial part – Moderate Uniqueness. I take it that the considerations against Credal Uniqueness are sufficiently strong to endorse Credal Permissivism. So, in the remainder of this chapter, I'll be solely concerned with arguing for Moderate Uniqueness.

2.2 An Argument for Moderate Uniqueness

In this section, I offer a novel argument for Moderate Uniqueness (MU, for short in this section) that improves on a similar style of argument due to Roger White (2014). As I've discussed in the introduction (Section 1.2.1), White's argument is built around the principle that the support relation between evidence and a hypothesis is *objective*: so that if evidence E supports a hypothesis H , then E supports H for all agents. In this section, I construct a new argument for MU that appeals to a logically weaker, less demanding view about evidential support, *Relational Objectivity*: whether a body of evidence E is more likely if H is true than if H is false is an objective matter and does not depend on how any agent interprets the relationship between E and H . As its name suggests, Relational Objectivity is a relational or contrastive principle. It is solely concerned with the conditional probabilities called *likelihoods*, and does not put substantive constraints on prior and posterior credences of an

agent. For this reason, the presented argument avoids the standard permissivist criticism levelled against White's argument.

I proceed as follows. First, I provide a detailed discussion of White's argument in Section 2.2.1. This discussion will allow us to see in what respects the novel argument for MU, introduced in Section 2.2.2, departs from White's original argument. I conclude in section 2.2.3 that it is possible to endorse MU and still retain an important permissivist idea that subjective, agent-relative factors rationally influence an agent's doxastic states.

2.2.1 White's Argument for Moderate Uniqueness

As noted in Section 1.2.2, at the heart of White's arguments is the idea I call the *Objectivity of Evidential Support* (Objectivity, for short): the view that the support relation between evidence and a proposition (or hypothesis) is absolute and does not change from agent to agent. White (2014) has specified Objectivity in modal terms, as the thesis that evidential support relations hold *necessarily*: that is, if E supports H then *necessarily* E supports H .

The claim that the evidential support relation holds necessarily may sound unobvious, and even trivially false to some. To take White's (ibid., 313-314) example that illustrates the worry about the necessity claim:

That the gas gauge reads Full supports the conclusion that the tank is full. But it need not. Suppose we know that the gauge is stuck on Full, or even that the wiring is switched so that it tends to read Full only when the tank is empty. In these cases the gauge's reading Full seems to support no conclusion or the opposite conclusion.

So, according to the above example, the evidence g : "The gas gauge reads Full" may support different conclusions, depending on what else we know about the gas gauge. To this example, White responds that g , in itself, does not support any conclusion about the tank. It is only when we combine g with our background evidence that we can meaningfully talk about what the evidence supports. For instance, if our background evidence is that the gauge is typically

reliable, then g unequivocally supports the conclusion that the tank is full. As White (ibid. 314) concludes, when our background evidence is sufficiently specified, “it is hard to make sense of the idea that all of that information might have supported a different conclusion.”

Now, using some additional premises, White gives the following argument for MU from Objectivity (2014, 314):²⁷

The Evidential Support Argument

- (P1) If E supports H then necessarily E supports H .
- (P2) It cannot be that E supports H and E supports $\neg H$.
- (P3) Necessarily, it is rational for S to believe that H iff S 's total evidence supports H .

Therefore:

- (C1) If an agent whose total evidence is E is rational in believing H , then it is impossible for an agent with total evidence E to rationally believe $\neg H$.

The first premise is White's version of Objectivity. Regarding the two other premises: P2 is what might be called the *Univocity Principle* (*Univocity*, for short), the view that “evidence speaks univocally, not equivocally” (Weisberg 2020, 2). So, according to Univocity, if evidence points to H it cannot also point to $\neg H$. The last premise, P3, is a bridge principle connecting evidential support with justified/rational belief. And the conclusion of White's argument, C1, is equivalent to Moderate Uniqueness (MU).

Permissivists have found the argument unconvincing. The most popular criticism of the argument is centred around White's account of evidential support. Several authors have

²⁷ The argument is quoted verbatim, but the order of premises and the original formalism is changed for the uniformity of reading.

argued that the relation of support is always relative to a third relatum. To quote Kopec and Titelbaum (2019, 208):

... support facts obtain only relative to a third relatum; absent the specification of that third relatum, there simply is no matter of fact about whether the evidence justifies the hypothesis.

Permissivists have developed a couple of different interpretations of this “third relatum” (see Kopec and Titelbaum 2016, 194); the most common interpretation is in terms of epistemic standards (Schoenfield 2014). Epistemic standards are the norms of evaluating and reasoning about evidence deemed reliable or truth conducive.

A popular and elegant way of representing epistemic standards is in terms of Bayesian *credence functions*. According to the standard Bayesian position, the degree to which an agent ought to believe a hypothesis, H , depends on (at least) two factors: (i) her (total body of) evidence and (ii) her prior probability in H . Prior probabilities (or priors) encode an agent’s degree of belief in H before receiving evidence E . An agent’s priors may reflect her epistemic standards: say, how much an agent values the simplicity of a hypothesis compared to its explanatory power. So, equally rational agents may adopt non-trivially different priors, depending on how much weight they give to the simplicity considerations over the explanatory considerations; and different priors may lead to non-trivially different posteriors.

Hence, permissivists contend that two individuals can rationally respond to the same body of evidence differently if they endorse different epistemic standards.

In the next section, I’ll state a novel argument for MU that avoids the standard criticism of White’s argument. I will substitute White’s Objectivity with a relatively undemanding principle about evidential support, which I call *Relational Objectivity*: the view that whether a body of evidence is more likely if a hypothesis is true than if the hypothesis is false is an objective matter. On White’s argument, the objective support relation has the

belief-guiding role (see the third premise, P3; more on this in the next section). By contrast, the presented argument won't assume that relational facts about support (fully) determine what an agent should believe and to what degree. For this reason, Relational Objectivity won't be susceptible to the standard permissivist objections; or so I will argue.

2.2.2 A New Argument for Moderate Uniqueness

My argument for MU is a *reductio* argument: it starts with an assumption that MU is false and shows that this assumption contradicts a couple of very plausible epistemic principles. As before, I shall call the negation of MU – *Extreme Permissivism*:

Extreme Permissivism (EP): There are some bodies of evidence E , such that E rationally permits believing that H and believing that $\neg H$.

The argument consists of three premises and a theorem of the probability theory. The first premise is a conditional that states that EP implies the existence of a certain type of permissive cases, and the other two premises are epistemic principles which I call “Moderate Principle” and “Relational Objectivity”. In what follows, first, I'll state the argument in a premise-conclusion form and then discuss these premises one at a time.

The Relational Objectivity Argument

- (1) If EP is true, then two equally informed agents who rationally suspend judgment about H can rationally come to adopt opposing doxastic attitudes about H upon learning some new evidence E : one agent may believe H and the other agent - $\neg H$.
- (2) The Moderate Principle: If evidence E justifies you in believing that H and prior to learning that E , you were not justified in believing H , then E makes it rational to increase your probability in H ; i.e., $P(H|E) > P(H)$, where P represents your credence function and P is a rational credence function for you to have.

- (3) Relational Objectivity: Whether evidence E is more likely on H than on $\neg H$, depends on the evidence and hypotheses themselves and not on how any agent interprets the relationship between these evidence and hypotheses; i.e., for any two equally informed agents with rational credence functions P and P^* , it cannot be the case that $P(E|H) > P(E|\neg H)$ and $P^*(E|H) \leq P^*(E|\neg H)$.
- (4) Theorem: For any H , E and credence function P , E confirms H iff $P(E|H) > P(E|\neg H)$.²⁸

Therefore:

- (5) Moderate Uniqueness (MU): For any hypothesis H and evidence E , it is not the case that E justifies belief that H and belief that $\neg H$.

The argument is valid. To see this, suppose, for reductio, that EP is true. Given the first premise and the Moderate Principle, EP entails that a body of evidence E could confirm H for one agent and $\neg H$ for some other agent. Now, it is a theorem of the probability theory, that, for any H and probability function P , E confirms H iff $P(E|H) > P(E|\neg H)$. And by Relational Objectivity, if the inequality $P(E|H) > P(E|\neg H)$ is true for some agent, then is true for all equally informed agents. Therefore, it cannot be the case that E confirms H for one agent and $\neg H$ for some other agent; contrary to our assumption.

Now that we established that the argument's validity, let us proceed to discuss each of its premises, one at a time.

The first premise does not follow from the definition of EP, but it articulates the key idea behind EP; that some bodies of evidence, in themselves, are radically permissive: so the reason why two individuals can adopt opposing doxastic attitudes towards H in light of their

²⁸ If not otherwise noted, I always assume that for all hypotheses x , $0 < P(x) < 1$.

shared evidence is not because of their prior convictions about H but because of their different evaluation of the same evidence.²⁹ Let me provide an example: consider two open-minded agents who, at some time, share the same (background) evidence K about the existence of God of traditional theism, G , and these agents rationally suspend judgment on G , in light of K . Now, according to premise (1), if EP is true, it may be possible that upon learning some new evidence E these agents rationally come to opposing conclusions about God's existence; so that one agent rationally believes G , while the other rationally believes $\neg G$.

I should note that premise (1) does not imply any substantive constraints on an agent's degrees of belief or credences. If two agents are agnostic about God's existence, this does not imply that their credences in God are the same. For instance, two agents may agree that a necessary condition for believing G is that it has a high probability of being true, say higher than 0.7. So, on this assumption, both agents may suspend judgment on G , even if one is, say, 0.6 confident in G and the other is 0.4 confident in G .

As premise (1) is highly plausible, I expect that it won't be a controversial step in my argument.

The second premise, the Moderate Principle, as its name suggests, is a moderate, uncontroversial thesis. It does not say that you are rational in believing H based on E whenever E rationally increases your probability or rationally *confirms* H . It only states a

²⁹ Many permissivists like Kelly (2014), Schoenfield (2014) and Titelbaum and Kopec (2019) explicitly argue that Permissivism is true because rational individuals can evaluate the same evidence in different ways, and not because they already have opposing attitudes towards H without any evidence. See, for instance, Kopec and Titelbaum's (2019., Sect. 4) *Reasoning Room* example. So given the published defences of EP, a version of EP that is incompatible with premise (1) does not seem to be an appealing view even for permissivists.

necessary (and not sufficient) condition on when it is rational to *start* believing H based on E ; and this necessary condition is that E rationally confirms H . For instance, let H be the hypothesis that Smith did the crime, let E be some new body of evidence; say, the evidence that Smith's fingerprints were found in the crime scene. Now, if prior to receiving evidence E you were not rational in believing H , and if E makes you rational in believing that H , then E must, at least, make it rational to be more confident in H than before.³⁰

What makes an agent's credence function rational (or rationally permissible)? Subjective Bayesians hold that the only rationality requirement on an agent's credence function P is that P is a probability function (that is, P satisfies the standard axioms of probability). This requirement is called *Probabilism*. As we will see, via premise (3), I'll defend an additional constraint on P that goes beyond (probabilistic) coherence;³¹ but I won't appeal to any set of conditions that are jointly sufficient for P to be rational (for an agent).

³⁰ One may object to the Moderate Principle for a reason related to the debate about the permissibility to form a belief in light of "*mere statistical evidence*". For instance, suppose that new evidence reduces your probability for a hypothesis but gives you non-statistical evidence for it, where you previously had only statistical evidence for it. Now, if we think that mere statistical evidence cannot suffice for rational belief, then we'll get cases where gaining evidence can justify moving from suspension of judgment to belief, despite reducing the probability of the believed proposition. Such alleged counterexamples against the Moderate Principle are irrelevant to the argument of this paper. Even if mere statistical evidence is insufficient for rational belief (which is a controversial assumption), we can restrict the argument against EP to the cases that do not involve a transfer from statistical to non-statistical evidence. After all, there is no reason whatsoever to think that EP is only true when an agent's evidence changes from statistical to non-statistical evidence.

³¹ Of course, it is not surprising that any argument against EP should go beyond a purely subjective Bayesian account of confirmation. I should also note that, while subjective Bayesianism is a popular view, many (e.g., see Maher 1996; Hawthorne 2005) have argued that any purely subjective account of confirmation faces some

Footnote continued on the next page.

It is important to note that whether E confirms H for an agent depends on what else the agent knows or takes for granted. So it is useful to divide an agent's total body of evidence into two parts: new evidence E and a body of *background evidence*, denoted by K . What counts as new evidence E and what counts as background evidence K is largely an arbitrary matter and depends on an agent in question and her context of reasoning. For instance, suppose you are particularly interested in how a piece of evidence E bears on the hypothesis, H , that Smith did the crime. E may be the evidence that Smith's fingerprints were found at the crime scene. In evaluating evidence E , your background evidence will include every relevant proposition that you take for granted at that time: such as common-sense propositions about how the world works (e.g., people leave fingerprints and that the fingerprint matching technology is highly accurate) and the assumption that the evidence has not been planted, etc. So given your background evidence K , it is clear that E confirms H .

If we make an agent's background evidence explicit,³² the Moderate Principle can be stated more fully as follows (for simplicity, I assume that both E and K are sets of propositions):

The Moderate Principle: Suppose your total body of evidence is $E \cup K$. If evidence E justifies you in believing that H and prior to learning that E , K alone did not justify you in believing H , then $P(H|E \wedge K) > P(H|K)$, where P represents your credence function and P is a rational credence function for you to have.

serious problems, most notably the so-called problem(s) of old evidence. See Maher (1996) for a detailed argument that a purely subjective account of confirmation is unattainable.

³² Sometimes, for the sake of readability, I won't explicitly mention an agent's background evidence here. But it should be remembered that the talk of confirmation only makes sense relative to a given background evidence. I will make background evidence explicit only when necessary.

The Moderate Principle is as plausible as an abstract epistemic principle can be. It is neutral between permissivist and impermissivist epistemologies. For instance, within subjective Bayesianism, the Moderate Principle is obviously right: after all, if relative to your credence function P , a new evidence E does not increase your credence in H , then why start believing H on E ? If E does not add to your credence in H , then E cannot be a part of your reason for believing H .

Therefore, I also do not expect the Moderate Principle to be a controversial premise in my argument.

By contrast, the third premise, Relational Objectivity, is a controversial premise from a permissivist perspective. So, it requires a more detailed discussion and motivation, compared to the previous two premises.

Relational Objectivity is a comparative principle that is solely concerned with how likely evidence E is if a hypothesis H is true *than* if H is false. Relational Objectivity is *not* concerned with either a *prior probability* of a hypothesis H , denoted by $P(H)$, or a *posterior probability* of H , $P(H|E)$. Instead, Relational Objectivity is about the conditional probabilities of the following type – $P(\text{evidence}|\text{hypothesis})$, called *likelihoods*. A likelihood encodes what a hypothesis, H , says about evidence E : that is, how likely E is on the supposition that H is true.

As with confirmation relation, whether $P(E|H) > P(E|\neg H)$ depends on a body of background evidence K . So, taking background evidence into account, Relational Objectivity can be stated more precisely as follows:

Let P and P^* be rational credence functions of two agents who share the same background evidence K ; then for any evidence E it cannot be the case that $P(E|H \wedge K) > P(E|\neg H \wedge K)$ and $P^*(E|H \wedge K) \leq P^*(E|\neg H \wedge K)$.

The relevance of Relational Objectivity for our argument is made explicit by the theorem of probability theory:

Theorem: For any H , E , K , and probability function P , E confirms H relative to K iff $P(E|H \wedge K) > P(E|\neg H \wedge K)$.

So, given this theorem, Relational Objectivity is equivalent to the thesis that if a piece of evidence E (relative to the fixed background evidence) rationally confirms a hypothesis, then it rationally confirms the hypotheses for all (equally informed) agents.

It may be useful to note that, instead of Relational Objectivity, I could have used a similar principle that has been defended by Maher (1996, 163). Maher has argued that the following, more objectivist analysis of confirmation should substitute the subjective Bayesian analysis of confirmation:

Let $R(K)$ denote the set of all probability functions that are rationally permissible on background evidence K ; then E confirms H relative to K , iff for all $P \in R(K)$, $P(H|E \wedge K) > P(H|K)$.

The gist of Relational Objectivity and Maher's principle is the same: on both principles, whether evidence E rationally confirms H does not depend on how an agent subjectively evaluates the relationship between E and H (relative to K). Clearly, my argument would remain valid if we substitute Relational Objectivity with Maher's principle.

But, unlike Maher's principle, Relational Objectivity makes explicit that the objectivity of confirmation is due to, what Hawthorne (2005, 278) has called "the objectivity or "publicness" of likelihoods that occur in Bayes' theorem." (More on this in section 4.1). So, Relational Objectivity is stated in a way that emphasises this publicness or objectivity of likelihoods.

Why accept Relational Objectivity? Firstly, Relational Objectivity is logically weaker than White's Objectivity: the former is entailed by the latter but not the other way around. So, any reason for accepting Objectivity is also a reason for accepting Relational Objectivity. Let me elaborate on this.

Objectivity is concerned with the traditional notion of evidential support which is closely related to the notion of rational belief. This is made explicit by the third premise of White's argument:

- (P3) Necessarily, it is rational for S to believe that H iff S 's total evidence supports H .

By contrast, as Theorem makes explicit, Relational Objectivity is concerned with the notion of *confirmation*. And the confirmatory relation between E and H is necessary but often insufficient for an agent to rationally believe H on E , even if S 's total body of evidence is E (I give an example shortly). An alternative way of explaining the difference between White's Objectivity and Relational Objectivity is by invoking Carnap's (1962, Preface to the Second Edition) well-known distinction between "concepts of firmness" and "concepts of increase in firmness". White's Objectivity concerns the firmness of a hypothesis; it says that whether a hypothesis is sufficiently firm or probable (for belief) is an objective matter. By contrast, Relational Objective concerns whether the evidence increases the firmness or confirms the hypothesis. And, as it is well-known, the evidence may increase the firmness of a hypothesis without making the hypothesis firm (or sufficiently firm). For instance, consider a detective who received reliable testimony that a suspect, John, was seen near the crime scene. Suppose that this piece of evidence, T , is the detective's total body of evidence that John committed the crime (denoted by J). Now, even if T rationally confirms or increases the probability of J , it is clearly irrational to believe that J solely on the basis of T . In Carnap's terms, T increases

the firmness of J but does not make J firm enough (for the detective). Hence, while confirmation is necessary for rational belief, it is often insufficient.

To sum up: Relational Objectivity is motivated by the same core idea as Objectivity, that the evidential support relation is objective, at least to some extent. But, unlike White's Objectivity, Relational Objectivity only commits us to a moderate view about the extent to which the support relation is objective.

Certainly, permissivists may call Relational Objectivity into question. But, as we will see, to call Relational Objectivity into question requires more than the appeal to the standard permissivists claims: that subjective, agent-relative factors such as epistemic standards, goals, and personal credence functions have a rational influence on an agent's doxastic states. So, I'll be happy to concede to permissivists that there is no objective support relation in White's sense: where the objective support relation fully determines what an agent ought to believe. However, as I argue next, Relational Objectivity won't commit us to such a demanding view about objective support.

2.2.3 Relational Objectivity

As I've already explained, Relational Objectivity is solely concerned with the type of conditional probabilities called likelihoods. Unlike prior and posterior probabilities, likelihoods are widely considered to be the most objective part of Bayesian inference. To illustrate this, let us consider one of the most common forms of Bayes' theorem:

$$(1) P(H|E) = \frac{P(E|H) * P(H)}{P(E|H) * P(H) + P(E|\neg H) * P(\neg H)}$$

Equation (1) enables us to calculate posterior probability, $P(H|E)$, in terms of the prior probability of H , $P(H)$, and two likelihoods: $P(E|H)$ and $P(E|\neg H)$.³³ Notice that (by the law of total probability) the denominator in Bayes' theorem - $P(E|H) * P(H) + P(E|\neg H) * P(\neg H)$ - equals to the *expectedness of evidence*, $P(E)$. Hence, $P(E)$ is nothing else than the *probability-weighted average of likelihoods*. So, (1) can be simplified to:

$$(2) \quad P(H|E) = \frac{P(E|H) * P(H)}{P(E)}$$

And by using equation (2), if we take the *ratio* of $P(H|E)$ and $P(\neg H|E)$ we get the *ratio form of Bayes' theorem*:

$$(3) \quad \frac{P(H|E)}{P(\neg H|E)} = \frac{P(E|H) * P(H) / P(E)}{P(E|\neg H) * P(\neg H) / P(E)}$$

By simplifying, we get:

$$(4) \quad \frac{P(H|E)}{P(\neg H|E)} = \frac{P(E|H)}{P(E|\neg H)} * \frac{P(H)}{P(\neg H)}$$

Finally, let R_{Post} be the ratio of *posteriors*, R_L the ratio of *likelihoods*, and R_{Prior} the ratio of *priors*, then, the ratio form of Bayes' theorem can be summarised succinctly as:

$$(5) \quad R_{Post} = R_L * R_{Prior}$$

As equation (5) makes explicit, the *impact* of evidence on any pair of priors is completely exhausted by R_L , the ratio of likelihoods.³⁴

³³ Sometimes, likelihoods written as $P(E|\neg H)$ denote the likelihoods of a special kind known as *catchall* likelihoods. Catchall likelihoods are discussed at the end of this section.

³⁴ The claim that the ratio of likelihoods is the only factor that impacts how the evidence changes the ratio of priors (which, as equation (5) illustrates, is a fact of probability theory) should not be conflated with a different claim that the ratio of likelihoods provides the adequate *measure* of the degree to which the evidence confirms a

Footnote continued on the next page.

Relational Objectivity is solely concerned with the value of R_L and not at all concerned with R_{prior} and R_{post} . This is an important selling point of Relational Objectivity as prior probabilities are unanimously acknowledged as the most subjective and problematic part of Bayesian inference. And what makes Relational Objectivity more appealing is that it is not a *quantitative* but a *comparative* principle: Relational Objectivity is not concerned with the precise numerical values of likelihoods, but only with their comparative probabilities. As $R_L = P(E|H)/P(E|\neg H)$, it follows that $P(E|H) > P(E|\neg H)$ iff $R_L > 1$. So, on Relational Objectivity, the exact value of R_L is unimportant; what is important is whether R_L is greater than 1.

Now, even the so-called subjective Bayesians – that is, Bayesians who allow the multitude of *coherent* prior distributions as rationally permissible – accept that R_L is the most objective part of Bayesian inference (Hawthorne 2005, 283). The objective status of R_L is due to the fact that, in many cases, an agent’s evidence defines an objective (or inter-subjectively justified) probability distribution over a set of competing hypotheses *without presupposing any prior probability distribution over these hypotheses*.³⁵

To illustrate the independence of likelihoods from prior probabilities, consider the following diagnostic example (a more philosophical example is considered shortly):

You are a physician who assesses a patient on whether she has some skin disease D .

Based on the extensive medical records, you know that the symptoms S , a peculiar

hypothesis. While some have argued that likelihoods are sufficient to adequately measure confirmation, not everyone accepts this. See Festa and Cevolani (2017) for a relevant discussion and references.

³⁵ See also Bandyopadhyay et al. (2016) for a detailed discussion about the special status and role of likelihoods in Bayesian inference.

rash on her hands, is 90% likely if she has D and only 10% likely if she does not have

D . So, you know that $P(S|D) = 0.9 > P(S|\neg D) = 0.1$.

Based on this information, you already know that evidence S confirms D : $P(D|S) > P(D)$.

And as the likelihood ratio is quite high, $P(S|D)/P(S|\neg D) = 0.9/0.1 = 9$, we know that the

evidence S makes the posterior ratio, $P(D|S)/P(\neg D|S)$, nine times greater than the prior

ratio, $P(D)/P(\neg D)$. So it is clear that S provides quite good evidence for D . But, this being

said, the new evidence, S , is *insufficient* to conclude that the posterior probability of D is high

(say, higher than 0.5). This is so, because the prior of D may be quite low. So, suppose that D

is a rare disease and only 1 in 1000 have it. And if your prior in D is 1/1000, then,

unintuitively, simple calculations show that your posterior probability in D should be less

than 1%: $P(D|S) \approx 0.009$. And in general, even if you do not have sufficient information to

provide an objective, uncontentious estimate of the prior of D , you can still rationally

conclude that S is more likely on D than on $\neg D$: hence you can rationally conclude that S

confirms D .

As this diagnostic example illustrates, likelihoods may be independent of priors and in many contexts have “objective or inter-subjectively agreed values” (Hawthorne 2005, 283).

For this reason, even subjective Bayesians accept the special status of likelihoods; for

instance, Edwards et al. (1963, 199) called likelihoods *public* because “[i]n many applications

practically all concerned find themselves in substantial agreement with regard to

[likelihoods]”.

Certainly, fixing the precise numerical values of likelihoods is not always as easy and objective as in the above diagnostic example. However, Relational Objectivity is *not*

concerned with precise numerical values of likelihoods but merely with their comparative

plausibilities. So, on Relational Objectivity, whether $P(E|H)$ is greater than $P(E|\neg H)$ is an

objective matter, even if $P(E|H)$ and $P(E|\neg H)$ do not always have objective numerical values. This makes Relational Objectivity a modest and appealing thesis even from a subjectivist perspective; as there are many cases where the exact *numerical values* of likelihoods are highly debatable, but we may still be in a position to know comparative claims about these likelihoods. To illustrate this, consider the following, more philosophically interesting example. Suppose two agents agree that the existence of evil constitutes evidence against God's existence, in the sense that the existence of evil is less likely if God exists than if God does not exist: $P(\text{evil exists}|\text{God}) < P(\text{evil exists}|\text{no God})$. And the agreement about the comparative plausibilities of these likelihoods requires neither the agreement about the priors, nor the agreement about the precise numerical values of these likelihoods.

While in many scientific and philosophical settings comparative claims about likelihoods are objective (or intersubjectively justified), it is unrealistic to suppose that this is always the case. Essentially, the problem is that sometimes it is not possible to approximate in a non-subjective manner the values of the so-called *catchall* likelihoods: the likelihoods that contain catchall (or composite) hypotheses. A catchall hypothesis is a disjunction of simple (or non-composite) hypotheses. To take a simplistic example: the hypothesis H_1 : "the coin is fair" is simple while the hypothesis $\neg H_1$: "the coin is not fair" is a catchall, as $\neg H_1$ is the disjunction of all the specific alternatives to H_1 (there are many specific ways in which the coin fails to be fair, if not assumed otherwise). Now, suppose the coin is tossed ten times and eight heads are obtained (denote this observation as "e"). The likelihood of e on the supposition that the coin is fair, H_1 , is completely objective and does not require the specification of a prior distribution over the competing hypotheses (no matter what the prior distribution is, $P(e|H_1) = 45 * 0.5^{10} \approx 0.04$). By contrast, the corresponding catchall likelihood, $P(e|\neg H_1)$ is sensitive to the prior distribution. To calculate the value of

$P(e|\neg H_1)$ we must know the values of *ordinary likelihoods* of e on each specific (mutually inconsistent) alternative to H_1 and the *prior distribution* over these alternatives. In symbols:

$$P(e|\neg H_1) = \frac{\sum_{i \neq 1} P(e|H_i) * P(H_i)}{P(\neg H_1)}$$

So, mathematically, catchall likelihoods are reducible to priors and ordinary likelihoods: if we know the values of priors and likelihoods, then we can calculate the value of any catchall likelihood. Thus, in many important settings, the values of catchall likelihoods cannot be neatly separated from the values of priors. The reader may worry that this feature of catchall likelihoods calls the argument of this section into question, as the argument relies on the independence of likelihoods and priors.

But the sensitivity of catchall likelihoods on priors is wholly consistent with my argument. For a start, there is an important *asymmetry* between catchall likelihoods and priors (Fitelson 2007, Section 5). If one knows the values of priors and ordinary likelihoods, then one can calculate the value of the corresponding catchall likelihoods; but not the other way around. Knowing the value of catchall likelihoods and ordinary likelihoods does not determine the prior distribution.³⁶ So, less information is required to fix the values of catchalls than to fix the values of priors. Because of this, the subjectivity of priors does not necessarily translate to the subjectivity of catchall likelihoods: we could have an objective approximation of the value of a catchall likelihood, but not the value of the corresponding priors.

³⁶ To take the coin example again, we could be justified to think that the likelihood of getting eight heads out of ten tosses is higher on the supposition that coin is *not* fair than on the supposition that the coin is fair, even if it is not possible to estimate the value of priors in a non-subjective manner.

I expect the following objection at this point: “But what if two equally rational agents have different estimates of priors and, due to this, they disagree about comparative claims involving catchall likelihoods? Does not this show that Relational Objectivity is false?” The answer is “No”. Relational Objectivity does not entail the strong and unobvious claim that the inequalities between likelihoods are always objectively well-defined. It is compatible with Relational Objectivity that in some cases, the available evidence and hypotheses do not justify even comparative claims about likelihoods (or confirmation). And if such cases obtain, i.e., if the available evidence does not justify objective claims about comparative likelihoods, then, according to my argument against EP, it is irrational to believe (or disbelieve) the relevant hypothesis on that evidence.³⁷ Instead, we should conclude that there is no fact of the matter whether E is more likely on H than on $\neg H$. This conclusion is entirely compatible with Relational Objectivity, as it does not require that the inequalities between likelihoods are objectively well-defined on any evidence.

Hopefully, this discussion convinces even those sympathetic with permissivist epistemologies that Relational Objectivity is a plausible, moderate principle, especially compared to such principles as White’s Objectivity that imposes very strong constraints on rational belief.

³⁷ For instance, Sober (2008, 26-30) has argued that when we deal with “deep and general” theories, such as the general theory of relativity, then some comparative claims about likelihoods cannot be objectively justified. Sober’s overall view agrees with our conclusion that if evidence and hypotheses themselves do not justify the comparative claims about likelihoods, then an agent should abstain from forming beliefs on such evidence.

I return to this topic in detail in Chapter 7. There I discuss in much more detail why in many important scientific settings the values of some catchall likelihoods cannot be approximated in an objective, intersubjectively justified manner. I also put forward a view on how to think about comparative belief in such settings.

2.2.4 Conclusion

To wrap up Section 2.2: I've presented a novel argument for Moderate Uniqueness (MU). According to this argument, MU is true because two equally informed agents who suspend judgement about a proposition H cannot adopt opposing attitudes towards H upon receiving the same new evidence; as adopting opposing attitudes towards H requires that E confirms H for one agent, but $\neg H$ for the other agent; but given a plausible and relatively weak principle about likelihoods – which I've called Relational Objectivity – this cannot be the case.

Since Relational Objectivity is only concerned with relational probabilities of likelihoods, the presented argument for MU is wholly compatible with a plausible permissivist idea that non-evidential factors – such as epistemic standards, goals, and credence functions – have (some) rational influence on what an agent ought to believe (and to what degree).

To illustrate this, suppose that Credal Permissivism is true (that is, equally informed agents can adopt different credences to a proposition). Now, consider two equally informed detectives, Salome and Naomi, who adopt non-identical credence functions P_S and P_N and suspend judgment on whether John committed the crime (denoted by J). Salome may be more sceptical about John's crime and attach lower prior in J ; for simplicity, assume that relative to their shared background evidence K , $P_S(J) = 0.25$ and $P_N(J) = 0.4$. Now, further suppose that they learn a new piece of evidence E that rationally confirms J : $P(E|J) > P(E|\neg J)$. Since Salome had a lower prior in J , her posterior in J , $P_S(J|E)$, may not be high enough for her to believe that J . By contrast, from Naomi's point of view, $P_N(J|E)$ maybe sufficiently high to believe that J .

This example is wholly consistent with the presented argument for MU. So, even if Relational Objectivity is true, two agents may have non-trivially differing credences towards

a proposition and rationally disagree about whether to believe or suspend judgment on that proposition.

Certainly, permissivists may find Relational Objectivity too demanding and advance some novel objections against it. But, as I've argued, these objections must go beyond the standard permissivist claims that subjective, non-evidential factors have a rational influence on what an agent ought to believe and to what degree.

Next, I provide a summary of Chapter 2.

2.3 Summary

In this chapter, I've motivated Hybrid Impermissivism. First, I've argued the intuitive appeal of Uniqueness and Permissivism is sensitive to whether we focus on belief or credence. Hence, Hybrid Impermissivism endorses the most intuitively plausible versions of Uniqueness and Permissivism: Moderate Uniqueness and Credal Permissivism.

Second, I've put forward a novel argument for a more controversial part of Hybrid Impermissivism, Moderate Uniqueness. This argument is built around a moderate principle about evidential support, Relational Objectivity. I showed that Relational Objectivity is wholly consistent with one of the guiding ideas of Permissivism that agent-relative, subjective factors have a rational influence on what we ought to believe.

I submit that, taken together, these considerations in favour of Hybrid Impermissivism provide sufficient justification for developing a hybrid theory of doxastic rationality.

In the next chapter, I will lay the conceptual and formal groundwork and discuss some preliminaries for developing such a theory.

3 Belief, Credence, and their Relationship

This chapter concerns the doxastic attitude-types which figure in the hybrid theories of doxastic rationality developed in this dissertation: (categorical, qualitative) belief and (quantitative, numerical) credence. Belief is understood as a categorical, nongraded attitude of holding a proposition to be true, and credence – a numerically graded attitude of having certain confidence in the truth of a proposition. In section 3.1, after discussing some preliminaries about belief and credence, I will introduce the framework for talking and modelling the objects of these doxastic attitudes: the possible-worlds account of propositions. Section 3.2 gives a detailed statement of the standard logical model of belief, and Section 3.3 – the standard Bayesian model of degrees of belief.³⁸ In Section 3.4, I discuss two important worries with the Bayesian account. One worry has to do with what might be seen as the excessive idealisation of the Bayesian model, according to which rational degrees of belief must be hyper-precise, numerically graded credences (discussed in Section 3.4.1). And the second worry has to do with the Bayesian account of credal revision (discussed in Section 3.4.2). As I explain, the main arguments of this dissertation do not require the assumption that the Bayesian account is fully adequate. In Section 3.5, I'll discuss the issues surrounding the relationship between belief and credence. And in Section 3.6, I summarise.

3.1 The Setting

Contemporary work on rational belief is characterised by two approaches. On a more traditional approach, belief is seen as a coarse-grained, all-or-nothing attitude, where one

³⁸ Some parts of Sections 3.1, 3.2, and 3.3 may be too introductory and slow-moving for some readers. But I believe that being explicit about each relevant aspect of the introduced models will have significant payoffs in terms of facilitating the reading of more difficult chapters, Chapters 5 and 6.

either believes or fails to believe a proposition. By contrast, from the second half of the 20th century, an increasing number of epistemologists have focused on belief as a finely graded attitude, where believing comes in various degrees.

These two approaches to rational belief have a solid grounding in common-sense psychology. Each (neurotypical) human believes a vast number of propositions. Yet, it is a common-sense fact about humans that all believed propositions are not believed with the same strength. For instance, I believe that I locked the door of my apartment. I also believe that the Georgian football team won't win the next World Cup. But I am more confident in the latter proposition than the former. As sometimes I falsely believe that I locked the door, there is roughly 1 in 10 chance that my current belief is false. By contrast, given the competitiveness of the World Cup in football and the relative weakness of the Georgia team, I'm practically certain that Georgia won't be the next world cup winner.

Common examples like this strongly indicate that beliefs differ in strength: an agent may believe proposition *A* and proposition *B*, while, at the same time, believe *A* more strongly than *B*.

It is a standard way of thinking that there is only one fundamental qualitative or coarse-grained attitude, belief.³⁹ I also follow the majority view here, and assume that belief is the fundamental qualitative attitude. Interestingly, the best-known account of fine-grained doxastic attitudes, the Bayesian account, also recognises just one fundamental type of fine-grained attitude: numerically graded belief or credence. On this account, if I have a fine-grained doxastic attitude towards a proposition, *A*, then I land some (numerical) credence to *A*. I will call the numerical notion of belief which is at the heart of Bayesianism – *credence*.

³⁹ Though, some (e.g., Pettigrew 2015), dispute this.

This dissertation will be primarily concerned with these two attitude types: belief and credence. Next section, I'll provide the well-known setting for discussing and analysing these doxastic attitudes in a precise and systematic way: the possible worlds account of propositions.

3.1.1 The Possible Worlds Account of Propositions

This dissertation makes a couple of standard assumptions about belief and credence. Both belief and credence are understood to be in the same business, so to say: they are *propositional attitudes* that aim to *represent* (some portion of) the world. An agent believes a proposition if and only if she holds the proposition to be true. Similarly, an agent's credence is her degree of confidence in the truth of the proposition.

Propositions – the objects of doxastic attitudes – are represented (or modelled) by sets of *possibilities*, or *possible worlds*. More precisely, let $W = \{w_1, w_2, \dots, w_n\}$ be a finite set of mutually exclusive and jointly exhaustive possible worlds. Proposition X (over W) is the set of worlds in which X is true. For instance, X can be set $\{w_1, w_3, w_7\}$. Hence a proposition over W is nothing but a subset of W .

Modelling propositions as sets of possible worlds is formally precise and fruitful because, on this approach, propositions fully inherit a *set-theoretic structure*. The conjunction of two propositions, $X \wedge Y$, is equivalent to a set-theoretic intersection, $X \cap Y$. The disjunction, $X \vee Y$, is equivalent to the union, $X \cup Y$. While the negation, $\neg X$, is the *complement* of X with respect to W , $W \setminus X$. I will sometimes use set-theoretic, instead of logical connectives, when I explicitly think or represent propositions as sets of possible worlds; so, for instance, sometimes I will write $W \setminus X$ instead of $\neg X$ and $X \cup Y$ instead of $X \vee Y$.

The logical relationships between propositions are equivalent to the set-theoretic relationships. Propositions A and B are inconsistent *if and only if* (iff, for short) they have an empty intersection: $A \cap B = \emptyset$. So, A and B are consistent iff they have a non-empty intersection: $A \cap B \neq \emptyset$. Proposition X entails proposition Y iff X is a subset of Y : $X \subseteq Y$.

Proposition X is true at a world w iff $w \in X$; In general, X is true at worlds $w_1, w_2 \dots w_n$ iff each $w_1, w_2 \dots w_n$ is a member of X .

A set-theoretic structure of propositions enables us to state and prove precise theorems about coarse- and fine-grained doxastic attitudes and their (possible) interaction.

On the possible worlds account, propositions are quite coarse-grained objects. Any tautology is true in all possible worlds, each tautology expresses the same proposition: the set of all possible worlds (denoted by W). So, on this account, the sentences “ $2 + 2 = 5$ ” and “Bachelors are unmarried” express the same proposition, W . Similarly, all contradictions are equivalent to the empty set, \emptyset , as there is no possible world in which a contradiction is true.

For the purposes of this dissertation, such a coarse-grained picture of propositions is adequate. No conclusions that I will reach in this dissertation presupposes that propositions are necessarily coarse-grained.

So objects of beliefs – propositions – are subsets of set W of possible worlds. But how should we think about W ? Which possible worlds should go into it? The extension of W (i.e., the possible worlds it contains) is not fixed and depends on an agent in question and her context of reasoning. For instance, consider an agent, Salome, who wonders whether her garden has been watered. She is certain that her garden would be watered if she left the sprinkler on or if it rained. So, in her context of reasoning, she is focusing on the following two propositions:

S: The Sprinkler on.

R : It has rained.

We call set $\{S, R\}$ Salome's *opinion set*. There are four logical possibilities with respect to propositions S and R considered together. Each of these possibilities is associated with a possible world, w_i . As there are four logical possibilities with respect to S and R , Salome needs to attend to four possible worlds:

w_1 : S is true and R is true.

w_2 : S is true and R is false.

w_3 : S is false and R is true.

w_4 : S is false and R is false.

Now, the possibilities that Salome needs to attend to are coarsely individuated. This is entirely appropriate. As there is no need to dissect w_1 into worlds where, say, (i) S and R are true, and Obama is a president, and (ii) S and R are true, and Obama is not a president. It is already assumed that w_1 contains all possibilities where S and R are true. So, in this context, W is the set of all logically possible worlds relative to opinion set $\{S, R\}$.

Any proposition over W that Salome needs to attend to in her context of reasoning is simply a set of worlds in which this proposition is true. For instance, $S = \{w_1, w_2\}$ and $R = \{w_1, w_3\}$.

As in the above example, we will be concerned with an agent who thinks and forms beliefs with respect to set W of possible worlds. W itself is obtained by the agent's opinion set: the set of propositions that the agent is focused on. And as I've already noted, an agent's opinion set is not fixed and depends on the agent's context of reasoning.

Following Lin and Kelly (2021), I will call a proposition that is true in just one possible world a *basic* proposition. In the above example proposition $S \wedge R$, which is

equivalent to proposition $\{w_1\}$, is a basic proposition. In general, a singleton subset of W is a basic proposition. We let Π to denote a set of all basic propositions. Π is a *partition* of propositions, such that: (i) any two propositions in Π are mutually exclusive, and (ii) exactly one proposition in Π is true.

Sometimes, it will be convenient to represent an agent's opinion set as a *partition* of basic propositions, Π . Set Π is a set of all singleton subsets of W . For instance, if W is a set of three possible worlds, $\{w_1, w_2, w_3\}$, then Π is all singleton subsets of W :

$$\Pi = \{\{w_1\}, \{w_2\}, \{w_3\}\}$$

Π and W have the same cardinality, but different types of members: W is a set of worlds while Π is a set of sets of worlds or propositions.

We will call a proposition *relevant* over Π iff (i) it is a basic proposition or (ii) it can be obtained by a *conjunction* of basic propositions. For instance, suppose we have a partition consisting of three propositions, $\Pi = \{A, B, C\}$. There are seven logically possible relevant propositions and one impossible (or contradictory) relevant proposition over Π . These propositions and logical relationships between them are depicted in Figure 3.1. The entailment relationship between the relevant propositions is represented by lines upward (i.e., the entailment have a direction from bottom to top).




Figure 3. 1: Logical relationships between three basic propositions

As the example illustrates, each non-contradictory (or non-empty) proposition over Π is either a basic proposition or is logically equivalent to a conjunction of basic propositions. For instance, $\neg C$ is entailed by A , and it is entailed by B . Hence, $\neg C$ is logically equivalent to $A \cup B$.

In general, each relevant proposition over Π corresponds to some subset of W and vice versa. Hence, the number of relevant propositions over Π is the same as the number of subsets of W , and equals to 2^n , where n is the number of worlds (or cardinality) of W (or the number of basic propositions in Π). For this reason, whenever convenient, we can define the agent's beliefs over a set of propositions Π , instead of a set of possible worlds W .

3.2 Belief

An agent's (categorical) beliefs will be represented by Bel , which is a set of all propositions believed by the agent. So, for any proposition X , X is believed iff $X \in Bel$. For convenience, I will write $Bel(X)$ instead of $X \in Bel$.

Belief comes in three types: one can either believe a proposition, disbelieve it, or suspend judgement. In our model, the attitude of disbelief in a proposition is identical to the attitude of belief in the proposition's negation. So, disbelief in a proposition, X , is written as $Bel(\neg X)$. Regarding the attitude of suspension: if neither X nor $\neg X$ is a member of Bel , then our agent would suspend judgment on X . So, suspension of judgement is logically equivalent to neither believing nor disbelieving a proposition.

