
Imprecise Probability in Epistemology

Lee Elkin



München 2017

Imprecise Probability in Epistemology

Lee Elkin

Inauguraldissertation
zur Erlangung des Doktorgrades der Philosophie
an der Ludwig–Maximilians–Universität München

vorgelegt von
Lee Elkin
aus Busan

2017

Erstgutachter: Prof. Dr. Stephan Hartmann
Zweitgutachter: Prof. Dr. Richard Pettigrew
Tag der mündlichen Prüfung: 28. Juni 2017

Contents

Acknowledgements	vii
Published Work	viii
Abstract	x
1 Introduction	1
1.1 A Brief History of a Formal Epistemology	2
1.1.1 Motivation	4
1.2 Outline of the Intended Project	7
2 Subjective Probability: The Bayesian Paradigm	11
2.1 Pragmatic Justification of Probabilism	15
2.2 A Non-Pragmatic Justification of Probabilism	16
2.3 Imprecision in Belief	19
3 Ambiguity Induced by Peer Disagreements	25
3.1 Splitting the Difference	27
3.1.1 A Positive Consequence of Linear Pooling	30
3.1.2 Irrational Consequences of Linear Pooling	33
3.2 Set-Based Credences	37
3.2.1 Why Set-Based Credences?	44
3.3 Summary Thus Far	49
3.4 A Qualitative Account	51
4 Complete Ignorance in Imprecise Probability Theory	59
4.1 Exiling Ignorance	60
4.2 Vacuous Priors	65
4.2.1 Interpreting Vacuous Priors	68
4.2.2 A Behavioral Interpretation	76
4.3 Belief Inertia	81
4.4 Summary	84
5 Probabilistic Confirmation Theory with Imprecise Probabilities	87
5.1 Imprecise Probability Revisited	88

5.1.1	Rejection of Indifference From an Aversion to Regret . . .	91
5.2	Bayesian Confirmation Theory	94
5.3	Ways of Describing Confirmation in Imprecise Probability Theory	97
5.3.1	Confirmational Extremity	97
5.3.2	A Behavioral Confirmation Theory	103
5.3.3	Confirmational Sensitivity	105
5.3.4	Interval Dominance Confirmation	108
5.4	Summary	110
	Conclusion	113
	Bibliography	116

Acknowledgements

There are a lot of people deserving gratitude for helping make this dissertation happen. First and foremost, I am indebted to my advisor, Stephan Hartmann, for giving me the opportunity to join the Munich Center for Mathematical Philosophy and providing guidance so that I might make a positive contribution to the field of epistemology. Second, I would like to thank Gregory Wheeler who acted as a mentor during my time in Munich. Much of what I know about imprecise probability today is largely due to Greg patiently explaining the ins and outs. In addition, Richard Pettigrew, Branden Fitelson, and John Norton were all kind enough to be involved with my research, offering encouragement and thought-provoking challenges along the way. For this, I am grateful. I would also like to thank Martin Kocher for his willingness to take part in the dissertation defense.

Aside from the aforementioned, a number of other people provided helpful discussion over the past few years and should be recognized. They are: Yann Benétreau-Dupin, Seamus Bradley, Peter Brössel, Catrin Campbell-Moore, Justin Caouette, Jennifer Carr, Jake Chandler, Matteo Colombo, Richard Dawid, Malte Doehne, Leon Geerdink, Remco Heesen, Catherine Herfeld, Jason Konek, Matthew Kotzen, Hannes Leitgeb, Aidan Lyon, Conor Mayo-Wilson, Ignacio Ojea Quintana, Arthur Paul Pedersen, Clayton Peterson, Hans Rott, Jan Sprenger, Rush Stewart, Scott Sturgeon, Naftali Weinberger, Kevin Zollman and audiences at the 2014 British Society for the Philosophy of Science Annual Meeting, 2014 Pittsburgh Area Philosophy Colloquium, University of Calgary's 4th Annual Philosophy Graduate Conference, Tilburg University's Research Seminar in Epistemology and Philosophy of Science, 2015 American Philosophical Association Pacific Division Meeting, 2016 American Philosophical Association Eastern Division Meeting, and 11th Cologne Summer School in Philosophy. I would also like to thank the Alexander von Humboldt Foundation for generously supporting my doctoral research financially through a doctoral fellowship.

Finally, I am thankful to have had support from family members and friends throughout the long academic journey. I am especially fortunate to have a wonderful and caring partner, Kimberly, who stood by my side, struggled with me in each step, listened to every word, and put up with the same Bayesian babble for years. Without her continued support, this dissertation would not have been possible.