There are some obvious objections against the view that suspension is simply a matter of non-belief (see Friedman 2013). For instance, consider a typical nine-year-old child who never considered or entertained the theory of natural selection. Clearly, this child does not believe nor disbelieve the theory of natural selection. Does this mean that the child suspends judgement on the question? Certainly not. The child simply lacks any doxastic attitude towards the theory of natural selection. Such a problem is not confined to children. Adult humans also lack belief about a vast number of propositions because they never considered them.

We can set aside such problems with the “non-belief” view of suspension for the following reason. As I've already explained in the preceding section, we shall discuss and analyse the agent's beliefs relative to her opinion set (or a set of possible worlds obtained from the opinion set). The agent's opinion set is just a set of propositions that she considers in some context of reasoning. So, on this picture, suspension of judgement about H is equivalent to neither believing nor disbelieving H only if H is in the agent's opinion set; or, in other words, only when the agent considers H .⁴⁰

⁴⁰ The issue here is far more complicated. Most people never considered the proposition that salt is made of gold atoms, but I submit that they nevertheless believe the proposition to be false. So, considering a proposition does not seem necessary for belief. To avoid such complications, we need to focus on the so-called *occurrent beliefs*

Footnote continued on the next page.

Friedman (2013, 2017) has argued that some aspects of suspension cannot be captured by a kind of non-belief view that I've sketched above. For instance, one may drive a wedge between the kind of mental states described in the following situations:

- (1) An agent considers *H* but so far have *not reached a conclusion* about *H*; hence she neither believes nor disbelieves *H*.
- (2) An agent considers *H* and *concludes* that neither belief nor disbelief is appropriate in *H*; hence she neither believes nor disbelieves *H*.

Friedman thinks that (2) describes the suspension of judgement but not (1).⁴¹ I don't think so. But whether (1) should be qualified as a suspension of judgement is inconsequential to the arguments of this dissertation. So, we will assume the qualified non-belief account of suspension is adequate for this dissertation.

Let's now turn to the normative side of belief. What norms should rational belief satisfy? Alternatively: what norms should rational belief set *Bel* satisfy?

The standard, popular requirements on *Bel* are the requirements of *consistency* and *deductive closure*:

(Peels 2017): i.e., the beliefs concerning propositions that the agent considers at some given time. And believing a proposition, in the occurrent sense, by definition, requires considering the proposition.

⁴¹ Here is one of the examples that Friedman (2013, 170) considers and analysis:

“Let's say that S deliberates about p at 15:00 for 5 min (until 15:05) and is in a state of p-non-belief throughout. It is perfectly appropriate to describe S at 15:02 as considering or having considered p. But at 15:02 “mid-wondering”, it does not look appropriate to describe S as having suspended judgment about p. All S has done by 15:02 is begun to think about p. But that alone does not turn his state of mere non-belief into one of suspended judgment. That you have simply started thinking about whether de Gaulle took that walk along the Seine does not on its own (or in combination with non-belief) mean you've suspended on the matter.”

Consistency: *Bel* should be consistent: it is not the case that $Bel(\emptyset)$.

Closure: *Bel* should be closed under conjunction and logical entailment: if $Bel(X)$ and $Bel(Y)$, then $Bel(X \cap Y)$; and if $Bel(X)$ and $X \subseteq Y$, then $Bel(Y)$.

If *Bel* satisfies both Consistency and Closure, then we will say that *Bel* is *deductively cogent*.

All deductive cogent belief sets have the following nice formal feature: if *Bel* is deductively cogent, then there is some non-empty proposition B_W in *Bel* and all other propositions in *Bel* logically follow from B_W . Proposition B_W can be regarded as the logically strongest proposition believed by the agent or, equivalently, the conjunction of all propositions believed by the agent. So, *Deductive Cogency reduces complexity*: any finite deductively cogent belief set can be represented with one unique proposition, B_W (or, in set-theoretic terms, we get a reduction of complexity from a *set of sets of worlds* to a single *set of worlds*). This feature of any deductively cogent belief set can be stated as the following theorem:

Theorem 3.1 (Deductive Cogency reduces complexity):

Bel over W is deductively cogent iff for some non-empty proposition B_W over W ,
 $Bel(X)$ iff $B_W \subseteq X$.⁴²

To illustrate the theorem, consider set W of four possible worlds, $W = \{w_1, w_2, w_3, w_4\}$, visualised in Figure 3.2:

⁴² General remark: I omit the proofs of theorems which are common knowledge within (formal) epistemology. Lesser known or novel theorems will be proved in the footnote at the end of the respective theorems. Theorems are numbered by two numbers; the first indicates the chapter in which the theorem is stated, and the second – the placement of the theorem within the chapter. Some theorems, like the one above, have a caption whenever informative or helpful for referring to it.

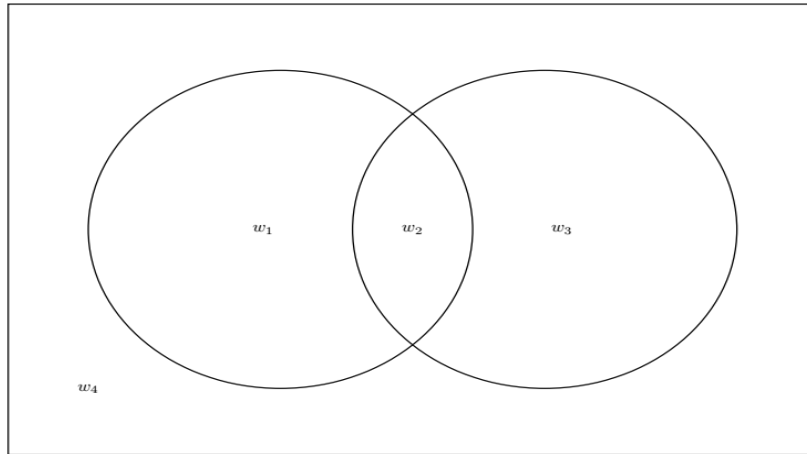


Figure 3. 2: Possible worlds over $W = \{w_1, w_2, w_3, w_4\}$

Four possible worlds partition the logical space, W , into four regions. And if an agent's belief set Bel is deductively cogent, then Bel can be obtained from some non-empty region, B_W over W ; so that any believed proposition must contain B_W as a subset. For instance, if $B_W = \{w_1, w_3\}$, then Bel is equivalent to the set containing the following sets (as each of these sets contain B_W as a subset; or equivalently, each of these sets is a *superset* of B_W):

$$\{w_1, w_3\}$$

$$\{w_1, w_2, w_3\}$$

$$\{w_1, w_3, w_4\}$$

$$\{w_1, w_2, w_3, w_4\}$$

To visualise the relationship between $B_W = \{w_1, w_3\}$ and Bel , consider Figure 3.3, where the shaded region represents B_W :




Figure 3. 3: $B_W = \{w_1, w_3\}$ for $W = \{w_1, w_2, w_3, w_4\}$

The shaded region represents the strongest believed proposition, B_W . And on the supposition of Deductive Cogency, any region in the Venn diagram that fully contains this shaded region is also believed. So, for instance, region $\{w_1, w_2, w_3\}$ contains B_W ; hence $Bel(\{w_1, w_2, w_3\})$. By contrast, if a region, X , does not contain any region of B_W , then, set-theoretically, the complement of this region, $\neg X$, must fully contain B_W ; hence, $Bel(\neg X)$. For instance, region $\{w_4\}$ does not contain any region of B_W , so the agent *disbelieves* $\{w_4\}$: $Bel(\neg\{w_4\})$, or, equivalently, the agent believes the negation or complement of $\{w_4\}$: $Bel(\{w_1, w_2, w_3\})$.

What if a region, X , contains some but not all regions of B_W ? In this case, the agent does not believe either X or $\neg X$: $\neg Bel(X)$ and $\neg Bel(\neg X)$. Hence, by the definition of the suspension of judgement, the agent suspends judgment on X .

To summarise the implications of Theorem 3.1: if the agent's belief set Bel over W is deductively cogent, then Bel can be obtained from a non-empty (i.e., non-contradictory) proposition B_W , the strongest proposition that the agent believes (B_W must be non-empty because, by consistency, Bel does not contain an empty proposition). Because B_W is some subset of W , for any subset X over W , one of the following must obtain:

- (i) B_W is a subset of X .

(ii) B_W and X have an empty intersection.

(iii) B_W is not a subset of X but B_W and X have a non-empty intersection.

If (i) is the case, the agent believes X , if (ii) – disbelieves X , and if (iii) – suspends judgement on X . So, the standard tripartite picture of belief is entailed by the possible worlds account of propositions and Deductive Cogency.

Granted that deductively cogent beliefs have nice formal features, it is still an open question whether Deductive Cogency is a norm of rationality. After all, even though Deductive Cogency is a popular, widely used norm, some explicitly reject it (e.g., Kyburg 1961, 1970).

The obvious worry with Deductive Cogency is that it is *overly demanding*. For instance, the logical consequences of some of our beliefs are often extremely hard or even practically impossible to trace. We may believe a set of propositions but disbelieve some of their consequences because we cannot recognise the highly complex entailment relationships between them. So Deductive Cogency faces the so-called problem of *logical omniscience*.

Do such considerations show that humans cannot satisfy Deductive Cogency? If so, why consider Deductive Cogency as an interesting norm on rational belief?

Such worries are avoided by the version of Deductive Cogency that we will be concerned with. As it may be clear from the previous section, we are not interested in Deductive Cogency as a requirement on a set of all propositions the agent believes. Instead, we are interested in Deductive Cogency as applying to a precise number of coarsely individuated basic propositions. For instance, suppose you are focused on whether your garage door is locked. In this context, you consider three “basic” answers to the question (i.e., the answers that you treat to be mutually inconsistent and jointly exhaustive): (i) you locked the door; (ii) your partner locked the door; (iii) nobody locked the door. In this context of

reasoning, there are just eight relevant propositions that you must attend to, with one necessary proposition and one contradictory proposition. Even ordinary humans can be expected to be deductively cogent in this context of reasoning, if they are rational.

To my knowledge, such an understanding of Deductive Cogency has been first explicitly proposed by Levi (1967, 41). As he put it:

To be sure, if deductive cogency is supposed to apply to all the beliefs held by a person over his entire career, without any restrictions, cogency clearly deserves to be abandoned. But it is highly doubtful whether anyone ever seriously took the principle to be understood in that sense. In any event, as it is understood here, the principle of deductive cogency has its scope of application restricted to beliefs that are held relative to a fixed body of evidence and a given set of relevant answers.

In more recent years, Lin and Kelly (2012, 2021) and Leitgeb (2013, 2014, 2017), among others, endorsed Levi's approach to Deductive Cogency. As Leitgeb (2017, 35) puts it:

... [in] coarse-grained contexts, the usual worries concerning epistemic logic's assumption of 'logical omniscience' lose much of their bite: there is just one logical truth to be rationally believed, and checking for logical implication amounts to a mere test for subsethood, neither of which is particularly delicate given a reasonably small number of coarse-grained possibilities.

So, overall, Deductive Cogency seems quite plausible in coarse-grained contexts that we shall be dealing with in this dissertation.

I will use Deductive Cogency to develop two different hybrid theories of doxastic rationality in Chapters 5 and 6, respectively. Though, as I will argue in Section 6.5, Deductive Cogency is not indispensable for a viable hybrid theory.

Deductive Cogency imposes a *synchronic* coherence requirement on *Bel*: it is solely concerned with how *Bel* ought to be at any given time. By contrast, *diachronic* requirements concern how *Bel* ought to be revised over time in response to new evidence. In most of this dissertation, I won't impose any explicit diachronic requirements on *Bel*. This is not because

there are no such requirements; but because imposing the diachronic requirements on *Bel* would complicate the discussion unnecessarily. And as I'll explain in detail in Section 5.3, belief change will simply piggyback on degree-of-belief change. In any case, in one of the appendices to chapter 5, Section 5.5.2, I'll show how to articulate a hybrid theory that explicitly appeals to the norms about belief change (in terms of the well-known AGM theory of belief revision).

Now I move to discuss another central attitude-type within epistemology: the fine-grained, numerical attitude called credence.

3.3 Credence

Some of the most significant work in contemporary epistemology is centred on the Bayesian notion of credence, according to which rational degrees of belief are numerically graded and have the structure of mathematical probabilities. The Bayesian framework recognises infinitely many credal attitudes towards a proposition, where each credal attitude is represented by a real number in the unit interval.

The terms “degree of belief” and “credence” are often used interchangeably. In this dissertation, by “degree of belief”, I mean fine-grained doxastic attitude in general, while by “credence” – the Bayesian conception or model of degree of belief. As I discuss in the next section, Section 3.3.1, there are good reasons for thinking that not all types of degrees of belief are credences. Still, this dissertation primarily focuses on the Bayesian model of degrees of belief for two reasons.

First, credence is the best-known and most widely discussed fine-grained attitude in contemporary epistemology in general and the Uniqueness debate in particular, and for good reasons, as the notion of credence, in many contexts, provides a useful, fruitful explication of the ordinary notion of confidence or degree of belief. Secondly, and most importantly, the

main arguments of this dissertation work without the presupposition that degrees of belief are always numerically graded. So, the hybrid theories that I develop do not require the assumption that all degrees of belief are credences.

Now that the above qualification is in place, let's proceed with discussing the Bayesian notion of credence.

As I've already noted, on the Bayesian model, rational degrees of belief has the structure of mathematical probabilities. The standard mathematics of probability is axiomatised and involves only three simple axioms. There are various formally equivalent ways of stating these axioms. In this dissertation, we state these axioms in terms of the possible worlds account of propositions and show that these axioms are jointly equivalent to three simple, easy to use conditions involving the assignment of numbers to worlds.

Let P an agent's credence function that assigns each proposition that the agent considers (in her context of reasoning) some real number in the unit interval $[0,1]$. More precisely, let W be a set of worlds that the agent attends to and let P be a function from all propositions (or subsets) over W to interval $[0,1]$. P is defined to be a probability function iff it satisfies the following three conditions:

- (1) For all propositions X over W , the probability of X is a real number in the unit interval. In symbols, $P(X) = c$, where $c \in [0,1]$.
- (2) The probability of a tautological proposition is 1. That is, $P(W) = 1$.
- (3) For any propositions X and Y over W , if X and Y are mutually exclusive propositions, then $P(X \text{ or } Y) = P(X) + P(Y)$.

Now, given the possible world account of propositions, the above definition of a probability function is logically equivalent to three simple and easy to use conditions. This is made explicit by the following theorem:

Theorem 3.2

Function P over set of possible worlds W is a probability function iff:

(C1) For each $w \in W$, P assigns some non-negative real number to singleton set $\{w\}$.

(C2) The sum of all $P(\{w\})$ equals to unity:

$$\sum_{w \in W} P(\{w\}) = 1$$

(C3) For any X , where $X \subseteq W$, $P(X)$ is the sum of all the probability values of worlds that are in X .⁴³

$$P(X) = \sum_{w \in X} P(\{w\})$$

The first condition, (C1), says that P assigns every singleton subset (or basic proposition) over W some non-negative real number. (C2) says that the sum of all singleton subsets is 1. And (C3) says that, for every proposition X over W , its probability can be obtained by summing the probabilities of each world in which X is true.

The following simple example illustrates the application and usefulness of the theorem. Consider set W of three possible worlds. Each possible world in W corresponds to a *basic proposition*: a proposition that is true in exactly one possible world. Denote these propositions as A , B , and C , respectively. To define a complete probability distribution over W (or, over a set of basic propositions $\Pi = \{A, B, C\}$) it is sufficient to assign probabilities to just these three basic propositions. For instance, your credence function P may be such that $P(A) = 0.5$, $P(B) = 0.3$, and $P(C) = 0.2$. If this is the case, then conditions (C1) and (C2)

⁴³ Note that, by the phrase “probability of a world, w ”, I mean the probability of the corresponding singleton set, $P(\{w\})$.

are satisfied. And, if X is a non-basic proposition, then, by condition (C3), $P(X)$ is obtained by summing the probabilities of the worlds in which X is true. Or, equivalently, $P(X)$ is the sum of the probabilities of basic propositions that entail it. For instance, if $X = B \cup C$, then $P(X) = P(B) + P(C) = 0.5$. Hence, we have a simple method for determining a complete probability distribution over any set of propositions (or possible worlds): we simply assign probabilities to basic propositions that sum to 1 and then use condition (C3).

In general, any probability function over a set of three basic propositions may be equated with the triple (a, b, c) , where $a = P(A)$, $b = P(B)$, and $c = P(C)$. The set of all such triples can be represented geometrically via an equilateral triangle, as depicted in Figure 3.4:




Figure 3. 4: Geometric representation of all coherent credence functions

The black dot in the middle of the triangle denotes the uniform probability distribution. The A vertex corresponds to the probability function that assigns the maximal credence to A and zero credence to B and C . Similarly, the vertices B and C represent the probability functions where B and C have probability 1. The further the point is from the A vertex, the less probability A gets on that point. Hence, the opposite side from the A vertex represents the set of credence distributions where $P(A) = 0$.

If a probability distribution over three propositions is not located on this triangle, this distribution violates some condition(s) of Theorem 3.2.

According to the standard Bayesian view, rational degrees of belief have a structure of mathematical probabilities, and hence should satisfy the conditions listed in Theorem 3.2. This Bayesian norm is called *Probabilism* (which I first introduced in Section 1.1.2). Degrees of beliefs that satisfy probabilities are called *probabilistically coherent*. In this dissertation, I generally assume that rational degrees of belief are probabilistically coherent.⁴⁴ More precisely, I assume that if an agent's degrees of belief are numerically graded (i.e., representable by real numbers in the unit interval), then the degrees of belief must be probabilistically coherent.

Probabilism is the synchronic coherence norm on rational degrees of belief in the same way as Deductive Cogency is the synchronic coherence norm on categorical belief. That is, Probabilism puts a constrain on an agent's credences at any given time. In addition to Probabilism, I'll also assume a diachronic norm called *Conditionalisation* about how the

⁴⁴ There are many kinds of arguments for Probabilism; the most popular is the *Dutch Book Argument* which traces back to Ramsey (1926) and de Finetti (1937). I do not consider or evaluate them in this dissertation. Probabilism is the standard assumption within the Uniqueness debate, and I see no benefit in calling it into question in our context.

agent's credence ought to be revised after acquiring some new information.

Conditionalisation can be stated as follows:

Conditionalisation:

For any agent with credence function P and any proposition X over W , if the agent learns that some proposition E over W is true, then her new credence function should be P_E where:

$$P_E(H) = P(H|E)$$

In plain English, Conditionalisation says that if the agent learns some proposition E to be true, meaning that she is certain that E is the case, then her new credence in H should be equal to her old conditional credence in H on the supposition that E . So, Conditionalisation explicates the notion of learning in terms of conditional probability.

Within the standard probability theory, conditional probability is defined as follows:⁴⁵

$$P(H|E) =_{Def} \frac{P(H \cap E)}{P(E)}$$

Therefore, by this definition, Conditionalisation is identical to the following:

$$P_E(H) = \frac{P(H \cap E)}{P(E)}$$

Equivalently, we can use the *Bayes' theorem* instead of the definition of conditional probability (as the two are logically equivalent under the axioms of probability) to state Conditionalisation:

$$P_E(H) = \frac{P(E|H) * P(H)}{P(E)}$$

Now, there is a very useful, simple way to think about Conditionalisation within our possible worlds account of propositions. Consider an agent with credence function P defined

⁴⁵ If not otherwise noted, I assume that for all x , $P(x) > 0$.

over W of four possible worlds, $W = \{w_1, w_2, w_3, w_4\}$. Now, suppose the agent learned some proposition E over W . For any world w_i in W , E is either true at w_i or not. Or, in other words, E either has an empty or non-empty intersection with $\{w_i\}$. If E is not true at w_i , then $P(E|\{w_i\}) = 0$; therefore $P_E(\{w_i\}) = 0$. And if E is true at w_i , then $P(E|\{w_i\}) = 1$; thus, if E is true at w_i :

$$P_E(\{w_i\}) = \frac{1}{P(E)} * P(H)$$

So her new credence function P_E is simply obtained by multiplying the probabilities of each world in W compatible with E by the same ratio, $1/P(E)$. For instance, suppose that $P(\{w_1\}) = 0.35, P(\{w_2\}) = 0.3, P(\{w_3\}) = 0.2, P(\{w_4\}) = 0.15$ and $E = \{w_1, w_4\}$. Now, $P(E) = 0.5$, hence $\frac{1}{P(E)} = 2$. So, the probabilities of all compatible worlds are multiplied by the same constant, 2: $P_E(\{w_1\}) = 0.35 * 2 = 0.7$ and $P_E(\{w_4\}) = 0.15 * 2 = 0.3$.

Because P_E is obtained by multiplying all worlds in W compatible with E with the same constant, this means that the ratio of all probabilities of basic proposition over W remains the same. Hence, we have the following highly useful theorem:

Theorem 3.3 (Conditioning preserves ratios of worlds/basic propositions):

For any basic propositions A and B , and any probability function P over W , if E is compatible with A and B , then revising or conditioning on E preserves the probability ratios of A and B :

$$\frac{P(A)}{P(B)} = \frac{P_E(A)}{P_E(B)}$$

This theorem will be handy in Chapters 5 and 6 (for proving, what I call, the *diachronic coordination theorems* for belief and credence).

According to Conditionalisation, the agent should revise her credences by E only when she learns E for certain. But this seems overly restrictive. I may learn that the

probability of E is, say, higher than I previously thought, and such learning seems to require me to update all my relevant credences.

There are diachronic norms that allow credence revision in light of uncertain evidence, the most popular of which is *Jeffrey Conditionalisation* (due to Richard Jeffrey 1965, Chapter 11):

Jeffrey Conditionalisation: Let $P_{E,c}(H)$ denote an agent's new credence in H after her new probability in E is c ; then:

$$P_{E,c}(H) = P(H|E) * c + P(H|\neg E) * (1 - c)$$

I discuss Jeffrey Conditionalisation and its impact on the main arguments of this dissertation in detail in Section 5.3.3. But for the most part, I assume Conditionalisation as a general diachronic norm on credal revision.

So, to sum up: according to the standard Bayesian model that is widely used in this dissertation, degrees of belief are represented by credences: numerically graded doxastic attitudes that obey the axioms of probability. And on the diachronic side, we mostly assume that credences are revised by Conditionalisation.

In the next section, I will critically discuss both the synchronic and diachronic parts of the standard Bayesian account of degrees of belief. As I explain, even if the standard Bayesian model is problematic, this does not affect the main arguments of this dissertation.

3.4 Worries with the Bayesian Model of Degree of Belief

3.4.1 Probabilism and Doxastic Precision

According to the traditional Bayesian view extensively used in this dissertation, rational degrees of belief come in precise numerical strengths and are represented by a single probability function. Hence, the standard Bayesian view accepts the following norm I call *Doxastic Precision*:

Doxastic Precision: For any perfectly rational agent and proposition X , there is some numerical strength to which the agent believes X .

Many have argued that Doxastic Precision is false. The most obvious worry with Doxastic Precision has to do with the situations where an agent has rough, qualitative evidence. For instance, consider the following example by Sturgeon (2020, 68), which, as he argues, points against Doxastic Precision:

Consider whether it will be sunny in Tucson tomorrow. If you are like me, you know that Tucson is located in the Sonoran Desert and is almost always sunny. But you will have no precise statistics about how often it is sunny in Tucson, and nor will you have evidence about the objective chance of sun there tomorrow. All you will have is qualitative evidence to the effect that it is extremely likely to be sunny in Tucson tomorrow...common sense underwrites the idea that you should be confident that it will be sunny in Tucson tomorrow, given your evidence; but it does not insist that you should lend a specific credence to that hypothesis. And nor does common sense see you as picking one just to get on with things rationally. No. The ordinary description of a case like this involves high confidence without credence.

Sturgeon's example, as well as many similar examples discussed in the literature, aims to establish a general point: in most cases where an agent has limited, qualitative evidence, it does not seem rational for the agent to have Bayesian hyper-precise credences.⁴⁶

⁴⁶ For instance, as Dorling (1977), among others, have noted, in many scientific contexts, even the doxastic states of highly trained experts seldom have the fine-grained structure of Bayesian credences. In Dorling's words (*ibid.*, 180):

... scientists always conducted their serious scientific debates in terms of finite qualitative subjective probability assignments to scientific hypotheses ... and they still do regularly do so in their philosophically unguarded moments. The probabilities they thus assign are often large: 'more probable than not', 'very probable', 'almost certainly correct', 'so probable as to be almost necessary' and so on.

Unless it can be shown that the behaviour of all these scientists (e.g. their betting behaviour) was such

Footnote continued on the next page.

The obvious counterexamples with Doxastic Precision have led many theoreticians to the view (more precisely, to the family of views) called *imprecise Bayesianism*. Roughly, this is the thesis that an agent's degrees of belief should be represented with ranges or intervals instead of point-valued probabilities.⁴⁷ For instance, consider the above-discussed example given by Sturgeon. In the example, you are very confident that it will be sunny in Tucson tomorrow, based on your belief that Tucson is located in the Sonoran Desert and is almost always sunny there. Certainly, it is highly unrealistic to insist that your fine-grained belief has a unique numerical strength. Instead, according to imprecise Bayesianism, in such cases, your belief should be represented by some range, say $[0.9, 0.99]$, meaning that you are at least 0.9 and at most 0.99 confident that it will be sunny in Tucson tomorrow.

I agree that, in some cases, an agent's degrees of belief seems to be better modelled via ranges of credence distributions. However, a moment's reflection shows that imprecise Bayesianism cannot fully mitigate the worry with the Bayesian model of degree of belief. The problem is that the agent's degree of belief often lacks the hyper-precise upper and lower bounds required by imprecise Bayesianism. Again, considering Sturgeon's example, there is no non-arbitrary way of representing your strength of belief in proposition T: "It will be sunny in Tucson tomorrow" by a range with some lower and upper bounds. For instance, is it more correct to say that $P(T) = [0.9, 0.99]$ rather than $P(T) = (0.9, 0.999]$, or $P(T) = [0.87, 0.96]$? Given the information present in Sturgeon's example, there is simply no way of

as to be irreconcilable with the truth of these first person statements about their own beliefs, I cannot see how we can do other than regard these statements as the best evidence we have about their beliefs.

⁴⁷ Within philosophy, the idea of imprecise probabilities has been developed and defended by Levi (1974, 1985), Kyburg (1983), Joyce (2011), Sturgeon (2020), among many others.

choosing one unique range-credence representing your high confidence in T . Therefore, imprecise Probabilism cannot fully mitigate the worry with Doxastic Precision.

A natural conclusion to draw from this is that not all degrees of belief are Bayesian credences; as neither the point-valued nor range-valued Bayesian models are generally adequate.

Fortunately, even if the above conclusion is correct, it would not affect the main arguments of this dissertation. While I will use Bayesian precise probabilities to develop a hybrid account of doxastic rationality, none of my arguments presupposes Doxastic Precision. Instead, as we will see in Chapter 5, our hybrid approach to doxastic rationality will be built on a principle about *doxastic rankings of worlds* (or basic propositions), and such doxastic rankings do not presuppose that the worlds or basic propositions are ordered on the numerical scale. And many uses of the point-valued probabilities in this dissertation are for the sake of simplicity and convenience and can be discarded, if necessary.

Now that we have set aside the worry about the hyper-precise credences, next, I shall critically discuss Conditionalisation as a general rule for revising credences.

3.4.2 Conditionalisation and Conceptual Innovations

As we've discussed, the diachronic part of the standard Bayesianism consists of one norm:

Conditionalisation

For any agent with credence function P and any proposition X over W , if the agent learns that some proposition E over W is true, then her new credence function should be P_E where:

$$P_E(H) = P(H|E)$$

Our statement of Conditionalisation makes it clear that Conditionalisation only applies to the cases where an agent already has the probability for both H and E , prior to

learning that E . Thus, if the agent's new evidence E contains some new concepts or theories that the agent was unaware of before learning that E , then her prior probability of $P(H|E)$ (over her (prior) set of possibilities W) is undefined. In other words, E may not be a proposition over W ; hence neither $P(H|E)$ nor $P(E)$ is well-defined. Similarly, an agent may learn about a new hypothesis H_{New} which she previously has not entertained, hence her (prior) conditional probability $P(H_{New}|E)$ is not well-defined. So, Conditionalisation does not cover *conceptual innovation* (Sober 2008, 12), which is an essential part of many kinds of learning in both scientific and ordinary contexts.⁴⁸

As an example, consider a pre-Darwinian biologist who never entertained the theory of natural selection (NT). Suppose she is interested in accounting for evidence E that biological organisms are adapted to their environment (meaning they are “built” to survive and reproduce in their environment). Suppose this biologist entertains just two specific hypotheses to account for the evidence:

God: Biological adaptation is due to divine design (i.e., God consciously designed biological organisms).

Chance: Biological adaption is due to chance.

She has priors for *God* and *Chance* that sum to 1 (so she treats these hypotheses as mutually exhaustive). But now suppose that this biologist reads Darwin's book on natural selection and starts considering NT as a potential explanation of the evidence, E . Now, if she is a good Bayesian, her credence in NT should be non-negative (otherwise, E won't have an impact on NT). But if the biologist's prior in NT is non-negative, she would violate Probabilism: as the sum of her priors in all three hypotheses would exceed unity.

⁴⁸ This problem for Conditionalisation is sometimes called the *problem of new theories* (see Earman 1992, Chapter 5). There is also the closely related *problem of old evidence* (Eells 1985).

A standard way out of this difficulty is to stipulate that the agent should always leave some probability for a *catchall hypothesis*, which says that none of the currently available specific hypotheses is true.⁴⁹ In this example, call the catchall hypothesis *None*: “neither *God* nor *Chance* is true”. This obvious suggestion is not satisfactory. Even if the agent leaves some probability to *None*, why think that $P(NS)$ should be derived from *None*? After all, before learning about *NT*, the biologist may have been pessimistic of any explanation of *E* besides *God* and *Chance* and assigning a very small prior to *None*; but now, after reading Darwin’s book, she may well change her mind and see *NT* as much more plausible than, say, *Chance*. There is nothing wrong with such a change of mind; and if Bayesianism finds changing one’s mind problematic, then so much worse for Bayesianism. So, simply reserving probabilities for some possible hypotheses that may be articulated in the future does not solve the problem.

Instead, I endorse a less orthodox approach to the problem of conceptual innovations (which is similar to the one suggested by Strevens 2017, 11.4). I understand conceptual learning as requiring an agent to change her old set of possibilities, W , to a new set W_{New} that fully represents her current total evidence and adopt a new credence function P_{New} over W_{New} . What should the relationship between P and P_{New} be? This depends on the particular case in question. For instance, consider our biologist again: if she thinks that $P(NT)$ is, say, 0.1, then she can change her credence in *God* and *Chance* by multiplying them with the same constant so that $P(God) + P(Chance) = 1 - P(NT) = 0.9$. Interestingly, this procedure

⁴⁹ Another response is to appeal to an agent’s *counterfactual credence function*: the function she would have had prior to learning that *FT*, on the supposition that her stock of concepts (or propositions) involves the possibility that *FT*. Most Bayesians are suspicious of the appeals to counterfactual credence functions (see Earman 1992, Chapter 5, for a classical discussion), so I don’t consider it here.

follows Jeffrey Conditionalisation (Jeffrey Conditionalisation is discussed in detail in Section 5.3.3):

$$P_{New}(God) = P_{New}(NT) * P(God|NT) + (1 - P_{New}(God)) * P(God|God \cup Chance)$$

This approach has a very nice formal feature: if we assume that a new hypothesis H_{New} is incompatible with some old competing hypothesis H , then the above procedure simplifies to:

$$P_{New}(H) = P(H) * (1 - P_{New}(H_{New}))$$

Therefore, if we assume that $P(God)$ is, say, 0.7, and $P(Chance)$ is 0.3, then:

$$P_{New}(God) = 0.7 * 0.9 = 0.63$$

$$P_{New}(Chane) = 0.3 * 0.9 = 0.27$$

So, $P_{New}(God \cup Chance \cup NT) = 0.63 + 0.27 + 0.1 = 1$. Conclusion: there exists a mechanical procedure to derive P to P_{New} via Jeffrey Conditionalisation, which always satisfies Probabilism. But this mechanical procedure depends on an agent's new probability in the newly learned hypothesis, H_{New} ; and there may well not be any precise, mechanical procedure for determining the value of $P_{New}(H_{New})$ from P .

So, overall, the problem involving conceptual evidence could be dealt with within a broadly Bayesian framework. But this broadly Bayesian framework is importantly different from, what we may call, the *clockwork Bayesianism* (Strevens 2017, 125), which endorses an inadequate conception of credal revision or learning: where an agent has a *fixed* stock of concepts or propositions, represented by her set of all possibilities W and she updates her opinions in an orderly, formally precise way by learning certain propositions from W . Instead, contra to the clockwork Bayesianism, we must accept that in many cases, new evidence requires the agent to revise her old set of possibilities and change her priors to “make room” for new hypotheses which she was unaware of before. Such revisions may not follow any pre-established mechanical procedure and may require the agent to fix new

probabilities “by hand”, so to speak, in a way that best fits her currently available evidence (including a set of competing hypotheses available to her).

While we won’t be discussing the cases of conceptual learning in this dissertation in any detail, it is worth pointing out that such cases could be dealt with in a way that I’ve sketched above.

3.5 The Relationship between Belief and Credence

In section 3.2, we have discussed the standard logical account of belief, according to which belief is a tripartite attitude towards a proposition and should satisfy the requirement of Deductive Cogency. And in Section 3.3, we have discussed the probabilistic account of degrees of belief according to which rational degrees of belief, or credences, are numerically graded attitudes towards a proposition and should satisfy the standard axioms of probability theory and the principle of conditionalisation.

While beliefs and credences have been discussed independently from each other, they are too similar for not considering the question of their relationship. After all, both an agent’s beliefs and credences comprise the agent’s representation of the world. In Anscombe’s (1963) terms, both have the *mind-to-world* direction of fit. Belief aims at representing the world truly, while credence aims at representing the world accurately. Besides the directedness towards truth and accuracy, rational beliefs and credences share other similarities: both seem to be constrained by evidence (e.g., an agent’s beliefs and credences, if rational, must be supported by the agent’s evidence), and both play the similar role in *decision making* (i.e., both an agent’s beliefs and credence, when paired with the agent’s desires or her utility function, commit the agent to certain actions).

So, how should we understand the relationship between belief and credence? This question has two prominent readings: *metaphysical* and *normative*. The metaphysical reading of the question is essentially this:

Metaphysical Question: On the supposition that beliefs and credences are both real mental states, are they, at bottom, the same or different mental states?

There are three basic alternative answers to this question:

- (1) *The credence-as-belief theory:* Belief is the fundamental doxastic attitude, and credence is just a type of belief; i.e., credence is reducible to belief.
- (2) *The belief-as-credence theory:* Credence is the fundamental doxastic attitude, and belief is just a type of credence; i.e., belief is reducible to credence.
- (3) *Belief-credence dualism:* Both belief and credence are fundamental doxastic attitudes; i.e., neither is reducible to the other.

While this metaphysical question is certainly interesting, it has no clear bearing on the main topics of this dissertation. Independent of which of these answers is correct, a hybrid theory of doxastic rationality can still be correct. To see this, suppose that answer (1), the credence-as-belief theory, is correct: hence, to have a credence c in H is nothing but to have the belief that the probability of H is c . Even if the credence-as-belief theory is true, an agent's evidence can still rationally require to believe that H , but permit to believe or disbelieve that the probability of H is c . Hence, if the credence-as-belief theory is correct, the requirements of evidence would be permissive for beliefs with *numerically graded contents* and impermissive with *non-graded contents*.

Or take a much more plausible answer to the metaphysical question, the belief-as-credence theory, according to which belief that H is nothing but sufficiently high credence in H . If this theory is true, then the evidence could rationally permit believing H and could

permit any credence in H which is sufficiently high (what count as “sufficiently high” would depend on a version of the belief-as-credence theory).

And trivially, belief-credence dualism is compatible with a hybrid account of doxastic rationality.

So, we can set aside the metaphysical side of the belief-credence relationship and focus on the following normative question:

Normative Question: On the supposition that an agent has both beliefs and credences, how should her beliefs and credences interact (or cohere) with each other?

This simple-looking question has proved to be one of the most difficult questions of contemporary epistemology. The problem is that it proved to be tremendously hard to provide a plausible theory of belief-credence relationship that can satisfy the two fundamental requirements on belief and credence that we have considered in Sections 3.2 and 3.3: Deductive Cogency and Probabilism.

To illustrate the problem, consider what seems to be the most obvious normative principle that imposes joint constraints on rational belief and credence: the so-called *Lockean thesis* (see Foley 1993, 140-41).

The Lockean thesis: There is threshold r greater than $1/2$ and less than 1 , and, for all agents, it is rational to believe X iff $P(X) \geq r$.

This Lockean thesis does not *define* or *reduce* belief to credence or vice versa. Instead, it states the joint normative constraints on these doxastic attitudes. Following familiar usage, I call any such principle imposing joint constraints on belief and credence a *bridge principle*.

Unfortunately, this Lockean thesis conflicts with Deductive Cogency and Probabilism. This can be illustrated by the famous Lottery Paradox (Kyburg 1961). By the Lockean thesis, threshold r should be less than 1 (and greater than $1/2$). Now, consider a fair

lottery consisting of more than $1/(1 - r)$ tickets, where *exactly one* ticket will win; the lottery is fair in the sense that each ticket is equally likely to win. So, if $r = 0.9$, then we will consider a fair lottery consisting of 11 tickets (as $11 > 1/(1 - 0.9)$), where the probability of each ticket winning is $1/11$. By the axioms of probability, the probability of each ticket losing is $10/11$; and as $10/11 > 0.9$, the Lockean Thesis entails that it is rational to believe of each ticket that it will lose. And, by assumption, it is also rational to believe that *no ticket* will win. However, as the probability that *some* ticket will win is 1, we get the contradiction: it is rational to believe that *no ticket* will win, and *some* ticket will win.

In response to the Lottery Paradox (and other similar style paradoxes like the Preface Paradox), many epistemologists have abandoned either Deductive Cogency or the very idea behind the Lockean thesis that belief corresponds to sufficiently high credence in a believed proposition.⁵⁰ For instance, Foley (1993, 2009), following Kyburg (1970), has argued that one of the key norms of Deductive Cogency, the closure of belief under conjunction, must be rejected to avoid the absurd consequences from the Lottery Paradox and other similar style paradoxes. As he (2009., 42) explains:

... any theory of rational belief must either reject the conjunction rule or face absurd consequences. I conclude that we ought to reject the conjunction rule, which in any event lacks initial plausibility. After all, a conjunction can be no more probable than its individual conjuncts, and often is considerably less probable.

I will return to the Lottery Paradox in the subsequent chapters. As we shall see in Chapter 5 (Section 5.2.1), there is a plausible way of avoiding the Lottery Paradox without abandoning either Deductive Cogency or the key idea behind the Lockean thesis.

⁵⁰ See Leitgeb (2017, Section 1.7) for a detailed discussion and references.

3.6 Summary

In this chapter, I have laid the conceptual and formal groundwork for developing a hybrid theory of doxastic rationality. I've discussed the standard logical account of belief, the Bayesian account of degrees of belief, and the relationship between the two.

The logical account of belief is defined by the norm of Deductive Cogency, according to which rational beliefs are consistent and closed under logical consequence. Deductive Cogency entails that a rational agent's system of beliefs can be recovered by a single proposition, B_W – the strongest proposition believed by the agent (as articulated by Theorem 3.1). And the Bayesian account of credence is defined by the norms of Probabilism, according to which rational degrees of belief are mathematically graded and obey the standard laws of probability theory. Probabilism entails that a rational agent's complete credence distribution can be recovered by the credence distribution over the basic propositions that the agent considers (as articulated by Theorem 3.2). Both theorems will come in handy in developing precise hybrid accounts of doxastic rationality in Chapters 5 and 6.

In the next chapter, Chapter 4, I will put forward the first, simple hybrid theory of doxastic rationality and present a general problem that any hybrid theory faces. The presented problem will preoccupy us in Chapters 5 and 6.

4 Hybrid Theories and the Coordination Problem

According to the terminology used in this dissertation, a hybrid theory of doxastic rationality consists of (i) Uniqueness about one (doxastic) attitude-type, (ii) Permissivism about some other attitude-type, and (iii) a bridge principle connecting these attitude-types. All hybrid theories that we shall study in this dissertation endorse Hybrid Impermissivism, which is a conjunction of two theses:

Moderate Uniqueness: For any hypothesis H and evidence E , it is not the case that E rationally permits belief that H and belief that $\neg H$.

Credal Permissivism: For some evidence, E , and proposition, H , E rationally permits more than one credence towards H .

Hence, the sole difference between the considered hybrid theories concerns the bridge principles for belief and credence.

In this chapter, I will study a simple, precise hybrid theory, which endorses a version of the well-known bridge principle called the Lockean thesis: the view that categorical belief corresponds to a degree of belief that reaches a certain threshold. The version of the Lockean thesis that this hybrid theory endorses is logically weaker than the standard Lockean thesis I've introduced in the previous chapter (in Section 3.5).

As I shall explain in detail, the presented hybrid theory is attractive in several respects: it captures the plausible aspects of Uniqueness and Permissivism and avoids some important problems associated with them. However, we shall also see that the theory faces a special problem that has to do with coordinating an agent's beliefs and credences over time. I call this the *diachronic coordination problem*. The overall conclusion of this chapter is that, the success of any hybrid theory of doxastic rationality is conditional on solving the diachronic coordination problem.

The chapter is structured as follows: in section 4.1, I put forward the first precise hybrid theory, which I call the *Lockean Hybrid Impermissivism* (L_{HI} , for short): it combines (i) a weak version of Belief Uniqueness, (ii) Credal Permissivism, and (iii) a version of the so-called *Lockean thesis* – the bridge principle connecting an agent’s beliefs and credences. In Section 4.2, I will motivate L_{HI} by arguing that it avoids some of the standard objections against impermissivist as well as permissivist epistemologies: on the impermissivist side, L_{HI} avoids the *epistemic standards objection* and on the permissivist side – the *arbitrariness objection*. In section 4.3, I identify the coordination problem for L_{HI} which shows that L_{HI} , as it stands, cannot coordinate beliefs and credences over time without violating either Moderate Uniqueness or Credal Permissivism. In Section 4.4, I consider two initial approaches to the coordination problem and argue that they are unsuccessful.

4.1 The Lockean Hybrid Impermissivism

The distinguishing feature of the Lockean Hybrid Impermissivism (L_{HI}), differentiating it from all the other hybrid theories we will consider, is the bridge principle that it endorses. As its name suggests, this bridge principle is related to the Lockean thesis, which roughly says that categorical belief corresponds to a degree of belief that reaches a certain threshold r , where r is greater than $1/2$ and at most 1. Following Leitgeb (2014, 2017), We distinguish two versions of the Lockean thesis:

Strong Lockean Thesis (SLT): There is a threshold r greater than $1/2$ and at most 1, and any agent’s belief set Bel and credence function P should be such that, $Bel(X)$ iff $P(X) \geq r$.

Weak Lockean Thesis (WLT): An agent’s belief set Bel and credence function P should be such that, there a threshold r greater than $1/2$ and at most 1, and $Bel(X)$ iff $P(X) \geq r$.

SLT is equivalent to the Lockean thesis discussed in Section 3.5. It says that there exists one unique threshold r , and all agents should believe a proposition iff their credence in that proposition is equal to or above that threshold.

By contrast, on WLT (which has been endorsed by Leitgeb 2014, 2017), there need not be a unique universally correct Lockean threshold for every agent and context of reasoning. Instead, WLT permits an agent's Lockean threshold to change, depending on her evidential situation and other agent-relative factors.

Why favour WLT over SLT? Overall, the choice between SLT and WLT is inconsequential to the main arguments of this chapter. However, the strength and weakness of hybrid theories are best illustrated by the logically weaker version of the Lockean thesis, WLT. There are also general reasons for favouring WLT over SLT. For a start, assuming that the threshold r is less than 1, SLT (unlike WLT) conflicts with our two main assumptions about rational belief and credence: Deductive Cogency and Probabilism. This is illustrated by the Lottery Paradox that we have considered in Section 3.5. As, on SLT, for any value r where $r < 1$, we can construct a fair lottery setting where the probability of each ticket losing is greater or equal to r (simply construct a fair lottery consisting of more than $1/(1 - r)$ tickets). For instance, suppose that the universally correct Lockean threshold $r = 0.999$. Now, construct a fair lottery consisting of, say, 1000 tickets. The probability of each ticket losing is 0.999. Hence, in such settings, SLT licenses the belief about each individual ticket that it will lose and also the belief that some ticket will win; and these beliefs violate Deductive Cogency.

By contrast, on WLT, there is no threshold r that must be applied in all contexts of reasoning. So, if an agent considers a fair lottery consisting of 1000 tickets, then the agent's threshold r can be greater than 999/1000. With this higher Lockean threshold, the agent will

only believe the tautological proposition that some ticket will win. Hence both Deductive Cogency and Probabilism will be satisfied. And, in general, with respect to any set of propositions that the agent considers, it is always possible to choose a Lockean threshold (less than 1) that won't violate Deductive Cogency.

While WLT does not violate Deductive Cogency, it is completely silent on which Lockean thresholds are permissible in any given evidential situation. So, WLT does not provide a full account of how rational belief and credence ought to interact with each other. Fortunately, as we will see in the next chapter, it is possible to transform WLT into a precise bridge principle for belief and credence. And, at this point, WLT is fully adequate as is. We will simply choose the appropriate Lockean thresholds on a case-by-case basis, in such a way that neither Deductive Cogency nor any other plausible norm of rational belief is violated.

Now, even bracketing the worries about the Lottery Paradox, SLT is very implausible for another reason. It is unclear why some unique Lockean threshold is valid for all agents in all possible epistemic circumstances. To see that this requirement is very implausible, consider the following situation. Suppose I sincerely believe that I will send a postcard to a friend in the next few days. But, I'm not absolutely certain that I'll do so; my credence in the proposition is somewhere around 0.75. Now, does this mean that I'm rationally required to believe every proposition that I deem to be more probable than 0.75? This seems very implausible as there might be epistemic situations where my threshold for rational belief is considerably higher. For instance, I would be unwilling to believe that my lottery ticket won't win if the probability of losing is only 0.75. So, it seems that one's standard of rational belief can be different in different circumstances. In some evidential situations, it may be rational to believe a proposition whose probability is slightly greater than 0.5, but in some other situations, the probabilistic threshold for belief may be very close to 1.

Hence, as WLT is logically weaker and considerably more plausible than SLT, it will be our favoured interpretation of the Lockean thesis.⁵¹ So, from now on, I'll sometimes refer to WLT as the Lockean thesis, without the modifier "Weak".

Now, we are ready to precisely state our first hybrid theory, the *Lockean Hybrid Impermissivism* (L_{HI} , for short),

L_{HI} : Moderate Uniqueness and Credal Permissivism are true, and beliefs and credences are coordinated via the Lockean Thesis (with a context-dependent threshold r).

At first blush, L_{HI} seems a rather plausible and moderate view. It combines two appealing versions of Uniqueness and Permissivism: Moderate Uniqueness and Credal Permissivism. Moreover, as I will show next, L_{HI} avoids some of the standard objections against both impermissivist and permissivist epistemologies: the epistemic standards objection (discussed in Section 4.2.1) and the arbitrariness objection (discussed in Section 4.2.2). The positive features of L_{HI} generalise to more sophisticated hybrid theories that I consider and motivate in this dissertation (in Chapters 5 and 6).

4.2 The Lockean Hybrid Impermissivism and the Uniqueness Debate

4.2.1 The Epistemic Standards Objection Reconsidered

Among all published criticisms of Uniqueness, the ones that appeal to the notion of epistemic standards have been the most popular. At the heart of these criticisms is the following idea that I will call *Permissivism's Core*:

⁵¹ We shall also see in the next chapter that WLT is a logical consequence of a plausible theory of belief-credence interaction, the *stability theory*.

Permissivism's Core: Agents can rationally hold different epistemic standards which influence what they are permitted to believe and to what degree.

There have been a couple of different interpretations of Permissivism's Core. In this section, I'll illustrate how L_{HI} can capture one of the most popular and well-articulated interpretations of Permissivism's Core, which I call *Jamesian Permissivism* (because this interpretation is inspired by the ideas proposed by William James in "The Will to Believe" 1896):

Jamesian Permissivism: Different agents can rationally attach different values to the two fundamental epistemic goals, "Believe Truth!" and "Avoid Error!" and how much value the agents attach to these goals influences what the agents are permitted to believe.

Within the Uniqueness debate, the most detailed discussion and defence of Jamesian Permissivism is due to Kelly (2014).⁵² As he writes (ibid., 301):

... the more value one gives to not believing what's false about some issue, the more it behooves one to be relatively cautious or conservative in forming beliefs about that issue. That is, the more weight one gives to not believing something false, the more it makes sense to hold out until there is a great deal of evidence that H is true before taking up the belief that H . On the other hand, the more one values not missing out on believing the truth, the more it makes sense to take a somewhat more liberal attitude about how much evidence one expects before taking up the relevant belief.

So, on this Jamesian view, two equally informed people can rationally adopt different doxastic attitudes towards a proposition, depending on how much value they attach to the

⁵² Many other permissivists, e.g., Schoenfield (2014), Peels and Booth (2014), Titelbaum (forthcoming, Section 4.3), also positively discussed the ideas similar to Jamesian Permissivism.

epistemic goals of “Believe Truth” and “Avoid Error”. Kelly (ibid., 301-302) illustrates and motivates this Jamesian Permissivism by analysing the following example:

Suppose that the evidence that you and I have that bears on some hypothesis H is E . Although it’s clear enough that E supports H over not- H , it’s not as though E is overwhelming evidence that H is true. Indeed, let’s suppose that this is a marginal case, in that E is just barely sufficient to justify believing H ... Recognizing that E suffices to justify belief in H , I take up the belief in response. I notice, however, that you don’t take up the same belief, despite having the same evidence. ...

In these circumstances, am I committed to thinking that you’re guilty of making some kind of mistake, that you’ve misjudged the probative force of our shared evidence? Before attempting to answer these questions, let’s add one further detail to the story. With respect to the question at hand, you’re a bit more concerned than I am to avoid believing what’s false, while I’m a bit more concerned than you are to not miss out on believing what’s true in virtue of suspending judgment. That is, there is a subtle difference in our cognitive goals, or rather, in the relative weights that we give to the two cognitive goals with respect to the question at hand.

Once this further stipulation is added, your not believing H on the basis of evidence that is only marginally sufficient to justify such belief seems eminently reasonable. ... The upshot: 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.

Now, Kelly’s example and the conclusion he reaches are wholly compatible with L_{HI} . Here is why. On L_{HI} , an agent’s Lockean threshold r is not universally fixed and depends on various context-sensitive factors. And if we accept Jamesian Permissivism, then one of the factors determining the agent’s Lockean threshold can be the weights the agent gives to the fundamental epistemic goals. For instance, if an agent attaches more value to the goal “Believe truth!” over the goal “Avoid error!”, her Lockean threshold can be low, say a bit above 0.5. By contrast, if the agent attaches more value to the “Avoid error!” goal, her Lockean threshold can be higher, say 0.8. So, suppose that in Kelly’s example, both you and I agree that the probability of H is roughly 0.7. As I give more weight to the “Believe truth!”

goal, my Lockean threshold can be a bit above 0.5. By contrast, as you give more weight to the “Avoid error!” goal, your Lockean threshold can be considerably higher, say 0.8. So, given our different Lockean thresholds, I am rationally required to believe H while you – rationally required to withhold judgement.

So, our hybrid theory, L_{HI} , can easily accommodate the Jamesian view that epistemic goals influence what the agents ought to believe. And as we will see in the subsequent chapters, this conclusion about the compatibility of Jamesian Permissivism and L_{HI} generalises to all hybrid theories that we will consider in this dissertation. We will also show how different epistemic goals influence an agent’s *entire system of beliefs*, and not just isolated, individual beliefs (as in Kelly’s example).

4.2.2 The Arbitrariness Argument Reconsidered

As I have already discussed in section 1.2.1, White’s (2005) arbitrariness argument against Permissivism is the most widely debated argument within the literature. The argument can be stated succinctly as follows: if you think that, say, your evidence equally permits belief that H and belief that $\neg H$, then believing either H or $\neg H$ is arbitrary. Hence, Permissivism entails that it is rationally permissible to form beliefs arbitrarily. But this consequence of Permissivism seems absurd; Therefore, Permissivism must be false.

In a premise-conclusion form, the argument can be stated as follows:

- (1) Assumption: Let e be some permissive evidence.
- (2) If an agent’s (total) evidence is e , then the agent cannot determine what to believe by examining the evidence alone.
- (3) If the agent cannot determine what to believe by examining the evidence alone, then it can be rationally permissible for the agent to form beliefs arbitrarily.
- (4) It is not rationally permissible to form beliefs arbitrarily.

Therefore:

(5) Evidence cannot be permissive.

White's arbitrariness argument is quite general and aims to show that all versions of Permissivism, including Credal Permissivism, are false. Hence, the arbitrariness argument conflicts with L_{HI} and any hybrid theory in general that endorses some version of Credal Permissivism.

In Section 1.2.1, we have seen that the most popular response to the arbitrariness objection appeals to an agent's epistemic standards that, together with the evidence, fully determine what the agent is rationally required to believe. So, on this response, the third premise of White's argument is false: even if an agent's evidence does not fully determine what she ought to believe, this does not entail that the agent is permitted to form beliefs arbitrarily. This is so because the agent's evidence, together with her epistemic standards, determines what she ought to believe in any given evidential situation.

We have also seen that this response from epistemic standards faces its own arbitrariness problem. The worry is well-summarised by Simpson (2017, 529):

Where before our problem was arbitrarily favoring one among two rationally permissible doxastic attitudes, now our problem is arbitrarily favoring one among two rationally permissible sets of epistemic standards.

In this section, I articulate a response to this problem that fits well with our hybrid approach to doxastic rationality in general and with L_{HI} in particular. This response grants that the choice (or favouring) between different standards is arbitrary but argues that such arbitrariness is unproblematic. When I use the term "choice" here, I do not want to suggest that an agent deliberately chooses (or is even capable of choosing) one epistemic standard over the other. Take the Jamesian pragmatist view again. Relative weights an agent gives to the two fundamental epistemic goals, "Believe Truth" and "Avoid Error", is one popular way

of characterising the agent's epistemic standards. Call an agent *epistemically brave* if she attaches greater weight to the Believe Truth norm and *epistemically cautious* if she attaches greater weight to the Avoid Error norm. Certainly, an agent can be epistemically brave without ever deliberately choosing to be so; as her level of epistemic braveness may be an outcome of her early education, academic training or overall environment. Still, given that different levels of epistemic braveness could be equally permissible for a single agent, the influences on belief (or credence) due to this agent's specific level of epistemic braveness count as arbitrary.

Now, my response to White's argument rejects its fourth premise, (4), and contends that rational belief may depend on some arbitrary factors. But because on L_{HI} arbitrary factors may have a relatively *minor* influence on what we ought to believe – they could only make the difference between believing and suspending judgement on a proposition – the arbitrariness argument loses much of its appeal against L_{HI} . Let me elaborate on this.

Suppose you and I agree that God's existence is very unlikely on our shared evidence. But, unlike you, I have interests in evaluating new arguments for or against God's existence. And I think that a fair evaluation of these arguments requires me to have some non-negligible credence in God's existence. For this reason, I'm roughly 0.05 confident that God exists. By contrast, you may have no interest in wasting more of your time thinking about God's existence. Hence, the God question is closed for you, and your credence in God is close to 0.

Whether we are interested in the God question has no bearing on God's existence. Hence, both of our credences in God are influenced by arbitrary factors. But are such arbitrary influences worrisome? I submit that there are not. On the supposition that there is no non-subjective basis for deciding what exact credence we should have in God, it does not seem irrational to align our credences with our epistemic interests and goals. After all, even if

our interests and goals are different, we share the same core belief that God's existence is very improbable. For this reason, a relatively minor difference between our credences should not worry us at all.

But if such minor arbitrary influences on belief and credences are not problematic, then the arbitrariness objection is not effective against L_{HI} . Here is why. L_{HI} endorses both Credal Permissivism and Moderate Uniqueness, together with the (Weak) Lockean thesis. This means that according to L_{HI} , the evidence puts significant constraints on rational credences: for instance, no matter what evidence an agent has, it cannot be permissible for the agent to be confident in H and confident in $\neg H$. This is so because if evidence permitted radically different confidences in H , say any credence in the interval $[0.1, 0.9]$, then it would be rationally permissible both to believe H and its negation; and this contradicts Moderate Uniqueness. Hence, on L_{HI} arbitrary factors have a relatively minor influence on belief and credence. And as the above God example aims to illustrate, in cases where the available evidence does not fully determine what to believe, minor arbitrary influences on belief are not as problematic as White's argument suggests.

After all, arbitrariness is a matter of degree; and not all degree of arbitrariness is equally problematic. For instance, both subjective Bayesianism and modern objective Bayesianism accept that rational credence involves arbitrary choices. But it is misleading to consider these theories in the same category as far as arbitrariness is concerned. After all, the degree to which rational credence is arbitrary on subjective Bayesianism is far greater than on objective Bayesianism. As Williamson (2010, 158), an objective Bayesian, puts it:

... all the [modern] Bayesian positions—strict subjectivism, empirically based subjective probability and objective Bayesianism— accept the fact that selection of degrees of belief can be a matter of arbitrary choice, they just draw the line in different places as to the extent of subjectivity. Strict subjectivists allow most choice, drawing the line at infringements of the axioms of probability.

Proponents of empirically based subjective probability occupy a halfway house, allowing extensive choice but insisting that evidence of physical probabilities as well as the axioms of probability constrain degrees of belief. Objective Bayesians go furthest by also using equivocation considerations to narrow down the class of acceptable degrees of belief.

So, as we see, both subjective and objective Bayesianans accept that arbitrary factors influence rational credences. However, it is bizarre to claim that these epistemologies are equally susceptible to White's argument. As, on objective Bayesianism, arbitrary factors have considerably less influence on rational credence than on subjective Bayesianism.

Therefore, I conclude that White's arbitrariness argument is ineffective against epistemologies that only permit minor arbitrary influences on belief.

4.3 The Coordination Problem for Hybrid Impermissivism

I have argued that L_{HI} is a plausible, moderate view that avoids the epistemic standards objection against Uniqueness as well as the arbitrariness objection against Permissivism. Unfortunately, though, L_{HI} runs into a problem once we start looking at the view from the *diachronic* point of view; that is, once we look at how an agent's beliefs and credences change over time, due to learning new information. I call this new problem the *diachronic coordination problem* (coordination problem, for short). As we shall see, the coordination problem is a general problem for hybrid theories of doxastic rationality and is not restricted to L_{HI} .

First, I will illustrate the coordination problem in a relatively informal way, by considering an example related to the hotly debated topic of cosmological fine-tuning. After discussing this example, I'll state the coordination problem within a more formal setting (Section 4.3.1), where we assume complete probability distributions over a set of propositions.

So, here is an example:

Fine-Tuning

Cathy and Julien are colleagues who often discuss various topics in philosophy and religion. On Monday, they had a lengthy discussion about the existence of God. They both concluded that, on the available evidence, it is rational to believe that the God of traditional theism does not exist. Now, while their categorical attitudes about God's existence are the same, Julien is more confident that God does not exist than Cathy is. Their levels of confidence can be represented as follows, where "God" denotes the proposition that God exists:

$$P_{Cathy}(God) = 0.1$$

$$P_{Julien}(God) = 0.02$$

On Sunday, Cathy and Julien meet each other again to discuss a recent paper about the fine-tuning argument for the existence of God. The paper argues that the new evidence that the so-called cosmological constants are *finely tuned* supports the hypothesis that God exists.⁵³ Somehow, both Cathy and Julien are convinced that fine-tuning provides strong evidence for God. The paper estimates that the fine-tuning

⁵³ The basic idea behind the fine-tuning argument is as follows: according to contemporary physics, the fact that life exists in the universe depends on the very precise values that the so-called fundamental constants of physics take. For instance, if the mass of proton had been slightly different from its actual value, then the complex structures that we find in the universe would not have existed; and hence, life would not have existed. But, life does exist in our universe. So, some think that the fine-tuning evidence speaks in favour of the God hypothesis.

data is approximately 25 times more likely on the supposition that God exists than on the supposition that God does not exist.⁵⁴ In symbols:

$$\frac{P(FT|God)}{P(FT|\neg God)} = 25$$

Cathy and Julien think that this estimate is correct.

Now, suppose that on Monday, Cathy and Julien are rational in believing that God does not exist. Further, suppose that their corresponding credences are also rational. For simplicity, let's assume that their beliefs and credences are related via the Lockean Thesis with a threshold of 0.6 (but, as I explain shortly, the choice of this threshold is inconsequential to my overall argument). So, it is rational for Cathy and Julien to have the following combinations of beliefs and credences:

On Monday: $Bel_{Cathy}(\neg God)$ and $P_{Cathy}(God) = 0.1$

$Bel_{Julien}(\neg God)$ and $P_{Julien}(God) = 0.02$

Now, let's suppose that after they receive and analyse the new evidence on Sunday, Cathy and Julien are rational to believe that fine-tuning is approximately 25 times more likely on the supposition that God exists than on the supposition that God does not exist. Probabilities of the form $P(evidence|hypothesis)$ are called *likelihoods*. So, Cathy and Julien are rational in believing that the ratio of likelihoods is approximately 25. Now (as I've already discussed in Chapter 2, Section 2.2.3), there is a theorem of probability calculus that enables us to calculate the ratio of posterior probabilities, given the ratio of likelihoods and ratio of priors. The theorem is usually called the *ratio form of Bayes' theorem*:

⁵⁴ See Hawthorne and Isaacs (2018, 161) for a discussion on estimating the relevant likelihood ratio. It is utterly inconsequential whether the likelihood ratio estimate is correct. The example is taken solely for illustrative purposes.

$$\frac{P(H|E)}{P(\neg H|E)} = \frac{P(E|H)}{P(E|\neg H)} * \frac{P(H)}{P(\neg H)}$$

If we let R_{Post} be the ratio of *posteriors*, R_L the ratio of *likelihoods*, and R_{Prior} the ratio of *priors*, then the theorem can be summarised succinctly as:

$$R_{Post} = R_L * R_{Prior}$$

Now, in the fine-tuning example, we know the corresponding values of R_L and R_{Prior} for Cathy and Julien. And the simple calculations show that, upon learning the new information about fine-tuning (denoted as “FT”), Cathy’s and Julien’s posteriors in God’s existence should be:

$$P_{Cathy}(God|FT) \approx 0.73$$

$$P_{Julien}(God|FT) \approx 0.34$$

And via the Lockean Thesis with a threshold of 0.6, we conclude that $Bel_{Cathy}(God)$ and $Bel_{Julien}(\neg God)$. So, even if Cathy and Julien’s relevant beliefs are the same on Monday, and even if they receive the same evidence which they interpret in the same way, their corresponding beliefs on Sunday are conflicting. To summerise their beliefs:

On Monday: $Bel_{Cathy}(\neg God)$ and $P_{Cathy}(God) = 0.1$

$Bel_{Julien}(\neg God)$ and $P_{Julien}(God) = 0.02$

On Sunday: $Bel_{Cathy}(God)$ and $P_{Cathy}(God) \approx 0.73$

$Bel_{Julien}(\neg God)$ and $P_{Julien}(God) \approx 0.34$

What this example shows is that, even if two agents are rational in their beliefs and credences, and even if they receive the same body of evidence which they interpret in exactly the same way, their newly formed beliefs may still be different.

This is a serious problem for our hybrid theory, L_{HI} . The assumptions that we've made about Cathy and Julien's doxastic states on Monday do not contradict any postulates of L_{HI} . Cathy and Julien start with the same categorical beliefs; their credences and beliefs are related via the Lockean Thesis. However, upon learning new information, their categorical beliefs are mutually inconsistent: Cathy believes that God exists, while Julien retains his old belief that God does not exist. So, as it stands, two fully rational individuals who do not violate any postulates of L_{HI} can come to violate Moderate Uniqueness after learning new information.

One straightforward way of responding to the above problem is as follows: Cathy and Julien adopt opposing beliefs on Sunday only on the supposition that their Lockean thresholds are relatively low: given the fixed value of the ratio of likelihoods, $R_L = 25$, their respective Lockean thresholds should not exceed 0.66. So, one might think that the coordination problem may be avoided if we assume a higher Lockean threshold, say, 0.9.

But this suggestion is too quick. For a start, why think that Cathy and Julien are rationally required to have such high Lockean thresholds for belief in God? In general, it seems at least rationally permissible for agents to have relatively low standards for belief. So, some explanation needs to be given why their Lockean thresholds need to be higher in this case.

But even if there is some good reason for them to have a higher Lockean threshold, this still does not solve the problem. It is true that if we assume the Lockean threshold of 0.9, and if we also suppose that all permissible priors in God for Cathy and Julien are within some specific range, say $[0.07, 0.1]$, then there is no *single* value of the ratio of likelihoods, R_L , such that it will make Cathy and Julien adopt opposing beliefs about God's existence. But, the assumption that Cathy and Julien agree on the value of R_L is an inessential, simplifying

assumption that can be easily dropped. Within a more idealised setting, where the agents' credence functions are fully specified, we do not need to assume that the agents agree about the value of R_L to derive the same coordination problem. This is what I'm going to show next

4.3.1 The Coordination Problem in a Formal Setting

Consider set W of four possible worlds: $W = \{w_1, w_2, w_3, w_4\}$. We can think about these possible worlds as corresponding to all logical possibilities associated with the two propositions in our fine-tuning example:

God: God exists.

FT: Our universe is fine-tuned for life.

w_1 corresponds to $God \wedge FT$, w_2 to $God \wedge \neg FT$, w_3 to $\neg God \wedge FT$, and w_4 to $\neg God \wedge \neg FT$.

Now, let's define two probability distributions, P_1 and P_2 over W , represented by the table below:

Possible worlds	P_1	P_2
w_1	$P(\{w_1\}) = 0.04$	$P(\{w_1\}) = 0.09$
w_2	$P(\{w_2\}) = 0.03$	$P(\{w_2\}) = 0.01$
w_3	$P(\{w_3\}) = 0.37$	$P(\{w_3\}) = 0.01$
w_4	$P(\{w_4\}) = 0.56$	$P(\{w_4\}) = 0.89$

Table 4.1

As we see, God's existence is very improbable on these credence functions, $P_1(God) = 0.07$; $P_2(God) = 0.1$ and fall within the range $[0.07, 0.1]$.

Now, suppose that relative to these credence functions, we have the same high Lockean threshold of 0.9. In that case, we get the same belief set Bel relative to both P_1 and P_2 , which contains $\{w_3, w_4\}$ (i.e., proposition *God*) and everything that follows from $\{w_3, w_4\}$: as on both of these credence functions, $\{w_3, w_4\}$ has a probability greater than or

equal to 0.9, and, trivially, only propositions that follow from $\{w_4\}$ has a probability greater than 0.7.

Now, suppose that we revise these credence functions with new information, FT , or $\{w_1, w_3\}$. That is, we assume that $\{w_1, w_3\}$ is true and calculate new credences in *God* via Bayes' theorem:

$$P_1(God|FT) = \frac{P_1(God \text{ and } FT)}{P_1(FT)} = \frac{P_1(\{w_1\})}{P_1(\{w_1, w_3\})} = \frac{0.04}{0.41} \approx 0.09$$

$$P_2(God|FT) = \frac{P_2(\{w_1\})}{P_2(\{w_1, w_3\})} = \frac{0.09}{0.1} = 0.9$$

So, while prior to receiving the new evidence FT , God's existence was very improbable on these credence functions, $P_1(God) = 0.07$; $P_2(God) = 0.1$, the new information changes what it is rational to believe relative to these credence functions. Given the Lockean threshold of 0.9, it is rational to believe $\neg God$ relative to P_1 and believe *God* relative to P_2 .

In response to this example, one may suggest that if credence functions P_1 and P_2 were more similar, then this coordination problem could have been avoided. This suggestion raises two distinct questions. The first question is more formal or technical, and the second is more philosophical. On the formal side, we should ask:

Given some plausible bridge principle for belief and credence, is it possible to restrict a set of rationally permissible credence functions in a way that these credence functions never license opposing beliefs, no matter on which possible evidence we condition these credence functions?

If the answer to this question is negative, then, trivially, the coordination problem is fatal for Hybrid Impermissivism. But, even if the answer is positive, there is still an outstanding, more philosophical question that needs to be answered:

Is it possible to restrict a set of permissible credence functions in a way that Hybrid Impermissivism avoids the coordination problem but does not commit to an overly strong and implausible view about the evidential constraints on credence?

This second question concerns whether there is a *plausible* solution to the coordination problem; i.e., a solution that does not impose overly strong and implausible constraints on rational credences. After all, even if it is possible to impose constraints on a set of permissible credence functions that block the coordination problem (given some plausible bridge principle for belief and credence), these constraints may be overly strong and demanding. And in such a case, the version of Credal Permissivism that Hybrid Impermissivism would require (for avoiding the coordination problem) won't be significantly more plausible than Credal Uniqueness. This, in itself, would contradict one of the central guiding motivations behind Hybrid Impermissivism: to avoid overly strong and implausible views about credences. After all, if the version of Credal Permissivism that Hybrid Impermissivism requires is still an extremely strong and unobvious view, then Hybrid Impermissivism loses much of its original appeal.

In Chapter 5 (Section 5.3), I will argue that there is a formally precise way of restricting a set of permissible credence functions without endorsing an overly strong and implausible view about the evidential constraints on credence. So, I shall defend the positive answers to both the formal and philosophical questions raised by the coordination problem.

But, before I develop this approach to the coordination problem, the remainder of this chapter will propose and evaluate two different, more straightforward responses to the coordination problem. I will argue that both of these responses are unsuccessful.

4.4 Two Initial Responses to the Coordination Problem

Arguably, the most straightforward response to the (diachronic) coordination problem is as follows:

Response 1

In the fine-tuning example, Cathy or Julien are rationally required to *change* their old Lockean thresholds. So, for instance, if their new Lockean thresholds are 0.8, then both of them would be rationally required to suspend judgment on the existence of God.

This response to the coordination problem is compatible with our favoured version of the Lockean thesis. This is because the (weak) Lockean thesis does not fix one unique threshold that applies to all contexts.

But why should Cathy and Julien change their standards of categorical belief in this case? There seems to be no motivation behind this proposal except shielding our hybrid theory from the diachronic coordination problem. I should emphasise that there are cases where a change in an agent's Lockean threshold is permissible, if not rationally required. For instance, if an agent's old Lockean threshold license her to adopt internally inconsistent beliefs, then the agent might well be required to revise this threshold. But in the fine-tuning example, there is no internal inconsistency in either Cathy or Julien's beliefs. The conflict is entirely *interpersonal*: Cathy's belief in God conflicts with Julien's corresponding belief. Therefore, it is rather unconvincing to respond that Cathy should change her old standard of rational belief (i.e., her Lockean threshold) only because her new belief contradicts Julien's belief.

By contrast, the second initial response to the diachronic coordination problem rejects the Lockean thesis. This response can be summarised as follows:

Response 2

The diachronic coordination problem shows that the following three theses, under prima facie plausible assumptions, are in tension: Moderate Uniqueness, Credal Permissivism and the Lockean thesis. So something has to give. As the first two theses are non-negotiable for any Hybrid Impermissivist view, we should reject the Lockean thesis and endorse a more sophisticated bridge principle that would avoid the diachronic coordination problem.

One such bridge principle that avoids the diachronic coordination problem makes an agent's rational beliefs sensitive not just to the agent's credence function but also to a *set* of all permissible credence functions. I call this bridge principle the Interpersonal Lockean thesis:

The Interpersonal Lockean thesis (ILT): For all agents and propositions H , there is threshold r greater than $1/2$ and at most 1 , and $Bel(H)$ iff for *all* permissible credence functions P , $P(H) \geq r$.

Some clarifications are in order. According to ILT, an agent has a Lockean threshold r , and she categorically believes a proposition, H , iff for all rationally permissible credence functions P , $P(H) \geq r$. Now, ILT is equivalent to the Lockean thesis when there is a unique credence function that is justified by the evidence. But ILT and the Lockean thesis diverge when the available evidence permits more than one credence towards a proposition. For instance, suppose that with respect to an agent's credence function P , $P(H) \geq r$, where r is the agent's Lockean threshold; but, if, there is some other credence function P' such that, P' is rationally permissible on the agent's evidence and $P'(H) \leq r$, then the agent is no longer permitted to believe H .

If we substitute the Lockean thesis with ILT, then hybrid impermissivism becomes immune to the diachronic coordination problem. To illustrate this, let's reconsider the fine-tuning example. We would assume that both Cathy and Julien's Lockean thresholds do not change upon learning the new information. Now, on ILT, Cathy is justified to disbelieve God's existence on Monday, because relative to her and Julien's credence functions (both rationally permitted on the evidence), the probability of God's existence is greater than 0.5. By contrast, on Sunday, Cathy's credence in God's existence is greater than 0.5, while Julien's credence is below that threshold. Therefore, on ILT, Cathy is *not* rational in believing that God exists.

Notice that, ILT makes Moderate Uniqueness and Credal Permissivism coherent without introducing any change in Cathy's *epistemic standards* from Monday to Sunday. The same threshold fixes her standard of belief both on Monday and Sunday. And the reason why she is no longer rational to have her old belief on Sunday is not due to change in her standards, but due to change in her evidence: as, unlike Monday, on Sunday, she knows that evidence justifies credence in God's existence that is below the 0.5 threshold.

While ILT allows us to avoid the diachronic coordination problem for Hybrid Impermissivism, it does not seem to give a plausible account of rational belief. For a start, ILT treats an agent's categorical and credal doxastic states quite differently. On this view, it is permissible for an agent to adopt one among equally justified credal states, but her categorical beliefs should reflect *all* rational credence functions that the evidence justifies. Such a discrepancy between a rational agent's credal and categorical beliefs looks objectionable.

In defence of ILT, one might suggest that the sensitivity of categorical belief to all rationally permissible credence functions is a selling point and not a bug. One way to

motivate this suggestion is by citing different functional roles of categorical beliefs and credences. For instance, credences can be more prominent in decision making and hence more influenced by an agent's *practical interests* and other non-evidential factors. For instance, if failing to do some action, *A*, has great negative consequences for an agent, then the agent might attach lower credences to the propositions that minimise the *expected* disutility of not doing *A*. In our fine-tuning example, we can say that Cathy, who attaches higher credence in God's existence, is relatively more concerned about the possibility of not believing in God, if God exists. By contrast, Julien is less concerned about such a possibility. Hence, non-epistemic factors, such as, how much practical importance an agent attaches to a proposition, can influence the agent's credence. By contrast, one might claim that an agent's categorical beliefs should be *epistemically pure* and independent of any non-evidential factors. Hence, the combination of *pragmatism* about credences and *purism* about belief fits well with the Interpersonal Lockean thesis: so that, if an agent's categorical belief is sensitive to *all* rationally permissible credence functions, then her beliefs would be *shielded* from non-evidential influences that permeate any *single* credence function.⁵⁵

The above defence of ILT is unsuccessful. Even if we grant that beliefs and credences have different functional roles, ILT is still problematic. According to ILT, a rational agent's categorical beliefs should be as sensitive to the agent's credence function as to any other rationally permissible credence functions. And this is extremely odd. To illustrate the problem, consider an agent with evidence *E* and credence function *P*. Suppose *E* is permissive, and there is a credence function *P'* which is also rational on *E*. Now, on ILT, even if the agent adopts credence function *P* instead of *P'*, her categorical beliefs should be equally sensitive to *P* and *P'*. So, the fact that *P* is the agent's credence function does not

⁵⁵ For a discussion of Belief/Credal Pragmatism and Belief/Credal Purism see Jackosn (2019, Sections 4.1, 4.2).

entail that the agent gives *more weight* to P over P' , as far as her categorical belief are concerned. And this is an extremely odd feature of ILT. For this reason, I don't think that substituting the Lockean thesis with ILT provides a plausible solution to the coordination problem.⁵⁶

So, overall, neither of these initial strategies for responding to the coordination problem is successful.

4.5 Summary

In this chapter, we have defined the first precise, simple hybrid theory of doxastic rationality, L_{HI} . While L_{HI} is an attractive view in several respects, we have seen that it faces the (diachronic) *coordination problem*; the problem about how to coordinate beliefs and credences over time without violating the required combination of Uniqueness and Permissivism.

In this chapter, we have only reached negative conclusions with respect to the coordination problem: we have considered two initial responses to the problem and found them unsuccessful.

In the next chapter, I'll propose an alternative, more plausible approach to the coordination problem; or so I argue. This approach utilises Leitgeb's *stability theory* about

⁵⁶ It is possible to modify ILT so that it gives more weight to the agent's credence function compared to other permissible credence functions. For instance, the agent could give different weights to equally rational credence functions depending on how similar they are to her credence function; and due to this weighting, more similar credence functions would have more influence on what the agent ought to believe. I am not sure whether this version of ILT can deal with the coordination problem. But even if it could do this, this approach will involve controversial assumptions about how to attach weights to different credence functions. For this reason, I don't think this is a promising way of approaching the coordination problem for a hybrid theorist.

how belief and credence ought to interact. The stability theory is not conveniently chosen because it can shield Hybrid Impermissivism against the coordination problem. Instead, as I'll discuss, the stability theory is the only theory that satisfies some of the most plausible coherence conditions on belief and credence. So, the proposed hybrid theory will avoid the coordination problem by appealing to an independently plausible theory of belief-credence interaction.

5 Hybrid Impermissivism and the Stability Theory

In this chapter, I state a solution to the coordination problem by using a formally precise and specific bridge principle for belief and credence: the *stability theory* (Leitgeb, 2017). The stability theory roughly says that a rational agent believes a proposition if and only if she assigns a *stably high* degree of belief to the proposition. Leitgeb also calls this view the *Humean thesis*, because, according to some Hume scholars (e.g., Loeb 2002, 2010), for Hume, the distinguishing feature of belief is that belief (unlike a mere “idea”) is steady, firm and resilient and not easily abandoned by new experiences or reasoning.

I call a hybrid theory that endorses Hybrid Impermissivism and the stability theory, the *Humean Hybrid Impermissivism* (H_{HI} , for short). As we shall see, the stability theory entails our favoured interpretation of the Lockean thesis, the weak Lockean thesis. So, H_{HI} can be thought of as a specification of the hybrid theory from the previous chapter, the Lockean Hybrid Impermissivism (L_{HI}).

The stability theory is not conveniently chosen because it can shield Hybrid Impermissivism from the coordination problem. Instead, as I’ll show, the stability theory is the *only* bridge principle that can accommodate some of the most plausible normative requirements on belief and credence. So, there are good independent reasons for endorsing the stability theory.

I also show in Chapter 6 that the proposed solution to the coordination problem can be implemented by using a plausible alternative to the stability theory that I’ve developed.⁵⁷ So,

⁵⁷ As I’ve discovered, the alternative bridge principle that I’ve developed is logically equivalent to Lin and Kelly’s (2012, 2021) tracking theory, which is the main published alternative to the Humean thesis.

the solution to the coordination problem that I offer in this chapter is quite general and does not need to presuppose the stability theory.

Regarding the structure of Chapter 5: It is useful to think of this chapter as consisting of two overall parts. The first part (Sections 5.1 and 5.2) discusses and motivates the stability theory. In Section 5.1, I rehearse the well-known difficulties with providing a satisfactory account of the belief-credence interaction. Against this background, I'll introduce and provide an in-depth analysis of the stability theory in Section 5.2.

The second part of the chapter (Section 5.3) uses the stability theory to articulate a solution to the coordination problem. I will identify a condition under which the Humean Hybrid Impermissivism (H_{HI}) avoids the coordination problem. I call this condition *Order Uniqueness*, which roughly says that for any evidence and proposition, the evidence justifies the unique plausibility order of relevant possibilities (or possible worlds) associated with this proposition (a detailed summary of Section 5.3 is at the end of Section 5.2).

5.1 Introduction

In Section 3.5, we have discussed the difficulty in combining the following individually plausible norms into a precise, systematic view on how rational belief and credence ought to interact:

Deductive Cogency: A rational agent's beliefs should be consistent and closed under logical consequence. More fully:

Consistency: It is not the case that $Bel(\emptyset)$.

Closure: If $Bel(X)$ and $Bel(Y)$, then $Bel(X \cap Y)$.

If $Bel(X)$ and $X \subseteq Y$, then $Bel(Y)$.

Probabilism: A rational agent's credences should satisfy the axioms of probability theory. More fully:

For any P over W (i) the probability of each world in W is non-negative, (ii) the sum of all worlds is 1, and (iii) for any proposition X over W , $P(X)$ is the sum of all worlds in which X is true.

As discussed (in Section 3.5), Deductive Cogency and Probabilism are inconsistent with the so-called strong Lockean thesis (SLT) and the highly plausible assumption that absolute certainty is not necessary for rational belief:

SLT: There is a threshold r greater than $1/2$ and at most 1, and any agent's belief set Bel and credence function P should be such that, $Bel(X)$ iff $P(X) \geq r$.

Non-certainty: An agent can believe a proposition, even if she is not certain that the proposition is true.

The inconsistency of Deductive Cogency, Probabilism, SLT and Non-certainty is illustrated via the Lottery Paradox. Consider a partition of propositions t_1, t_2, \dots, t_n and define the uniform probability distribution over this partition, so that, for any t_i $P(t_i) = 1/n$. The uniform probability distribution over a partition is the standard probabilistic model of a *fair lottery*, where each ticket is equally likely to win, and one and only one ticket will win. Given Non-certainty, there always exists a uniform distribution over a partition where it is rational to believe of each proposition that it is false. And given Deductive Cogency, this entails that it is rational to believe that no proposition in the partition is true. But, by definition, a partition of propositions includes one true proposition. Contradiction. So, we have the following general result:

The following principles are inconsistent:

1. *Deductive Cogency,*
2. *Probabilism,*

3. *Non-certainty*,
4. *The strong Lockean thesis (SLT)*.

Because of such paradoxes, many epistemologists have concluded that, given Deductive Cogency, Probabilism, and Non-certainty, having a high credence in a proposition is not sufficient for believing the proposition. As Levi (1967, 41) put it (presupposing Probabilism and Non-certainty): "... either [Deductive] Cogency or the requirement of high probability as necessary and sufficient for acceptance must be abandoned."

But Hannes Leitgeb (2014, 2017) has shown that the above conclusion is too quick. Understood in a certain way, believing a proposition can be equivalent to assigning sufficiently high credence in the proposition. How so? The trick is not to fix some unique, universally correct Lockean threshold, but instead to allow different Lockean thresholds in different reasoning contexts. More precisely, Leitgeb has proposed to substitute SLT with the following weaker principle that we have called the weak Lockean thesis (WLT):

WLT: An agent's belief set *Bel* and credence function *P* should be such that, there a threshold *r* greater than 1/2 and at most 1, and *Bel*(*X*) iff $P(X) \geq r$.

As we can see, WLT does not demand that all *Bel* and *P* should be related via a unique threshold *r*. Instead, it only says that, given any *Bel* and *P*, there is some threshold *r* that connects *Bel* with *P*. And this threshold *r* could be different for different *Bel* and *P*. But how do we determine what Lockean threshold *r* an agent should have?

According to Leitgeb, the threshold *r* in the Lockean thesis depends on the agent's *context of reasoning*. This context is characterised by the familiar contextual factors like the agent's interests and her level of epistemic cautiousness or braveness, as well as a more unfamiliar factor of the agent's credence function *P* itself. I discuss in detail how these factors determine the threshold *r* in sections 5.2.3, But first, let's illustrate how choosing

certain, very specific Lockean thresholds “by hand” makes it possible to harmonise Deductive Cogency, Probabilism, and Non-certainty with a key idea behind the Lockean thesis that belief corresponds to sufficiently high credence.

Suppose an agent is concerned with just four mutually exclusive and exhaustive possibilities or possible worlds, denoted by W , where $W = \{w_1, w_2, w_3, w_4\}$. Think about these possible worlds as corresponding to all logical possibilities associated with the two propositions in our fine-tuning example from the previous chapter:

God: God exists.

FT: Our universe is fine-tuned for life.

w_1 corresponds to $God \wedge FT$, w_2 to $God \wedge \neg FT$, w_3 to $\neg God \wedge FT$, and w_4 to $\neg God \wedge \neg FT$.

The agent’s credences over W are as follows:

Possible worlds	P
w_1	$P(\{w_1\}) = 0.03$
w_2	$P(\{w_2\}) = 0.02$
w_3	$P(\{w_3\}) = 0.18$
w_4	$P(\{w_4\}) = 0.77$

Table 5.1

Now, suppose that the agent’s Lockean threshold r , in this context, equals the probability of the world w_4 : $r = P(\{w_4\})$. Trivially, each proposition X over this W has the probability of at least 0.77 only if X follows for $\{w_4\}$; or, equivalently, if X contains $\{w_4\}$ as a subset. This is so because if X does not contain w_4 , then it can have the probability of at most $1 - P(\{w_4\}) = 0.33$.

So, if the agent’s threshold $r = P(\{w_4\})$, then the agent’s beliefs will be deductively cogent.

By contrast, if the agent's threshold r is a bit higher and equals 0.79, then the corresponding *Bel* would not be deductively cogent. Here is why. The "first" proposition that has the probability of 0.79 is $\{w_2, w_4\}$:

$$P(\{w_2, w_4\}) = P(\{w_2\}) + P(\{w_4\}) = 0.02 + 0.77 = 0.79.$$

Now, the proposition $\{w_3, w_4\}$ also has a probability greater than 0.79:

$$P(\{w_3, w_4\}) = 0.18 + 0.77 = 0.95$$

However, the intersection of these two propositions does not reach the probability of 0.79:

$$P(\{w_2, w_4\} \cap \{w_3, w_4\}) = P(\{w_4\}) = 0.77$$

Therefore, given $r = 0.79$, we have $Bel(\{w_2, w_4\})$, $Bel(\{w_3, w_4\})$ but not $Bel(\{w_2, w_4\} \cap \{w_3, w_4\})$. And this contradicts Deductive Cogency.