Published Work

The work presented in this dissertation has yet to appear in print as a result of publication at this point in time. However, Chapter 3 of the dissertation draws on a forthcoming article, “Resolving Peer Disagreements Through Imprecise Probabilities”, co-written with Dr. Gregory Wheeler (MCMP, LMU Munich) and is expected to appear in the journal, *Noûs*. Specifically, the content of pages 25-27, 31-47, and 51-52 is taken from the forthcoming article. Of course, the work has been re-written in my own words to better situate it within the chapter, and I take full responsibility for any errors stemming from the re-write.

Abstract

There is a growing interest in the foundations as well as the application of imprecise probability in contemporary epistemology. This dissertation is concerned with the application. In particular, the research presented concerns ways in which imprecise probability, i.e. sets of probability measures, may helpfully address certain philosophical problems pertaining to rational belief. The issues I consider are disagreement among epistemic peers, complete ignorance, and inductive reasoning with imprecise priors. For each of these topics, it is assumed that belief can be modeled with imprecise probability, and thus there is a non-classical solution to be given to each problem. I argue that this is the case for peer disagreement and complete ignorance. However, I discovered that the approach has its shortcomings, too, specifically in regard to inductive reasoning with imprecise priors. Nevertheless, the dissertation ultimately illustrates that imprecise probability as a model of rational belief has a lot of promise, but one should be aware of its limitations also.

List of Tables

2.1	Ellsberg Experiment: Three-Color	20
3.1	Forecasters p_1 and p_2 and Equally Weighted Average p^*	35
3.2	Forecaster One and Forecaster Two's Payoff	36
3.3	Reasonable Ranges and Loss of Independence	40
3.4	Buy-side Payoffs (left) and Sell-side Payoffs (right)	48
3.5	Checklist	50
4.1	Truth-tables for NOT, AND, and OR in \mathcal{L}_3	71

Chapter 1

Introduction

The present dissertation concerns the use of imprecise probability, or generalized Bayes, as a formal tool in an attempt at addressing a class of philosophical problems relating to rational belief. Of course, it would be practically impossible to offer up anything near a comprehensive study covering every epistemological problem of interest in contemporary circles. With that said, the following philosophical questions have been chosen as the main focus of the dissertation.

- *How should equally competent peers respond to a disagreement?*
- *Can a state of complete ignorance be represented probabilistically?*
- *When is a theory or hypothesis confirmed by evidence if prior opinion is imprecise?*

While the respective chapter devoted to each of these questions may constitute a stand-alone essay, the dissertation is unified by a recurring application of imprecise probability for modeling rational belief, thus resulting in a cohesive project.

The formal nature of analysis to be given on each subject matter places the philosophical work under the heading of formal epistemology, a small yet growing and lively field in philosophy. Its growth has resulted from an increasing number of philosophers who regard mathematical theories as invaluable tools for addressing contemporary philosophical issues, especially in epistemology. The non-standard attitude is embraced in this collection of essays by focusing particularly on ways a theory of imprecise probability can or at least attempt to lend a helping hand in the process of engineering solutions to various epistemic challenges.

1.1 A Brief History of a Formal Epistemology

Formal methods in 20th century analytical philosophy were primarily confined to the fields of logic, philosophy of language, philosophy of mathematics, philosophy of science(s), and to some extent analytical metaphysics. Epistemology, on the other hand, proceeded with a fixation on Cartesianism, which resulted in refutations of skepticism and conceptual accounts of knowledge that were largely influenced by G.E. Moore (1939). To this day, conceptual analysis remains the dominant method of epistemology celebrated in almost all Western analytic philosophy departments and typically involves little to no engagement with formal methods.

The closest attempt at developing a formal epistemology arose in mid-20th century philosophy of science led by Carnap, Hempel, and Popper who put to use deductive logic and probabilistic methods in studying scientific reasoning (see Horsten & Douven 2008). While Hempel's (1945) notable logic of confirmation relied on a set of deductive principles for assessing the plausibility of scientific theories and hypotheses, Carnap (1962) focused his attention on an inductive logic involving logical probability in which the logical relation between a statement and evidence is the *degree* to which the evidence (objectively) confirms the statement.¹ Looking back, Carnap seemed to be on the right track given the many difficulties that would soon appear with Hempel's deductive method.² What is more, a probabilistic rather than deductive theory of confirmation boldly made an attempt at resolving Hume's (1888) problem of induction, which has worried many for so long. But despite the program's boldness, logical probability received little endorsement, even though it appeared to be in the right arena (see Hájek 2011).