So, as the above example illustrates, for *Bel* and P to satisfy Deductive Cogency and Probabilism, the Lockean threshold should be chosen in a very specific way. Is there a systemic way for determining appropriate Lockean thresholds (that preserve Deductive Cogency) in any context of reasoning?

The answer is "Yes". As Leitgeb has shown, there is a general, formally precise theory of how any *Bel* and P should be related, and this theory fully determines which Lockean thresholds are appropriate for an agent to have, given that the agent ought to satisfy Deductive Cogency and Probabilism. This is what I will explain next.

5.2 The Stability Theory

Fortunately, Leitgeb's theory can be simply explained by using his notion of a *stable proposition*:

Stable Proposition (Definition): For any proposition X and credence function P over W , X is a stable proposition (relative to P) iff $P(X) = 1$ or for all worlds w_i where X is true, $P(\{w_i\}) > P(\neg X)$.

To put it simply, X is a stable proposition iff $P(X) = 1$ or each world in which X is true is more probable than $\neg X$. Equivalently, in set-theoretic terms:

X is a stable proposition iff $P(X) = 1$ or for all w_i , such that $w_i \in X$, w_i is more probable than the sum of all worlds in which X is false: $P(\{w_i\}) > P(W \setminus X)$.

Let us illustrate this new definition with an example. Consider again a set W of four possible worlds: $W = \{w_1, w_2, w_3, w_4\}$ and the same probability distribution as in Table 5.1:

Possible worlds	P
w_1	$P(\{w_1\}) = 0.03$
w_2	$P(\{w_2\}) = 0.02$
w_3	$P(\{w_3\}) = 0.18$
w_4	$P(\{w_4\}) = 0.77$

Table 5.1

As it can be easily verified (via the definition of a stable proposition), there are four different stable propositions over P : $\{w_4\}$, $\{w_3, w_4\}$, $\{w_1, w_3, w_4\}$, and W (the tautological proposition, which is always stable because $P(W) = 1$). As we can see, the first stable proposition, $\{w_4\}$, is logically strongest as it entails all the other stable propositions. And the “second” stable proposition, $\{w_3, w_4\}$ is logically stronger than each of the remaining stable propositions, and so forth. W is the weakest or last stable proposition, and trivially so, as W , by definition, is stable with respect to any P .

In set-theoretic terms, stable propositions are well-ordered with respect to the subset relation: the first (or the least/smallest) stable proposition, $\{w_4\}$, is a subset of the “second” stable proposition $\{w_3, w_4\}$ (i.e., the stable proposition that is logically weaker than $\{w_4\}$, but logically stronger than each of the other stable propositions) and the second stable proposition is a subset of the third stable proposition and so on, where W is always the

weakest or largest stable proposition. Figure 5.1 represents how all non-trivial stable propositions in our example are nested like nested spheres:




Figure 5. 1: Nested spheres representing the nested stable propositions

The smallest sphere at the centre represents the first or the smallest stable proposition $\{w_4\}$, and the next sphere, which covers the first sphere, represents the next stable proposition $\{w_3, w_4\}$ and so on. For any finite W , the stable propositions can always be represented as the nested spheres with the smallest sphere at the centre.

Now given this definition, Leitgeb’s theory can be stated as follows (for simplicity, I assume that each world in W has a non-zero probability):

The Stability Theory: For an agent with belief set Bel and credence function P , $Bel(X)$ if and only if there is a stable proposition Y in Bel and $Y \subseteq X$.⁵⁸

⁵⁸ If some worlds in W have zero (personal) probability, then we should add the following proviso to the stability theory: “If $P(Y) = 1$, then Y is the *least* proposition over W with the probability 1”. The reason we need this proviso is related to the Lockean thesis. If Y is not the least (or the logically strongest) proposition with probability 1, then the stability theory won’t satisfy the Lockean thesis. For instance, say X is logically weaker than Y and $P(X) = P(Y) = 1$. And if Y is the least believed proposition, then X won’t be believed (as Y does not entail X). But, by assumption, the agent assigns the same probability to both X and Y . Therefore, the Footnote continued on the next page.

So, according to the stability theory, a rational agent's belief set Bel must include a stable proposition, such that everything else that the agent believes follows deductively from this stable proposition. Following Leitgeb, we denote a proposition that entails all of the agent's beliefs as " B_W ". B_W is the strongest proposition believed by the agent. So, on Leitgeb's theory B_W should be a stable proposition.

To illustrate the stability theory, let's consider the probability distribution in Table 5.1. Given the stability theory, there are four choices for the strongest believed proposition, B_W : $\{w_4\}$, $\{w_3, w_4\}$, $\{w_1, w_3, w_4\}$, and W . So, for instance, if $B_W = \{w_3, w_4\}$, then $Bel = \{\{w_3, w_4\}, \{w_2, w_3, w_4\}, W\}$. As far as the stability theory is concerned, any of these stable propositions could be the least believed stable proposition B_W . Intuitively, which stable proposition is the agent's least believed proposition depends on her level of epistemic cautiousness or braveness. If in this context of reasoning, $B_W = \{w_4\}$, then the agent is maximally "brave", as this choice of B_W enables the agent to have the "largest" permissible belief set. On the other extreme, if $B_W = W$, then the agent is maximally cautious and suspends judgement on all propositions except the trivial proposition. And in between these two extremes, there are two more moderate choices of B_W : $\{w_3, w_4\}$ or $\{w_1, w_3, w_4\}$.

While the choice between different B_W makes the difference for what the agent's belief set Bel is, it does not make the difference between believing a proposition and believing its negation. In the above example, $\{w_4\}$ is contained as a subset in any stable

agent's beliefs violate the Lockean thesis. Such cases only arise when W contains a world with zero probability (the reader should easily convince herself of this. If this is not obvious, just consider a simple example involving, say, four worlds, where one world has zero probability). But, in all our examples, we assume W to include only the worlds with non-zero probabilities.

proposition. So, any believed proposition X and Y should be true at w_4 . Hence, X and Y cannot be contradictory.

The rest of Section 5.2 discusses the details of the stability theory, its important strengths and overall weakness. First, in Section 5.2.1, I show that the stability theory is a highly fruitful view: it entails (i) Deductive Cogency, (ii) the (weak) Lockean thesis, (iii) and the Humean view that rational belief is not easily defeated by new evidence. In section 5.2.2, I state an important representation theorem that, under some standard assumptions, the stability theory is the *only theory* that achieves *perfect coordination* between an agent's beliefs and credences. By perfect coordination, I mean that each believed proposition is more probable for the agent than each non-believed proposition. And in Section 5.2.3, I consider and respond to the most widely discussed problematic consequence of the stability theory: a strong form of context-sensitivity of belief.

Because of the above-explained features of the stability theory, a hybrid theory that I shall study from Section 5.3 is the only possible hybrid theory that satisfies the standard logical and probabilistic constraints on belief and achieves perfect coordination between belief and credence.

5.2.1 Three Consequences of the Stability Theory

Consequence 1: Deductive Cogency

The stability theory entails that the agent's belief set Bel is deductively cogent. This is evident from Theorem 3.1 (from Section 3.2):

Theorem 3.1:

Bel over W is deductively cogent iff:

For some non-empty proposition B_W over W , $Bel(X)$ iff $B_W \subseteq X$.

On the stability theory, every believed proposition is entailed by the least stable proposition believed by the agent. Hence, on the stability theory, non-empty proposition B_W exists, and it is some stable proposition (over the given credence distribution).

Consequence 2: The Lockean Thesis with a context-sensitive threshold

The fact that the stability theory entails deductive cogency is easy to see. What may be more surprising is that the stability theory entails the weak Lockean thesis:

Weak Lockean Thesis: An agent's belief set Bel and credence function P should be such that, there a threshold r greater than $1/2$ and at most 1 , and $Bel(X)$ iff $P(X) \geq r$.⁵⁹

⁵⁹ Here is how we can show this. Suppose in line with the stability theory, B_W is some stable proposition over W . Trivially, by the axioms of the probability theory, for any proposition X , if B_W entails X , then $P(X) \geq P(B_W)$. So, any proposition in Bel should have the probability that is equal to or greater than the probability of B_W . Now, take any proposition Y which is not in Bel ; i.e., the agent either suspends judgement or disbelieves Y . Suppose the agent suspends judgement on Y : hence B_W and Y have a non-empty intersection: $B_W \cap Y \neq \emptyset$ and $B_W \cap \neg Y \neq \emptyset$. By the axioms of probability:

$$P(Y) = P(Y \cap B_W) + P(Y \cap \neg B_W)$$

$$P(B_W) = P(B_W \cap Y) + P(B_W \cap \neg Y)$$

So, $P(Y) \geq P(B_W)$ iff $P(Y) = P(Y \cap \neg B_W) \geq P(B_W \cap \neg Y)$. Because the agent suspends judgement on Y , we have $Y \cap \neg B_W \notin B_W$ and $B_W \cap \neg Y \in B_W$. And by the definition of a stable proposition, any world in B_W is more probable than the sum of all worlds not in B_W . Hence, $P(Y \cap \neg B_W) < P(B_W \cap \neg Y)$, contradicting the assumption that $P(Y) \geq P(B_W)$. Therefore, if an agent suspends judgment on Y , then for any proposition X , such that $Bel(X)$, $P(X) > P(Y)$.

And if Y is disbelieved then $B_W \cap Y = \emptyset$. And, by definition, any world in B_W should be have a higher probability than that of Y .

Footnote continued on the next page.

Because the stability theory entails both Deductive Cogency and the weak Lockean thesis, and permits believing non-trivial propositions, i.e., the propositions that don't have the credence of 1, we have the following general result:

The following principles are consistent:

1. *Deductive Cogency,*
2. *Probabilism,*
3. *Non-certainty,*
4. *The weak Lockean thesis.*

This result implies that the stability theory avoids the Lottery paradox. To see how, consider a fair lottery consisting of, say 100 tickets. Given this, we have the uniform credence function P defined over set W of 100 possible worlds. Each world, w_i , is associated with the proposition that ticket i wins. Now, it is easy to check that no proposition over W , except the tautological proposition W , is stable. Take any proposition X that is true at some worlds but false in some other(s). And as each world in W is equally probable, X cannot be stable: as for any world w in X , there is a world w' outside X , such that $P(\{w\}) = P(\{w'\})$. In general:

If P is a uniform probability distribution, then Bel satisfies the stability theory iff Bel is trivial; i.e., $Bel=W$.

The above conclusion does not entail that, on the stability theory, a lottery proposition like "ticket 1 will not win" is never rationally believable. To see why this is so, we need to

Therefore, any proposition in Bel is more probable than any proposition outside Bel . And this establishes that, for any Bel and P that satisfies the stability theory, there is some threshold r , which equals to the probability of the least believed stable proposition B_W , and for all X , $Bel(X)$ iff $P(X) \geq P(B_W)$. As required.

approach the above-described lottery situation from a different perspective. Consider again the fair lottery consisting of 100 tickets. But this time, suppose that the agent is solely concerned with the question: “Will ticket 1 win?”. In this context, the agent may attend to a coarse-grained partitioning of possibilities consisting of just two worlds – one in which ticket 1 wins and the other where ticket 1 loses: $W' = \{w_1, v\}$, where v represents a set of all worlds where ticket 1 loses; i.e., the set $\{w_2, w_3 \dots w_{100}\}$. So, W' can also be represented as $\{w_1, \{w_2, w_3 \dots w_{100}\}\}$, where the large set inside W' is treated as a single unit or world. In this context, the agent does not differentiate the specific ways in which ticket 1 loses; instead, all these specific possibilities are lumped into one bulky possibility, v , or represented by the *catchall* hypothesis (i.e., the hypothesis that $\{w_1\}$ is false). Now, relative to W' and P , proposition $\{v\}$: “the ticket 1 does not win” is a stable proposition: as $P(\{v\}) = 0.99$ and as v is treated as a single world – the world in which ticket 1 does not win – $\{v\}$ is stable. So, if the threshold $r = P(\{v\}) = 0.99$, then the stability theory requires to believe $\{v\}$.

This example illustrates all context-sensitive parameters that, on the stability theory, determine the threshold r , and hence determine what it is rationally permissible to believe. These context-sensitive parameters are (1) the partitioning of possibilities, W , (2) the agent’s level of cautiousness or braveness, and (3) the agent’s credence function P . So, on the stability theory, an agent’s context of reasoning is a *triple* consisting of the above three parameters.

Leitgeb (2017, Section 3.4) has argued that the stability theory’s context-sensitive account of belief “works to the theory’s advantage” in the lottery-type cases. His reasoning is as follows. Suppose the agent is focused on the question “Which ticket from the hundred fair tickets will win?”, and, hence, her partitioning of possibilities consisting of 100 possible worlds. In this context, the agent should not believe about any individual ticket that it will

lose. After all, each ticket is equally likely to lose. But, in a more coarsely grained context of reasoning where $W' = \{w_i, v\}$, if the agent's level of epistemic cautiousness is not maximal, then she could rationally believe that the ticket i will not win.

In response to Leitgeb, many have argued that the type of context-sensitivity that the stability theory entails is unacceptable. We consider this debate in Section 5.2.3

Consequence 3: The Humean Thesis

Another important consequence of the stability theory is, what Leitgeb has called, the *Humean thesis* on belief. The Humean thesis says that a rational agent believes a proposition X if and only if, for any proposition Y that the agent does not disbelieve, her conditional credence in X given Y is high (at least, higher than 0.5). More precisely:

The Humean thesis: For all X , $Bel(X)$ iff for all Y such that if Y is not disbelieved by the agent and $P(Y) > 0$, then $P(X|Y) > c$, where $0.5 \leq c < 1$.⁶⁰

Some clarifications are in order. It is immediately obvious that the Humean thesis, as stated, entails that the tautological proposition W is believed, and the contradictory proposition \emptyset is not believed. This is because for any Y , if $P(Y) > 0$, the $P(W|X) = 1$. Further, Because W is believed and the negation of W , \emptyset , is not believed, each believed proposition must have an unconditional probability greater than 0.5, as $P(X|W) = P(X)$. So, the Humean thesis entails that for any $Bel(X)$, X has an unconditional probability greater than 0.5 as well as a conditional probability greater than 0.5, if we condition X on any proposition Y which is not disbelieved by the agent. Here is how we can interpret this. Following the Bayesian principle of *conditionalisation*, think about conditional probability

⁶⁰ The Humean threshold c should not be conflated with the Lockean threshold r , as these two thresholds may, and often do, come apart. More on this shortly.

$P(X|Y)$ as an agent's estimate of her future probability in X , if she learns Y and nothing else. So, given the Humean thesis, if an agent believes X and if the agent learns a new piece of information Y that is compatible with X , then the agent's credence in X will remain high. In other words, the agent's belief that X should be *resilient* or have a *stably high probability* when the agent learns new evidence that does not contradict X .

Let $Poss(Y)$ be short for “ Y is not disbelieved” or “ $\neg Bel(\neg Y)$ ”. $Poss(X)$ can be read as “ X is doxastically possible for the agent”. So, besides an agent's belief set Bel , we can define another set, $Poss$, which includes the propositions that are doxastically possible from the agent's point of view. And the Humean thesis can be stated as the thesis that an agent's beliefs are stable with respect to set $Poss$. More precisely:

For all X , $Bel(X)$ iff for all Y such that if $Poss(Y)$ and $P(Y) > 0$, then $P(X|Y) > c$, where $0.5 \leq c < 1$.

This resiliency or stability of rational belief that the Humean thesis entails is the rationale behind the label “Humean”. As, according to some Hume scholars (e.g., Loeb 2002, 2010), Hume characterises belief not only in terms of *force* and *vivacity* but also in terms of “solidity, or firmness, or steadiness” (Hume 1978, Sect. VII, Pt. III, Bk. I).

Now, we can prove that the stability theory is logically equivalent to the Humean thesis. The proof is delegated to the footnote.⁶¹

⁶¹ First, let's prove that the stability theory entails the Humean thesis.

On the stability theory, Bel should be obtained via some (non-empty) stable proposition B_W : for all X , $Bel(X)$ iff $B_W \subseteq X$. Now, suppose that Y is not disbelieved by the agent and: $Y \cap B_W = w_n$, where w_n is some world over W . Now, by the axioms of probability:

$$P(B_W|Y) = \frac{P(Y \cap B_W)}{P(Y)} = \frac{P(w_n)}{P(Y)} = \frac{P(w_n)}{P(w_n) + P(Y \setminus B_W)}$$

Now, by the definition of a stable proposition, $P(w_n) > P(Y \setminus B_W)$, therefore:

Footnote continued on the next page.

Because the two theories are logically equivalent, I'll sometimes state the stability theory in terms of the Humean thesis, if this helps to facilitate the discussion.

To summarise the section: the stability theory, which is equivalent to the Humean thesis, unifies Deductive Cogency, Probabilism, Non-certainty, and the important idea behind the Lockean thesis (that each believed proposition should be more probable than each non-believed proposition) into one, formally precise, coherent view.

All these may sound too good to be true. Is there a catch? As we have noted, the stability theory achieves all these by making rational belief sensitive to the agent's context of reasoning. And many have argued that such context-sensitivity of rational belief is highly problematic. I disagree. I consider the context-sensitivity problem in detail in Section 5.2.3.

While the above discussed three consequences of the stability theory provide strong motivation for the view, in the next section, I explain what I consider the strongest motivation for endorsing the stability theory. As I show, the stability theory is *the only theory* that achieves *perfect coordination* between belief and credence, without violating either

$$\frac{P(w_n)}{P(w_n) + P(Y \setminus B_w)} > 0.5$$

As required.

Now, let's prove that the Humean thesis entails the stability theory. To prove this, we will use a lemma, proved by Leitgeb (2017, 3.21), that the Humean thesis entails Deductive Cogency.

Suppose for reductio that *Bel* and *P* satisfy the Humean thesis, but not the stability theory. Now, given Deductive Cogency, there must be some non-empty proposition *X* that entails each proposition in *Bel*. Now, *X* cannot be a stable proposition, otherwise, *Bel* and *P* would satisfy the stability theory. Now, by definition, there should be some world w_n in *X* which is not more probable than some world w_j outside *X*. By definition, such worlds must exist. Now define proposition *Y* to be $\{w_n, w_j\}$. As *Y* and *X* have a non-empty intersection, we have $Poss(Y)$. And, by the axioms of probability theory as, $P(\{w_n\}) \leq P(\{w_j\})$, $P(X|Y) \leq 0.5$; contradicting our assumption that *Bel* and *P* satisfy the Humean thesis.

Deductive Cogency or Probabilism. By perfect coordination, I mean that each believed proposition is more probable than each non-believed proposition. In other words, perfectly coordinated beliefs and credences satisfy the following norm that, following Leitgeb, I call the *Monotonicity principle* (Monotonicity):

Monotonicity: For any rational agent with credence function P , if she believes X and $P(Y) \geq P(X)$, then she also believes Y .

And we shall see that, if one accepts Deductive Cogency, Probabilism, and a highly plausible, weak norm that each believed proposition is more probable than not, then it can be proved that Monotonicity is *logically equivalent* to the stability theory. So, given these popular norms, if one accepts Monotonicity, then there is no choice but to accept the stability theory.

5.2.2 From Monotonicity to the Stability Theory

Monotonicity is one of the most plausible (if not the most plausible) coherence principle between rational belief and credence.

Monotonicity: For any rational agent with credence function P , if $Bel(X)$ and $P(Y) \geq P(X)$, then $Bel(Y)$.

Monotonicity articulates a condition under which an agent's categorical doxastic attitudes perfectly reflect the strengths of her numerical doxastic attitudes. So, if Bel and P satisfy Monotonicity, then any proposition inside Bel is more probable for an agent than any proposition outside Bel .

Monotonicity and the stability theory seem very distinct. But, surprisingly, it can be proven that given some standard assumptions about rational belief and credence, Monotonicity and the stability theory are logically equivalent (this theorem is implicit in one of Leitgeb's (2017) theorems, Theorem 7).

Theorem 5.1

Let Bel be the set of all propositions believed by an agent and let P be her credence function. We assume that:

Deductive Cogency: Bel is consistent and closed under logical entailment.

Probabilism: P satisfies the axioms of probability.

High Probability: All believed propositions have a probability greater than 0.5 (relative to P).

Given these assumptions, the following two theses are logically equivalent:

- (1) Monotonicity: For all X and Y , if $Bel(X)$ and $P(Y) \geq P(X)$, then $Bel(Y)$.
- (2) The stability theory (stated as the Humean thesis): For all X , $Bel(X)$ iff for all Y such that $Poss(Y)$, $P(X|Y) > c$, where $0.5 \leq c < 1$.⁶²

⁶² To prove this theorem, we will use a simple theorem from Chapter 3, Theorems 3.1, that if an agent's beliefs are deductively cogent, then there must be the *least believed proposition*, B_W , such that, for all propositions X , $Bel(X)$ iff B_W is a subset of X .

First, let's prove that the stability theory entails Monotonicity. By the stability theory, we know that the probability of the least believed proposition is greater than 0.5: $P(B_W) > 0.5$. Now, suppose for reductio that there is some Y , $P(Y) \geq P(B_W)$, but $\neg Bel(Y)$. Given that $P(Y) > 0.5$, we know that Y has non-empty intersection with B_W . Hence, we know that $\neg Y$ is not believed: $Poss(Y)$. Now, by the axioms of probability, $P(Y) = P(Y \cap B_W) + P(Y \cap \neg B_W)$ and $B_W = P(Y \cap B_W) + P(\neg Y \cap B_W)$. Therefore, if $P(Y) \geq P(B_W)$, then $P(Y \cap \neg B_W) \geq P(\neg Y \cap B_W)$. But, by the definition of $Poss$, we have $Poss((Y \cap \neg B_W) \cup (\neg Y \cap B_W))$. And by the Humean thesis, $P(B_W | (Y \cap \neg B_W) \cup (\neg Y \cap B_W)) > 0.5$, which, by the probability theory, entails that $P(\neg Y \cap B_W) > P(Y \cap \neg B_W)$. But, as we have seen, by $P(Y) \geq P(B_W)$, we also have $P(Y \cap \neg B_W) \geq P(\neg Y \cap B_W)$. Contradiction.

Now, we need to prove that Monotonicity entails the Humean thesis. Assume for reductio that $Bel(X)$, $Poss(Y)$ but not $P(X|Y) \leq 0.5$. There are two cases here: either $Bel(Y)$ or $\neg Bel(Y)$. First, assume that $Bel(Y)$:

Footnote continued on the next page.

The only new assumption in this theorem is High Probability: a weak, ubiquitously accepted assumption that an agent should not believe a proposition unless she is more confident in the proposition than in its negation.

So, the stability theory is the only bridge principle for belief and credence that satisfies Deductive Cogency, Probabilism, and High Probability and achieves perfect coordination between belief and credence: meaning that if P and Bel satisfy the stability theory, then each proposition inside Bel is more probable than each proposition outside Bel .

5.2.3 The Problem of Context-Sensitivity

The stability theory of belief is truly remarkable in terms of its power, scope, and elegance. But over the past several years, some important criticisms have been levelled against it.

Leitgeb himself acknowledges that the most worrisome feature of the stability theory is “a strong form of sensitivity of belief to a context” (2017, 157). Given Leitgeb’s theory, we can fully represent an agent’s context of reasoning by a triple: (C1) the partitioning of possibilities, W , (C2) the agent’s level of cautiousness and braveness, and (C3) the agent’s credence function P . (C2) and (C3) pose no special problem for the stability theory.

by P3, $P(X|Y) = P(X \cap Y)/P(Y)$ and by P1 and P3, $P(Y) \geq P(X \cap Y) \geq P(B_W) > 0.5$. Hence, $P(X \cap Y)/P(Y) > 0.5$, which contradicts our assumption that $P(X|Y) \leq 0.5$. Now, assume that $\neg Bel(Y)$. $Poss(Y)$ entails that Y and B_W have a non-empty intersection. Hence, $(Y \cap B_W) \neq \emptyset$. By P3, $Y = ((Y \cap B_W) \cup (Y \cap \neg B_W))$. And by the definition of B_W , we have $Poss(Y \cap B_W)$ and $\neg Poss(Y \cap \neg B_W)$. Therefore, $P(Y \cap B_W) > P(Y \cap \neg B_W)$. This implies that $P(Y \cap B_W)/P(Y) > 0.5$. And as $P(X \cap Y) \geq (Y \cap B_W)$, by the axioms of probability theory, we have $P(X \cap Y)/P(Y) > 0.5$; contradicting our original assumption that $P(X|Y) \leq 0.5$. This completes my proof of the theorem.

Concerning (C2), as we have seen, many epistemologies take on board some version of the Jamesian pragmatist view that an agent's levels of epistemic cautiousness and braveness have some influence on what she ought to believe.⁶³ (C3) also does not seem troublesome (as acknowledged by critics such as Titelbaum 2020). We saw that given Deductive Cogency, the probabilities that P assigns to the relevant propositions influence which Lockean thresholds r are permissible to have in a given context of reasoning. So, the sensitivity of r to P can be seen as an innocuous consequence of Deductive Cogency.

By contrast, the sensitivity of belief to a partitioning of possibilities, W , is unanimously acknowledged as the controversial part of the stability theory. Hence, from now on, when I speak about the problem of context-sensitivity, I mean the problem associated (C1).

For many (Pettigrew 2015; Staffel 2016; Douven 2019; Titelbaum 2020), accepting a form of strong sensitivity of belief to a partitioning of possibilities is too high a price to pay, even for a theory that has much going for it. Titelbaum (ibid., 11) summarised the context-sensitivity worry with the stability theory as follows:

It's particularly unfortunate that Leitgeb's stability theory of belief is sensitive to partitions in this way. We're used to the idea that an agent's beliefs will change when she receives new evidence. One of the innovations of Leitgeb's theory is to suggest that for a credence to support belief, it must be resilient in the face of evidential change. Yet even when an agent's evidence remains constant, Leitgeb allows her beliefs to crumble in the face of partitional change. He (2017, 146) writes, "In this way, the agent is

⁶³ However, if needed, it is possible to re-define the stability theory in a way that does not give any weight to an agent's subjective levels of cautiousness and braveness. To see how this can be done, see Leitgeb's (2013), which is his first presentation of the stability theory, albeit a very technical one. The main idea is quite simple: just equate an agent's belief set with the least stable proposition over her set of possibilities W . This way, any credence function P will be associated with a unique belief set Bel .

able to maintain the logic of belief, the axioms of probability, and the Lockean thesis simultaneously. The price to be paid is this very dependency of belief on contexts.” I’m not sure the payoff is worth the price.

Now, in response, one may argue, as Leitgeb does, that the context-sensitivity of belief may be a feature rather than a bug of his theory. For instance, we saw (in Section 3.2.1) that the stability theory offers a way of reconciling the prima facie conflicting intuitions in the Lottery paradox by appealing to shifts in context. Further, some philosophers – most recently Yalcin (2018) – have articulated independent arguments for thinking that belief is partition-sensitive.

Leitgeb (2017, 141) also appeals to proof by Lin and Kelly (2012, Sections 13-14), that given Deductive Coherence and some weak assumptions, the partition-sensitivity of rational belief is *unavoidable*. As he (ibid.,) explains:

Stability theory’s partition-dependence of belief is just a special case of the general phenomenon that Lin and Kelly refer to as lack of “question-invariance” of acceptance. Roughly, what they prove is that given some pretty general background assumptions on belief and degree of belief, assuming the logical closure of belief will always necessitate belief to be partition-sensitive.

But, even if belief is context-sensitive to some degree, Titelbaum’s point still rings true: at least in many cases, beliefs should not “crumble in the face of partition change”. To drive this point home, consider an example from Staffel (2016) (the example is drawn from unpublished works by Fitelson and Schurz). An agent attends to a coarse partition of possibilities with respect to a proposition H , when she is only concerned with the possibility w_1 where H is true and the possibility w_2 where H is false. Hence $W = \{w_1, w_2\}$. Her credence function P over W is: $P(H) = 0.66$, $P(\neg H) = 0.34$. And as H is a stable proposition over W , it is permissible for her to believe H . Now, suppose that the agent fine-grains her possibility and takes into account the proposition C : “A coin flip landed heads.” The coin is fair and has nothing to do with whether H is true. Now, if the agent changes her

context of reasoning, and attends to four possibilities, associated with propositions H and C , then her credences over these new partitioning, $W' = \{w_1, w_2, w_3, w_4\}$ is as follows:

$$P(\{w_1\}) = P(H \cap C) = 0.66 * 0.5 = 0.33$$

$$P(\{w_2\}) = 0.33$$

$$P(\{w_3\}) = 0.17$$

$$P(\{w_4\}) = 0.17$$

And over W' , H is no longer a stable proposition (as $P(\{w_2\}) < P(\{w_3, w_4\})$); therefore, it is impermissible to believe H once the agent takes into account the probabilistically irrelevant proposition C . And this certainly looks bad for the stability theory.

Leitgeb has not given a response to such counterexamples. However, as we shall see, a plausible response can be given on behalf of the stability theorists that builds on the thought that there is something wrong for the agent to shift her context of reasoning from W to W' . More fully, the response agrees with the stability theory that belief is partition-sensitive, but it imposes constraints on which partitioning of possibilities the agent can rationally attend to and which shifts of context are rationally permissible.

The constraint on the choice and shifts of the partitioning of possibilities that I develop is similar to David Lewis's (1996) rules governing what possibilities an agent may or may not rationally ignore in knowledge attribution; though, of course, we won't be concerned with knowledge, but with belief. This response takes into account Douven's (2019, 372) criticism that the stability theory must provide some account of what we should and should not attend to. As he explains:

... the stability theory of belief is meant as a theory of rational belief: it is about the beliefs of a perfectly rational agent. But then we would like to hear more about which shifts in context are rationally permissible and which are not. What should we and should we not be paying attention to? ...

For all we are told, anything goes as far as these matters are concerned. One would have liked to see

something akin to what David K. Lewis (1996) does in his contextualist account of knowledge attributions. Lewis makes at least an attempt to lay out the normative principles governing what we may and may not properly ignore.

In line with Douven's criticism and suggestion, I state a general rule that puts an important constraint on which contexts or partitioning of possibilities are relevant for rational belief. To simplify the exposition, we will think about an agent's partitioning of possibilities W in terms of her *opinion set* \mathcal{O} , a set of propositions that the agent attends to (as we did in Chapter 3, Sections 3.1 and 3.2). If \mathcal{O} contains one proposition, X , then the corresponding W contains all logical possibilities associated with X : namely one coarse possibility where X is true, and the other coarse possibility where X is false; similarly, if \mathcal{O} contains two propositions, then W contains four worlds and so on. So, shifts in contexts will be understood as shifts in opinion sets.

I call the proposed rule the *rule of case reasoning*. Roughly, it says that if presupposing either E or $\neg E$ does not affect whether the agent should believe H , then the opinion sets relevant to whether an agent is permitted to believe a proposition H must not include E . To explain the rule in detail, we need a new definition. We will say that a proposition E is *case irrelevant* to an agent with respect to a proposition H iff presupposing either E or $\neg E$ does not affect whether she is permitted to believe H relative to the opinion set \mathcal{O} that includes just H . For instance, let $\mathcal{O} = \{H\}$, and suppose that the agent is rational in believing H relative to \mathcal{O} because $P(X) \geq r$, where r is the agent's threshold for the belief that satisfies the stability theory relative \mathcal{O} . Now, we shall say that a proposition E is *case irrelevant* for H for the agent iff the agent's doxastic attitude towards H is not affected by learning (or presupposing) either E or $\neg E$: $P_H(E) \geq r$ and $P_{\neg H}(E) \geq r$. Given this new definition, the rule of case reasoning can be stated as follows:

The rule of case reasoning: For any rational agent, if E is case irrelevant for H for that agent, then the opinion sets relevant to whether an agent is permitted to believe H must not include E .⁶⁴

The guiding idea behind the rule is very plausible: if learning either E or $\neg E$ does not affect your belief about H , then E cannot be relevant to whether you can believe H before learning anything about E .

Now, as we shall see, the rule of case reasoning blocks the common counterexamples against the stability theory involving “strange” shifts of contexts.

Trivially, the rule prohibits an agent from adding any probabilistically irrelevant proposition to her relevant opinion set; as, if E is probabilistically irrelevant to H , then learning either E or $\neg E$ has no impact on H . Hence, Staffel’s counterexample is avoided: if the agent is interested in whether to believe H , then her partitioning of possibilities should not include a probabilistically irrelevant proposition.

⁶⁴ This rule about permissible shifts in context is built from Lin and Kelly’s (2012, 2021) principle about rational belief – *Case Reasoning*. As they (2021) explain:

If you believe that A given information E and also given information $\neg E$, then it is a foregone conclusion that you will end up believing A , so you may as well believe that A already.

In their earlier paper (2012, Footnote 1), Lin and Kelly point out the similarity between Case Reasoning and van Fraassen’s famous *reflection principle*, which is concerned with rational credence:

The principle is analogous in spirit to the reflection principle (van Fraassen 1984), which, in this context, might be expressed by saying that if you know that you will accept a proposition regardless what you learn, you should accept it already.

The rule of case reasoning has a wider scope than simply prohibiting taking into account probabilistically irrelevant propositions. It also prohibits an agent to *unnecessarily fine grain* her partitioning of possibilities even with a probabilistically relevant proposition if the proposition has no *overall effect* on what the agent ought to believe. Let me give an example. Consider the following distribution involving two propositions: $P(H \cap E) = 0.47$, $P(H \cap \neg E) = 0.21$, $P(\neg H \cap E) = 0.22$, $P(\neg H \cap \neg E) = 0.1$. Suppose that an agent who considers these propositions has a Lockean threshold of 0.6. Then, it can be verified that E is case irrelevant to H : $P(H|E) = \frac{0.47}{0.47+0.21} \approx 0.69$ and $P(H|\neg E) = \frac{0.21}{0.21+0.1} \approx 0.67$; therefore, given the agent's standard of belief (represented by her Lockean threshold), it is impermissible to fine grain her partitioning of possibility with E . So, while on the stability theory, H is not a stable proposition over $\mathcal{O} = \{H, E\}$ (this is easy to verify: $P(H \cap \neg E) < (\neg H \cap E)$), the agent is still rationally permitted to believe H and properly ignore the possibilities associated with E .⁶⁵

The rule of case reasoning also explains why fine graining of possibilities seems permissible in the *lottery* settings. Consider a fair lottery consisting of 100 tickets; define the uniform probability distribution P over W of 100 possible worlds, representing the lottery case. Now suppose the agent's opinion set \mathcal{O} is grained coarsely and contains just one proposition – $\{w_1\}$. We know that, on the stability theory, $\neg\{w_1\}$ is permissible to believe over the opinion set \mathcal{O} . Now suppose we fine-grain \mathcal{O} with just one additional proposition, $\{w_2\}$, and get a new opinion set $\mathcal{O}' = (\{w_1\}, \{w_2\})$. It is simple to verify that $\neg\{w_1\}$ is no

⁶⁵ This being said, the agent can attend to the opinion set $\mathcal{O} = \{H, E\}$ if she is interested in some question or topic involving both H and E . However, if she is only concerned with whether to believe H , her permissible opinion sets should not include case irrelevant proposition E .

longer permissible to believe, over \mathcal{O}' : let $W' = \{w_1, w_2, n\}$, where $n = \{w_3, w_4 \dots w_{100}\}$. Now, $\neg\{w_1\} = \{w_2, n\}$; hence, as $P(\{w_2\}) = P(\{w_1\})$, we know that $\{w_2, n\}$ is not stable over \mathcal{O}' (or, over the corresponding set of possibilities W'). We can also verify that $\{w_2\}$ is not a case irrelevant proposition: $P(\neg\{w_1\}|\{w_2\}) = 1$ and $P(\neg\{w_1\}|\neg\{w_2\}) = 98/99$; hence given the old Lockean threshold of $r = P(\neg\{w_1\}) = 99/100$; $Bel_{\{w_2\}}(\neg\{w_1\})$ and $\neg Bel_{\neg\{w_2\}}(\neg\{w_1\})$.

In summary: the proposed rule is sufficient to block Staffel's and other similar counterexamples against the stability theory, and it also explains why some prima facie plausible shifts in context are permissible. One may object to the rule for the following reason. Suppose that two pieces of (potential) evidence E_1 and E_2 are individually not case relevant to H ; However, taken together, $E_1 \cap E_2$ is case relevant to H . The rule does not permit an agent to attend to the possibility involving E_1 alone, but allows attending the possibility involving $E_1 \cap E_2$. And this may seem odd.

In response, I don't see this consequence of the rule as problematic. Here is why. We should not conflate (i) a rule about what possibilities are relevant to whether an agent is permitted to believe H at a given time, with (ii) a view about what counts as actual or potential evidence for H . The rule of case reasoning concerns only (i) and not (ii). Of course, E_1 and E_2 are both individually and jointly relevant to H in the sense that learning, say, E_1 , impacts the probability of H . Still, considered individually, E_1 or E_2 are not relevant to whether the agent is permitted to believe H at a given time (before learning either E_1 or E_2). And this is perfectly in line with the proposed rule.

So, overall, the rule of case reasoning is independently plausible – it is motivated by the plausible thought that if learning either E or $\neg E$ does not affect your belief about H , then E cannot be relevant to whether you are permitted to believe H at a given time – and explains

which choices and shifts of partitions are permissible. For these reasons, I submit that the partition sensitivity of the stability theory is not as problematic as it may initially seem, on the supposition that we add the rule of case reasoning to the stability theory.

Let us take stock. So far in this chapter, I have explained, motivated and defended the stability theory. As we have seen, the stability theory not only satisfies some popular, plausible requirements on belief and credence, but is the only theory that can satisfy these requirements.

In the remaining sections of this chapter, I'll use the stability theory to articulate a solution to the diachronic coordination problem that we have identified in the previous chapter. I will conclude that the *Humean Hybrid Impermissivism* (H_{HI}), a hybrid theory that endorses Hybrid Impermissivism together with the stability theory, can successfully avoid the coordination problem.

I proceed as follows. After discussing some preliminaries, Section 5.3.1 states a solution to the coordination problem. I show that H_{HI} avoids the coordination problem on the supposition of Order Uniqueness: the view that evidence justifies a unique order of relevant worlds associated with an agent's context of reasoning. In Section 5.3.2, I argue that Order Uniqueness is a relatively undemanding and plausible view. In Section 5.3.3, I discuss the coordination problem within a setting where an agent could update her credences in response to uncertain evidence (i.e., within the setting where we substitute Conditionalisation with *Jeffrey Conditionalisation*).

Section 5.3 has an appendix that consists of three sections. The first section of the appendix, 5.5.1, argues that the coordination problem is fatal for a hybrid theory that substitutes Moderate Uniqueness with Extreme Uniqueness. In Section 5.5.2, I provide a solution to the coordination problem for Hybrid Impermissivism that uses independent

diachronic norms on belief; namely, the AGM theory of belief revision (the best-known model of belief revision). In Section 5.5.3, I discuss the relationship between Order Uniqueness and Relational Objectivity (the principle about evidential support from Chapter 2).

5.3 The Humean Solution to the Coordination Problem

In this section, I develop a solution to the coordination problem. I will show that, given the stability theory, it is possible to restrict a set of permissible credence functions over a set of possibilities W in a way that these credence functions never license opposing beliefs over W , no matter on which possible evidence we condition these credence functions. And as I will argue, the proposed solution to the coordination problem does not commit to a strong and implausible view about the evidential constraints on credence.

As the coordination problem concerns how beliefs and credences change over time, first, we need to make the dynamics of the stability theory explicit.

As before, we assume an agent revises her credences via Conditionalisation.

Conditionalisation

For any agent with credence function P and any proposition X over W , if the agent learns that some proposition E over W is true, then her new credence function should be P_E where:

$$P_E(H) = P(H|E)$$

Conditionalisation specifies the *unique policy* of credal revision: there is only one unique way to revise P by E .

With respect to belief, we only need to assume that whenever an agent revises her belief set Bel by any evidence E , her new belief set Bel_E is such that P_E and Bel_E satisfy the stability theory. So, we need not assume that there are separate diachronic norms on belief; as

the belief revision can piggyback on credal revision: first, we update P by E , and then we determine the set of all permissible Bel_E , such that Bel_E and P_E satisfy the stability thesis. This piggybacking of belief revision on credal revision is depicted in Figure 5.2. The horizontal arrow from P to P_E indicates that P_E has been obtained by conditioning P by E ; while there is no horizontal arrow from Bel and Bel_E . There are only vertical double arrows between Bel and P and between P_E and Bel_E denoting the synchronic joint constraints on belief and credence at any given time (in this chapter, the synchronic joint constraints is given by the stability theory).




Figure 5. 2: Piggybacking of belief revision on credal revision

But can we also approach the coordination problem via the separate diachronic norms on belief? The answer is *yes*. We will see this in Section 5.4.1, where I utilise the AGM belief revision theory – the most well-known and extensively studied theory of (qualitative) belief revision – to provide another solution to the coordination problem.

The solution to the coordination problem that I will develop is based on a type of theorem that I call the *coordination theorem*. The coordination theorem uses a bridge principle for belief and credence and identifies the conditions under which different agents with non-identical credence functions won't adopt opposing attitudes towards a proposition, no matter which possible evidence they learn. In this section, I state such a coordination theorem and show that H_{HI} (i.e., the hybrid theory that endorses Hybrid Impermissivism together with the stability theory), avoids the (diachronic) coordination problem under certain assumptions. To state the theorem, we need a new definition.

We will say that two credence functions, P and P' , defined over the same set of possible worlds W , are *order equivalent* relative to W iff P and P' determine the same *ordering* of the worlds over W . More precisely:

Order Equivalence: For any P and P' defined over the same W , P and P' are order equivalent iff:

For any w and w' over W , $P(\{w\}) \geq P(\{w'\})$ iff $P'(\{w\}) \geq P'(\{w'\})$.

So, order equivalent credence functions agree with respect to the at-least-as-probable relation over the worlds in W . That is, for any order equivalent functions P and P' over some fixed W , if w is at least as probable as w' according to P , then w is at least as probable as w' according to P' , and vice versa. Sometimes we won't mention an agent's credence function explicitly and instead write that " $w_n \geq w_j$ " meaning that the agent considers w_n to be at least as probable as w_j ; or, alternatively, the agent *doxastically ranks* w_n at least as high as w_j . Similarly, " $w_n = w_j$ " means that the agent's doxastic ranking of w_n and w_j is the same and " $w_n > w_j$ " means that the agent ranks w_n over w_j . The doxastic rankings that we are concerned with here are purely epistemic: they only concern the agent's evaluation of relative probabilities of two worlds, and not their relative desirability or utility.

Order equivalence should not be conflated with what is usually called *ordinal equivalence*. P and P' are ordinally equivalent relative to W iff for any propositions X and Y (over W), $P(X) \geq P(Y)$ iff $P'(X) \geq P'(Y)$. So, if P and P' are ordinally equivalent then they are order equivalent as well; but not the other way around. For instance, consider the two probability distributions below:

Possible worlds	P_1	P_2
w_1	$P_1(\{w_1\}) = 0.03$	$P_2(\{w_1\}) = 0.08$
w_2	$P_1(\{w_2\}) = 0.02$	$P_2(\{w_2\}) = 0.02$
w_3	$P_1(\{w_3\}) = 0.18$	$P_2(\{w_3\}) = 0.09$
w_4	$P_1(\{w_4\}) = 0.77$	$P_2(\{w_4\}) = 0.81$

Table 5.2

P_1 and P_2 are order equivalent but not ordinally equivalent: for instance, $P_1(\{w_3\}) > P_1(\{w_1, w_2\})$, but $P_2(\{w_3\}) < P_2(\{w_1, w_2\})$.

Now, we are ready to state a novel theorem that shows that H_{HI} , under certain assumptions, avoid the coordination problem.

5.3.1 The Coordination Theorem

The theorem goes as follows:

Theorem 5.2: *The Coordination Theorem (CT)*

For any two agents with prior credence functions P and P' and prior belief sets Bel and Bel' defined over the same partition W , if these agents (i) satisfy the stability theory, (ii) update their credence functions via Conditionalisation, and (iii) their credence functions are order equivalent, then the following obtains:

For any evidence (proposition) E and proposition X over W and for any permissible posterior belief sets Bel_E and Bel'_E over W_E (i.e., the set of worlds in W compatible with E), it is not the case that $Bel_E(X)$ and $Bel'_E(\neg X)$.⁶⁶

⁶⁶ To prove this theorem, we need to know that conditioning on E preserves all ratios of worlds compatible with E (This has been established via Theorem 3.3, in Section 3.3). That is, for any w and w' over W and for any P , if E is compatible with each w and w' , then $P(\{w\}) > P(\{w'\})$ iff $P_E(\{w\}) > P_E(\{w'\})$. So, conditioning preserves all the relevant ratios.

Footnote continued on the next page.

CT includes a proviso that the agents' beliefs are defined with respect to the same partition. To see why this proviso is needed, consider again our fine-tuning example, where the agents are solely focused on two propositions:

God: God exists.

FT: Our universe is fine-tuned for life.

The set of all possible worlds that they need to attend is $W = \{w_1, w_2, w_3, w_4\}$. Now, consider these two credence distributions defined over W :

Possible worlds	P_1	P_2
w_1	$P_1(\{w_1\}) = 0.01$	$P_2(\{w_1\}) = 0.04$
w_2	$P_1(\{w_2\}) = 0.2$	$P_2(\{w_2\}) = 0.2$
w_3	$P_1(\{w_3\}) = 0.3$	$P_2(\{w_3\}) = 0.25$
w_4	$P_1(\{w_4\}) = 0.49$	$P_2(\{w_4\}) = 0.51$

Table 5.3

Trivially, P_1 and P_2 are order equivalent relative to W . But suppose that the agents are interested in a more coarse-grained partition of possibilities than W . For instance, suppose that these agents were concerned with just one proposition: $\{w_4\}$. Hence, there are just two possibilities these agents need to attend, which we can represent as $W' = \{\{w_1, w_2, w_3\}, w_4\}$.

Now, assume that P and P' over W are order equivalent. Let W_E be the set of worlds compatible with E . We know that P_E and P'_E must agree with respect to the orderings of these worlds (as conditioning preserves the ratios, and hence, the orderings of these worlds). Define w_{Max} to be a member of W_E which is at least as probable as any world in W_E . By definition, w_{Max} must exist in any W_E . And, by definition of a stable proposition, for any stable propositions over W_E relative to either P_E and P'_E , this stable proposition must contain w_{Max} . So, relative to both P_E and P'_E , there is no stable proposition over W_E that does not contain w_{Max} . Because of this, for all propositions X and Y , such that $Bel_E(X)$ and $Bel'_E(Y)$, X must contain w and Y must contain w . Therefore, X and Y cannot be contradictory. And in general, for any X over W , it cannot be the case that $Bel_E(X)$ and $Bel'_E(\neg X)$. As required.

Now, given W' , proposition $\{w_4\}$ is stable relative P_2 and its negation, $\{w_1, w_2, w_3\}$ is stable relative to P_1 . Therefore, given W' , on the stability theory, it is permissible to believe $\{w_4\}$ relative to P_2 , and permissible to believe $\neg\{w_4\}$ relative to P_1 . The example illustrates a general point: two credence functions can be order equivalent relative to W , but still permit opposing beliefs relative to a more coarse-grained partitioning of possibilities.

So, CT would be useful to provide a *general* solution to the coordination problem for H_{HI} only if set W over which the agents' beliefs are defined has some special status or role; otherwise, the theorem only shows that two agents won't adopt opposing beliefs relative to *some* partition of possibilities (one among equally appropriate partitions of possibilities that these agents could attend to). This qualified (or special) result could still be useful. However, I will argue that we can identify a partition that, given the agents' shared evidence and context of reasoning, has a special role in what the agents could rationally believe.

So how should we fix this special W ? The short answer is: in a way that best represents the agents' shared evidence concerning a proposition or topic they are attending to. Here is how we can precisify this answer. Consider two agents concerned with whether to believe a hypothesis H in light of their shared evidence E . Call W their *relevant partition of possibilities* with respect to H and E when W includes the possibilities representing the logical combinations of H and E as well as all relevant alternatives to H available to the agents.⁶⁷ For instance, suppose that the agents are concerned with how the evidence of fine-

⁶⁷ Think about relevant alternatives to H as case relevant alternatives to H . As discussed in Section 5.2.3, a proposition X is *not* case relevant to H for an agent iff presupposing either X or $\neg X$ does not affect what the agent is permitted to believe about H . So, case relevant alternatives to H need to be sufficiently probabilistically relevant to H to be included in the relevant partition (where the meaning of "sufficiently probabilistically relevant" depends on the agent's Lockean threshold).

tuning – represented by proposition *FT*: “Our universe is fine-tuned for life” – bears on the following two hypotheses:

God: God exists.

M: There are a vast number of universes, and most (maybe all) possible values of cosmological constants are actualised in some universe (so, the majority of universes are *not* fine-tuned. We just happen to inhabit the universe which is fine-tuned).

In this case, the agents’ relevant partition *W* is a set of eight worlds representing the logically possible combinations of these three propositions: *FT*, *God*, *M*. The more relevant alternatives to *God* the agents consider, the more fine-grained the relevant partitioning.

In Sections 1.1.2-3 and the appendix to the first chapter (Section 1.4), I have argued that the set of alternative hypotheses available to the agent is itself a part of the agent’s total evidence. So, we can think of an agent’s relevant partition as representing her *total body of evidence*. It is not important whether there is a precise, mechanical procedure for determining an agent’s relevant partition, given her interest and total evidence. We only assume that, given an agent’s context of reasoning, we can identify a special partition that accurately represents her total relevant body of evidence (relative to a proposition/topic she is attending to).

So, given the notion of relevant partition, the coordination theorem, CT, can be stated as follows:

For any two agents with prior credence functions *P* and *P'* and prior belief sets *Bel* and *Bel'* defined over the same relevant partition *W*, if these agents (i) satisfy the stability theory, (ii) update their credence functions via Conditionalisation, and (iii) their credence functions are order equivalent, then these agents won’t adopt opposing beliefs no matter which new evidence from *W* they learn.

The theorem establishes a *condition* under which the stability theory, Moderate Uniqueness, and Credal Permissivism are consistent over the relevant partition. I call this condition *Order Uniqueness*:

Order Uniqueness: For any two equally informed agents whose credence functions P_1 and P_2 are defined over the same relevant partition W , there is a unique order of worlds that their evidence justifies over W .

From now on, if I suppress the reference to a partition of possibilities and say that two credence functions are order equivalent, then I mean that these credence functions are order equivalent with respect to the same relevant partition. So, Order Uniqueness can be restated as: for any two equally informed agents, their respective credence functions should be order equivalent.

Now, via CT, we have established the following compatibility result:

For any agents who form their beliefs with respect to the same relevant partition and revise their credences via Conditionalisation, the following theses are consistent:

1. *The Stability Theory,*
2. *Moderate Uniqueness,*
3. *Credal Permissivism,*
4. *Order Uniqueness.*

The combination of (1), (2), and (3) is what we have called the *Humean Hybrid Impermissivism* (H_{HI}):

H_{HI} : Moderate Uniqueness and Credal Permissivism are true, and beliefs and credences are coordinated via the stability theory.

Hence, the theorem shows that to solve the coordination problem for H_{HI} for the agents who form their beliefs relative to the shared relevant partition, we only need to endorse Order Uniqueness.⁶⁸

I should emphasise an important limitation of the coordination theorem, CT. As I've already shown (see Table 5.3), two order equivalent credence functions (relative to a given relevant partition) could permit opposing beliefs relative to some coarser partition. So, CT alone is insufficient to establish that Order Uniqueness (with some additional, stated assumptions) entails Moderate Uniqueness.

We should not worry about this. It is certainly true that, for some order equivalent credence functions defined over a relevant partition W , if we gerrymander the space of possibilities, then these credence functions would license (some) opposing beliefs. But I don't think this is a significant result for radical permissivists. A convincing criticism of Moderate Uniqueness should show that, relative to a context of reasoning where agents attend to a hypothesis and some evidence (or distinct pieces of evidence), they can adopt opposing beliefs towards the hypothesis because of the shared evidence; and not because of artificial gerrymandering of space of possibilities. And, as the coordination theorem, CT, demonstrates, such clear-cut (radically) permissive cases cannot obtain if Order Uniqueness

⁶⁸ CT, as stated, only deals with cases where the evidence E is a proposition over the agents' (prior) set of possibilities W . So, CT is silent about the cases where the agents' new evidence involves, what we've called, *conceptual innovation*: i.e., when new evidence involves concepts and propositions that the agents were unaware of. We can extend CT to cover conceptual learning. As discussed in Section 3.4.2, if an agent learns new conceptual information, then she ought to revise her old set of possibilities W and adopt new set of possibilities W_{New} (that represents her current conceptual evidence); and given Order Uniqueness, all equally informed agents should adopt new credence functions over W_{New} which are order equivalent.

and some other plausible assumptions are correct (as I will show in Chapter 6, the stability theory is not essential for this result).

So, I submit that the central question regarding whether CT provides an adequate solution to the coordination problem depends on the plausibility of its key supposition – Order Uniqueness. In the next section, I argue that Order Uniqueness is sufficiently plausible as it does not commit to an overly strong and demanding view of rational credence.

5.3.2 Order Uniqueness, Credal Uniqueness, and Ordinal Uniqueness

At the end of Section 4.3.1, I discussed that a successful solution to the coordination problem should not commit to an overly strong and demanding view about the evidential constraints on credence. As we saw, my proposed solution depends crucially on the thesis about the evidential constraints on credence, Order Uniqueness: the view that equally informed agents should agree on the plausibility ordering of possible worlds in their relevant partition. In this section, I show that Order Uniqueness is not an overly strong and demanding thesis.

For a start, Order Uniqueness is logically weaker and significantly less demanding than Credal Uniqueness. According to Credal Uniqueness, *evidence parses finely* (Staplefor 2019): even when one has vague, qualitative evidence, there is still a unique credence function that the evidence justifies over any proposition. By contrast, according to Order Uniqueness, *evidence may parse coarsely*: in some cases, evidence may not justify a unique credence function over a set of propositions, but only the unique order of worlds associated with these propositions. Hence, the requirements of evidence on credence are far less demanding on Order Uniqueness compared to Credal Uniqueness.

To illustrate this in more detail, consider set W of three possible worlds. Each possible world in W corresponds to a *basic proposition*: a proposition that is true in exactly one possible world. Here we have three *basic propositions* A , B , and C . As discussed in

section 3.3, a complete probability distribution over W can be recovered from the probability assignment to these three basic propositions. We depicted all such probability assignments via an equilateral triangle (Figure 3.4).




Figure 3.4: Geometric representation of all possible coherent credence functions

Now, on Credal Uniqueness, any evidence uniquely determines one rational credence function over W , out of a set of *infinitely many probability functions*. Geometrically, on any evidence E , there is one point in the above triangle, and this point represents maximally rational credence function on E .

By contrast, on Order Uniqueness, evidence need not always justify a single credence function but, instead, multiple credence functions that are members of some set of order

equivalent credence functions. And there are *finitely* many order equivalent credence functions over any finite W . More specifically, there are only 13 distinct order equivalent sets over a set of three possible worlds. These sets can be represented geometrically (Figure 5.3). As we see, there are six *bulky* sets of probability distribution in the triangle. These bulky sets represent distinct sets of order equivalent credence functions. For instance, the lower triangle adjacent to the A vertex represents the set of probability functions where $A > B > C$. In general, the probability distributions near the A vertex have A as the most probable proposition. And, as each point in the lower triangle adjacent to the A vertex is closer to the B vertex than to the C vertex, on these probability distributions B is more probable than C . Similarly, in the upper triangle adjacent to the A vertex, C is the second most plausible basic proposition, as each region in the triangle is closer to the C vertex than to the B vertex.




Figure 5. 3: Geometric representation of order equivalent probability distributions

Besides these six bulky sets, there are seven *thin* sets of order equivalent credence distributions. First, we have a singleton set consisting of the equiprobable distribution at the centre of the triangle (on the equiprobable distribution $A = B = C$). Next, we have three sets represented by the lines from the vertices to the midpoint. For instance, the line from the A vertice to the midpoint represents the set of all distributions where $A > B = C$. Similarly, the lines from the other two vertices represent the distributions where $B > A = C$ and $C > A = B$, respectively. And finally, we have the lines from the midpoints of each side to the

midpoint of the triangle. Such a line from the side $A-C$, represents the distributions where $A = C > B$. Similarly, the other two lines represent the distributions where $B = C > A$ and $C = B > A$.

So, there are 13 distinct possible sets of order equivalent distributions over three basic propositions. And on Order Uniqueness, evidence uniquely determines one set of order equivalent distributions out of these *13 possibilities*.

Hence, the requirements of evidence on belief are far less demanding on Order Uniqueness compared to Credal Uniqueness. On Order Uniqueness, evidence may parse rather coarsely, while on Credal Uniqueness, evidence always parses extremely finely.

Order Uniqueness is also compatible with the cases where two equally informed agents disagree about the comparative probabilities of two propositions. In other words, Order Uniqueness does not commit to, what may be called, *Ordinal Uniqueness*:

Ordinal Uniqueness: For any two rationally permissible credence functions P_1 and P_2 defined over the set of possible worlds W and evidence E , and for any propositions X and Y over W , P_1 and P_2 agree on whether X is at least as probable as Y on the evidence E .

To compare these two theses, consider again Figure 5.2. The lower triangle adjacent to the A vertex represents the set of probability functions where $A > B > C$. This region can be partitioned into three further regions corresponding to these three possibilities: (i) $A > (A \cup B)$, (ii) $A = (A \cup B)$, (iii) $A < (A \cup B)$. So, with respect to 6 bulky order equivalent sets, we have $3 * 6 = 18$ distinct ordinally equivalent credence functions (and there are additional ordinally equivalent functions within the *thin* order equivalent sets, excluding the equiprobable distribution). So, Ordinal Uniqueness requires evidence to parse significantly more finely compared to Order Uniqueness.

For these reasons, I submit that we do not commit ourselves to an overly demanding view about rational credence by endorsing Order Uniqueness. After all, Order Uniqueness is not only compatible with Credal Permissivism but also with *Ordinal Permissivism*: the view that two equally informed agents could rationally disagree about the comparative plausibilities of some propositions.

Therefore, the proposed solution satisfies the desideratum that I've set at the end of Section 4.3.1: it avoids an overly demanding view about rational credence.

5.3.3 The Coordination Problem and Jeffrey Conditionalisation

The solution to the coordination problem that I've defended assumes that the agents in question update their credences via *Conditionalisation*: that is, the agents update their credences only when they learn a piece of information for *certain*. More fully, an agent whose credence function P is defined over a set of possible worlds W , should revise or update P by a proposition X over W only if her new evidence makes X certain for her.

But this understanding of credal revision (or learning from evidence) is not generally satisfactory. New evidence does not need to make any proposition over W certain to influence the agent's beliefs and actions. Suppose that you are wondering whether your lawn is wet (denoted as L). Given your evidence, your original credence in L is 0.6. Now, suppose you hear rain-like noises from a room where the window curtains are shut. You know that the wind rustling tree leaves can make similar rain-like noises. So, this new evidence does not make you certain that it is raining; but your credence in the rain increases significantly: say your new credence in the rain ($Rain$) is 0.7. Now, while the proposition $Rain$ is not learned for certain, it is obvious that your credence in L should change, now that the probability of $Rain$ is quite high. But how exactly? We cannot use Conditionalisation to revise P by $Rain$. So, we need to go beyond Conditionalisation.

The most famous and intuitive rule of credal change that does not require evidential certainty is the so-called *Jeffrey Conditionalisation* (due to Richard Jeffrey 1965). On Jeffrey Conditionalisation an agent should revise her beliefs on the pair (E, c) , where c denotes the new credence in E that the agent has due to the new evidence/observation:

Jeffrey Conditionalisation: For any proposition H and evidence E :

$$P_{E,c}(H) = P(H|E) * c + P(H|\neg E) * (1 - c)$$

Trivially, Jeffrey Conditionalisation is equivalent to Conditionalisation when $c = 1$; that is, when the agent updates on certain evidence.

Now, we can apply Jeffrey Conditionalisation to determine your new credence in L (the lawn is wet) in light of the new uncertain evidence *Rain*:

$$P_{Rain,c}(L) = 0.7 * P(L|Rain) + 0.3 * P(L|\neg Rain)$$

Now, as *Rain* entails L , $P(L|Rain) = 1$; and we can assume that $P(L|\neg Rain) = 0.4$, because I could have left the sprinkler on. Plugging the numbers:

$$P_{Rain,c}(L) = 0.7 + 0.3 * 0.4 = 0.8$$

As we shall see, Jeffrey Conditionalisation, like Conditionalisation, holds fix the ratios of all worlds consistent with the new evidence. However, because on Jeffrey Conditionalisation, a world in W could change its probability without becoming certain, this may upset the old order of the worlds in W .

To illustrate this, consider Table 5.4, representing two order equivalent credence functions:

Possible worlds	P_1	P_2
w_1	$P_1(\{w_1\}) = 0.03$	$P_2(\{w_1\}) = 0.08$
w_2	$P_1(\{w_2\}) = 0.02$	$P_2(\{w_2\}) = 0.02$
w_3	$P_1(\{w_3\}) = 0.38$	$P_2(\{w_3\}) = 0.11$
w_4	$P_1(\{w_4\}) = 0.57$	$P_2(\{w_4\}) = 0.79$

Table 5.4

Now, suppose that both agents learn the same new information according to which the probability of w_2 is 0.12. To update these credence functions by evidence $(\{w_2\}, 0.12)$ via Jeffrey Conditionalisation, we simply change the old probability of w_2 to its new probability, and multiply all the remaining worlds with the same constant, so that these worlds sum to $1 - 0.12 = 0.88$. So, with respect to P_1 , we have $c(0.03 + 0.38 + 0.57) = 0.88$, where c denotes the ratio by which the probabilities of all worlds compatible with the new evidence (except w_2 itself) change. By simple calculations: $c = 0.88/0.98$. The table below gives the revision of P_1 and P_2 by $(\{w_2\}, 0.12)$:

Possible worlds	$P_{1,New}; (\{w_2\}, 0.12)$:	$P_{2,New}; (\{w_2\}, 0.12)$
w_1	≈ 0.03	≈ 0.07
w_2	$= 0.12$	$= 0.12$
w_3	≈ 0.34	≈ 0.1
w_4	≈ 0.51	≈ 0.71

Table 5.5

As we can see, when we update these credence functions with the new evidence, via Jeffrey Conditionalisation, they are no longer order equivalent: $P_{1,New}(\{w_2\}) < P_{1,New}(\{w_3\})$ and $P_{2,New}(\{w_2\}) > P_{2,New}(\{w_3\})$. So, Jeffrey Conditionalisation can change the overall order of worlds in W .

This tells us that the coordination theorem we have proved in Sections 5.3.1 does not work if we substitute Conditionalisation with Jeffrey Conditionalisation. So, unfortunately, we have a distinct coordination problem for Hybrid Impermissivism involving Jeffrey Conditionalisation that cannot be solved solely by our coordinations theorems. Call this new problem the *J-coordination problem*.

How should we respond to the J-coordination problem? For a start, adding more stringent requirements on the agents' credence functions won't solve the problem: as in the above example, the agents' credence functions are not only ordered equivalent but ordinally equivalent. So, we should approach the J-coordination problem in some other way.

I don't think the J-coordination problem is as damning as it looks. Suppose Table 5.4 above represents two agents' credences over the same two propositions: *God* and *FT*. Hence, new evidence $(\{w_2\}, 0, 12)$ represents their change of opinion about the probability of $w_2 = \text{God} \wedge \neg \text{FT}$. But why think that this new information fixes the same new probability in w_2 for both agents? After all, these agents update on *uncertain evidence*; they don't learn w_2 for certain, but simply receive the information that makes w_2 more probable than before. And, in this case, it seems highly plausible that their new probability in w_2 would not be solely determined by their new information, but also by their prior probability functions. Let me explain.

Suppose these agents read a paper that argues that God may well create the universe which supports life without finely-tuned cosmological constants (for instance, God may prefer the laws of nature that support life under most values of the cosmological constants). The agents find the paper quite convincing and, in response to it, increase their confidence in $w_2 = \text{God} \wedge \neg \text{FT}$. Plausibly, their new confidence in w_2 would depend on their prior credence functions. In any case, it is contrary to the spirit of Credal Permissivism (that we are assuming, given Hybrid Impermissivism) that this new information fixes unique new credence in w_2 . Now, in our example represented by Table 5.4, the agents have different prior credence functions; hence, it is unwarranted to suppose that these agents have the exact same estimate for the new probability of w_2 . So, when $P_1 \neq P_2$, it is not justified to assume that when P_1 and P_2 are revised by some uncertain information, they are necessarily revised

with the same pair (E, c) . And once we allow P_1 and P_2 to update on the same uncertain information but on *different pairs* of evidence-proposition and a real number, the J-coordination problem could be avoided. For instance, if P_1 revised by $(\{w_2\}, 0.12)$ and P_2 by $(\{w_2\}, 0.1)$, then the new credence functions would remain order equivalent (as shown in the table below):

Possible worlds	$P_{1,New}; (\{w_2\}, 0.12):$	$P_{2,New}; (\{w_2\}, 0.1)$
w_1	≈ 0.03	≈ 0.07
w_2	$= 0.12$	$= 0.1$
w_3	≈ 0.34	≈ 0.1
w_4	≈ 0.51	≈ 0.73

Table 5.6

But what if the agent's new information fixes the unique value of c for the pair (E, c) ? In such cases, the J-coordination problem only arises if the evidence available to the agents does not justify a unique credence distribution for them (otherwise, the agents would have the identical credence functions and, hence, the identical doxastic ordering of possibilities). I don't think that such cases are realistic. But let's assume that they could obtain: so, suppose two agents with non-equivalent but order equivalent prior credences update on the same pair (E, c) ; but, due to their differing prior credences, their doxastic orderings diverge (Tables 5.4 and 5.5 represent the formal features of such an example). Does not this show that the J-coordination problem could not be avoided in general? Not necessarily. The above example assumes that the agents in question have different prior probabilities in light of the same evidence. If these agents revise their credences by the same pair (E, c) , then Order Uniqueness entails that at least one of their credence functions are irrational. So, if Order Uniqueness is true, then some credal revision must be called for: so, the agents must either suspend judgement on the disputed ordering of worlds or re-calibrate

their credences to make them order equivalent again. It is a task of the epistemology of disagreement to tell us how exactly the credences need to be revised in such cases. Here, it suffices to say that such a revision must be rationally required, given Order Uniqueness.

This approach to the J-coordination problem that I've sketched above is a long way from a more clockwork, mechanical solution of the coordination problem that I've given with respect to (ordinary) Conditionalisation (where no matter what evidence the agents learn, they can mechanically update their credences via Conditionalisation, without any risk of upsetting the old ordering of possibilities). But this departure from the mechanical solution to the coordination problem should not trouble us. For a start, as we have seen in Section 3.4.2 (when discussing the problem of *conceptual innovations*), there are good general reasons to think that in response to some type of new evidence, it is *irrational* for the agent to slavishly follow the old update policy, as determined by Conditionalisation. For instance, when the agent discovers some new hypothesis, H_{New} , such that her old credence function does not leave room for a non-zero probability of H_{New} , then it is plainly wrong for the agent to slavishly follow Conditionalisation. Instead, the agent should re-calibrate her priors in light of the new evidence to make room for the non-zero probability of H_{New} . Similarly, Jeffrey updates that break the order equivalence of two credence functions may well require the respective agents to re-calibrate their credences, to make them order equivalent again.

Given that Hybrid Impermissivism is an attractive position in many respects (as I've argued in Chapters 2 and 4) and it mechanically avoids the coordination problem when the agents update on certain evidence (via Conditionalisation), admitting some re-calibration of credences in cases of *some* Jeffrey updates does not seem too problematic.

5.4 Summary

In this chapter, I have used the stability theory to articulate a solution to the coordination problem. In the first part of the chapter (Sections 5.2), I introduced, analysed, and motivated the stability theory. As we have seen, the stability theory is the only bridge principle that satisfies the standard logical and probabilistic requirements on belief and credence and achieves the perfect correspondence between an agent's belief and credence; where each believed proposition is more probable for the agent than each non-believed proposition.

In the second part, in Section 5.3, I used the coordination theorem to articulate a solution to the coordination problem for Hybrid Impermissivism. I've identified a relatively undemanding constrain on rational credence, Order Uniqueness, and showed that if two agents satisfy this constraint, they won't adopt opposing beliefs, no matter what possible evidence they learn in the future. For this reason, I conclude that the proposed solution to the coordination problem is adequate, as it does not commit to an overly demanding view about rational credence.

5.5 Appendix to Chapter 5

As already noted, Chapter 5 has an appendix where I discuss three distinct topics related to the proposed solution to the coordination problem. The relative placement of these sections does not convey an orderly continuity between them. Each section deals with a distinct topic related to Section 5.3 and can be read in any order.

In Section 5.5.1, I argue the coordination problem is fatal for a stronger hybrid theory that substitutes Moderate Uniqueness with Extreme Uniqueness. In Section 5.5.2, I provide a solution to the coordination problem for Hybrid Impermissivism by using independent diachronic norms on belief; namely, the AGM theory of belief revision (the best-known model of belief revision). And finally, I discuss the relationship between Order Uniqueness and Relational Objectivity (the principle about evidential support from Chapter 2).

5.5.1 The Coordination Problem for Strong Hybrid Impermissivism

In Section 5.3, I've proposed and defended a solution to the coordination problem for the Humean Hybrid Impermissivism (H_{Hi}). As I've shown, if rational agents satisfy the stability theory, and if their credence functions are order equivalent, then no matter which evidence these agents learn, they won't adopt opposing beliefs in light of the same new evidence.

This section examines whether a similar strategy can be used to solve the coordination problem for a logically stronger version of Hybrid Impermissivism, *Strong Hybrid Impermissivism* (SHI). SHI substitutes Moderate Uniqueness with Extreme Uniqueness. Hence, SHI combines the following two theses:

Extreme Uniqueness: Given any body of evidence, E , and proposition, H , there is a unique belief-attitude (either belief, disbelief, or suspension) that any agent should adopt toward H .⁶⁹

Credal Permissivism: For some evidence, E , and proposition, H , E rationally permits more than one credence towards H .

I defend the negative conclusion about the prospects of providing a satisfactory solution to the coordination problem for SHI. The severity of the coordination problem for SHI can already be demonstrated within a relatively informal setting, without assuming that the agents adopt complete probability distributions over the relevant propositions.

Consider again our fine-tuning example from Chapter 4 (Section 4.3): on Monday, Cathy and Julien are rational in believing that God does not exist. Julien is a bit more confident that God does not exist than Cathy is: $P_{Julien}(God) = 0.02$, $P_{Cathy}(God) = 0.1$. For simplicity, let's assume that their beliefs and credences are related via the Lockean thesis with a threshold of 0.7 (but, as I explain shortly, the choice of this threshold is inconsequential to my argument).

Now suppose that they receive and analyse the new evidence on Sunday that the fine-tuning data is approximately 25 times more likely on the supposition that God exists than on the supposition that God does not exist.

Given their prior probabilities, we can use the ratio form of Bayes' theorem to determine their posterior probabilities:

$$\frac{P(H|E)}{P(\neg H|E)} = \frac{P(E|H)}{P(E|\neg H)} * \frac{P(H)}{P(\neg H)}$$

⁶⁹ To remind the reader, Moderate Uniqueness, unlike Extreme Uniqueness, is compatible with the cases where the evidence permits both believing a proposition and suspending judgement.

As before, we let R_{Post} be the ratio of *posteriors*, R_L the ratio of *likelihoods*, and R_{Prior} the ratio of *priors*. So, the theorem can be summarised succinctly as: $R_{Post} = R_L * R_{Prior}$.

Using this theorem, the simple calculations show that, upon learning the new information about fine-tuning (denoted as “FT”), Cathy’s and Julien’s posteriors in God’s existence should be:

$$P_{Cathy}(God|FT) \approx 0.73$$

$$P_{Julien}(God|FT) \approx 0.34$$

And via the Lockean thesis with a threshold of 0.7, we conclude that $Bel_{Cathy}(God)$ and $\neg Bel_{Julien}(God)$; Contradicting Extreme Uniqueness.

In the above example, we have assumed that Cathy and Julien’s Lockean thresholds are relatively low. But this assumption is inessential. We can easily show that within the setting where the agents agree on the value of the ratio of likelihoods, no matter how we specify the Lockean threshold r (or if we specify different thresholds r for different agents), if two agents have different credences towards a proposition, it is always possible for them to learn some new information that would require one agent to believe a proposition and the other agent – to disbelieve or suspend judgement about the proposition. For instance, suppose that Cathy and Julien’s Lockean thresholds are quite high, say, $r = 0.9$. By algebra, for any proposition H :

$$P(H) > 0.9 \text{ iff } \frac{P(H)}{P(\neg H)} > 9$$

Now, if Cathy’s and Julien’s credence functions P_{Cathy} and P_{Julien} are not identical, then there is some value of R_L such that:

$$\frac{P_{Cathy}(H)}{P_{Cathy}(\neg H)} * R_L > 9$$

$$\frac{P_{Julien}(H)}{P_{Julien}(\neg H)} * R_L < 9$$

Therefore, to guarantee that Cathy and Julien won't adopt different doxastic attitudes towards H , we need to assume that they have identical credences towards H . And this assumption directly contradicts Credal Permissivism.

Notice that the above development of the coordination problem does *not* apply to Hybrid Impermissivism. This is because, if we make two credence functions sufficiently similar, say $P_{Cathy}(God) = 0.05$, $P_{Julien}(God) = 0.04$, and if we assume the Lockean threshold of 0.9, then there is no *single* value of R_L that will permit Cathy to believe God and Julien to believe $\neg God$. Therefore, even within this informal setting, the coordination problem is more serious for SHI than for Hybrid Impermissivism.

Regarding the setting where we work with complete probability distributions: it can be easily verified that if two credence functions P and P' are order equivalent, it is still possible that P permits believing H while P' – *prohibits* believing H . To illustrate this, consider Table 5.6:

Possible worlds	P_1	P_2
w_1	$P_1(\{w_1\}) \approx 0.13$	$P_2(\{w_1\}) \approx 0.42$
w_2	$P_1(\{w_2\}) \approx 0.09$	$P_2(\{w_2\}) \approx 0.1$
w_3	$P_1(\{w_3\}) \approx 0.78$	$P_2(\{w_3\}) \approx 0.48$

Table 5.6

As we see, proposition $\{w_3\}$ is stable with respect to P_1 but not with respect to P_2 . Hence, given the stability theory, it is rationally permissible to believe $\{w_3\}$ relative to P_1 but not relative to P_2 . The same conclusion can be reached by assuming a more straightforward Lockean view (without appealing to the stability of believed propositions). No matter how we choose the Lockean threshold r , two credence functions can be ordered equivalent but license

different beliefs. For instance, suppose $r = 0.99$; then define P_1 as a triple $(0.06, 0.04, 0.99)$, corresponding to three possible worlds (in that order), and leave P_2 as in the above table. These credence functions are order equivalent. And given the Lockean threshold of 0.99, it is rationally permissible to believe $\{w_3\}$ relative to P_1 but not relative to P_2 .

So, even on the supposition of Order Uniqueness, SHI, unlike Hybrid Impermissivism, is still open to the coordination problem.⁷⁰

We can show more: if we assume that two credence functions are *ordinally equivalent*,⁷¹ it still seems possible that these credence functions could permit both believing a proposition and suspending judgement. For instance, consider the following table representing two ordinally equivalent credence functions:

Possible worlds	P_1	P_2
w_1	$P_1(\{w_1\}) = 0.03$	$P_2(\{w_1\}) = 0.08$
w_2	$P_1(\{w_2\}) = 0.02$	$P_2(\{w_2\}) = 0.02$
w_3	$P_1(\{w_3\}) = 0.38$	$P_2(\{w_3\}) = 0.11$
w_4	$P_1(\{w_4\}) = 0.57$	$P_2(\{w_4\}) = 0.79$

Table 5.7

Suppose these credence functions permit believing a proposition only if $P(X)$ is very high: at least higher than 0.9. Hence they agree that $\{w_3, w_4\}$ should be believed. But now, update these credence functions by new evidence $\{w_2, w_3\}$. By Bayes' theorem:

$$P_1^{New}(\{w_3\}) = 0.95$$

⁷⁰ As we see, we do not even need to look at how beliefs and credences change over time to verify that given the stability theory, different order equivalent credence functions can both permit and prohibit believing the same proposition.

⁷¹ P and P' are ordinally equivalent relative to W iff for any propositions X and Y (over W), $P(X) \geq P(Y)$ iff $P'(X) \geq P'(Y)$.

$$P_2^{New}(\{w_3\}) \approx 0.85$$

And given the same Lockean threshold of 0.9, it is permissible to believe $\{w_3\}$ relative to P_1^{New} but not relative to P_2^{New} .

Hence even on the supposition of Ordinal Uniqueness – the view that for any two equally informed agents and propositions X and Y , these agents must agree on whether X is at least as probable as Y on their evidence – SHI is still open to the coordination problem.

In summary, the coordination problem is significantly more challenging within the considered informal setting for SHI than for Hybrid Impermissivism; and the strategy of restricting the set of permissible credence functions (via Order Uniqueness, or even via Ordinal Uniqueness) does not solve the coordination problem for SHI (but the same strategy is successful with respect to Hybrid Impermissivism). For these reasons, I submit that the coordination problem is fatal for SHI.

Of course, this conclusion is not certain, as it is impossible to exclude that some alternative approach would be more successful for SHI. What is certain is that, even if there is a plausible approach to the coordination problem for Hybrid Impermissivism, this approach may not be extendible to defend SHI from the same problem.

5.5.2 A Coordination Theorem within AGM Belief Revision Theory

In this section, I provide another approach to the coordination problem, which, as I show, coheres with the *AGM theory of belief revision*. AGM is the most well-known and extensively studied (normative) theory of (qualitative) belief revision. It was articulated by Carlos Alchourron, Peter Gardenfors, and David Makinson (Alchourron et al. 1985).

I divide AGM into two parts: one part deals with the revision of beliefs by evidence that is consistent with each of the beliefs being revised; and the other deals with the revision of beliefs by evidence that is not consistent in that way. AGM can be presented in a variety of

equivalent ways. I present the first part of AGM axiomatically and the second part via the so-called *Levi identity* (I find such a presentation of AGM most intuitive).

So, let's start with the first part of AGM. As before, I assume that an agent's belief set is consistent at any given time and, hence, represented by her strongest believed proposition B_W . A new evidence E is consistent with Bel iff the negation of E is not believed: $\neg Bel(\neg E)$. And given the requirement of consistency, it follows that E is consistent with Bel iff E is consistent with the agent's strongest believed proposition, B_W :

Consistent Evidence (Definition): New evidence E is consistent with Bel iff $B_W \cap E \neq \emptyset$.

The central AGM axiom for belief revision by consistent evidence is called *Preservation*:

Preservation: If an agent revises (her belief set Bel) by E and E is consistent with Bel , then the agent should retain (or preserve) all of her beliefs after revising Bel by E .

In other words, Preservation states that the agent should not *lose* any of her old beliefs when revising on any proposition that she previously did not disbelieve. For instance, suppose that you initially believed only the following two propositions:

B : I bought a Bottle of sparkling water

$K \cup R$: I either put the bottle on the Kitchen table (K) or in the Refrigerator (R).

Because you are fully rational, your belief set Bel is deductively cogent. Hence, your strongest believed proposition is B_W : $B \cap (K \cup R)$. Now, suppose that you go to the kitchen and see that the water is *not* on the kitchen table. How should you revise your beliefs? Because the new evidence, $\neg K$, is consistent with B_W , Preservation tells you that you should not lose any of your old beliefs. In this case, this seems plausible. Absent additional relevant

information, it seems that you can rationally believe B and rationally believe R even on the supposition that K is false. And as $K \cup R$ follows logically from R , and given that you still believe R , then $K \cup R$ should also be in your revised or posterior belief set.

Some think that Preservation is not a norm of rationality (see Lin 2019 for a discussion). But as Preservation is an important part of AGM, this section assumes it to be true (though, in the next chapter, I develop a bridge principle for belief and credence that violates Preservation).

More formally and precisely, Preservation can be stated as follows, via the following definitions:

For any agent with belief set Bel and evidence E , an agent's posterior belief Bel_E is equivalent to the revision of Bel by E , denoted by $Bel * E$, where $*$ is the belief revision operator: $Bel_E = Bel * E$.

Preservation: If E is compatible with Bel , then Bel should be a subset of $Bel * E$: $Bel \subseteq Bel * E$.

We need one additional axiom, besides Preservation, to fully characterise AGM theory of how Bel should be revised by consistent evidence E . This axiom is called *Inclusion*:

Inclusion: If an agent revises Bel by E , then her posterior belief set should not contain more propositions than the logical consequences of $Bel \cup E$, denoted by $Cn(Bel \cup \{E\})$. In symbols: $Bel * E \subseteq Cn(Bel \cup \{E\})$

Inclusion places an upper-bound on $Bel * E$: $Bel * E$ cannot contain X unless $X \in Cn(Bel \cup \{E\})$. Inclusion is trivially true if Bel and E are inconsistent because inconsistent sets of propositions logically entail everything. And in non-trivial cases, Inclusion is highly intuitive: after all, how could $Bel * E$ contain more proposition than logical consequences of $Bel \cup E$?

Taken together, Preservation and Inclusion are equivalent to the following norm that I call, *Belief Revision Rule 1* (R1):

R1: If Bel and E are consistent, then $Bel * E$ is equivalent to $Cn(Bel \cup \{E\})$

R1 says that, when Bel and E are consistent, then an agent's posterior belief set $Bel * E$ is obtained by adding E to Bel , $\{E\} \cup Bel$, and closing the obtained set deductively. There is a very useful reformulation of R1 in terms of the notion of the least believed proposition B_W .

R1: Let B_W be the least believed proposition with respect to the prior belief set Bel .

Then for all E , if $Bel \cap E \neq \emptyset$, then $X \in Bel * E$ iff $B_W \cap E \subseteq X$.

So, according to R1, if E is compatible with Bel , then any proposition X is included in the agent's posterior belief set $Bel * E$ iff X follows from the intersection of B_W and E . So, R1 simply says that, when E and Bel are consistent, the posterior belief set, Bel_E , is equivalent to the set of all logical consequences of $B_W \cap E$: $Bel_E = Cn(B_W \cap E)$.

Now we can show that the *stability theory satisfies R1* (equivalently, the stability theory satisfies Preservation and Inclusion; this has been already shown by Leitgeb 2017, Section 4.2). More fully:

Theorem 5.3

Let Bel_E denote the posterior belief set, and P_E denote the posterior credence function obtained by conditioning P by E . If (i) Bel satisfies R1, i.e., if $Bel \cap E \neq \emptyset$, then $Bel_E = Cn(B_W \cap E)$, and (ii) Bel and P satisfy the stability theory, then for any consistent evidence E , Bel_E and P_E satisfy the stability theory.⁷²

⁷² *Proof:* on the stability theory, any Bel and P is rational only if the strongest believed proposition B_W is a stable proposition (relative to P). So, to prove that the stability theory satisfies R1, we need to show that Bel_E , which is equivalent to $Cn(B_W \cap E)$, is a stable proposition. Trivially, if B_W entails E , then $B_W \cap E = B_W$; and (because Conditionalisation preserves the ratios of worlds) B_W will remain stable, as each world in B_W will be

Footnote continued on the next page.

We can show more. Given a weak assumption of *High Probability* (that each believed proposition is more probable than not), we can already deduce that if *Bel* satisfies R1, then Bel_E and P_E satisfy the stability theory. More precisely:

Theorem 5.4

On the supposition that $Bel \cap E \neq \emptyset$, if *Bel* satisfies R1, and *Bel* and *P* satisfy High Probability, then Bel_E and P_E satisfy the stability theory.⁷³

Let us now turn to belief revision where new evidence *E* is *incompatible* with *Bel*. According to AGM theory, in such cases, an agent first should “remove” some proposition(s) from *Bel* in such a way the resulting set is no longer incompatible with *Bel*. This new set is called the *contracted belief set* and is denoted by $Bel \div \neg E$, where \div is the *contraction* operator. Then, the agent can *add* new evidence *E* to the contracted belief set and close everything under logical consequences. The resulting set is denoted as $(Bel \div \neg E) + E$, which is the agent’s posterior belief set. This method is called *Levi Identity* (due to Levi 1977):

$$\text{Levi Identity: } Bel * E = (Bel \div \neg E) + E$$

Now, when *Bel* and *E* are compatible, it is evident that Levi Identity is equivalent to R1 (i.e., Preservation + Inclusion). So, Levi Identity adds content to R1 only when *E* is incompatible with *Bel*.

more probable than $P_E(\neg B_W)$. Now, suppose B_W does not entail *E*, hence Bel_E is a subset of B_W . In that case, B_E will be a stable proposition relative to P_E , because each world in B_E will be more probable than $P_E(\neg B_W)$. As required.

⁷³ Proof: R1 and High Probability entail that, for any *E*, if $B_W \cap E \neq \emptyset$, then $P(B_W|E) > 0.5$. Hence, by the definition of a stable proposition, B_W is stable. As required.

Now, if E is incompatible with Bel , then there could be more than one contraction of Bel that is consistent with E : in other words, there can be multiple candidates for $Bel \div \neg E$. To illustrate this consider a set W of six possible worlds, $W = \{w_1 \dots w_6\}$ and an agent's belief set Bel defined over W . Suppose that the agent's strongest believed proposition, B_W , is $\{w_1\}$. Further, suppose that the agent learns that $\neg\{w_1\}$. Now, there are multiple contractions of B_W which are consistent with $\neg\{w_1\}$. For instance, if the new strongest believed proposition, B_W^{New} , is $\{w_2\}$, then the agent's posterior belief set $Bel * \neg\{w_1\}$ will satisfy Levi's Identity, as $\neg\{w_1\} \in Cn(\{w_2\})$ and $\{w_1\} \notin Cn(\{w_2\})$. But Levi's Identity is also satisfied if B_W^{New} were, say, $\{w_2, w_3\}$. So, the question is: how to choose among competing contractions?