Meanwhile, decision theorists in the post-war era were working on von Neumann-Morgenstern (1944) expected utility theory together with a subjective theory of probability developed earlier by Frank Ramsey (1926) and Bruno de Finetti (1931/1989). This effort ultimately led to what is known as *Bayesian Decision Theory*. The most notable explication was that of Savage's (1954) in *The Foundations of Statistics*. Although statisticians and economists were those mainly interested in the mature subjective Bayesian method at the time, some philosophers also had taken notice of its virtues, particularly in providing an inductive logic that avoids interpretational issues associated with logical probability and the neglect of prior opinion in a frequentist account of probability (though, not everyone thought

¹An earlier development of logical probability is found in Keynes (1921).

²See Crupi (2015) for a general discussion on the problems with Hempelian confirmation.

the absence of prior opinion was such a bad thing).

Moreover, early endorsement of Bayes appeared in Isaac Levi's (1961) "Decision Theory and Confirmation" where he advanced a skeptical attitude towards there being a non-pragmatic account of "accept/reject" in inductive inference, which aimed at softening the resistance against subjectivity in the Bayesian method. A few short years later, a more comprehensive Bayesian view materialized in Richard Jeffrey's (1965) *The Logic of Decision* that was much inspired by Ramsey and de Finetti. In it, Jeffrey gave a philosophical theory of Bayesian belief, decision, and induction that has had a long-lasting effect on the subsequent generations of philosophers of science and decision theorists.

As Levi and Jeffrey continued promoting Bayesianism quite generally for decades, the mainstream tended specifically toward its application in the logic of confirmation. In fact, an entire industry devoted to Bayesian confirmation theory (a topic that will later be picked up in the dissertation) emerged and attracted much support, but the theory had also faced some tough challenges culminating from critics like Clark Glymour (1980) and Deborah Mayo (1996).³ Despite the problems raised against probabilistic confirmation theory, however, the Bayesian method remained alive and well in the philosophy of science and decision theory, but it still had made relatively little impact on epistemology proper even through the '90s, decades after Jeffrey's book was published. It was not until the turn of the century that Bayesianism successfully infiltrated epistemology proper.

At the turning point, much of the inspiration for the movement in the current century, at least I think, emerged from Luc Bovens and Stephan Hartmann's (2003) *Bayesian Epistemology*. The significance of the book lies in the demonstration of how Bayesian probability can be successfully employed in the study of epistemological problems relating to the mainstream interests including coherence (see e.g. Bonjour 1985), reliability (see e.g. Goldman 1979), and testimony (see e.g. Graham 1997; Goldman 1999; Lackey 1999; Goldberg 2001). Shortly after its publication, many had recognized that a formal epistemology—that is, an epistemology employing formal methods in the broadest sense—could very well be a successful field of research, and the newly developed interest led to a flood of articles in top ranking journals tackling new and old problems using formal techniques. (The empirical claim can be verified by searching digital archives.⁴) Moreover, the field of

³See Norton (2011) for a recent and substantial critical analysis of a Bayesian confirmation theory.

⁴To get a sense of how influential Bayes is in philosophy nowadays, the recent philpapers.org archive returned 1000+ results for a query with an exact phrase match "Bayesian" and dates ranging

formal epistemology has garnered much support in recent years, and by the looks of things, it is not going away anytime soon.

1.1.1 Motivation

Skipping ahead now to more recent times and the point where my story begins. I was brought to Munich by one of the movement's architects, Stephan Hartmann, who has guided me every step of the way in completing this dissertation. When I first arrived in Munich, Stephan suggested early on to undertake a project that would further the field. The suggestion was quite intimidating at first since I had come to Munich with a background in traditional epistemology. But at the same time, I was excited to have the opportunity to learn a different and fascinating way of doing philosophy in a recently established center specializing in mathematical and scientific approaches, the Munich Center for Mathematical Philosophy.

Once I got started, the learning of formal methods quickly led to exploration in research outside of philosophy, which provided new opportunities to attend conferences and engage with academics in other fields such as computer science, economics, and statistics. Beforehand, I had very little connection to such fields since 'skepticism', 'infallibilism', 'epistemic luck', and the like attracted very little interest beyond the philosophy seminar room. The lack of interest from researchers in other fields was neither surprising nor unreasonable provided that such concepts fail to be well-defined, and some may even think that contemplation might do more harm than good by hindering scientific progress. However, I quickly discovered that a lack of interest in mainstream epistemology does not prevent philosophical inquiry from having a place in other disciplines. It does indeed have a place.