A plausible answer, which is at the heart of AGM theory, is to choose the contraction that makes as few changes in Bel as possible.⁷⁴ How can we make this idea precise? AGM theory has a standard answer to this question (via the so-called *Partial Meet Contraction*). But, here, I give an alternative answer in terms of the stability theory, which closely follows the standard AGM answer (and is much easier to explain).

As we see, given R1 and High Probability, we must assume that an agent's belief set Bel is obtained via a stable proposition B_W (stable relative to her credence function P). Now, if the agent's new evidence E (which is assumed to be a non-empty proposition over W) is incompatible with B_W , then the agent's new least believed proposition should be the *weakest stable proposition* over W compatible with E . More fully, we have the following rule:

⁷⁴ In other words, AGM revision endorses, what Shear and Fitelson (2019) call, the *principle of conservativity* (Conservativity, for short):

Conservativity: For any agent with a prior belief set Bel , when the agent learns E , her posterior belief set Bel_E (i.e., Bel revised by E) should be *as similar as possible* to her prior belief set Bel .

Belief Revision Rule 2 (R2): For any agent with a belief set obtained by B_W , and for any evidence E over W , if E is *inconsistent* with B_W , then there must be some *least strong stable proposition* B'_W that is compatible with E , and B'_W should be the new strongest believed proposition: $Cn(\{B'_W\}) = Bel * E$.

Here is an example to illustrate R2: take a set W of six possible worlds, $W = \{w_1 \dots w_6\}$, and the following probability distribution over W : $P(\{w_1\}) = 0.51$, $P(\{w_2\}) = 0.19$, $P(\{w_3\}) = 0.17$, $P(\{w_4\}) = 0.07$, $P(\{w_5\}) = 0.05$, $P(\{w_6\}) = 0.01$.

Now, suppose the agent's belief set Bel over this W is given by $\{w_1\}$. Trivially, $\{w_1\}$ is stable relative to P and also the strongest stable proposition over W . Beside $\{w_1\}$, we have three other non-trivial stable propositions, listed below in terms of their logical strength (where the weakest stable proposition is listed last):

$$\{w_1, w_2, w_3\}$$

$$\{w_1, w_2, w_3, w_4\}$$

$$\{w_1, w_2, w_3, w_4, w_5\}$$

Now, suppose that the agent learns new evidence $\neg\{w_1\}$, which is, of course, inconsistent with her B_W . Now, the logically weakest stable proposition consistent with $\neg\{w_1\}$ is $\{w_1, w_2, w_3\}$. Hence according to the rule, $Bel * \neg\{w_1\} = Cn(\{w_1, w_2, w_3\})$. So, within the stability theory framework, revision by evidence inconsistent with Bel is quite simple: once you know the set of all stable propositions, it is easy to check which stable propositions are consistent with new evidence; next, determine the logically weakest stable proposition consistent with the evidence, obtain all of its logical consequences, and you will have the new posterior belief set.

Now that we stated the rules that fully specify how an agent should revise her beliefs, we can state the following coordination theorem:

Theorem 5.5 (The AGM coordination theorem)

For any two agents with belief sets Bel_1 and Bel_2 and credence function P_1 and P_2 , if:

- (1) These agents' prior beliefs and credences satisfy the stability theory,
- (2) Bel_1 and Bel_2 are revised via rules R1 and R2, and P_1 and P_2 are revised via Conditionalisation, and
- (3) P_1 and P_2 are ordered equivalent,

Then, the following obtains:

For all propositions X over W , it is not the case that $Bel_1 * E \in X$ and $Bel_2 * E \in \neg X$.⁷⁵

So, this theorem shows that, given the stability theory, Conditionalisation, and the AGM revision rules R1 and R2, two agents whose credence functions are order equivalent (but not identical), would not adopt opposing beliefs over W , no matter which new evidence they learn from W . So, we have derived another coordination result for Hybrid Impermissivism, but this time, we use separate diachronic norms on belief from the AGM theory.

Therefore, it is possible to solve the coordination problem for Hybrid Impermissivism by using the separate diachronic norms on belief and credence.

⁷⁵ Proof: we have seen that for any Bel and P (over W) that satisfy the stability theory, if we revise Bel by E (over W) via rules R1, R2, and if we revise P by E via Conditionalisation, then Bel_E and P_E satisfy the stability theory. Hence, given the stability theory, for any proposition $X \in Bel_1 * E$, X should be true at the most probable world w_{Max} relative to P_1 . And because P_1 and P_2 are ordered equivalent, w_{Max} is the same relative to these credence functions. Hence, for no propositions X and Y , where $X \in Bel_1 * E$ and $Y \in Bel_2 * E$, $X \cap Y = \emptyset$. As required.

5.5.3 Order Uniqueness and Relational Objectivity

In Section 2.2, I have put forward an argument for Moderate Uniqueness, which is built on the following principle about rational confirmation:

Relational Objectivity: Whether evidence E is more likely on H than on $\neg H$, depends on the evidence and hypotheses themselves and not on how any agent interprets the relationship between the evidence and hypotheses; i.e., for any two equally informed agents with rational credence functions P and P^* , it cannot be the case that $P(E|H) > P(E|\neg H)$ and $P^*(E|H) \leq P^*(E|\neg H)$.

As we have discussed, Relational Objectivity does not require that, for any evidence-hypotheses pair (and a body of *background* evidence/auxiliary hypotheses), there is a unique numerical value for their likelihood. Instead, Relational Objectivity only requires that all equally rational credence function P and P^* agree on *comparative* judgements about likelihoods; so, it cannot be the case that $P(E|H) > P(E|\neg H)$ and $P^*(E|H) \leq P^*(E|\neg H)$, if P and P^* are equally rational. We will say that P and P^* *agree in direction* with respect to likelihoods iff they agree on comparative judgements about likelihoods (This definition is equivalent to Hawthorne’s 2018 definition, Section 5).⁷⁶

In Section 5.3, I have defended Hybrid Impermissivism by appealing to the following thesis about evidential constraints on credence, Order Uniqueness:

⁷⁶ It is worth noting that Relational Objectivity (though, not in name) has been very favourably discussed by Hawthorne (2018), who derives the so-called *Likelihood Ratio Convergence Theorem* from it; as he (ibid., 28) claims, his version of the theorem “overcomes many of the objections raised by critics of Bayesian convergence results”. Hawthorn also notes that his result only presupposes agreement on *comparative* claims about likelihoods (or, as he calls it, agreement in direction), and not agreement about the numerical values.

Order Uniqueness: For any two equally informed agents whose credence functions P_1 and P_2 are defined over the same relevant partition W , there is a unique order of worlds that their evidence justifies over W .

The question that I will consider in this section is: what is the relationship between Order Uniqueness and Relational Objectivity?⁷⁷

For a start, Order Uniqueness and Relational Objectivity don't entail each other. To illustrate this, consider once again set W of four possible worlds: $W = \{w_1, w_2, w_3, w_4\}$ representing all logical possibilities associated with the two propositions in our fine-tuning example: *God* – *God* exists, and *FT* – Our universe is fine-tuned for life. w_1 corresponds to $God \wedge FT$, w_2 to $God \wedge \neg FT$, w_3 to $\neg God \wedge FT$, and w_4 to $\neg God \wedge \neg FT$. By the axioms of the probability theory, the probability of each world w_i equals to a probability of a likelihood *weighted* by the prior probability of a hypothesis which is true at that world:

$$P(\{w_1\}) = P(God \cap FT) = P(God) * P(FT|God)$$

⁷⁷ Another epistemic norm I used in my argument for Moderate Uniqueness is, what I've called *Moderate Principle*:

If evidence E justifies you in believing that H and prior to learning that E , you were not justified in believing H , then E makes it rational to increase your probability in H ; i.e., $P(H|E) > P(H)$, where P represents your credence function and P is a rational credence function for you to have.

Moderate Principle is trivially true within the framework of the stability theory on the supposition that we are concerned with mutually exclusive competing hypotheses or basic propositions. For instance, say an agent's opinion set consists of three basic propositions $\{A, B, C\}$; if before learning evidence E , it was rationally impermissible for an agent to believe A (according to the stability theory), and if learning E makes it permissible to believe A , then E must increase the probability of A . And Moderate Principle, within my argument for Moderate Uniqueness, was concerned with such cases alone (where an agent thinks about belief in terms of mutually exclusive hypotheses, H and $\neg H$).

$$P(\{w_2\}) = P(God \cap \neg FT) = P(God) * P(\neg FT|God)$$

$$P(\{w_3\}) = P(\neg God \cap FT) = P(\neg God) * P(FT|\neg God)$$

$$P(\{w_4\}) = P(\neg God \cap \neg FT) = P(\neg God) * P(\neg FT|\neg God)$$

So, in other terms, $P(\{w_i\})$ equals to the product of the prior probability of a hypothesis (which is true at w_i) and its likelihood.

P and P^* can agree in direction without being order equivalent. For instance, consider the following two credence distributions over W , defined in terms of the likelihood and prior distributions:

P_1	P_2
$P_1(God) = 0.3$	$P_2(God) = 0.8$
$P_1(FT God) = 0.8$	$P_2(FT God) = 0.6$
$P_1(FT \neg God) = 0.2$	$P_2(FT \neg God) = 0.3$

Table 5.8

As we see, P_1 and P_2 agree in direction (with respect to likelihoods). But as it can be easily verified, these credence functions are not order equivalent. The most probable world according to P_1 is w_4 : $P_1(\{w_4\}) = 0.7 * 0.8 = 0.56$; and the most probable world according to P_2 is w_1 : $P_2(\{w_1\}) = 0.8 * 0.6 = 0.48$.

We can also show that order equivalent credence functions do not necessarily agree in direction with respect to all relevant likelihoods. Consider the following probability distributions over the same proposition God and FT :

Possible worlds	P_1	P_2
w_1	$P_1(\{w_1\}) = 0.15$	$P_2(\{w_1\}) = 0.15$
w_2	$P_1(\{w_2\}) = 0.16$	$P_2(\{w_2\}) = 0.16$
w_3	$P_1(\{w_3\}) = 0.34$	$P_2(\{w_3\}) = 0.17$
w_4	$P_1(\{w_4\}) = 0.35$	$P_2(\{w_4\}) = 0.52$

Table 5.9

These credence functions are order equivalent, but they disagree about whether FT confirms God :

$$P_1(FT|God) \approx 0.3 > P_1(God) = 0.31;$$

$$P_2(FT|God) \approx 0.47 > P_2(God) = 0.31.$$

So, these credence functions do not agree in direction.

To foreground the relevant details, we represent these credence functions in terms of their prior and likelihood distributions:

P_1	P_2
$P_1(God) = 0.31$	$P_2(God) = 0.31$
$P_1(FT God) \approx 0.48$	$P_2(FT God) \approx 0.48$
$P_1(FT \neg God) \approx 0.49$	$P_2(FT \neg God) \approx 0.25$

Table 5.10

So, as we see, on P_1 the likelihood of FT is slightly higher on the supposition that $\neg God$ than on the supposition of God . By contrast, on P_2 the likelihood of FT is considerably higher on God than on $\neg God$.

But, while there is no purely logical connection between Order Uniqueness and Relational Objectivity, the two are still closely related. As we have seen, the ordering of worlds over W depends on two factors: (i) the prior probabilities of hypotheses and (ii) their respective likelihoods. For this reason, we can think about the probabilities of possible worlds as being *epistemically composite*; i.e., being composed of more basic types of probabilities: priors and likelihoods. And while the likelihood distribution alone is not sufficient to determine an ordering of worlds over W , the credence distributions that agree in direction will also agree on the doxastic orderings of worlds, given that their prior distributions are not widely different (where what “widely different” means depends on a context).

For instance, in our first example, represented by Table 5.8, two credence functions agree in direction but are not order equivalent. This is due to wide differences in their respective prior distributions: $P_1(God) = 0.3$ and $P_2(God) = 0.8$. By contrast, if we make their prior distributions similar, say $P_1(God) = 0.3$ and $P_2(God) = 0.4$, then, given the same likelihood distribution as in Table 5.8, these credence functions will be order equivalent. The new distribution is represented by the table below:

P_1	P_2
$P_1(God) = 0.3$	$P_2(God) = 0.4$
$P_1(FT God) = 0.8$	$P_2(FT God) = 0.6$
$P_1(FT \neg God) = 0.2$	$P_2(FT \neg God) = 0.3$

Table 5.11

The most probable world, according to both credence functions, is w_1 :

$$P_1(\{w_4\}) = 0.7 * 0.8 = 0.56; P_2(\{w_4\}) = 0.6 * 0.7 = 0.42.$$

Regarding the other worlds:

$$P_1(\{w_1\}) = 0.3 * 0.8 = 0.24; P_2(\{w_1\}) = 0.4 * 0.6 = 0.24.$$

$$P_1(\{w_2\}) = 0.3 * 0.2 = 0.06; P_2(\{w_2\}) = 0.44 * 0.4 = 0.16.$$

$$P_1(\{w_3\}) = 0.14; P_2(\{w_3\}) = 0.18$$

So, Table 5.11 represents credence functions that are order equivalent and agree in direction. Generally, credence functions that agree in direction would impose the same ordering of worlds on many different prior distributions.

For this reason, Relational Objectivity provides an important motivation for Order Uniqueness. The objective ordering of worlds can be seen as a derivative of objective likelihoods and a range of prior distributions that are rationally permissible in a given evidential situation. This way of looking at Order Uniqueness is entirely consistent with Credal Permissivism in particular and impermissible epistemologies in general.

6 Hybrid Impermissivism without the Stability Theory

The primary goal of this chapter is to develop an alternative hybrid theory that avoids the coordination problem but does not endorse the stability theory. Achieving this goal is important for two reasons. First, it shows that the hybrid approach to the Uniqueness debate is flexible, as it does not depend on a very specific understanding of how belief and credence ought to interact. And second, as we will discuss in detail, there is a serious worry with the stability theory that has to do with its being an overly strict view of rational belief. For these reasons, it is important to put forward a hybrid theory that does not rely on the stability theory.

The hybrid theory that emerges from this chapter substitutes the stability theory with, what I call, the *dominant core theory*. The resulting hybrid theory – i.e., the combination of Hybrid Impermissivism and the dominant core theory – is dubbed the *Dominant Hybrid Impermissivism* (D_{HI} , for short). D_{HI} solves the coordination problem in exactly the same way as the Humean Hybrid Impermissivism (H_{HI}). Hence, the only reason for favouring D_{HI} over H_{HI} would be to consider the dominant core theory as a superior alternative to the stability theory. And as I will argue in this chapter, all things considered, the dominant core theory is a superior alternative to the stability theory. But, because both hybrid theories provide the same solution to the coordination problem, the choice between the stability theory and the dominant core theory will be inconsequential to the main argument of this dissertation. Still, it is important to establish that there is an alternative, plausible bridge principle that does not have the same weakness as the stability theory and avoids the coordination problem for Hybrid Impermissivism in the same way.

The chapter proceeds as follows. In Section 6.1, I argue that the stability theory entails an overly restrictive view of rational belief. As we have seen in Section 5.2.2, the

stability theory, under the standard logical and probabilistic assumptions about belief and credence, is logically equivalent to Monotonicity: the view that each believed proposition should be more probable for the agent than each non-believed proposition. And, as already discussed, Monotonicity enjoys a great deal of intuitive, pre-theoretical appeal. So any plausible criticism of the stability theory should provide an explanation of why, in some cases, it is rational to violate Monotonicity. To the best of my knowledge, Section 6.2 provides the first such detailed explanation, without appealing to the alleged impermissibility of forming beliefs on probabilistic considerations (or the *mere statistical evidence*) alone. As I argue, Monotonicity may be violated but only with respect to, what I call, *inferentially trivial disjunctions*. Roughly, a disjunction is inferentially trivial for an agent when the agent cannot reliably apply the rule of disjunctive syllogism to the disjunction (to infer either of its disjuncts). Building on this discussion, in Section 6.3, I will articulate a new bridge principle for rational belief and credence, the *dominant core theory*, that is less demanding than the stability theory but supports the same solution to the coordination problem for Hybrid Impermissivism. In Section 6.4, I provide an in-depth comparison between the stability theory and the dominant core theory and conclude that the latter is a superior alternative to the former. I provide the concluding discussion in Section 6.5, where I also show that the main strengths of the proposed hybrid theories do not require the assumption of Deductive Cogency.

6.1 Is the Stability Theory Overly Demanding?

In Section 5.2.3, I have discussed the most glaring problem with the stability theory: a strong form of context-sensitivity of belief to an agent's partitioning of possibilities. I've argued that this problem can be avoided by invoking a rule (the Rule of Case Reasoning) about which

possibilities an agent can rationally ignore in her reasoning context and which shifts of contexts are rationally permissible.

But there is another important worry with the stability theory, not concerned with context-sensitivity: that the stability theory is *overly demanding*.⁷⁸ In what follows, first, I explain how this worry has been developed in the relevant literature so far, and then I explain my preferred version of it.

Rott (2017), Douven and Rott (2018), and Schurz (2018) have argued that the stability theory is too close to the so-called *Certainly Proposal*:

Certainty Proposal: For any agent with belief set *Bel* and credence function *P*,
 $Bel(X)$ iff $P(X) = 1$

For instance, Douven and Rott (ibid.) have shown that the stability theory is *likely* to permit exactly the same belief as the Certainty Proposal. They showed this by calculating the chance of a randomly chosen probability distribution (over some *W*) to permit different beliefs on the stability theory and the Certainty Proposal (their calculations use the technique of *Monte Carlo sampling*). For instance, they calculated (via Monte Carlo integration) that if *W* contains eight possibilities and given the Humean threshold of 0.8, then the chance that a randomly chosen probability distribution over *W* permits more beliefs than the Certainty proposal is roughly 30%. If we increase the Humean threshold and the number of possible

⁷⁸ There is another type of criticism against the stability theory, due to Staffel (2016), that points out the obvious fact about the stability theory: that it permits believing a proposition solely on statistical or probabilistic evidence. Now, many, including Leitgeb and everyone else who is engaged in providing a formal theory of how belief and credence ought to interact, accept that belief can be formed in response to statistical evidence. In any case, given the topics and aims of this dissertation, the debate about the permissibility of believing something on statistical evidence is irrelevant.

worlds in W , this chance decreases significantly. For instance, with the Humean threshold of 0.95 and W of 50 possible worlds, the chance of non-certain beliefs is roughly 5% (See their Table 4 for full details).

In response to Douven and Rott, Leitgeb (2021) has commented that, in certain contexts, where an agent attends to a large number of possibilities and when her Humean threshold is high, it is not bad at all that the stability theory licenses the same beliefs as the Certainty Proposal.⁷⁹ After all, if the agent's standard of belief is quite demanding, and if she attends to lots of possibilities, it is not surprising that she won't have any non-certain beliefs.

In defence of the stability theory, one may also question Douven and Rott's assumption that it is an advantage for a bridge principle to render it probable for an agent to have non-trivial beliefs if her credence function is chosen randomly. After all, neurotypical

⁷⁹ Here is Leitgeb's comment in full:

I do not regard the Certainty Proposal “for all X , $Bel(X)$ iff $P(X) = 1$ ” as a “bad” bridge principle for rational all-or-nothing and graded belief at all, which is why I am less worried by ... the Humean thesis approximating the Certainty proposal on fine-grained partitions and/or close-to-uniform probabilities. It is just that in many everyday contexts, one should be able to believe in a proposition X without assigning to X the maximal degree of belief 1. The ... Humean thesis allows an agent to do so, and it even allows an agent to assign to X a “realistic” degree of belief of 0.9 or so, as long as the partitioning of possibilities is coarse-grained enough. In that sense, the thesis improves the Certainty proposal. I doubt that in real-world contexts in which a great many possible cases are considered at once – e.g., when a great many consequences of a scientific theory are compared with a great many scientific data – the appropriate attitude towards such sets of fine-grained possible cases is that of all-or-nothing belief. For reasons analogous to those in the lottery paradox, longish conjunctions of “acceptable” statements that are sufficiently probabilistically independent of each other will always have low probabilities that should not be assigned to propositions that are believed in the all-or-nothing sense.

humans don't adopt random credences over a set of possibilities. So, why think that random probability distributions are likely to license non-trivial beliefs?

In any case, I think there is a better way of articulating the restrictiveness worry (i.e., the worry that the stability theory is too demanding/restrictive). This strategy directly challenges the idea that rational beliefs should be stable to the extent required by the stability theory. Let me elaborate. As we have seen, the stability theory entails the following: if $Bel(X)$, then for *all* doxastically possible propositions Y , $P(X|Y) > 0.5$. So, assuming the Bayesian principle of *conditionalisation*, the stability theory entails that a rational belief is *stable* under learning new information. In other words, a rational agent should not consider it to be likely to learn a new piece of information that turns her believed proposition improbable. Or so the stability theory demands.

While some degree of stability seems necessary for rational belief, does the stability theory require too much stability? For instance, it seems reasonable to think that a scientist may rationally believe a theory, even if it is a *real possibility* for her to learn new information that would render the theory improbable. And even in ordinary contexts, it seems to be a manifestation of epistemic modesty to think that some of our beliefs may realistically be defeated by future evidence.

To develop the restrictiveness worry, the next section provides an in-depth analysis of an example that shows that the stability theory is overly demanding.

6.1.1 A Counterexample

I begin by stating an example:

Citizen Hannes

Richard believes that Hannes is either a German citizen (G) or was born in Austria (A), but not both: $Bel(G \vee A)$. (where \vee is the exclusive disjunction, $G \vee A = (G \wedge \neg A) \vee (\neg G \wedge A)$).

$G \vee A$ is the *strongest proposition* that Richard believes, meaning that Richard does not believe anything more specific than $G \vee A$. Further, he considers $G \wedge \neg A$ to be more probable than $\neg G \wedge A$. Because of this, his degrees of belief are distributed such that, on the supposition that a more probable disjunct, $G \wedge \neg A$, is false, his credence in the disjunction, $G \vee A$, is less than 0.5. In symbols:

$$P(G \vee A | \neg(G \wedge \neg A)) < 0.5$$

It is easy to verify that the above-described combination of beliefs and credences violate the stability theory. Because $G \vee A$ is the strongest proposition that Richard believes, we know that $\neg Bel(G \wedge \neg A)$. And, by the definition of *Poss*, $\neg Bel(G \wedge \neg A)$ is equivalent to $Poss \neg(G \wedge \neg A)$. Hence, on the stability theory, Richard's beliefs and credences should be such that $P(G \vee A | \neg(G \wedge \neg A)) > 0.5$. And this contradicts our example.

But while the Citizen Hannes example contradicts the stability theory, there seems to be nothing wrong with Richard's beliefs and credences. After all, why should it be rationally *required* to have a high credence in a believed disjunction, $X \vee Y$, if one assumes that its more probable disjunct is false? To illustrate the point, suppose that Richard's main reason for accepting disjunction $(G \wedge \neg A) \vee (\neg G \wedge A)$ is its first disjunct. So, if Richard assumes that $(G \wedge \neg A)$ is false, then Richard's main reason for accepting the disjunction is invalidated; and, for this reason, Richard conditional credence $P(G \vee A | \neg(G \wedge \neg A))$ does not exceed 0.5. This sounds perfectly correct but contradicts the stability theory. Therefore, the example seems to show that the stability theory demands too much stability. It should be at least rationally permissible for Richard to believe $G \vee A$, even if the conditional

probability of $G \vee A$ is relatively low under the supposition of a proposition which Richard does not disbelieve.

How can a stability theorist respond to such an example? In his book, Leitgeb offers a response to a counterexample similar to Citizen Hannes. Leitgeb's strategy is as follows: first, he correctly shows such counterexamples contradicts the following highly plausible principle, which I label the *Monotonicity of Poss*:

The Monotonicity of *Poss*: For all X , if $Poss(X)$, then for all Y , if $P(Y) \geq P(X)$, then $Poss(Y)$.

The Monotonicity of *Poss* says that if an agent does not disbelieve a proposition X , then she should also not disbelieve a proposition that is at least as probable as X . It is easy to verify that Richard's beliefs contradict the Monotonicity of *Poss*. By assumption $P(G \vee A | \neg(G \wedge \neg A)) < 0.5$. And by the axioms of the probability theory, this entails $P(\neg G \wedge A) / P(\neg(G \wedge \neg A)) < 0.5$; which simplifies to: $P(\neg G \wedge A) < P(\neg(G \vee A))$. However, by assumption, we also have $\neg Poss \neg(G \vee A)$ and $Poss(\neg G \wedge A)$; contradicting the Monotonicity of *Poss*.

While Leitgeb does not point this out, the Monotonicity of *Poss* is logically equivalent to Monotonicity. Here is a simple proof:

Monotonicity = Monotonicity of Poss

By definition and logic, Monotonicity is equivalent to the following:

(1) If $\neg Bel(Y)$ and $P(Y) \geq P(X)$, then $\neg Bel(X)$.

Now, let's define propositions A and B as follows:

(2) $A =_{df} \neg Y$; $B =_{df} \neg X$;

Given that X and Y are propositions, A and B are also propositions. Hence, we have:

(3) If $\neg Bel(A)$ and $P(A) \geq P(B)$, then $\neg Bel(B)$.

And by definition of *Poss* and the axioms of probability, we have:

(4) If $Poss(Y)$ and $P(X) \geq P(Y)$, then $Poss(X)$.

As X and Y can be any arbitrary propositions, (4) is equivalent to the Monotonicity of *Poss*. As required.

Now because the Monotonicity of *Poss* (or simply, Monotonicity) is highly intuitively plausible, Leitgeb (2017, 88) concludes the following:

Pre-theoretically, independently of considerations concerning the Humean thesis, it does not seem to be the case that a perfectly rational agent would regard a proposition as possible, another one as impossible, but assign the latter a higher degree of belief than the former. So the example does not seem to be a counterexample to the Humean thesis after all. At second glance, our pre-theoretic verdict coincides with that of the Humean thesis: a perfectly rational agent could not have the required combination of beliefs and degrees of belief.

How should we evaluate his response? One may retort that his response is *problematically circular*. After all, we know (from Theorem 5.1, Section 5.2.2) that Monotonicity, with some widely shared assumptions (which are satisfied in the Citizen Hannes example), is already logically equivalent to the stability theory. So, it should not be surprising at all that a counterexample to the stability theory would violate Monotonicity.

Still, Leitgeb's response will ring true for those who consider Monotonicity as "pre-theoretically" highly plausible. After all, the above example only targets the stability aspect of the stability theory and does not provide any direct reason for thinking that Monotonicity is false. So, if one accepts Monotonicity, then one may dismiss the Citizen Hannes example (and similar examples) as inconclusive for giving up the stability theory.

In the next section, I want to argue that rejecting Monotonicity is not as counter-intuitive as it may initially seem. I show there is a plausible alternative to Monotonicity,

Partial Monotonicity, that captures plausible aspects of Monotonicity but does not commit to the stability theory or a similar hyper stable view of rational belief.

6.2 An Alternative to Monotonicity

My proposed alternative to Monotonicity is defined in terms of a new notion of *inferentially trivial disjunction*.

Inferentially Trivial Disjunction: For any agent with belief set Bel and probability function P defined over W , a proposition $X \vee Y$ is inferentially trivial for the agent iff:

- (i) $X \vee Y$ is logically weaker than some proposition she believes, and
- (ii) The agent cannot reliably use the rule of *disjunctive syllogism* to $X \vee Y$:
i.e., $P(X \vee Y | \neg X) < 0.5$ or $P(X \vee Y | \neg Y) < 0.5$.

In simpler terms, the new notion says the following. Suppose that the agent believes X . Then $X \vee Y$ is inferentially trivial for the agent if she cannot reliably infer Y by assuming that $\neg X$. So $X \vee Y$ is inferentially trivial because it is both logically weaker than the believed proposition X and is an *unreliable premise* in the agent's reasoning. For a specific example, suppose that Richard believes R : "Rudolf owns a ford". But he has no clue about J : "Jones owns a ford". Disjunction $R \vee J$ is inferentially trivial for Richard because he cannot reliably infer J from $\neg R$.

Now, given this new notion, we define the following, logically weaker version of Monotonicity:

Partial Monotonicity: For any propositions X and Y , if $Bel(X)$ and $P(Y) \geq P(X)$, then $Bel(Y)$ *only if* Y is not an inferentially trivial disjunction.

So, Partial Monotonicity says that Monotonicity can be violated *but only with respect to inferentially trivial disjunctions*.

To illustrate Partial Monotonicity, let's re-examine the Citizen Hannes example. In the example, Richard attends to two propositions, G and A . To simplify the probabilistic calculations, we define a complete probability distribution over these propositions (this distribution respects all the assumptions about Richard's beliefs and credences that we've made):

$$P(\{w_1\}) = P(G \wedge A) = 0.26$$

$$P(\{w_2\}) = P(G \wedge \neg A) = 0.4$$

$$P(\{w_3\}) = P(\neg G \wedge A) = 0.28$$

$$P(\{w_4\}) = P(\neg G \wedge \neg A) = 0.06$$

On this credence distribution, we have $P(\{w_2, w_3\}) = 0.68 = P(G \vee A)$; hence the strongest believed proposition is more probable than not (as required). Further, $P(\{w_2\}) > P(\{w_3\})$, because Richard considers $(G \wedge \neg A)$ to be more probable than $(\neg G \wedge A)$; and, lastly, as $P(G \vee A | \neg(G \wedge \neg A)) < 0.5$:

$$P(\{w_2, w_3\} | \{w_1, w_3, w_4\}) = \frac{P(\{w_3\})}{P(\{w_1, w_3, w_4\})} = \frac{0.28}{0.6} \approx 0.46 < 0.5$$

Hence, all the assumptions of the Citizen Hannes example are satisfied. Now, on the supposition that Richard's beliefs are deductively cogent, we know that he does not believe $\{w_1, w_2, w_4\}$; as his strongest believed proposition $\{w_2, w_3\}$ does not entail $\{w_1, w_2, w_4\}$. Now, $P(\{w_1, w_2, w_4\}) = 0.72$, which is greater than $P(\{w_2, w_3\})$. Therefore, Richard violates Monotonicity:

$$Bel(\{w_2, w_3\}) \text{ and } \neg Bel(\{w_1, w_2, w_4\}).$$

$$P(\{w_2, w_3\}) = 0.68; P(\{w_1, w_2, w_4\}) = 0.72$$

But while Richard violates Monotonicity, he still satisfies Partial Monotonicity, as $\{w_1, w_2, w_4\}$ is an inferentially trivial disjunction for Richard. To see this, assume that proposition $\{w_1, w_2\}$ is false. Hence, we have:

$$P(\{w_1, w_2, w_4\} | W \setminus \{w_1, w_2\}) = \frac{P(\{w_4\})}{P(\{w_3, w_4\})} = \frac{0.06}{0.32} = 0.1875$$

It is easy to verify that if $\{w_2, w_3\}$ is Richard's strongest believed proposition, and if his beliefs are deductively cogent, then Monotonicity is violated only with respect to this inferentially trivial disjunction, $\{w_1, w_2, w_4\}$. To verify this, consider Richard's belief set, *Bel*:

$$Bel = \{\{w_2, w_3\}, \{w_1, w_2, w_3\}, \{w_2, w_3, w_4\}, W\}.$$

As we see, there is just one proposition over W , $\{w_1, w_2, w_4\}$, that has a greater probability than $\{w_2, w_3\}$, but $\neg Bel(\{w_1, w_2, w_4\})$. And we saw that $\{w_1, w_2, w_4\}$ is inferentially trivial for Richard. So, while $\{w_1, w_2, w_4\}$ has a relatively large probability, Richard does not lose any inferentially valuable belief by not believing $\{w_1, w_2, w_4\}$.

Hopefully, this convinces the reader that rejecting Monotonicity is not as problematic as it initially seemed. After all, if an agent's belief set satisfies Partial Monotonicity, then there are no inferentially valuable propositions that the agent does not believe. In other words, we have seen that violating Monotonicity may be *inferentially inconsequential*.

Certainly, my argument against Monotonicity is not that an agent should not believe an inferentially trivial disjunction. After all, it is a necessary consequence of the requirement of *deductive cogency* that an agent may believe *some* inferentially trivial disjunctions: the trivial disjunctions which *logically follow* from some believed proposition(s). Instead, I defend the claim that if an inferentially trivial disjunction, X , does not logically follow from an agent's beliefs, then the agent should not believe X , even if X 's probability is quite high.

While Partial Monotonicity is strictly logically weaker than Monotonicity, it is still a rather demanding thesis. To illustrate this, consider a weaker version of the Humean thesis, which Leitgeb calls the *Bel-variant of the Humean thesis* (HT_{Bel} , for short). The only difference between Leitgeb's favoured Humean thesis, and HT_{Bel} is as follows: on the former thesis, a believed proposition should be stable with respect to the set of all *doxastically possible* propositions (i.e., set $Poss$); while on HT_{Bel} – with respect to the set of *believed* proposition (i.e., set Bel). So, HT_{Bel} is defined as follows:

HT_{Bel} : For all X , $Bel(X)$ iff for all Y such that $Bel(Y)$, $P(X|Y) > 1/2$.

It is easy to show that HT_{Bel} violates Partial Monotonicity. For instance, consider the same probability distribution representing Richard's degrees of belief in Citizen Hannes:

$$P(\{w_1\}) = P(G \wedge A) = 0.26$$

$$P(\{w_2\}) = P(G \wedge \neg A) = 0.4$$

$$P(\{w_3\}) = P(\neg G \wedge A) = 0.28$$

$$P(\{w_4\}) = P(\neg G \wedge \neg A) = 0.06$$

Now, by HT_{Bel} , it is permissible to have belief set $Bel^* = \{\{w_1, w_2\}, \{w_1, w_2, w_4\}, W\}$. This is so because each believed proposition has a high probability conditional on any believed proposition (higher than 0.5). For instance, $P(\{w_1, w_2\}|\{w_1, w_2, w_4\}) \approx 0.91$. However, Bel^* violates Partial Monotonicity: as $\{w_1, w_2, w_3\}$ is a non-trivial disjunction, $P(\{w_1, w_2, w_3\}) > P(\{w_1, w_2\})$, but $\neg Bel^*(\{w_1, w_2, w_3\})$.

This illustrates two important points. First, Partial Monotonicity is still a demanding thesis: as HT_{Bel} does not impose weak constraints on rational belief, but HT_{Bel} violates Partial Monotonicity. Second, it should be possible to articulate a bridge principle between rational belief and credence that is less demanding than the stability theory, but more

demanding than HT_{Bel} , such that this bridge principle satisfies Partial Monotonicity. In what follows, I shall put forward such a bridge principle.

6.3 An Alternative to the Stability Theory: The Dominant Core Theory

The bridge principle I shall develop is built around a new probabilistic concept of a *dominant proposition*. The intuitive idea behind this new concept is as follows: a proposition, X , is *dominant* from an agent's point of view iff her credence in X is 1 or any possible world in which X is true is more probable for the agent than any possible world in which X is false. More precisely, we define a dominant proposition as follows:

Dominant Proposition: X is a dominant proposition relative to a probability function P over W iff $P(X) = 1$ or for all $w \in X$ and for all $w' \in \neg X$, $P(\{w\}) > P(\{w'\})$.

Let us illustrate this new definition by finding dominant propositions over the probability distribution in Citizen Hannes:

$$P(\{w_1\}) = P(G \wedge A) = 0.26$$

$$P(\{w_2\}) = P(G \wedge \neg A) = 0.4$$

$$P(\{w_3\}) = P(\neg G \wedge A) = 0.28$$

$$P(\{w_4\}) = P(\neg G \wedge \neg A) = 0.06$$

Trivially, W is a dominant proposition relative to any probability distribution. Further, in Citizen Hannes, we have three other dominant propositions:

$$\{w_2\}$$

$$\{w_2, w_3\}$$

$$\{w_2, w_3, w_1\}$$

As we see, there is the “first time” at which each world over W enters into a hierarchy of dominant propositions. So, the set of all dominant propositions is *well-ordered with respect to the subset relation*: the first (or the least) dominant proposition, $\{w_1\}$, is a subset of the

next dominant proposition $\{w_1, w_2\}$ and so forth, where all dominant propositions of probability less than 1 are subsets of the last dominant proposition of probability less than 1.

Given a probability distribution over W , a non-trivial dominant proposition (i.e., a dominant proposition of probability less than 1) almost always exists. In fact, only the *uniform* or the so-called lottery-type probability distributions do not contain a non-trivial dominant proposition. In all other cases, W contains at least one non-trivial dominant proposition.

Now, we are ready to state a bridge principle that is logically weaker than the stability theory and satisfies Partial Monotonicity. The bridge principle, which I shall call the *dominant core theory of belief*, roughly says that an agent's belief set Bel and credence function P (over W) should be such that her least believed proposition B_W is some dominant proposition whose probability exceeds 0.5. B_W is sometimes called the *core* of Bel . So, the theory simply says that the core of Bel should be a probable dominant proposition (hence the name "the dominant core theory"). More fully and precisely:

The Dominant Core Theory: For any agent with a belief set Bel and credence function P , defined over a finite partition of possibilities W , Bel and P should be such that:

There is a dominant proposition Y over W , $P(Y) > 0.5$,

if $P(Y) = 1$, then Y is the least subset of W with probability 1; and

for all X , $Bel(X)$ iff Y entails X .

The dominant core theory is similar to the stability theory in several important respects:

1. The former, like the latter, satisfies Deductive Cogency. We know this from the fact that Bel is deductively cogent iff it has a core, B_W : the least believed proposition that

entails each and every proposition in *Bel*. And both the stability theory and the dominant core theory require each rational *Bel* to have a core. In other terms, both stable and dominant propositions are well-ordered with respect to the subset relation.

2. Any stable proposition over W is also a dominant proposition. If X is a non-certain stable proposition, then each world in X is more probable than $\neg X$. Hence, trivially, each world in X is more probable than each world outside X . But, as we discuss shortly, not all dominant propositions are stable propositions.
3. Both theories provide the same kind of solution to the Lottery Paradox. To illustrate this, suppose we have a fair lottery consisting of 100 tickets, represented by a set W of 100 possible worlds and the uniform probability function P over W . Then, on both theories, it is impermissible to believe any non-trivial proposition over W . This is so because, for any non-trivial X over W , there is at least one world outside X which is as probable as any world in X .

But if we coarse-grain W and focus on partitioning $W' = \{w_1, \{w_2 \dots w_{100}\}\}$, then, on both theories, it is permissible to believe $\neg\{w_1\}$ over W' .

4. On both theories, a probability distribution may rationally support more than one belief set *Bel*. Consider the credence function P in Citizen Hannes: $P(\{w_1\}) = 0.26$, $P(\{w_2\}) = 0.4$, $P(\{w_3\}) = 0.28$, $P(\{w_4\}) = 0.06$. Here we have two non-trivial dominant propositions with a probability greater than 0.5: $\{w_2, w_3\}$ and $\{w_2, w_3, w_1\}$. So, we may choose any of these propositions to be the dominant core proposition. If Richard's beliefs are obtained by $\{w_2, w_3\}$, then the corresponding belief set *Bel* would be maximally *brave*: Richard would believe every proposition that is rationally permissible given his credence function. By contrast, if $\{w_2, w_3, w_1\}$ is Richard's least believed proposition, then the corresponding *Bel* would be more

cautious. As far as the dominant core theory is concerned, both choices of B_W are equally permissible.

As we see, the two theories are quite similar. But there are some significant differences between the two. For a start, a dominant proposition with a probability greater than 0.5 is not always a stable proposition. For instance, in Citizen Hannes, $\{w_2, w_3\}$ is a dominant proposition (whose probability exceeds 0.5) but not a stable proposition because $P(\{w_3\}) = 0.28$, which is less than $P(\{w_1, w_4\}) = 0.26 + 0.06 = 0.32$. Therefore the dominant core theory permits believing $\{w_2, w_3\}$, while the stability theory prohibits believing it. For this reason, the dominant core theory agrees with the intuitive analysis of Citizen Hannes: Richard's strongest believed proposition may be the exclusive disjunction that Hannes is either a German citizen or was born in Austria, $Bel(\{w_2, w_3\})$; even if Richard's conditional credence in this disjunction is not greater than 0.5 on the supposition that its more probable disjunct is false: $P(\{w_2, w_3\} | \neg\{w_2\}) < 0.5$.

Because on the dominant core theory it is permissible to believe $\{w_2, w_3\}$, it is immediate that the theory violates Monotonicity: if Bel is obtained via $\{w_2, w_3\}$, then, as we have seen, Bel violates Monotonicity with respect to the inferentially trivial disjunction $\{w_1, w_2, w_4\}$. But what may be far less obvious is that the dominant core theory always satisfies Partial Monotonicity:

Theorem 6.1 (The dominant core theory satisfies Partial Monotonicity):

For any perfectly rational agent with belief set Bel and credence function P , if Bel and P satisfy the dominant core theory, then Bel and P satisfy Partial Monotonicity.⁸⁰

⁸⁰ Proof: suppose for reductio that Bel is obtained via a dominant proposition D , but for some inferentially non-trivial disjunction X , $P(X) \geq P(D)$ and $\neg Bel(X)$. Given that Bel satisfies the dominant core theory, $D \not\subseteq X$.

Footnote continued on the next page.

So, given the dominant core theory, if a rational agent violates Monotonicity with respect to proposition X , we know that X is an inferentially trivial disjunction. Hence, if an agent satisfies the dominant core theory, there won't be any inferentially useful proposition that the agent fails to believe.

Of course, it is true that if an agent's belief set Bel satisfies the dominant core theory, there could still be some inferentially trivial disjunctions in Bel . As I've noted in the previous section, any theory of belief that endorses Deductive Cogency permits believing *some* inferentially trivial disjunctions. Still, the dominant core theory makes it impermissible for an agent to believe any *additional* inferentially trivial disjunctions other than those that deductively follow from her other beliefs.⁸¹

Finally, and most importantly, we can easily verify that the dominant core theory supports the exact same solution to the (diachronic) coordination problem as the stability theory. This solution is articulated in Section 5.3.1 via Theorem 5.2, which establishes that, given the required assumptions, if two equally informed agents credences functions are order

And because the set of all dominant propositions is well-ordered with respect to the subset relation, we know that X is not a dominant proposition. Now, let Y be a disjunct of X (we know that X must a disjunction, otherwise X would not have a greater probability than D) such that $\neg Y = \{w_D, w_X\}$, where $w_D \in D \setminus X$ and $w_X \in X \setminus D$. By the assumptions that we have made, proposition Y must exist. Finally, if X is a non-trivial disjunction, then $P(X|\neg Y) > 0.5$. But, $P(X|\neg Y) = P(\{w_X\})/P(\{w_D, w_X\})$ and we know that $P(\{w_D\}) > P(\{w_X\})$, which entails that $P(X|\neg Y) < 0.5$. Contradiction. Therefore, if Bel satisfies the dominant core theory but violates Monotonicity with respect to X , then X must be a trivial disjunction. As required.