For instance, there are many outside of the philosophy profession who often admit to the conceptual and practical limitations of modeling, which has led to theorizing about extensions or new approaches altogether for solving complex problems. Behavioral economics is an exemplary field, and it became the focus of my minor study at LMU Munich. Behavioral economists recognized that the principles of classical decision theory often tend to be violated, so they invented new empirically-informed theories, e.g. prospect theory, regret theory, reference-dependent utility, etc. The latter theories are capable of accommodating preferences under the influence of cognitive biases that ordinary people regularly face.

from 2000 to 2016, while only 338 results were returned for the same exact phrase match but with dates ranging from 1950 to 1999. Granted, technological advancements may have contributed to such disparity, but the difference is quite significant regardless.

While studying some of these theoretical innovations in behavioral economics, a particular cognitive phenomenon caught my attention and would ultimately influence my PhD research, namely ambiguity aversion that was pointed out by Daniel Ellsberg (1961). From a pair of experiments, he concluded that the preferences of actual decision-makers tend to be inconsistent with Savage's axioms when facing ambiguous prospects. Decades later, Gilboa & Schmeidler (1989) developed an axiomatized theory explaining the results of Ellsberg's experiments through maximin expected utility and *imprecise probabilities*.

Before the reader becomes confused by my tangential discussion of behavioral economic theory, I bring it to attention, especially the part about Ellsberg, mainly because my early days exploring other fields with a newly cultivated understanding and appreciation of formal methods in philosophy led me to conclude early on that formal epistemology would benefit from approaches beyond Bayes. Exposure to behavioral economics, in particular, provided the realization that a variety of belief models could and should be deployed under different circumstances, and one that I found quite attractive for addressing a class of problems was an imprecise probability model similar to that described by Gilboa and Schmeidler in their representation of ambiguity aversion.

Luckily for me, I was not the only one at the Center with an interest in the framework. Seamus Bradley, Jake Chandler, and Greg Wheeler were working on philosophical problems relating to imprecise probability such as dilation and sequential decision-making while Stephan had previously done some work on imprecise probability in quantum physics. Needless to say, I had much support in pursuing the topic. But it became apparent not too long afterward that the idea of employing imprecise probability in philosophy was not novel. In retrospect, I was very late to the party since a lot of work had already been done by Isaac Levi (1974), Richard Jeffrey (1983), Bas van Fraassen (1990), Teddy Seidenfeld and Larry Wasserman (1993), James Joyce (2005; 2010), Scott Sturgeon (2008), Roger White (2009), Stephan (2010), Greg (2014), and Seamus (2014) among others.

However, what I noticed to be missing in all of the work done up to that point were specific applications of imprecise probability to the mainstream problems discussed in epistemology and philosophy of science. This revelation came as a surprise since orthodox Bayes has been extensively applied by philosophers to epistemological problems. Nevertheless, it was an opportune moment for me. Like an engineer, I had little interest in contributing to the foundations, but instead I aimed at discovering ways the formal theory might be applied. So I began thinking seri-

ously about difficult problems in epistemology that might be better addressed using imprecise probability rather than classical Bayes.

In the spring of 2014, the Center hosted a conference, *Imprecise Probabilities in Statistics and Philosophy*, which supplied me with a better understanding of the foundations and formal structures. However, my ideas on how the theory might be applied to philosophical problems were still very muddy at that time. On the traditional side, I was thinking about the problem of peer disagreement in social epistemology for quite some time, and then the “Ah, ha!” moment came when I realized that *imprecise credences* as a way of resolving disagreement made the most sense from an evidentialist standpoint. The idea was sharpened through many discussions with Greg and fully developed when we joined together in proposing an imprecise probability solution to peer disagreement, which has formed the basis of the third chapter of this dissertation.

The other substantive work making up the remainder of the dissertation seemed to come more naturally once I got going with the project on peer disagreement. Having interest in general philosophy of science, and scientific reasoning in particular, it seemed appropriate for me to look in that direction next. And since Bayesian confirmation theory still rates highly among the candidate theories in the literature on scientific inference, I had an opportunity to explore a generalized Bayesian confirmation theory with imprecise priors. What was interestingly learned early on in my research is that many confirmation theories could be constructed upon introducing sets of probabilities. After some helpful discussions and guidance from Stephan and Branden Fitelson, I went on to detail plausible candidates for a generalized Bayesian confirmation theory, but ultimately I arrived at the conclusion that each candidate theory suffers from substantive problems, which is discussed fully in the fifth chapter. Although imprecise probability has much to offer in philosophical analysis, it appears to have limitations like any other method.

As for the remaining work, a final project targeting the epistemic state of complete ignorance emerged from a discussion I had with John Norton after presenting an early version of the confirmation paper in 2015. Although showing enthusiasm for a non-classical Bayesian approach to confirmation at first, he was quick to reject the idea that imprecise probability would provide a suitable inductive logic. Pointing me towards a collection of his papers, I learned of John’s critical views on probability as a logic for inductive inference. While I grew sympathetic toward his criticisms aimed at the orthodox Bayesian method, I was not at all convinced that imprecise probability could do no better. So I took on some of the challenges

laid out by him in the series of papers, which led to the fourth chapter on complete ignorance. I am grateful for John pushing me in such a direction since in the end, I arrived at a view of ignorance that, at least in my mind, is the most compelling and no better represented by any model other than imprecise probability.

In summary, the described sequence is essentially how the present dissertation had come about, and I am indebted to those mentioned for helping me to develop the project and discover new and interesting things.

1.2 Outline of the Intended Project

With the background and motivation for the dissertation out of the way, let us turn now to the particular details of the research project. There are three philosophical topics of special interest: *peer disagreement*, *complete ignorance*, and *confirmation*. The first is considered a “newer” topic in epistemology while the latter two are old hat. What each has in common is a classical Bayesian solution. However, I recognize and hope to convince the reader that each epistemological problem might also be described, in some situations, in the language of imprecise probability when a belief or credence is imprecise. Here is what the reader has to look forward to in the subsequent chapters of the dissertation.

- The question of how epistemic peers—individuals who are equally competent and share the same information—should respond to a disagreement has recently invited *equal weight* (Christensen 2007; Elga 2007) and *steadfast* (Kelly 2011) responses. To simplify the problem, suppose that epistemic peers, 1 and 2, disagree about a proposition A such that $p_1(A) \neq p_2(A)$ where p_1 and p_2 are probability measures representing 1 and 2’s credences or beliefs, respectively. What is the rational reaction to their disagreement? An equal weight view seemingly suggests that 1 and 2 both should adopt a middle ground, $p^*(A) = \frac{1}{2} p_1(A) + \frac{1}{2} p_2(A)$. Alternatively, a steadfast view demands, at least in some instances, that each peer stands their ground such that $p_1^*(A) = p_1(A)$ and $p_2^*(A) = p_2(A)$.

In Chapter 3, both of the proposed views are challenged. Against an equal weight view, an argument from sure loss (in expectation) is given for when two or more disputes are held over propositions that are epistemically irrelevant to one another. A novel view is then detailed, which introduces *set-based credences* modeled by imprecise probability. Simply put, the common ground recommended by this account is the full set of peer opinions, \mathbb{P} , that induces lower and upper prob-

abilities. This alternative approach generates a strong argument against a steadfast view. It is based on an aversion to the risk of regret and exploits the fact that opposing opinions signal that each peer may have miscalculated the appropriate buying and selling rates for a gamble on the disputed proposition(s). Towards the end of the chapter, a qualitative account is given in which agreement, disagreement, or indeterminacy among the group opinions can similarly be modeled.

- It has become clear that Bayesians are troubled by the epistemic state of complete ignorance, which has been sufficiently demonstrated in a series of papers by John Norton (2007a; 2007b; 2008; 2010). In particular, Norton points out that a theory of additive measures representing belief and disbelief fails to satisfy a desirable duality principle relating to ignorance. The failure to satisfy the duality principle is what prevents a representation of the epistemic state in Bayesian epistemology.

The technical idea is that if p is interpreted as a belief measure and its dual M a disbelief measure, then belief and disbelief should be interchangeable in a similar vein as True and the dual False are interchangeable in Boolean algebra. But this is not the case, for the dual M does not obey the same axioms constraining p unlike how the dual of True does obey the axioms of Boolean logic. This technical flaw together with the additivity property of the measures is where the problem begins as they entail that an increase in belief entails a decrease in disbelief and *vice versa*. Additivity ultimately excludes ignorance such that either belief or disbelief in a proposition is had. After laying out concrete examples illustrating the consequences of additivity, Norton goes on to say that a generalized Bayes model does not do any better. Admitting that there are self-dual sets of probability measures, he points out that it remains unclear which set represents the state of complete ignorance and further claims that such representation is unintuitive.

In Chapter 4, I find myself in agreement about Bayes' failure. However, I disagree that imprecise probability suffers the same fate. I suggest that the epistemic state is best captured by a vacuous prior, $\{0, 1\}$, for such representation expresses no opinion at all. Next, I demonstrate that the set of measures is self-dual. Actually, it is a lower probability $\underline{P} = 0$ yielding duality given that it automatically defines an upper probability $\overline{P} = 1$ through conjugacy, relative to contingent propositions A and $\neg A$. For any contingent propositions A and $\neg A$, a lower probability 0 associated with both propositions trivially induces vacuous priors. Afterward, I illustrate that imprecise probability is an extension of an inductive logic that Norton envisages followed by an interpretation of the representation that is seemingly intuitive.