⁸¹ I may also add that, on the dominant core theory, believed trivial disjunctions are more inferentially reliable than non-believed trivial disjunctions: for *any* trivial disjunctions X , such that $\neg Bel(X)$, there is a trivial disjunction Y , $Bel(Y)$ and Y is *more inferentially reliable* than X ; meaning that, for any Z which is a disjunct of both X and Y , $P(X|\neg Z) < P(Y|\neg Z)$.

equivalent relative to relevant partition W , then these agents won't adopt opposing beliefs no matter which proposition from W they learn.⁸² The same coordination theorem can be proved, in exactly the same way, by substituting the stability theory with the dominant core theory. Here is why. The proof of the theorem (as given in Section 5.3.1) requires a bridge principle that prohibits believing a proposition that is not true at the *most probable world*. Trivially, both the stability and dominant core theories have this feature: by definition, stable and dominant propositions over W are true at the most probable world in W . And, in general, any theory with that feature will enable us to provide the same solution to the coordination problem.

This feature of the dominant core theory and the stability theory that support the same solution to the coordination problem is due to their more fundamental feature, which, following Dietrich and List (2018), I call *propositionwise dependence*. According to this principle (of propositionwise dependence), whether an agent believes a proposition is not the sole function of the proposition's probability but also the other propositions that the agent considers. Certainly, whether a world (or a possibility) is the agent's most probable world

⁸² More fully, the theorem goes as follows:

For any two agents with prior credence functions P and P' and prior belief sets Bel and Bel' defined over the same partition W , if these agents (i) satisfy the stability theory, (ii) update their credence functions via Conditionalisation, and (iii) their credence functions are order equivalent, then the following obtains:

For any evidence (proposition) E and proposition X over W and for any permissible posterior belief sets Bel'_E and Bel''_E over W_E (i.e., the set of worlds in W compatible with E), it is not the case that $Bel'_E(X)$ and $Bel''_E(\neg X)$.

depends not only on its probability but also on what other propositions (or possibilities) the agent attends to.⁸³

By contrast, some theories, like the standard Lockean thesis (what we have called the *strong Lockean thesis*), endorses *propositionwise independence*: the principle that the probability of a proposition alone – independent of any other proposition(s) – determines whether an agent ought to believe that proposition.

Because the stability theory and dominant core theory satisfy propositionwise dependence and make rational belief sensitive to an agent's most probable world, they enable us to give the same solution to the coordination problem. Conversely: no theory that violates propositionwise dependence can support the proposed solution to the coordination problem. Hence, given the discussions in Chapters 4 and 5, it seems that no hybrid theory can afford to reject propositionwise dependence (due to the coordination problem).

The reader may enquire whether propositionwise dependence is acceptable. In response, I think that propositionwise dependence requires no additional motivation from what I've already offered in support of the stability and dominant core theories. Propositionwise dependence closely relates to the context-sensitivity of these theories (see

⁸³ While propositionwise dependence is closely related to context-sensitivity (of belief), the two should not be conflated. One can endorse the view that belief is context-sensitive (i.e., belief depends on a set of possibilities the agent attends to) but hold that belief is a function of its probability in a context and nothing else. Such a view is defended by Clarke (2013), who argues that belief is equivalent to credence 1 and that degrees of belief change from context to context. In contrast to Clarke's theory, the stability theory and the dominant core theory endorse both context-sensitivity and propositionwise dependence.

While context-sensitivity does not entail propositionwise dependence, the converse relationship seems to hold. It seems that any theory that is propositionwise dependent is also context-sensitive. Though, nothing hangs on verifying this (alleged) conceptual connection.

the footnote above), and I've argued at length (Section 5.2.3) that context-sensitivity is not as problematic as it may initially seem. So, the reader should judge the plausibility of propositionwise dependence based on the plausibility and fruitfulness of the two theories of rational belief that I've considered: the stability theory and the dominant core theory.

To sum up: I've articulated a precise bridge principle for belief and credence, the dominant core theory that avoids the Citizen Hannes counterexample, satisfies Partial Monotonicity (but violates Monotonicity) and supports the exact same solution to the coordination problem as the stability theory. In the next section, I give reasons for favouring the dominant core theory to the stability theory.

6.4 The Dominant Core Theory vs The Stability Theory

We saw that while the stability theory and the dominant core theory share several important features, the latter is less restrictive than the former and avoids counterexamples like Citizen Hannes. For this reason, the dominant core theory has an important comparative advantage over the stability theory. In this section, I want to argue that no other considerations outweigh this comparative advantage that the considerations about restrictiveness bestow on the dominant core theory.

For a start, it is true that out of the two, only the stability theory satisfies Monotonicity, which is a highly pre-theoretically plausible coherence condition on rational belief and credence. But, as I've shown (in the previous section), the dominant core theory satisfies Partial Monotonicity; for this reason, the violation of Monotonicity that the dominant core theory allows is inferentially insignificant: an agent won't lose any inferentially useful belief by satisfying Partial Monotonicity but violating Monotonicity.

In defence of the stability theory, one may point out that it harmonises with well-known belief revision principles, as given by the AGM axioms (considered in Section 5.5.2).

And, at this point, we have not seen whether the dominant core theory harmonises with any independent principles of belief revision.

But it is far from obvious whether the harmonization of the stability theory with the AGM axioms is a feature rather than a bug. Here is why. As we saw, at the heart of AGM is the idea of *minimal change* or *doxastic conservatism*:

A rational agent should revise her prior belief set Bel by evidence E in a way that her posterior belief set Bel_E is as similar as possible to Bel .

If an agent learns new evidence E which is consistent with Bel , i.e., $\neg Bel(\neg E)$, then the AGM axioms say the following:

R1: If Bel and E are consistent, then $Bel * E$ is equivalent to $Cn(Bel \cup \{E\})$.⁸⁴

So, according to R1, when Bel and E are consistent, an agent's posterior belief set $Bel * E$ is obtained by *adding* E to Bel and closing the obtained set, $\{E\} \cup Bel$, deductively.

Now the Citizen Hannes example provides a counterexample to R1. In this example, Richard believes that Hannes is either a German citizen (G) or was born in Austria (A), but not both: $Bel(G \vee A)$. His main reason for believing $G \vee A$ is $G \wedge \neg A$. Now suppose that he learns new information that $G \wedge \neg A$ is false: $\neg(G \wedge \neg A)$. By assumption, Richard does not believe the negation of $\neg(G \wedge \neg A)$: $\neg Bel(G \wedge \neg A)$. Hence, by R1, Richard must continue believing $G \vee A$, even if he learns that his main reason for believing the conjunction is false. And this seems wrong.

So, Citizen Hannes is a counterexample for both the stability theory and AGM. This should not surprise us as the two theories are closely related, as we have seen.

⁸⁴ R1 is logically equivalent to the following AGM axioms: *Preservation* and *Inclusion*. See Section 5.5.2 for a detailed discussion.

Now, it is easy to verify that the dominant core theory violates R1. On the dominant core theory, Richard should no longer believe $G \vee A$ when he learns that $G \wedge \neg A$ is false (even if this new evidence is consistent with her prior belief set Bel). Hence, the dominant core theory violates AGM.

But does the dominant core theory harmonise with any independent principles of belief revision? The answer is *yes*. We know this from work by Lin and Kelly (2012, 2021). They articulated a bridge principle based on the pair-wise comparisons between *worlds* and *worlds with maximal probability*. They call the resulting bridge principle *the odds threshold rule*. We can explain the rule as follows: let w_{Max} be the most probable world over W , relative to a probability function P . Then, according to the odds threshold rule, a rational belief set Bel should be obtained by the strongest believed proposition B_W such that B_W is determined via a threshold s in the open interval $(0,1)$ in the following way:

$$B_W = \{w_i \in W : \frac{P(\{w_i\})}{P(\{w_{Max}\})} > 1 - s\}$$

As we see, if s is sufficiently close to zero, then only the most probable world w_{Max} will be in B_W . It is easy to see that the odds threshold rule reflects the order of worlds over W : each world in B_W must be more probable than each world outside it. Hence, on the odds threshold rule, B_W must be a *dominant* proposition. So, surprisingly, Lin and Kelly's rule is extensionally equivalent to the dominant core theory.⁸⁵ While extensionally equivalent, I think the dominant core theory has a much simpler form (I may be biased, as I articulated it).

⁸⁵ To my surprise, I've discovered the logical equivalence between the two theories after it was pointed to me (by an anonymous reviewer) that my definition of a dominant proposition was identical to Cantwell and Rott's (2019) notion of *coherence*. As Cantwell and Rott also present a theory logically equivalent to Lin and Kelly's theory, and hence logically equivalent to my theory.

Now, Lin and Kelly (2012, 2021) showed that the odds threshold rule harmonises with the so-called Shoham-Boutilier (SB) revision (due to Yoav Shoham 1987 and C. Boutilier 1996). We don't need to go into the details of SB. Suffices it say that SB is similar to AGM, except that it does not satisfy the axiom of Preservation (or R1), which, as we saw, is open to counterexamples like Citizen Hannes.

So, in summary, the dominant core theory captures important plausible aspects of the stability theory, coheres with the SB revision, and is less restrictive than the stability theory. For these reasons, I submit that the dominant core theory is a superior alternative to the stability theory.

6.5 Concluding Discussion

We have seen (in Section 6.3) that substituting the stability theory with the dominant core theory does not affect the coordination theorem proved in Chapter 5. This is because, on both theories, a proposition X is permissible to believe over W iff X is true at the most probable world in W . So, given that two credence functions are orderer equivalent over W , these functions won't license opposing beliefs over W_E (when we update them by E), for any proposition E over W .

So, the choice between the dominant core theory and the stability theory is inconsequential to the main arguments of this dissertation. So the hybrid approach to the Uniqueness problem does not need to assume a specific, contentious view of rational belief and credence ought to interact. We can derive all the key conclusions about Hybrid Impermissivism by using any of the two main theories of belief-credence interaction that satisfy Deductive Cogency: the stability theory and the dominant core theory (equivalent to Lin and Kelly's threshold odds rule).

Further, we can also show that the assumption of Deductive Cogency is unnecessary to derive the coordination theorems.

Consider the following version of the stability theory that violated Deductive Cogency:

*The Stability Theory**:

For any agent with a belief set Bel and credence function P , defined over a finite partition of possibilities W , Bel and P should be such that:

There is a stable proposition Y over W ,

if $P(Y) = 1$, then Y is the least subset of W with probability 1; and:

for all X , $Bel(X)$ iff $P(X) \geq P(Y)$ and X is not a trivial disjunction.

The stability theory* includes one extra proviso that no believed proposition is inferentially trivial. As Deductive Cogency permits believing a trivial disjunction, it is easy to see that the stability theory* violates Deductive Cogency. Let us illustrate with a simple example, involving probability distribution over the partitioning W of six possibilities:

Possible worlds	P
w_1	$P(\{w_1\}) = 0.30$
w_2	$P(\{w_2\}) = 0.26$
w_3	$P(\{w_3\}) = 0.25$
w_4	$P(\{w_4\}) = 0.09$
w_5	$P(\{w_5\}) = 0.08$
w_6	$P(\{w_6\}) = 0.02$

Table 6.1

There are two non-trivial stable propositions over this W : $\{w_1, w_2, w_3\}$, $\{w_1, w_2, w_3, w_4, w_5\}$. Suppose $B_W = \{w_1, w_2, w_3\}$. Now $B_W \subseteq \{w_1, w_2, w_3, w_6\}$, and $\{w_1, w_2, w_3, w_6\}$ is a trivial disjunction: $P(\{w_1, w_2, w_3, w_6\} | \neg\{w_1, w_2, w_3\}) \approx 0.1$. And because $\{w_1, w_2, w_3, w_6\}$ is a trivial disjunction, it is impermissible to believe it on the stability theory*, even if $B_W \subseteq$

$\{w_1, w_2, w_3, w_6\}$. So, the stability theory* violates deductive cogency. On this theory, given

$B_W = \{w_1, w_2, w_3\}$, the following is the uniquely rational belief set:

$$Bel = \{\{w_1, w_2, w_3\}, \{w_1, w_2, w_3, w_4, w_5\}, W\}$$

Hence, as we see, each believed proposition is a stable proposition over W .

This is not an incidental feature of the above-considered probability distribution. On the stability theory*, all and only stable propositions are permissible to believe. This is so because of the following theorem:

Theorem 6.2 (All and only stable propositions can reliably support disjunctive syllogism):

For any non-basic (i.e., non-singleton) proposition X over W , X reliably supports disjunctive syllogism iff X is a stable proposition.

The proof is quite simple: by definition, X is a stable proposition when all worlds in X are more probable than $\neg X$. So, if we exclude any world from X (i.e., if we assume that some disjunct(s) of X are false), then it will still be true that the remaining worlds in X will be more probable than \neg . So, for any proper subset Y of X , $P(X|Y) > 0.5$. To illustrate this, consider the above table again. We know that $\{w_1, w_2, w_3\}$ is stable. Now, assume that the most probable two worlds in X are false:

$$P(\{w_1, w_2, w_3\} | \neg\{w_1, w_2\}) = \frac{P(\{w_3\})}{P(\neg\{w_1, w_2\})} = \frac{0.25}{0.49} \approx 0.51$$

Therefore: a stable proposition can always reliably support disjunctive syllogism (or classical modus ponens or modus tollens inference). And, again, substituting the stability theory with the stability theory* won't affect the two coordination theorems from Chapter 5. This is because two order equivalent credence functions over any fixed W won't regard contradictory propositions as stable (as each stable proposition must be true in the most probable world in W).

The same conclusion can be derived if we restate the dominant core theory in the same way as we have restated the stability theory (this should be obvious for the reader, so that I won't do it).

In conclusion: there are several plausible bridge principles similar to the stability theory, such that the choice between them does not affect the key conclusions from Chapter 5 about Hybrid Impermissivism.

7 Belief Within Incomplete Doxastic Rankings

In Chapters 5 and 6, I have developed hybrid theories by working with complete probability distributions over a partition of possibilities or possible worlds. I also assumed that a rational agent's doxastic order of worlds is *total*, meaning that, for every world w_n and w_j in the agent's W , the agent either considers w_n at least as plausible as w_j or w_j at least as plausible as w_n . In other words, there are *no gaps* in the agent's ranking of words in terms of their plausibility.

But, in some evidential situations, it seems that an agent's doxastic ranking of worlds can be *incomplete*. I will illustrate this by discussing Arthur Edington's famous experiment about light deflection that tested the general theory of relativity (*GTR*) against Newtonian theory (*NT*). Consider the following four possibilities, where O denotes Edington's observations: w_1 corresponds to $GTR \cap O$, w_2 to $GTR \cap \neg O$, w_3 to $\neg GTR \cap O$, and w_4 to $\neg GTR \cap \neg O$.

We can provide a pairwise comparison of *some* worlds in a fully uncontentious, objective manner by presupposing the information available to an "ideal" (or perfectly rational) physicist at the time of Edington's experiment (represented by credence function P). Because O is one of the important predictions of *GTR*, we know that the likelihood of O on the supposition of *GTR* is higher than the likelihood of $\neg O$ on the supposition of *GTR*. Now, by the axioms of the probability theory:

$$P(\{w_1\}) = P(GTR \cap O) = P(GTR) * P(O|GTR)$$

$$P(\{w_2\}) = P(GTR \cap \neg O) = P(GTR) * P(\neg O|GTR)$$

And as $P(O|GTR) > P(\neg O|GTR)$, we know that the ideal physicist ranks w_1 higher than w_2 .

As we see, to determine the ranking between w_1 and w_2 we only need to know the comparative values of the *ordinary likelihoods*, i.e., the likelihoods of an observation on the

supposition of a single, specific hypothesis. By contrast, to determine the ranking between, say, w_1 and w_3 we need to know or, more realistically, approximate the value of the *catchall* likelihood: $P(O|\neg GTR)$. It is called a catchall likelihood because it includes a catchall hypothesis, $\neg GTR$, asserting that *some* alternative to GTR is true. More fully, to calculate the value of $\{w_3\}$ we need to estimate the prior probability of the catchall hypothesis $\neg GTR$ and the probability of the catchall likelihood $P(O|\neg GTR)$:

$$P(\{w_3\}) = P(\neg GTR \cap O) = P(\neg GTR) * P(O|\neg GTR)$$

For the sake of argument, suppose that the probability of $P(\neg GTR)$ could be estimated in a non-subjective manner (which, in itself, is a very substantive assumption). Still, there does not seem to be an objective, uncontentious way to estimate the likelihood of the observation O on the supposition that GTR is false.⁸⁶ Let me explain. From all we know, there may be many specific alternatives to GTR that physicists have not discovered or conceptualised yet (and, probably, they never will). And from all we know, some of these alternatives could explain or predict Edington's observation, O , just as well, or even better than GTR . So, even if $P(O|GTR)$ is greater than $P(O|NG)$, from this we cannot conclude that $P(O|GTR) > P(O|\neg GTR)$.

It is true that the value of a catchall likelihood, $P(E|\neg H)$, can be *calculated* by knowing the values of *ordinary likelihoods* of E on each (mutually inconsistent) specific alternatives to H and the *prior distribution* over these alternatives. In symbols, if $Alt_1, Alt_2, Alt_3 \dots Alt_n$ denote the specific alternatives to H , then

$$P(E|\neg H) = \frac{\sum_i P(E|Alt_i)P(Alt_i)}{P(\neg Alt_i)}$$

⁸⁶ In Section 2.2.4, we have already discussed a general difficulty in estimating the values of some catchall likelihoods in an objective or intersubjectively justified way.

So, mathematically, catchall likelihoods are reducible to priors and ordinary likelihoods: that is, if one knows the values of priors and likelihoods, then one can calculate the value of the catchall. But as we have seen, scientists very often do not know all specific alternatives to a general hypothesis, like *GTR* (and even if they knew all specific alternatives to *GTR*, it is still required to specify a prior distribution over these alternatives; and this, in itself, is highly problematic).⁸⁷

The cases where the values of catchall likelihoods cannot be estimated objectively abound in science and philosophy (we consider more of such examples in this chapter). Hence, in many cases, it does not seem to be possible to determine an objective or interpersonally justified complete order of worlds.

The failure to rank possibilities objectively can make it impossible to talk about belief, degrees of belief, and confirmation/evidential support in an objective manner. Take our Addington example again. Because it is impossible to pairwise rank w_1 and w_3 objectively, it cannot be decided whether the observation O makes w_1 more probable than not; hence it cannot be decided whether learning that O should make $\{w_1\}$ believable. After all, the ideal physicist has no idea how to estimate the probability of $w_3 = \neg GTR \cap O$ either in absolute or comparative terms. For this reason, no probability distribution over W can represent an agent's incomplete ranking of possibilities over W .

This chapter is concerned with how to think about rational belief, degrees of belief, and evidential support in the contexts where the doxastic order of worlds is incomplete.

In the next section, after discussing some preliminaries, I introduce the so-called *likelihoodist* framework, primarily designed to analyse the notion of evidential support. The

⁸⁷ While the so-called *problem of priors* is extensively discussed in philosophy, arguably, the *problem of catchall likelihoods* pose an even greater problem for any non-subjective epistemologies.

overall aim of this chapter is to argue that the likelihoodist framework can be extended to provide guidance for *comparative* belief in contexts where the available evidence does not justify a complete ranking of possibilities.

7.1 Introduction: Bayesianism and Likelihoodism

As we have seen, sometimes our evidence cannot determine even a coarse-grained doxastic order over a set of possibilities. In other words, an agent's evidence E in her context of reasoning may not decide whether a possibility, w_n , is at least as probable as some other possibility, w_m .

How should we think about rational belief and degrees of belief in such settings? The standard Bayesian strategy is to *transform* incomplete doxastic orders to complete doxastic orders. One way to do this is to appeal to an agent's *subjective confidences* about each world in W . This approach is associated with the so-called *subjective* (or personalist, or permissive) Bayesianism. According to subjective Bayesianism, scientists should assign probabilities to hypotheses according to their personal degrees of belief, as long as these probabilities do not violate the axioms of probability. For instance, in our Edington example, the ideal agent may subjectively consider $P(O|GTR)$ to be greater than $P(O|\neg GTR)$, and assign the former a greater probability than the latter.

Certainly, such a subjectivist approach to catchall likelihoods is inconsistent with the main arguments of this dissertation. If (equally informed) rational agents' doxastic orders over the same set of possibilities could differ, then, trivially, the hybrid impermissivist view is doomed.

But, independent of this dissertation, the subjectivist approach to catchall likelihoods seems highly problematic; more problematic than the subjectivist approach to priors. Here is why. When scientists convey that a particular piece of evidence supports or confirms a

hypothesis, such claims are commonly understood as objective and not a mere expression of their subjective opinions. As Sober (2008) articulates the point:

if we think of the likelihoods as merely reflecting subjective degrees of confidence, someone might assert, as an autobiographical remark, that the *GTR* has a higher likelihood than its negation; but someone else, with equal autobiographical sincerity, could assert the opposite. And both would be right if the probabilities involved were merely subjective. In science, we want more than this.

Certainly, some (e.g., Howson and Urbach 2006) have defended the subjective account of confirmation in the face of such difficulties. It goes beyond the dissertation's scope to evaluate this debate. It suffices to say that this subjectivist approach to confirmation is highly unintuitive and should be avoided if possible.

Another Bayesian strategy to transform an incomplete doxastic order to a complete order is to appeal to some version of the Principle of Indifference:

The Principle of Indifference for Doxastic Ranking: For any agent and a set of possibilities W , if the agent's evidence does not determine the comparative ranking between some w_n and w_j in W , then the agent should consider these possibilities equally probable.

This strategy is associated with *objective Bayesianism*, which appeals to such indifference considerations to derive inter-personally justified probability distributions over any W .

Many (including myself) find such principles of indifference problematic. Take Edington's example again: if the ideal agent's evidence is insufficient to decide whether $w_1 = GTR \cap O$ is more probable than $w_3 = \neg GTR \cap O$, then why should she consider these

possibilities equally probable? This verdict does not seem more objective than to (probabilistically) favour one possibility over the other.⁸⁸

So, both Bayesian strategies to “fill” incomplete doxastic orders seem problematic. But what is an alternative? The alternative I want to consider is associated with the *likelihoodism* framework within the philosophy of statistics. This alternative approach *concedes* that, in many cases, it is impossible to provide an objective ranking of worlds over a relevant set of possibilities. Still, it is argued that, even in such cases, we can form rational *comparative beliefs* without appealing to subjective credences or a principle of indifference. Hence, the impossibility of assigning objective numerical values to likelihoods and priors is not as detrimental for belief guidance as it may initially seem.

What is likelihoodism? Following Sober (2008), I characterise likelihoodism as a *fallback position*. On likelihoodism, if the available evidence provides an objective justification for assigning values to priors and (catchall) likelihoods, then we should think about degrees of belief and related notions, like belief and confirmation, in a standard Bayesian way. But if the evidence lacks this feature, then the talk of belief, degree of belief and confirmation is unjustified. For instance, because we cannot estimate in an objective, uncontentious manner the value of catchall likelihood $P(O|\neg GTR)$, we cannot conclude that O confirms or increases the probability of GTR . But, on likelihoodism, we can still conclude that O supports GTR over its specific alternative, NT : as the likelihood of O is higher if GTR is true than if NT is true. This conclusion about comparative support follows from the so-

⁸⁸ My aim here is not to argue or demonstrate that an objective Bayesian approach to the problem of catchall likelihoods is wholly unpromising. Rather, I only point out a well-known difficulty associated with this approach and, as the chapter unfolds, explore an alternative impermissivist approach that avoids this difficulty.

called *Law of Likelihood* according to which an outcome, E , is evidence for a hypothesis, A , over its competitor, B , when E is more likely if A is true than if B is true.

In the rest of this chapter, I argue that a broadly likelihoodist framework can provide substantive guidance for *comparative* belief in cases where there is no objective basis for ranking each possibility over an agent's (set of possibilities) W . By comparative belief, I mean a belief of the form “ A is more probable than B ” or “ A and B are equally probable.” So, I shall argue that rational comparative belief can still be formed in cases where the available evidence does not justify a complete ranking of relevant possibilities, without invoking Bayesian subjective probabilities or the principle of indifference. For simplicity and convenience, instead of incomplete doxastic rankings, I will be concerned with *incomplete credence functions*. So, I'll be concerned with cases where an agent's evidence does not justify objective or interpersonally justified values for some relevant hypotheses (and their respective likelihoods).⁸⁹

Following Salmon (1990), my main argumentative strategy is to move from *non-relational probabilities* of individual hypotheses to *comparative evaluations* of competing hypotheses. A non-relational probability of a hypothesis, A , is commonly represented by a point-valued probability: e.g., when A is assigned a probability of, say, 0.6 (I also allow non-relational probabilities to be represented by *intervals* or *sets* of probability distributions). By contrast, comparative evaluations of competing hypotheses A and B do not require us to assign any non-relational probabilities to them. For instance, scientists may not have an objective basis for assigning non-relational probabilities to A and B , but they can still rationally judge that A and B are roughly equally plausible.

⁸⁹ As we have seen, by estimating the values of priors and catchalls over W we impose certain ranking of worlds over W .

Section 7.2 gives a general, broad-brush overview of the main aspects of likelihoodism. Before I develop a likelihoodist account of comparative belief within a setting of incomplete credence distributions, I'll first address a general argument that likelihoodism cannot provide guidance for belief. Section 7.3 gives a precise statement of the distinct problem that likelihoodism faces concerning belief guidance. Section 7.3.1 shows that the most detailed and rigorous version of this criticism, as put forward by Gandenberger (2016), is unsuccessful. In Section 7.4, I put forward a positive, likelihoodist account of guidance for comparative belief. This account utilises the *ratio form of Bayes' theorem*. I will illustrate both the applicability and limits of likelihoodist belief guidance by analysing two examples: one from cognitive neuroscience and one from philosophy. I conclude in section 7.5.

7.2 Two Tenets of Likelihoodism

Likelihoodism can be seen as an attempt to overcome the frequentism-Bayesian controversy within the philosophy of statistics by paving the way between “the illogic of the frequentists and the subjectivity of the Bayesians” (Royall 1997, XIV). So, to explain likelihoodism adequately, we should contrast them to the *frequentist* and the *Bayesian* paradigm that have been the two dominant approaches to statistical inference for the past 60 years or so.

A central procedure of frequentist statistics is the so-called *Null Hypothesis Significance Testing* (NHST). A significance test starts with a hypothesis, called the “null hypothesis”, which is examined against some relevant outcome or data. Simply put, NHST says that if a null hypothesis renders certain outcomes as highly improbable and if such an improbable outcome occurs, then the null hypothesis should be rejected.

While the guiding idea behind NHST seems plausible, many have found the method fundamentally defective.⁹⁰ Moreover, in certain fields of science, primarily in the social, behavioural, and biomedical sciences, some of the important results that relied on NHST failed to be replicated.⁹¹ This new evidence showing the significant lack of replication in these fields puts an additional strain on frequentism, so much so that there is an increasing call for some kind of statistical reform.

Many critics of frequentism see Bayesianism as providing superior methods of data analysis (e.g., Dienes 2011; Wetzels et al. 2011; Kruschke 2013). The key characteristic of Bayesianism is the use of the so-called *prior probabilities*. Unlike frequentism, Bayesian theory requires a probability distribution over both the sample space and statistical hypotheses.⁹² A probability distribution over statistical hypotheses (prior to considering the relevant evidence) is called *prior distribution*. A prior distribution encodes how likely the competing hypotheses are before the relevant evidence comes in.

⁹⁰ An immediate problem with NHST is that it embodies a defective form of inductive reasoning; even if a hypothesis renders a certain observation as unlikely, the observation might still support the hypothesis; this is so because the observation might be *even more improbable* if the hypothesis is false. To demonstrate this, suppose two individuals share a copy of a rare allele; only 1 in 10 000 have it. As siblings share half of their alleles on average, $P(\text{rare allele}|\text{siblings}) = 0.5 * 0.0001$, which is a very small number. However, if the two individuals are unrelated, the probability is much lower: $P(\text{rare allele}|\text{unrelated}) = 0.0001 * 0.0001$. Hence, the data supports the sibling hypothesis, even if the hypothesis renders the data quite unlikely. See Sober (2008, 48-58) for an accessible discussion.

⁹¹ For a general philosophical discussion of the replication crisis see Romero (2019). For an explanation/diagnosis of the crisis see Bird (2018).

⁹² NHST only considers probability distribution over sample space (i.e., how likely an outcome is on the supposition of a statistical hypothesis) and not over statistical hypotheses themselves (i.e., how likely the relevant statistical hypotheses are).

Certainly, the indispensability of priors in data analysis is the Achilles heel of Bayesian methods. The problem is that, in many contexts, there seems to be no objective, uncontroversial way to fix priors. And due to this unmistakably subjectivist component in Bayesian methods, many are quite reluctant to give up on the traditional frequentist methods.⁹³ While some Bayesians have proposed various theories for grounding “objective” prior probabilities, none of these proposals has been generally accepted.⁹⁴ Hence the problem of the subjective priors continues to haunt Bayesian methods.

Like frequentism (and unlike Bayesianism), likelihoodism requires only a probability distribution over the sample space and not over hypotheses themselves. And like Bayesianism (and unlike frequentism), likelihoodism holds that the impact of evidence on any two hypotheses is wholly determined by the likelihoods of these hypotheses.⁹⁵ Hence, likelihoodism endorses some true-and-tried principles from both frequentism and Bayesianism, without relying on controversial NHST or subjective probabilities.⁹⁶

The core of likelihoodism consists of (i) a *comparative, relational* conception of evidential support and (ii) the likelihood ratio measure of the degree of relational support.

⁹³ There have been some attempts to marrying the two paradigms together, in a unified frequentist-Bayesian theory. But so far, the most serious disputes between the two approaches are still raging. See Mayo (2018) for a lengthy discussion.

⁹⁴ For a positive, systematic account of objective or impermissible Bayesianism see Williamson (2007, 2010). For a critical discussion, see Meacham (2014).

⁹⁵ This view on the impact of evidence is called the *Likelihood Principle*. Frequentism is in tension with the principle as it allows various non-likelihood related factors to influence the impact of evidence. For a detailed discussion of the Likelihood Principle see Berger and Wolpert (1987) and Ganderberger (2014).

⁹⁶ For an influential statement of the likelihoodist program see Royall (1997). For a more philosophically rich discussion, see Sober (2008) and Bandyopadhyay et al. (2016).

The first is qualitative and the second is a quantitative aspect of likelihoodism. In what follows, I will characterise and explicate each of these aspects, starting with the likelihoodist view of (evidential) support.

To explain the likelihoodist view of support, it's useful to contrast it with a more orthodox, non-relational view. A theory of support is non-relational when it defines support for an individual hypothesis, without contrasting the hypothesis to its alternative, competitor hypothesis. For instance, consider the standard Bayesian view, which I call *Support-IP* (IP for Increase in Probability):

Support-IP: For any hypothesis H , evidence E , and personal probability function P , E supports H relative to P iff $P(H|E) > P(H)$.

Support-IP defines support in terms of the increase-in-probability relation (or confirmation). And Support-IP is a non-relational view because support for a hypothesis is defined without appealing to any *competitor hypothesis*.

By contrast, the likelihoodist view of support is *inherently relational*, as it requires two competitor hypotheses to define the relation of evidential support. This view is expressed by the so-called *Law of Likelihood* (LL), which roughly says that for any two competitor hypotheses A and B , E supports A more strongly than B iff E is more likely on the supposition that A than on the supposition that B . More precisely:

LL: For any two competitor hypotheses A and B , E supports A over B iff A confers greater probability on E than B does: $P(E|A) > P(E|B)$.

Why accept LL over its Bayesian competitors? The main strength of LL – according to its supporters – is that it provides an objective, inter-personally justifiable criterion for evidential support. LL defines support in terms of two likelihoods; i.e., the probabilities of the following form: $P(\text{Evidence}|\text{Hypothesis})$. A likelihood encodes the empirical content of a

hypothesis; that is, what the hypothesis says about evidence. For instance, let h = “25% of philosophy undergraduates are introverts”, and let e = “randomly chosen philosophy undergraduate is an introvert”. There is a certain logico-conceptual relationship between h and e that is articulated by likelihood $P(e|h)$. And even if we have no clue about the prior probability of h and e , the likelihood of $P(e|h)$ is still objectively given: $P(e|h) = 0.25$.

The example illustrates what Hawthorne (2005, 278) has called the *publicness of likelihoods* (for a detailed discussion of likelihoods and their public character see Section 2.2.3-2.2.3). So, even if two agents disagree about the prior probability of h , they can still agree on the value of the likelihood, $P(e|h)$.⁹⁷

Fixing likelihoods is not always as easy as in the above example. But even the so-called *subjective* Bayesians – that is, Bayesians who allow the multitude of coherent prior distributions as rationally permissible – grant that likelihoods can be objectively well-grounded in many scientific contexts (Edwards et al. 1963).

In addition to LL, likelihoodists provide a measure of (comparative) evidential support, which quantifies the basic idea behind LL; so that the degree of evidential support between A and B is defined as the ratio of their respective likelihoods.

⁹⁷ By the standard definition of conditional probability, likelihoods are still mathematically related to priors: as $P(E|H) = P(E \text{ and } H)/P(H)$. But this mathematical connection between likelihoods and priors does not imply that we cannot make an independent sense of $P(E|H)$, without appealing to the prior probability of H . For one thing, there is an important logical asymmetry between $P(E|H)$ and $P(E \text{ and } H)$ and $P(H)$. Knowing the values of $P(E \text{ and } H)$ and $P(H)$ fixes the value of $P(E|H)$. But not the other way around. So we can make an independent sense of $P(E|H)$ without assuming that the prior probabilities are known or well-defined. For a more detailed discussion see Sober (2008, 38-41).

Relational Measure of Support: The degree to which evidence E supports a hypothesis A over its competitor B equals the ratio of their respective likelihoods. In symbols:

$$\frac{P(E|A)}{P(E|B)}$$

Ratios of likelihoods (R_L , for short) has useful mathematical properties. Whenever the data is more likely on A than on B , the R_L is always greater than 1.⁹⁸ And the better E fits A over B , the greater the ratio. Following Royall (1997), it is common to postulate an arbitrary cut-off point for characterising weak and strong evidence. For instance, we can say that if $1 < R_L < 8$, then E provides *weak* evidence for A . And if $R_L \geq 8$, then E provides *strong* evidence for A .⁹⁹

The combination of LL and the measure of relational support, R_L , comprises the core of likelihoodism. Many (e.g., Fitelson 2007, 2011; Mayo 1996, 2018) have criticised these core principles on various grounds. I will not address any potential difficulties with either LL or the likelihoodist measure of support. Rather, the focus is on the applicability of these principles to the question of *belief guidance*.

The next section gives a detailed statement of the problem of belief guidance for likelihoodism.

⁹⁸ Except when $P(E|B) = 0$.

⁹⁹ The reader should not attach too much significance to the cut-off point 8. Certainly, whether evidence E provides strong evidence for A over B is a context-sensitive matter and depends on the evidence and hypotheses in question (Bandyopadhyay et al. 2016, 24).

7.3 The Problem of Belief Guidance

It has long been recognised that likelihoodist methods for interpreting data as evidence do not, by itself, determine what one ought to believe. The point is well-illustrated by Royall (1997, 2-4) by distinguishing three types of questions regarding the analysis of evidence:

(Q1) What does the present evidence support?

(Q2) What should you believe in light of your present evidence?

(Q3) What should you do in light of your present evidence?

The core of likelihoodism only applies to the first question. By contrast, answering questions (Q2) and (Q3) require more than the information about the likelihoods. To illustrate this, consider a physician, “*you*”, who investigates whether a patient, “*Eve*”, has a skin disease. Eve has taken a test, and the result came up positive. The probability of a true-positive is quite high, 95%, and the probability of a false-positive is quite low, 5%. From this information, we can already answer Royall’s first question by using the Law of Likelihood (LL): the test result strongly supports the hypothesis that Eve has the skin condition over the hypothesis that she does not have it.¹⁰⁰

But this information is insufficient to answer either question (Q2) or (Q3). To answer (Q2), you also need to know the prior probability of the disease. If the disease is quite rare and only 1 in 10 000 people have it, then the test result does not license the belief that Eve

¹⁰⁰ $P(+ \text{ result} | \text{Eve has the disease}) = 0.95$ and $P(+ \text{ result} | \text{Eve doesn't have the disease}) = 0.05$. So, the likelihood ratio is $0.95/0.05 = 19$. Hence, the positive test result provides strong evidence that Eve has the disease.

has the disease. So, to answer (Q2), you need to know both the relevant likelihoods and prior probabilities.

Regarding (Q3): whether you should give any medication to Eve, based on the test result, depends not only on your probabilities but on relevant *utilities*. If the common medication against the disease is harmless, then you can reasonably prescribe it to Eve, even without knowing the exact prior probability of the disease. And, if the medication can be harmful to a healthy person, you would prescribe it only when you are quite certain that Eve has the disease.

To sum up then: even if E strongly supports A over B this won't imply that E justifies either categorical belief in A or comparative belief in A over B (more on this in the next section). Hence, the core of likelihoodism only applies to question (1) and not to questions (2) or (3).

But how does a likelihoodist answer the belief question? As one of the motivating ideas of likelihoodism is to avoid the subjectivity of Bayesianism, likelihoodists cannot rely on *subjective priors* to guide beliefs. So, instead, they should rely on *objectively well-grounded priors*.

Generally speaking, there are two broad strategies for grounding objective priors: by appealing to (i) empirical information about frequencies or (ii) some a priori principle (e.g., the so-called *Principle of Indifference*). However, as I discuss next, both strategies are problematic for likelihoodists (Gandenberger 2016).

Regarding the first strategy: it is widely accepted that frequency data can provide objective justification for fixing prior probabilities. Frequency information is often *out there*, independent of our knowledge, as when a certain disease has some objective incidence rate in the population. For instance, the incidence rate of TB in England is approximately 9.2 per

100,000. And we can estimate the prior probability of a randomly selected individual in England to have TB, based on this frequency data.

Using empirically informed priors to guide belief seems to meet the likelihoodist standard of objectivity. But what if such priors are unavailable? One, popular likelihoodist position is that, when empirically well-grounded priors are unavailable, the only rational doxastic response is a suspension of judgment. Such a view about rational belief is well-summarised and endorsed by Sober (2008, 32):

When prior probabilities can be defended empirically, ... you should be a Bayesian. When priors and likelihoods do not have this feature, you should *change the subject*. In terms of Royall's three questions ..., you should shift from question (2), which concerns what your degree of belief should be, to question (1), which asks what the evidence says.

Unfortunately, though, Sober's proposal is unsatisfactory. Sober himself points out that in many cases, empirically informed priors are simply unavailable. As he (2008, 26) articulates the point:

There is a world of difference between this quotidian case of medical diagnosis and the use of Bayes' theorem in testing a deep and general scientific theory, such as Darwin's theory of evolution or Einstein's general theory of relativity. ... When we assign prior probabilities to these theories, what evidence can we appeal to in justification? We have no frequency data as we do with respect to the question of whether *S* has tuberculosis. If God chose which theories to make true by drawing balls from an urn (each ball having a different theory written on it), the composition of the urn would provide an objective basis for assigning prior probabilities, if only we knew how the urn was composed. But we do not, and, in any event, no one thinks that these theories are made true or false by a process of this kind.

Sober's view about the scarcity of frequency data is the majority view in the philosophy of science and statistics; as most would agree that a prior probability assignment cannot be

defended empirically in many cases of interest. Hence, the proposal that the talk of belief is inappropriate in the absence of frequency data seems to lead to a sceptical view of science, where scientific theories and models are rarely useful for guiding belief.

So, can likelihoodists pursue an alternative strategy and appeal to some a priori principle(s) to ground objective prior probabilities? This strategy is also problematic for likelihoodists, as they are generally sceptical about the prospects of grounding priors on *a priori* principles.¹⁰¹

To illustrate this, let us consider the most prominent a priori rule for fixing priors, the so-called *Principle of Indifference* (PoI). PoI roughly says that if you have no reason to favour a proposition over its competitor, then you should assign equal probabilities to them. More generally and precisely:

PoI: Let U be a finite set of all mutually exclusive and exhaustive hypotheses; if an agent has no evidence that favours any member of U over any other, then for all x in U , $P(x) = \frac{1}{|U|}$, where $|U|$ is the cardinality of U .

Likelihoodists are sceptical towards PoI for two separate but interconnected reasons. First, it's a common practice for scientists to consider a handful of competitor hypotheses, at any given time. In most cases, no one thinks that the considered hypotheses exhaust the space of all serious possibilities. So, scientists rarely know all members of the set of realistic hypotheses. But the application of PoI depends on such a set, whose members and cardinality are explicitly known. For instance, consider the contemporary theories of quantum gravity. There are just a couple of well-articulated theories of gravity, and no working physicist would think that these hypotheses exhaust the space of all possible realistic hypotheses. Of

¹⁰¹ See Sober (ibid., pp. 27-28).

course, one can negate the disjunction of the competing hypotheses, and hence fill the space of all possibilities. But this manoeuvre leads to a second problem. The problem is that there is more than one way to carve this logical space. For instance, assume that scientists focus on only three *specific* competitor theories of quantum gravity: Q_1 , Q_2 , and Q_3 . So, in total, they must consider four competitor hypotheses: $\{Q_1, Q_2, Q_3, \text{ and } \neg(Q_1 \vee Q_2, \vee Q_3)\}$. Assuming that one is indifferent between these four hypotheses, PoI mandates to assign the probability of $1/4$ to each. But it is possible to carve the space of possibilities in a more coarse-grained or fine-grained manner (for instance, by introducing another specific theory of quantum gravity). Different carvings would have licensed different priors. PoI, in itself, does not settle which carvings should be favoured.¹⁰²

All such a priori rules for deriving priors are relative to the set of competitor hypotheses; hence the two problems I've mentioned are not restricted to *PoI* and apply to other *a priori* rules for deriving priors.

To wrap up the above: according to the standard likelihoodist position, frequency data, essentially, is the only admissible evidence for grounding priors for scientific hypotheses. And as such data is unavailable for most scientific hypotheses, the likelihoodist methods seem to be practically useless for science.

Of course, one can simply deny that the lack of belief guidance is a problem. To paraphrase Sober, when empirically grounded priors are unavailable, one must simply change

¹⁰² While likelihoodists think that PoI is problematic even with *discrete* cases, there is also a well-known Bertrand's paradox that poses problems for PoI with respect to *continuous* probabilities (where one cannot straightforwardly appeal to the "finest" partition of the space of possibilities). Some Bayesians (e.g., Williamson 2007, 2010) have provided novel, nuanced defences of PoI. It is beyond the scope of the chapter to evaluate these defences as I'm solely concerned with why likelihoodists think that PoI is wrong.

the subject and answer the evidence question instead of the belief question. I won't argue that such a response is illegitimate. But I don't expect that this response would convince the critics. Hence, I shall proceed by presupposing that belief guidance is a genuine problem for likelihoodism.

In the remainder of this chapter, I shall argue that this received view on the inapplicability of likelihoodist methods to the belief question is incorrect.

Before I defend my positive proposal, first, I need to address a general worry against the very possibility of likelihood-based guidance for belief. The worry has been articulated in a detailed, rigorous manner by Gandenberger (2016). The next section provides a detailed analysis and critique of Gandenberger's argument.