In the end, I respond to the challenge of updating vacuous priors and propose an alternative method of *credal set replacement* that circumvents the inductive learning problem or *belief inertia*.

- In the study of confirmation, Bayes has been placed at center stage and reigns supreme in the philosophy of science. With an ability to simply capture the confirmation relation between hypotheses/theories and evidence, and an ability to accommodate surprising new evidence, there is no mystery for why Bayesian confirmation theory has had much influence on philosophers of science. The final chapter explores an extension of the theory that addresses situations in which prior judgment regarding a scientific theory or hypothesis is imprecise as a result of limited or unspecific background information. Introducing imprecise priors to the game, however, radically changes our understanding of confirmation from what Bayesians have become so acquainted with. Confirmational relations are no longer based on a comparison of a single posterior probability and prior probability as they are in ordinal Bayesian confirmation theory. So what are the relations based on, then?

In Chapter 5, I give four possible answers. First, a theory or hypothesis H is confirmed by evidence E if every conditional probability in a set $\mathcal{P}(H|E)$ is larger than every corresponding unconditional probability in the set $\mathcal{P}(H)$. This view is referred to as *extremity*, which yields something like a supervaluationist theory of confirmation. Second, H is confirmed by E upon an individual's lower and upper conditional *previsions* exceeding the corresponding unconditional previsions. This view gives confirmation a behaviorist reading and is referred to as *previsions-based confirmation*. Third, a more complex theory may be needed, for the previsions theory leaves out alternate possibilities like an increase in upper probability and a decrease in lower probability, i.e. dilation. An all-encompassing theory of *confirmational sensitivity* accounts for each possible outcome in lower and upper probability. Finally, one might choose instead an absolute theory of confirmation to prevent confirmational relations obtaining when "belief intervals" overlap, e.g. $\mathcal{P}(H) = [0.4, 0.5]$ and $\mathcal{P}(H') = [0.45, 0.55]$. A theory or hypothesis H is confirmed by E just in case $\mathcal{P}(H|E)$ *interval-dominates* $\mathcal{P}(H'|E)$ for all H' .

I go on to discuss the details of each candidate for a confirmation theory with imprecise probabilities, but ultimately I arrive at the conclusion that they all suffer from substantive problems, which generates a skeptical outlook as to whether a plausible confirmation theory with imprecise probabilities is at all possible. Although the classical model (or special case in imprecise probability) is fairly simple

and intuitive, we learn that a non-singleton set of probability measures creates quite some difficulty in defining confirmation.

Chapter 2

Subjective Probability: The Bayesian Paradigm

Before turning to the analyses outlined in the first chapter, it will be helpful for the reader to have some background (or review) in Bayesian probability since the dissertation will revolve around a formal account of belief grounded in the subjective Bayesian tradition. This chapter will serve a purpose throughout, especially in thinking about how the orthodox model compares to a non-classical, imprecise probability model that is of primary interest. So let me take the time now to rehearse the probabilistic approach often adopted in formal epistemology.

The story begins with a subjective interpretation of probability due to Ramsey (1926) and de Finetti (1931/1989), which has given rise to a formalized image of belief and rationality that so many are now familiar with, especially in the domains of computer science, decision and game theory, philosophy, and statistics. In simple terms, subjective probability is a theory of “orderly opinion” (Edwards, Lindman, & Savage 1963) in which (prior) belief formation and inference are governed mathematically by a set of axioms and rule(s) for conditional reasoning, respectively. This particular theory of probability has led to what is now widely known as Bayesian epistemology.

Those unfamiliar with this tradition might wonder what probabilities and beliefs have to do with one another. On the de Finetti-Ramsey view, probabilities are reflections of a *rational* individual’s beliefs, or more specifically, grades of credence invested in a set of events or propositions.¹ We must be clear, though, that

¹I loosely switch between terms ‘belief’ and ‘credence’ throughout. While some may consider the oscillation to be confusing since the term ‘belief’ is typically reserved for an all-or-nothing epistemic attitude, I make no distinction here and do not wish to engage in the debate between full

credences need not be probabilities, for one can believe however they wish. But if credences are probabilities, those credences are considered to be optimal or rational. Starting with the assumption, then, that probabilities are rational credences, credences that are not probabilities must ultimately suffer from some defect. Credences are said to be defective if they violate at least one of the axioms of (finite) mathematical probability, hence the relation between rational credence and probability. To make the picture precise, we are in need of some basic notation.

Let \mathcal{F} be an algebra over a finite set of worlds $W = \{w_1, w_2, \dots, w_n\}$ closed under complementation, union, and intersection. A function p from \mathcal{F} into the reals of the unit interval $[0,1]$, i.e. $p : \mathcal{F} \rightarrow [0, 1]$, is a probability measure satisfying:

- $p(W) = 1$;
- $p(A) \geq 0$ for all $A \in \mathcal{F}$;
- $p(A \cup B) = p(A) + p(B)$ for all $A, B \in \mathcal{F}$ if $A \cap B = \emptyset$.

The first of these axioms states that W should be assigned maximum probability supposing that one of the worlds in W is the actual world. The second states that the value assigned to any element in \mathcal{F} is non-negative. The first and second axioms then entail that $p(A) \in [0, 1]$ for all $A \in \mathcal{F}$ since no set in \mathcal{F} is assigned a negative value and no set is more probable than W . The third, and a bit more controversial, is an additivity axiom (finite additivity axiom, to be precise). It states that the probability of the actual world being either in A or in B is the sum of their individual probabilities. The axiom should seem acceptable, though, since the union of any two sets is at least as probable as one of the sets individually. For instance, if A and B are disjoint and exhaustive, then $A \cup B = W$ and thus $p(A) \leq p(A \cup B)$ and $p(B) \leq p(A \cup B)$.

The basic axioms above give rise to some useful mathematical consequences that one should keep in mind. They include:

- $p(A \cap B) \leq p(A)$ for all $A, B \in \mathcal{F}$;
- $p(A) = p(B)$ if $A = B$, for all $A, B \in \mathcal{F}$;
- $p(\bar{A}) = 1 - p(A)$.

Opposite of union, the probability of intersecting sets, A and B , is no greater than the probability of either individual set. This is quite intuitive, logically speaking.

and partial beliefs. The reader may assume that ‘belief’ means the same thing as ‘credence’ or ‘degree of belief’ unless otherwise noted.

If $\varphi \wedge \psi$ is true, then φ is guaranteed to be true. Provided that φ is a deductive consequence of $\varphi \wedge \psi$, φ must be at least as likely to be true as the sentence $\varphi \wedge \psi$ entailing it. That is the idea expressed in the first consequence (but in set-theoretic terms). The second consequence straightforwardly says that equivalent sets should be treated the same and thus given the same probability. Finally, the last consequence defines the probability of complementary sets where $\bar{A} = W \setminus A$ is the set of worlds not in A , and its probability, $1 - p(A)$, is implied by the basic axioms.

The elementary details of mathematical probability suffice for providing us with machinery from which a formal theory of rational credence may be constructed. In our formal theory, we will say that the measure p represents an individual's belief or credal state that is relativized to a finite structure, (W, \mathcal{F}) .² Within the canonical language of subjective probability, an individual has beliefs or credences toward *events*, i.e. elements of \mathcal{F} , and a set of events typically under consideration is a partition Θ of W . Accordingly, if $A \in \Theta$, then an individual is opinionated with respect to A , i.e. $p(A)$. Keep in mind that a partition is dependent on W , which we will assume to be finite throughout for the sake of ease.

While the axioms of finitely additive probability and their consequences purportedly provide rationality constraints on credences, they only tell one how their credences should be at a fixed time. But of course, an individual will often learn new information in a dynamic world. To accommodate learning, many adopt a diachronic updating rule of *conditionalization*. First, one employs Thomas Bayes' (1764) celebrated rule (hence the name 'Bayesian')

$$p(A|B) = \frac{p(B|A) p(A)}{p(B)} \quad (2.1)$$

for determining the conditional probability of A given B for some $A, B \in \mathcal{F}$ where $p(B) > 0$, followed by the individual adopting a new level of credence $p'(A) = p(A|B)$. The procedure is continued upon learning the results of subsequent experiments until it turns out that $p'(A) = 1$ or $p'(A) = 0$.

In a nutshell, that is the Bayesian theory of credence in its most basic form. What constitutes a theory of Bayesian credence is considerably broad these days given that 'Bayesian' has become an umbrella term for probabilistic theories of credence in general. Over many decades, a variety of interpretations and rationality constraints have been imposed on Bayesian epistemologies.

²In case the context is clear, I will omit reference to the structure (W, \mathcal{F}) when talking about an individual's credences determined by p .

Regarding interpretation, the logical view from earlier was a featured contender for representing rational credence, but the program faced difficulty as noted in the introduction. The objective attempt, contra the view to be sketched and endorsed later, had received a liking, but in a different fashion by Jaynes (1957), Rosenkrantz (1977), and Williamson (2010) where they invoked objective criteria for rational credence to evade equating rational credence with pure opinion that has tended to be the Achilles heel of subjective Bayesianism. Despite an attempt to find a middle ground between Bayesians and frequentists, however, objectivists also receive a fair amount of criticism just the same as subjectivists. Still, to this day, there remain tensions between these two camps of Bayesians.