7.3.1 An Argument Against Likelihood-based Belief Guidance

Gandenberger (ibid.) has articulated an argument against the possibility of deriving belief guidance from a likelihoodist framework. The argument identifies a principle that, as he claims, all likelihoodists should endorse. He calls this principle "*minimal comparative proportionalism*" (MCP, for short). To quote Gandenberger (ibid., p. 7):

This principle [MCP] says that there is a real number $r > 1$ such that for any pair of hypotheses A and B , a rational agent believes A over B either in an absolute sense or at least to some degree, if its total evidence favours A over B to degree r or greater.

Now, accepting MCP, as he demonstrates, leads to various epistemic paradoxes. Hence, he concludes that one cannot derive rules of belief from the likelihood framework alone.

A simple, but representative counterexample, similar to the one that Gandenberger puts forward, is as follows:

Example:

There is an urn consisting of 10 tickets, labelled T_0, T_1, \dots, T_9 , and a machine that selects tickets from the urn, without replacement. For each ticket, the machine will either select the ticket or not: so, it could select all 10 tickets, only some tickets, or no tickets at all. You want to know whether the machine selects the tickets randomly or deterministically. The machine may be selecting the tickets by following a random process, where each ticket has a 50% chance of being drawn. Alternatively, the machine may be following some deterministic rule and select the same set of tickets in each experiment. You don't know which process underlies the selection.

You've decided to switch the machine on and see which tickets it would select. In the first round, the machine has selected tickets 0, 4, 6, and 8; let's denote the data as d_{0468} .

Now let h_{random} be the hypothesis that *the machine selects tickets randomly* and let h_{0468} be the hypothesis that *the tickets 0, 4, 6, and 8 were bound to be selected*.

Should you believe h_{random} over h_{0468} ?

Now, the likelihoods of the observed data, d_{0468} , on each competing hypothesis are as follows: $P(d_{0468}|h_{random}) = 1/1024$; $P(d_{0468}|h_{0468}) = 1$.¹⁰³ Thus, the degree of evidential support of h_{0468} over h_{random} is 1024. The data seems to support h_{0468} quite strongly. So, if we let the threshold value, r , in *MCP* to be less than 1024, then you ought to believe h_{0468} over h_{random} .

¹⁰³ On h_{random} , each ticket is equally likely to be selected; hence each possible outcome is equally probable. As each of the 10 tickets is either selected or not, there are 2^{10} or 1024 possible outcomes, the probability of d_{0468} , on the supposition of h_{random} , is $1/1024$. And, on the supposition of the deterministic hypothesis, h_{0468} , the probability of d_{0468} is 1.

But this is clearly absurd. The data does not make h_{0468} more believable than h_{random} . We know from the outset that, for some deterministic hypothesis h_x , the first trial would inevitably favour h_x over h_{random} . Therefore, the first experiment cannot be interpreted as making any deterministic hypothesis more believable than h_{random} .

Notice that the above-identified problem for MCP would remain intact if we had chosen a higher threshold value than 1024. For any finite value of r , a similar counterexample can easily be devised (by increasing the number of tickets in the urn). Therefore, it is tempting to conclude that there are no reasonable likelihood-based rules for belief.

As I show shortly, the above conclusion is premature. One can tease out two readings from the original MCP, depending on the position of the “there is a real number r ” quantifier in relation to the universally quantified sentence “over pair of hypotheses A and B ”. These two readings are as follows:¹⁰⁴

MCP_{weak}: For any pair of hypotheses A and B , there is a real number r such that a rational agent believes A over B if her total evidence favours A over B to degree r or greater.

MCP_{strong}: There is a real number r , such that for any pair of hypotheses A and B , a rational agent believes A over B if her total evidence favours A over B to degree r or greater.

Any likelihoodist account that endorses *MCP_{strong}* is susceptible to a type of counterexample identified by Gendenberger. But notice that accepting *MCP_{weak}* alone does not give rise to

¹⁰⁴ As the reader will see, my strategy for blocking Gendenberger’s objection is similar to Leitgeb’s (2017, Chapter 3) defence of the Lockean thesis that we considered in Chapter 5, Section 5.2.

the same problem. This is so because MCP_{weak} allows threshold r to vary across *contexts of reasoning*. For instance, if we set threshold r to be equal to 1025 (instead of, say, 1023), then the first experiment would not settle the question of which hypothesis should be believed. Of course, the second experiment can go either of the two ways: (i) the machine can select the same set of tickets as in the first experiment or (ii) it can select a different set of tickets. If the second possibility is actualised, then the data would conclusively settle the issue in favour of h_{random} . On the other hand, if the machine selects the same set of tickets, then this would provide overwhelming evidence for the deterministic selection process. The probability that the machine selected the same set of tickets in two trials, on the supposition of a random process is $\frac{1}{1024} * \frac{1}{1024} = \left(\frac{1}{2}\right)^{20}$. Therefore, if we set the threshold value in Example to be greater than 1024 and less than 2^{20} , we would have avoided the problem.

I will discuss at the end of the next section, which aspects of an agent's context determine the value of threshold r . But even at this point of argumentation, we have reached an important conclusion: once we dissect MCP into two principles, MCP_{strong} and MCP_{weak} , it becomes evident that Example is only problematic for MCP_{strong} . Hence, Gendenberger's overall argument is inapplicable to MCP_{weak} .

Now, it is fairly uncontroversial that likelihoodists should accept MCP_{weak} . After all, if there are normative principles that relate likelihood functions with belief, then there should be *some* value for the ratio of likelihoods that would make A more probable than B .¹⁰⁵ But

¹⁰⁵ More than that, assuming that A and B are mutually exclusive and have non-zero probabilities, it is a consequence of Bayes' Theorem that there is some value for the ratio of likelihoods, $P(E|A)/P(E|B)$, that would make A more probable than B . This is evident from the following theorem of probability calculus:

For any mutually exclusive hypotheses A and B :

Footnote continued on the next page.

MCP_{weak} does not entail MCP_{strong} . And it's not clear why likelihoodists should accept MCP_{strong} . After all, why believe that there is one unique threshold value that should fix beliefs in all reasoning contexts?¹⁰⁶ Even bracketing its paradoxical consequences, the existence of such a unique threshold is rather implausible on its own; and I don't see how likelihoodists can be forced to accept such a principle. Hence the argument against a likelihood-based account of belief is wanting.

Of course, my response here is solely negative, as MCP_{weak} , on its own, is insufficient to derive any belief guidance. It remains to be seen whether a broadly likelihoodist account of belief guidance is tenable.

7.4 The Case for Likelihood-based Belief Guidance

As MCP_{weak} , by itself, cannot guide belief, some other principle(s) is needed to connect likelihood functions with beliefs. Like Gendenberger, I also focus on belief guidance for comparative belief; that is, beliefs of the following form: “ A is more probable than B ”, where A and B are any two competitor propositions (hypotheses/theories).

There are two additional reasons for focusing on comparative belief. Firstly, as we have seen, likelihoodism endorses a comparative conception of evidential support. Hence, it's

$$\frac{P(A|E)}{P(B|E)} = \frac{P(E|A)}{P(E|B)} * \frac{P(A)}{P(B)}$$

So, for any given value for the ratio of priors, there is *some* value for the ratio of likelihoods that would make $P(A|E) > P(B|E)$. Hence, MCP_{weak} is not something that a Bayesian – or anyone who accepts the standard definition of conditional probability – can reject.

The above theorem will play a crucial role for deriving guidance for comparative belief in the next section.

¹⁰⁶ In section 4.1, I have discussed in more detail the same point with respect to two versions of the Lockean thesis.

to be expected that likelihoodist methods would be better suited to accommodate comparative belief rather than categorical belief.

Secondly, comparative judgements and evaluations are *indispensable* in science. Typical testing in science is contrastive, where rival hypotheses are assessed against relevant evidence. And scientists often do interpret comparative testing in doxastic terms; as, when biologists conclude that the change in allele frequencies in a population is *probably* due to genetic drift rather than due to selection. Such comparative judgements in science seem to be less problematic, from the epistemic point of view, than categorical or non-relational probabilistic judgements.

So, can likelihoodists provide guidance for belief without lapsing into subjective Bayesianism? To answer this question, we need to be more clear about what “lapsing into subjective Bayesianism” means. From Gandenberger remarks, it is clear that by “lapsing into subjective Bayesianism”, he means accepting this core subjective Bayesian principle, which I call *Subjectivity*:

Subjectivity: When objective, empirically grounded priors are unavailable, scientists can rationally assign prior probabilities to hypotheses that reflect their *subjective degrees of belief* in the hypotheses.

Now, from Gandenberger remarks, it is clear that by “prior probabilities” he means precise or point-valued prior probabilities. But, to make Subjectivity more appealing, I do not assume that probabilities are always point-valued. Instead, in some cases, prior probabilities may be represented with *ranges* or sets of probability functions. So, Subjectivity is assumed to be consistent with situations where scientists represent the probability of H by some range, say $[0.1, 0.6]$.

Now, independent from whether we represent priors with points or ranges, the core of Subjectivity is the view that it is rational for scientists to assign priors to H based on their subjective degree of belief in H .

In what follows, I show how likelihoodists can accommodate belief guidance without accepting Subjectivity. By arguing this, I grant the main premise of Gandenberger's criticism: that scientists cannot appeal to objective, empirically grounded priors in many relevant cases. However, even granted this, we can make sense of rational comparative belief. Let me explain how.

It has been pointed out by Wesley C. Salmon (1990), among others, that when the information about prior probabilities is unavailable, rational comparative belief can be guided via the so-called *ratio form of Bayes' theorem*:¹⁰⁷

$$\frac{P(A|E)}{P(B|E)} = \frac{P(E|A)}{P(E|B)} * \frac{P(A)}{P(B)}$$

If we let R_{Post} be the ratio of *posteriors*, R_L the ratio of *likelihoods*, and R_{Prior} the ratio of *priors*, then the theorem can be summarised succinctly as:

$$R_{Post} = R_L * R_{Prior}$$

Now, the ratio form of Bayes' Theorem frees us from the need of knowing the exact, or even approximate prior probability of either A or B to determine whether A is more probable than B . All we need to know is the value (or approximate value) of the *ratio* of priors, R_{Prior} , and not the value of priors themselves. And fixing the approximate value of R_{Prior} requires strictly less information than fixing the approximate value of priors. To explain this, we need to differentiate *non-relational* priors from *relational priors*.

¹⁰⁷ It is interesting to note that Earman (1992, Chapter 7, Section 3) has criticised Salmon's strategy as too restrictive for Bayesianism, for reasons similar to Gandenberger's criticism of likelihoodism.

Non-relational priors are priors of an individual hypothesis (or a set of hypotheses); when, for instance, we assign a prior of 0.6 to A , or a range of $[0.1, 0.6]$ to A , we attribute a non-relational probability to A . By contrast, relational priors have to do with the relationship between competing hypotheses, A and B . And we may be rational in believing that A and B do not differ significantly in their probabilities without knowing their non-relational probabilities. All we need to know is that the ratio of their priors is approximately 1, $P(A)/P(B) \approx 1$. This ratio can be approximated for many competing theories by appealing to such non-subjective characteristics as their overall predictive accuracy, simplicity, explanatory scope, fruitfulness, etc. So, we may have a good objective basis for concluding that hypotheses A and B are roughly equal in prior plausibility, without knowing their non-relational probabilities. Again, I emphasise that such relational judgments do not require the assignment of non-relational probabilities to the hypotheses in question. Therefore, even when *non-relational priors* cannot be objectively well-grounded, we can still guide comparative belief in a way that meets the likelihoodists standard of objectivity.

Let us illustrate this with an example from cognitive neuroscience. It involves the famous Trolley Problem, which essentially is about whether it is morally permissible/required to sacrifice one innocent life to save several.

First, we need to distinguish two types of Trolley cases: the *impersonal* cases (otherwise known as the *bystander case*), where one needs to hit a switch which diverts a runaway trolley that kills one person but saves five; and the *personal* cases (otherwise known as the *footbridge case*), where one needs to push someone from a bridge to stop a runaway trolley. It is well-known that people respond differently to the impersonal and personal versions of the Trolley Problem. When presented with the bystander case, people tend to

answer that one should hit a switch and save five. By contrast, when presented with the footbridge case, most object to pushing someone to divert the trolley.

Greene et al. (2001) used brain scanning techniques to study which brain regions were “activated” when people engaged with the impersonal and personal Trolley problems. They were primarily concerned with the following two hypotheses (I borrow the formulation of these hypotheses from Machery 2014, 258):

H_1 : People respond differently to moral-personal and moral-impersonal dilemmas because the former elicit more emotional processing than the latter.

H_2 : People respond differently to moral-personal and moral-impersonal dilemmas because the single moral rule that is applied to both kinds of dilemmas (for example, the doctrine of double effect) yields different permissibility judgments.

Now, Greene et al. (2001) found that the personal cases elicited relatively greater activation of brain regions associated with automatic emotional responses; while the impersonal cases elicited relatively greater activation of brain regions associated with conscious reasoning. Let’s denote this neuroimaging evidence as e_{new} .

How should we interpret this neuroimaging evidence? One relatively uncontroversial inference is that the likelihood of e_{new} is higher on the supposition of H_1 than on the supposition of H_2 . This is so because H_2 , unlike H_1 , cannot account for why the two cases elicit the activation of brain regions associated with two very different psychological processes. But it is unclear whether we can make an informed estimate of the posterior probability of H_1 (or H_2) on this evidence. First, it is unclear how we should estimate the prior probabilities for these hypotheses. And even if priors can be fixed in some non-arbitrary way, there may well be some alternative hypothesis that predicts the evidence far better than H_1 . As Machery (2014, 256) puts it:

...in many cases, cognitive neuroscientists have no sense of the probability of obtaining a particular pattern of brain activation if psychological process p is not recruited by experimental tasks and, as a result, they do not know whether the observed pattern of activation gives them a reason to conclude that the psychological process of interest was involved during the task under consideration.

However, notice that even if we cannot estimate the posterior probability of H_1 and H_2 , we can still rationally conclude the evidence renders H_1 more probable than H_2 . By the ratio form of Bayes' theorem, to make this comparative inference, the only required information is that the ratio of likelihoods, $P(e_{new}|H_1)/P(e_{new}|H_2)$, is *not less than* the ratio of priors, $P(H_1)/P(H_2)$. And it is reasonable to think that the available evidence licenses us to accept this inequality. Therefore, even if we have no clue about the non-relational prior and posterior probabilities of H_1 and H_2 , we can still conclude that the former is more probable than the latter, on the relevant evidence.

Of course, I grant that there are many cases where the ratio of priors cannot be estimated in a non-subjective manner. In such cases, judging that a hypothesis A is more probable than B may be problematically sensitive to some subjective factors. As an example, consider a hotly debated topic of cosmological fine-tuning. Some background would be required to explain this.

According to contemporary physics, the fact that life exists in our universe depends on the very precise values that the so-called *fundamental constants* of physics take. For instance, if the mass of proton had been slightly different from its actual value, then the complex and stable structures we find in the universe, like galaxies, stars and planets, would not have existed; hence, life would not have existed. So, given the laws that most contemporary physicists accept, the existence of stable structures and, specifically, the existence of life, is very improbable. But, life, as we have known for some time, does exist in our universe. How can we account for this puzzling evidence?

Some (e.g., Leslie 1989; Hawthorne and Isaacs 2018) think that the likelihood that our universe is fine-tuned for life (denoted as F) is roughly the same relative to these two very different hypotheses:

G : The cosmological constants of the universe have been consciously designed by the God of traditional theism (as, if such a God exists, she would create the universe that can support life).

M : There exist very many (maybe infinitely many) universes. And most (maybe all) possible values of cosmological constants are actualised in some universe(s). Therefore, it is to be expected that some universe(s) among this vast ensemble of universes is fine-tuned for life; and we inhabit such fine-tuned universe.

Let's suppose that the ratio of their likelihood with respect to the fine-tuning evidence, F , is around 1: $P(F|G)/P(F|M) \approx 1$. Now, on this supposition, it is not clear whether there is a relatively unbiased or uncontentious way to evaluate the relative plausibilities of G and M , given evidence F .¹⁰⁸ Some philosophers (e.g., Hawthorne and Isaacs 2018, Section 7.7.3) suggest that the prior of G should be greater than M , because "... [it is] quite strange indeed to suppose that we are living in a multiverse" (ibid., 160). Certainly, many reject this. For instance, one may argue that most non-theists should assign a far greater subjective probability to the multiverse hypothesis than to the God hypothesis. Because, for most non-theists, the universe with God in it is more "strange" than the universes without God.¹⁰⁹

¹⁰⁸ If we use the imprecise probability framework, we may say that on some rationally permissible probability distributions, G is more likely than M , but for some other permissible distributions – M is more likely than G . Hence, on this framework, it seems that the evidence supports *suspending judgement* on whether G is more probable than M .

¹⁰⁹ I have developed this type of criticism in detail (Tokhadze, forthcoming a).

Therefore, at least at first blush, there does not seem to be a non-subjective way of assessing the relative plausibilities of G and M . And cases like these are abundant in philosophy and science.

Now, it should be clear that the existence of such cases does not conflict with the main argument of this chapter. Likelihoodists are not committed to the claim that comparative beliefs can be formed in *all* evidential situations. By contrast, all we needed to show is that, in *many* cases, comparative beliefs can be freed from subjective priors. And this is exactly what I've argued here: comparative belief can be objectively well-grounded even when the empirical information about non-relational priors are unavailable.

Before concluding, let me briefly discuss the connection between the ratio form of Bayes' theorem and MCP_{weak} . To remind the reader, MCP_{weak} is the thesis that:

For any pair of hypotheses A and B , there is a real number r such that a rational agent believes A over B if her total evidence favours A over B to degree r or greater.

As I've already noted in the previous section, threshold r in MCP_{weak} may be fixed differently in different contexts of reasoning. But I have not elaborated on how an agent's context of reasoning fixes the relevant threshold for comparative belief. The context-sensitive factor that fixes an agent's threshold r , given competing hypotheses A and B , is her *prior comparative probability function* that provides an estimate of the value of the ratio of priors (for A and B). So, if an agent estimates that A and B 's prior ratio is some number c , then threshold r should be greater than the *reciprocal* of c . As before, the rationale behind this is

provided by the ratio form of Bayes' theorem. So, the threshold r would be sensitive to an agent's estimate for the relevant ratio of priors.¹¹⁰

In some cases, like in the considered example from cognitive neuroscience, the relevant values of threshold r can be estimated in a relatively uncontentious manner. However, in the fine-tuning example, different agents may have different estimates for threshold r . And as I've already discussed, this is perfectly consistent with the main argument of this chapter.

This concludes the positive argument of this chapter.

7.5 Summary

The chapter has dealt with the following question:

How should we think about rational belief in contexts where the available evidence does not justify a complete order (or a complete probability distribution) over relevant possibilities?

I have provided an answer to this question that fits well with an impermissive account of rational belief defended in this dissertation. The answer has been articulated within a broadly likelihoodist framework. As I've argued, it is possible to guide *comparative belief* in cases where the empirically justified prior probabilities are unavailable without endorsing either subjective or objective Bayesianism. My main argumentative strategy has been the move from the non-relational to relational or comparative probabilities: probabilities of the form “ A is more probable than B ”. And, as we have seen, sometimes our evidence may not justify

¹¹⁰ My proposal follows Leitgeb (2017, Chapter 3). As we have discussed in detail (in Section 5.2.3), within the stability theory, an agent's degree of belief function P should be considered as a part of her reasoning context.

absolute probabilities towards competing hypotheses, but it could still justify the comparative claims about them.

I conclude that a broadly likelihoodist, impermissivist account of rational belief is not as restrictive as it may initially seem: even when the available evidence does not support a complete probability distribution over a set of possibilities, rational comparative beliefs can still be formed without appealing to subjective credences or the principle of indifference.

8 Conclusion

In this dissertation, I have proposed and defended a hybrid view about the requirements of evidence on doxastic attitudes. This hybrid view combines Moderate Uniqueness, the thesis about the coarse-grained, binary doxastic attitude of belief, with Credal Permissivism, the thesis about the fine-grained, numerical doxastic attitude of credence:

Moderate Uniqueness: For any hypothesis H and evidence E , it is not the case that E rationally permits belief that H and belief that $\neg H$.

Credal Permissivism: For some evidence, E , and proposition, H , E rationally permits more than one credence towards H .

I have called the combination of these two theses Hybrid Impermissivism. The main goal of this dissertation has been to articulate a precise *hybrid theory* that endorses Hybrid Impermissivism together with the norms of how rational belief and credence ought to interact.

I conclude by reflecting on the core claims I've made in this dissertation (including their most unintuitive/problematic aspect) and touching upon an important, relevant topic I have not discussed – the epistemic significance of disagreement.

There are many claims made in this dissertation. Not all of them are central to the overall view that I've defended. While I do not believe the conjunction of all claims made in this dissertation, I believe the following conjunction that articulates the core proposition of the dissertation:

The Dissertation's Core (Core):

- Hybrid Impermissivism combines versions of Uniqueness and Permissivism that are individually plausible (much more plausible than other, radical versions of Uniqueness and Permissivism), and

- Hybrid Impermissivism can be transformed into a formally precise *hybrid theory* by endorsing either the stability theory or the dominant core theory (or the varieties of these theories, without Deductive Cogency). The resulting hybrid theories capture important plausible aspects of permissivist and impermissivist epistemologies;
 - On the permissivist side: each considered hybrid theory makes room for agent-relative, subjective influences on belief such as an agent’s level of epistemic cautiousness or braveness (4.2.2).
 - On the impermissivist side: each considered hybrid theory avoids the problem of fine distinctions and is consistent with a version of objectivism about evidential support (i.e., Relational Objectivity; Section 2.2).
- The resulting hybrid theories avoid the arbitrariness objection (Section 4.2.2) and the diachronic coordination problem (Chapters 5 and 6), and
- An impermissivist approach to rational belief defended in this dissertation is not overly demanding in the following sense: when the available evidence does not justify a complete ordering of relevant worlds, rational comparative belief can still be formed by appealing to objective likelihoods (Chapter 7).

Given the relevant reasoning context, I believe this conjunctive proposition Core and invest high credence in it; at least, higher than 0.5. I don’t have a more specific finely-grained doxastic attitude towards Core (alluding to Section 3.4).

A theory that provides a detailed answer to an important philosophical question is very likely to have some problematic or unintuitive aspect(s). I consider such an aspect of the presented hybrid theory a sensitivity of belief to a context. More specifically, while the proposed hybrid approach avoids the coordination problem without appealing to overly

strong constraints about rational credence, this is achieved via *propositionwise dependence*: the view that whether an agent believes a proposition is not the sole function of the proposition's probability but also depends on the other propositions that the agent considers. Propositionwise dependence makes it possible to solve the coordination problem via a relatively undemanding view about the evidential constraints on rational credence, *Order Uniqueness*: the view that for any evidence and proposition, the evidence justifies the unique plausibility order of *relevant* possibilities (or possible worlds) associated with this proposition.

While I have argued that the proposed approach to the coordination problem is plausible and avoids objections, I recognise that the reliance on the context-sensitive view of belief may be problematic for many. Still, the context-sensitive framework for belief is highly fruitful within epistemology in general, as it helps us to make sense of the norms like Deductive Cogency and Probabilism (as discussed in Sections 3.2), and, most importantly, enables us to unify these norms into a coherent view, via the stability theory or the dominant core theory. Hence, I submit that the unintuitive aspects of the context-sensitive approach to rational belief are outweighed by its fruits.

Finally, let me finish by mentioning an important topic that I have not covered in any detail in this dissertation and which I see as the work for the future – the epistemic significance of disagreement. As noted in the introduction, I believe Uniqueness raises more foundational issues in epistemology that need to be addressed before turning to more localized debates, like the disagreement debate.

Hybrid Impermissivism opens up new approaches to the disagreement debate that have not been explored in detail so far. The sole exception is Jackson's recent paper (2021), where she suggests that peer disagreement requires an agent to change her credences but not

her beliefs. Hybrid Impermissivism is compatible with Jackson's suggestion, but does not commit to it. On Hybrid Impermissivism, an agent, in certain cases, could be rationally required to change her belief if the required change in her credences is significant. Interestingly, on Hybrid Impermissivism, on the supposition that the disagreement requires some credal revision(s), an agent may or may not be required to change her beliefs, depending on the case in question (i.e., depending on her prior credences, the extent to which these prior credences need to be revised, and/or whether the required revision changes her doxastic order of possible worlds). So, the hybrid approach to Uniqueness enriches possible approaches to the epistemic significance of disagreement considerably.

Bibliography

- Alchourrón, C. E., Gärdenfors, P., & Makinson, D. (1985). On the Logic of Theory: Partial Meet Contraction and Revision Functions. *The journal of symbolic logic* 50.2 , 510-530.
- Anscombe, G. E. (1963). *Intention* (2nd ed.). Oxford: Blackwell.
- Ballantyne, N., & Coffman, E. J. (2011). Uniqueness, Evidence, and Rationality. *Philosophers' Imprint* 11, 1–13.
- Bandyopadhyay, P. S., Brittan, G. G., & Taper, M. L. (2016). *Belief, Evidence, and Uncertainty: Problems of Epistemic Inference*. Springer International Publishing.
- Berger, J., & Wolpert, R. (1988). *The Likelihood Principle*. Beachwood, OH: Institute of Mathematical Statistics, 2nd edition.
- Bird, A. (2017). Inference to the Best Explanation, Bayesianism, and Knowledge. In K. McCain, & T. Poston (Eds.), *Best explanations: New essays on inference to the best explanation* (pp. 97-120). Oxford.
- Bird, A. (2018). Understanding the Replication Crisis as a Base Rate Fallacy. *The British Journal for the Philosophy of Science*.
- Boutilier, C. (1996). Iterated Revision and Minimal Change of Conditional Beliefs. *Journal of Philosophical Logic* 25.3, 263-305.
- Brueckner, A., & Bundy, A. (2012). On “epistemic permissiveness. *Synthese* 188.2, 165-177.
- Cantwell, J., & Rott, H. (2019). Probability, Coherent Belief and Coherent Belief Changes. *Annals of Mathematics and Artificial Intelligence* 87.3, 259-291.

- Carnap, R. (1962). *Logical Foundations of Probability*. 2nd. Chicago: University of Chicago Press.
- Castro, C., & Hart, C. (2019). The Imprecise Impermissivist's Dilemma. *Synthese* 196.4, 1623-1640.
- Christensen, D. (2007). Epistemology of Disagreement: The Good News. *The Philosophical Review* 116.2, 187-217.
- Clarke, R. (2013). Belief Is Credence One (In Context). *Philosophers' Imprint Volume 13 Issue 11*, 1-18.
- Craig, E. (1990). *Knowledge and the State of Nature*. Oxford: Clarendon Press.
- Decker, J. (2012). Disagreement, Evidence, and Agnosticism. *Synthese* (187), 753–783.
- Dennett, D. (2001). In Darwin's Wake, Where am I? *Proceedings and Addresses of the American Philosophical Association*, 11-30.
- Dienes, Z. (2011). Bayesian Versus Orthodox Statistics: Which Side Are You On? *Perspectives on Psychological Science*, 6(3), 274–290.
- Dietrich, F., & List, C. (2018). From Degrees of Belief to Binary Beliefs: Lessons from Judgment-Aggregation Theory. *Journal of Philosophy* 115(5), 225–270.
- Dorling, J. (1979). Bayesian Personalism, the Methodology of Scientific Research Programmes and Duhem's Problem. *Studies in the History and Philosophy of Science Part A* 10(3), 177-187.
- Douven, I. (2009). Uniqueness Revisited. *American Philosophical Quarterly* 46, 347–61.

- Douven, I. (2019). The Stability of Belief: How Rational Belief Coheres with Probability. *Philosophical Review*, 371-375.
- Douven, I., & Rott, H. (2018). From Probabilities to Categorical Beliefs: Going Beyond Toy Models. *Journal of Logic and Computation* 28.6, 1099-1124.
- Earman, J. (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.
- Edwards, W., Lindman, H., & Savage, L. J. (1963, 70). Bayesian Statistical Inference for Psychological Research. *Psychological Review* , 193-242.
- Eells, E. (1985). Problems of Old Evidence. *Pacific Philosophical Quarterly* 66.3-4, 283-302.
- Festa, R., & Cevolani, G. (2017). Unfolding the Grammar of Bayesian Confirmation: Likelihood and Antilikelihood Principles. *Philosophy of Science* 84 (1), 56–81.
- Fitelson, B. (1999). The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity. *Philosophy of Science*, 66, 362–378.
- Fitelson, B. (2007). Likelihoodism, Bayesianism, and Relational Confirmation. *Synthese* 156, 473–489.
- Fitelson, B. (2011). Favoring, likelihoodism, and bayesianism. *Philosophy and Phenomenological Research* 83 (3), 666-672.
- Foley, R. (1993). *Working without a Net*. Oxford: Oxford University Press.
- Foley, R. (2009). Beliefs, Degrees of Belief, and the Lockean Thesis. In *Degrees of Delief* (pp. 37-47). Dordrecht: Springer.
- Friedman, J. (2013). Suspended Judgment. *Philosophical studies* 162.2, 165-181.

- Friedman, J. (2017). Why Suspend Judging? *Noûs* 51.2, 302-326.
- Gandenberger, G. (2015). A New Proof of the Likelihood Principle. *The British Journal for the Philosophy of Science*, 66, 475–503.
- Gandenberger, G. (2016). Why I am not a likelihoodist. *Philosophers' Imprint*, 16 (7), 1-22.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge: Harvard University Press.
- Good, I. J. (1950). *Probability and the weighing of evidence*. London: Griffin.
- Greco, D., & Hedden, B. (2016). Uniqueness and Metaepistemology. *Journal of Philosophy* 113(8), 365–395.
- Greene, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, 293(5537), 2105-2108.
- Hacking, I. (1965). *Logic of statistical inference*. Cambridge University Press.
- Harman, G. (1999). *Reasoning, Meaning, and Mind*. Oxford: Oxford University Press.
- Hawthorne, J. (2005). Degree-of-Belief and Degree-of-Support: Why Bayesians need Both Notions. *Mind* 114 (454), 277-320. doi:<https://doi.org/10.1093/mind/fzi277>
- Hawthorne, J. (2018). Inductive Logic. *The Stanford Encyclopedia of Philosophy (Spring 2021 Edition)*, Edward N. Zalta (ed.). Retrieved from <https://plato.stanford.edu/archives/spr2021/entries/logic-inductive/>
- Hawthorne, J., & Isaacs, Y. (2018). Fine-tuning fine-tuning. In M. A. Benton, J. Hawthorne, & D. Rabinowitz (Eds.), *Knowledge, Belief, and God: New Insights in Religious Epistemology* (pp. 136-168). Oxford University Press.

- Hedden, B. (2015). A Defense of Objectivism about Evidential Support. *Canadian Journal of Philosophy* 45.5-6, 716-743. doi:<https://doi.org/10.1080/00455091.2015.1123038>
- Howson, C., & Urbach, P. (2006). *Scientific Reasoning: The Bayesian Approach* (3rd ed.). La Salle: Open Court Publishing.
- Hume, D. (1739/1978). *A Treatise of Human Nature* (2nd ed.). (L. A. Selby-Bigge, Ed.) Oxford: Oxford University Press.
- Jackson, E. (2019a). A Defense of Intrapersonal Belief Permissivism. *Episteme*, 1-15.
- Jackson, E. (2019b). Belief and Credence: Why the Attitude-type Matters. *Philosophical Studies* 176.9, 2477-2496.
- Jackson, E. (2021). Dilemmas, Disagreement, and Dualism. In K. M. Scott Stapleford, & M. Steup (Eds.), *Epistemic Dilemmas: New Arguments, New Angles* (pp. 217–231). New York: Routledge.
- Jackson, E., & Turnbull, M. G. (forthcoming). Permissivism, Underdetermination, and Evidence. In C. Littlejohn, & M. Lasonen-Aarnio (Eds.), *The Routledge Handbook of the Philosophy of Evidence*.
- James, W. (1896). The Will to Believe. In *The Will to Believe and Other Essays in Popular Philosophy* (pp. 1-31). New York: Dover Publications, 1956.
- Jeffrey, R. (1965). *The Logic of Decision*. New York: McGraw-Hill.
- Joyce, J. M. (2011). A Defense of Imprecise Credence in Inference and Decision. *Philosophical Perspectives*, 24, 281–323.
- Kelly, T. (2010). Peer Disagreement and Higher Order Evidence. In R. Feldman, & T. A. Warfield (Eds.), *Disagreement* (pp. 111-175). New York.

- Kelly, T. (2014). Evidence Can be Permissive. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary Debates in Epistemology* (2nd ed., pp. 298-312). Wiley Blackwell.
- Kelly, T. (2016). Evidence. (E. N. Zalta, Ed.) *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition). Retrieved from <<https://plato.stanford.edu/archives/win2016/entries/evidence/>>.
- Kopec, M. (2015). A Counterexample to the Uniqueness Thesis. *Philosophia* 43, 403–9.
- Kopec, M., & Titelbaum, M. G. (2016). The Uniqueness Thesis. *Philosophy Compass* 11, 189-200.
- Kopec, M., & Titelbaum, M. G. (2019). When Rational Reasoners Reason Differently. In M. Balcerak Jackson, & J. Balcerak-Jackson (Eds.), *Reasoning: Essays on Theoretical and Practical Thinking* (pp. 205-231). Oxford: Oxford University Press.
- Kyburg, H. E. (1961). *Probability and the Logic of Rational Belief*. Middletown: Wesleyan University Press.
- Kyburg, H. E. (1970). Conjunctivitis. In *Induction, Acceptance, and Rational Belief* (Marshall Swain ed., pp. 55-82). Dordrecht: Reidel.
- Kyburg, H. E. (1983). Rational Belief. *The Brain and Behavioural Sciences*, 6, 231–273.
- Lee, M. (2013). Conciliationism without Uniqueness. *Grazer Philosophische Studien* 88.1, 161-188.
- Leitgeb, H. (2013). Reducing Belief Simpliciter to Degrees of Belief. *Annals of Pure and Applied Logic*, 164(12), 1338-1389.
- Leitgeb, H. (2014). The Stability Theory of Belief. *Philosophical Review* 123.2, 131-71.

- Leitgeb, H. (2017). *The Stability of Belief. How Rational Belief Coheres with Probability*. Oxford: Oxford University Press.
- Leitgeb, H. (2021). Stability and the Lottery Paradox. In I. Douven (Ed.), *Lotteries, knowledge and Rational Belief: Essays on the Lottery Paradox* (pp. 147-170). Oxford University Press.
- Leslie, J. (1989). *Universes, London*. London: Routledge.
- Levi, I. (1998). Pragmatism and Change of View. *Canadian Journal of Philosophy, Supplementary Volume, 24*, 177-201.
- Levi, I. (1967). *Gambling with Truth: An Essay on Induction and the Aims of Science*. Cambridge, MA: MIT Press.
- Levi, I. (1974). On Indeterminate Probabilities. *Journal of Philosophy, 71*, 391–418.
- Levi, I. (1977). Subjunctives, Dispositions and Chances. *Synthese, 34*, 423–455.
- Levi, I. (1985). Imprecision and Indeterminacy in Probability Judgment. *Philosophy of Science, 52*, 390–409.
- Levinstein, B. A. (2017). Permissive rationality and sensitivity. *Philosophy and Phenomenological Research 94.2*, 342-370.
- Lewis, D. (1996). Elusive Knowledge. *Australasian journal of Philosophy 74.4*, 549-567.
- Lin, H. (2019). Belief Revision Theory. In R. Pettigrew, & J. Weisberg (Eds.), *The Open Handbook of Formal Epistemology*.
- Lin, H., & Kelly, K. T. (2012). Propositional Reasoning that Tracks Probabilistic Reasoning. *Journal of philosophical logic 41.6*, 957-981.

- Lin, H., & Kelly, K. T. (2021). Beliefs, Probabilities, and their Coherent Correspondence. In I. Douven (Ed.), *Lotteries, knowledge and Rational Belief: Essays on the Lottery Paradox* (pp. 185-222). Cambridge University Press.
- Loeb, L. E. (2002). *Stability and Justification in Hume's Treatise*. Oxford: Oxford University.
- Loeb, L. E. (2010). *Reflection and the Stability of Belief: Essays on Descartes, Hume, and Reid*. Oxford: Oxford University Press.
- Machery, E. (2014). In Defense of Reverse Inference. *The British Journal for the Philosophy of Science* 65.2, 251-267.
- Maher, P. (1996). Subjective and Objective Confirmation. *Philosophy of Science*, 63 (2), 149-174.
- Makinson, D. C. (1965). The Paradox of the Preface. *Analysis*, 25(6), 205-207.
- Matheson, J. (2011). The Case for Rational Uniqueness. *Logos & Episteme* 2, Issue 3, 359–373. doi:<https://doi.org/10.5840/logos-episteme20112319>
- Mayo, D. (1996). *Error and the growth of experimental knowledge*. University of Chicago Press.
- Mayo, D. (2018). *Statistical inference as severe testing: How to get beyond the statistics wars*. Cambridge: Cambridge University Press.
- Meacham, C. (2014). Impermissive Bayesianism. *Erkenntnis* 79, 1185–217.
- Peels, R. (2017). *Responsible Belief – A Theory in Ethics and Epistemology*. Oxford University Press.

- Peels, R., & Booth, A. (2014). Why Responsible Belief is Permissible Belief. *Analytic Philosophy* 55, 75–88. doi: <https://doi.org/10.1111/phib.12036>
- Pettigrew, R. (2015). II—Pluralism about Belief States. *Aristotelian Society Supplementary Volume*. Vol. 89. No. 1, 187-204.
- Podgorski, A. (2016). Dynamic permissivism. *Philosophical Studies*, 173(7), 1923-1939.
- Raleigh, T. (2017). Another Argument Against Uniqueness. *Philosophical Quarterly*, 67 (267), 327–346.
- Romero, F. (2019). Philosophy of Science and the Replicability Crisis. *Philosophy Compass*, 14 (11).
- Rott, H. (2017). Stability and Scepticism in the Modelling of Doxastic States: Probabilities and Plain Beliefs. *Minds and Machines* 27.1, 167-197.
- Royall, R. (1997). *Scientific Evidence: A Likelihood Paradigm*. London: Chapman and Hall.
- Salmon, W. (1990). Rationality and objectivity in science or Tom Kuhn meets Tom Bayes. *Minnesota Studies in the Philosophy of Science* 14, 175-204.
- Schoenfield, M. (2014). Permission to Believe: Why Permissivism is True and What it Tells us About Irrelevant Influences on Belief. *Noûs* 48, 193–218.
- Schoenfield, M. (2019). Permissivism and the Value of Rationality: A Challenge to the Uniqueness Thesis. *Philosophy and Phenomenological Research*, 286-297.
- Schurz, G. (2019). Impossibility Results for Rational Belief. *Noûs* 53.1, 134-159.
- Shear, T., & Fitelson, B. (2019). Two Approaches to Belief Revision. *Erkenntnis*, 487–518.

- Shoham, Y. (1987). A Semantical Approach to Nonmonotonic Logics. In *Readings in Nonmonotonic Reasoning* (pp. 227–250). Morgan Kaufmann Publishers Inc.
- Simpson, R. (2017). Permissivism and the Arbitrariness Objection. *Episteme, Volume 14, Issue 4 D*, 519-538.
- Smith, J. J. (2020). Unacknowledged Permissivism. *Pacific Philosophical Quarterly*, 158-183.
- Sober, E. (2008). *Evidence and Evolution: The Logic Behind the Science*. Cambridge University Press.
- Staffel, J. (2016). Beliefs, Buses and Lotteries: Why Rational Belief can't be Stably High Credence. *Philosophical Studies 173.7*, 1721-1734.
- Stapleford, S. (2019). Intraspecies Impermissivism. *Episteme 16.3*, 340-356.
- Strevens, M. (2017). *Notes on Bayesian Confirmation Theory (Unpublished Lecture notes)*. Retrieved from <http://www.strevens.org/bct/BCT.pdf>
- Sturgeon, S. (2020). *The Rational Mind*. Oxford: Oxford University Press.
- Thorstad, D. (2019). Permissive Metaepistemology. *Mind 128.511*, 907-926.
- Titelbaum, M. G. (2010). Not Enough There There: Evidence, Reasons, and Language Independence. *Philosophical Perspectives 24 (1)*, 477-528.
- Titelbaum, M. G. (2021). The Stability of Belief: How Rational Belief Coheres with Probability, by Hannes Leitgeb. *Mind*, 1006-1017.
- Titelbaum, M. G. (forthcoming). *Fundamentals of Bayesian Epistemology*. Oxford: Oxford University Press.

- Tokhadze, T. (forthcoming a). Fine-tuning, Weird Sorts of Atheism and Evidential Favouring. *Analytic Philosophy*.
- Tokhadze, T. (forthcoming b). Hybrid Impermissivism and the Diachronic Coordination Problem. *Philosophical Topics*.
- Tokhadze, T. (forthcoming c). Extreme Permissivism Revisited. *The European Journal of Analytic Philosophy*.
- Tokhadze, T. (forthcoming d). Likelihoodism and Guidance for Belief. *Journal for General Philosophy of Science*.
- Tversky, A., & Kahneman, D. (1971). Belief in the "Law of Small Numbers". *Psychological Bulletin* 76,, 105-110.
- Tversky, A., & Kahneman, D. (1973). Availability: A Heuristic for Judging Frequency and Probability. *Cognitive Psychology* 5, 207-232.
- van Fraassen, B. (1984). Belief and the Will. *The Journal of Philosophy*, 81(5), 235-256.
- Van Fraassen, B. (1989). *Laws and Symmetry*. Oxford University Press.
- Weisberg, J. (2020). Could've Thought Otherwise. *Philosophers' Imprint* 20(12), 1-24.
doi:<http://hdl.handle.net/2027/spo.3521354.0020.012>
- Wetzels, R., Matzke, D., Lee, M. D., Rouder, J. N., Iverson, G. J., & Wagenmakers, E. J. (2011). Statistical evidence in experimental psychology: An empirical comparison using 855 t tests. *Perspectives on Psychological Science*, 6(3), 291-298.
- White, R. (2005). Epistemic Permissiveness. *Philosophical Perspectives* 19, 445-459. doi:
<https://doi.org/10.1111/j.1520-8583.2005.00069.x>

- White, R. (2014). Evidence cannot be Permissive. In M. Steup, J. Turri, & E. Sosa (Eds.), *Epistemology, Contemporary Debates* (2nd ed., pp. 312-323). Wiley-Blackwell.
- Willard-Kyle, C. (2017). Do Great Minds Really Think Alike? *Synthese* 194 (3), 989–1026.
- Williamson, J. (2007). Motivating objective Bayesianism: From empirical constraints to objective probabilities. In W. E. Harper, & G. R. Wheeler (Eds.), *Probability and Inference: Essays in Honour of Henry E. Kyburg, Jr.* Amsterdam: Elsevier.
- Williamson, J. (2010). *In Defence of Objective Bayesianism*. Oxford University Press.
doi:<https://doi.org/10.1093/acprof:oso/9780199228003.001.0001>
- Yalcin, S. (2018). Belief as Question-Sensitive. *Philosophy and Phenomenological Research* 97.1, 23-47.