As for rationality constraints imposed on credences (objective or evidential), the most notable include the *principle of indifference* (Keynes 1921) or MAXENT (Jaynes 1957), *principal principle* (Lewis 1986) or *calibration* (Williamson 2010), and the *reflection principle* (van Fraassen 1984), just to name a few. Each principle is an advisement stating what an individual's credences should look like when the individual is either in a state of ignorance (indifference and maxent) or possesses statistical information about physical phenomena (principal principle and calibration) or thinking about their future mental state (reflection). Depending on the author, there are different ways of justifying each principle, and the various justifications have been subjected to scrutiny in the philosophical literature.

Moreover, the belief updating rule that depends on a theorem derived by the person who the theory is named seems to be the only uncontroversial feature of the theory. But it turns out that conditionalization has not gone unchallenged. As an alternative to conditionalization, for example, probability kinematics or Jeffrey conditioning (Jeffrey 1965) was proposed in order to overcome the unrealistic assumption of an individual having credence one in an evidential statement through simple conditional probability. Minimizing Kullback-Leibler divergence between prior and posterior probability distributions has been proposed for valid reasoning with uncertain premises (Hartmann & Eva, ms.). And as the reader will learn later, a generalized belief updating method of *credal set replacement* is given. So we see that conditionalization might not be a pillar of the theory after all.

I leave it to the reader, however, to explore the discussed controversies surrounding "Bayesian" theories of belief as a comprehensive survey on Bayesianism is beyond the scope of the current project.

2.1 Pragmatic Justification of Probabilism

So far, I have maintained without justification that the axioms of finite probability (and their consequences) along with conditionalization are the core Bayesian rationality constraints on credences. But for what reason should one think that credences need to be constrained in such ways? Simply because an artificial system happens to nicely describe an individual believing to some degree that a particular event will occur? No. Subjective Bayesians have a much more compelling justification for the probabilistic view of credence.

The long-standing tradition has been to defend the view by illustrating that an individual's credences regarding a set of events should obey the axioms of probability and be updated via conditionalization or else the individual ought to be willing to face a synchronic and/or diachronic *Dutch book* (de Finetti (1974); see Teller (1973) for a diachronic Dutch book argument). What this means is that a clever party would be in a good position to take advantage of the individual (prior to and/or after learning new information) by using a system of bets on the relevant events, which the individual considers to be fair, but ensures a monetary loss come what may. Accepting a set of sure-loss bets is clearly irrational. The argument concludes with a recommendation that one should form probabilistically coherent credences and update them by means of conditionalization in order to avoid being booked in a sure loss.

The pragmatic justification of what some refer to as *probabilism*—rational credences are probabilities—nicely unifies behavioral dispositions with an individual's epistemic attitudes. In addition, the justification yields a method of practical value, namely a way to determine what an individual believes and predict how they might behave. Specifically, we can learn about an individual's credences regarding a partition $\Theta = \{A, \bar{A}\}$, for example, through their *previsions* for a special type of gamble on the events (or learn about an individual's committed previsions for special gambles through their credences). To see this, let us introduce an indicator $I_X(w)$ on a subset $X \subseteq W$ that takes a world $w \in W$ as its argument and returns 1 if $w \in X$ and 0 otherwise. We will let I_X denote a special gamble that pays \$1 if the event X obtains and \$0 otherwise. With respect to Θ , an individual is expected to announce fair prices, x and y , their previsions, for gambles I_A and $I_{\bar{A}}$.

In the de Finetti-Ramsey tradition, an individual's prevision or fair price for the gamble I_A is two-sided, meaning that they would be willing to take either side of the gamble at a price x . To illustrate, suppose that an individual is willing to pay

a maximum of \$.50 for the gamble I_A . Accordingly, they should also be willing to sell the gamble to another for as low as \$.50—that is what it means to take the other side. If the individual avoids a Dutch book, then their fair price for $I_{\bar{A}}$ is \$.50 as this is implied by the infimum selling price of I_A . Now, what are we able to infer from the stated prices? Knowing the individual’s fair prices, we infer that their credence in A is $1/2$ and likewise for \bar{A} . As we observe at this moment, the individual’s epistemic state is coherent, and more importantly we have demonstrated that the epistemic state is determinable through the individual’s behavioral dispositions: what they are disposed to risk on uncertain events.

The previsions game just described is the subjective Bayesian’s *belief elicitation* method. Using this approach, we may learn whether or not an individual is rational in what they believe by how they are disposed to act, which leads us toward an operational epistemology. The method makes clear why the study of epistemic and practical rationality is a worthy endeavor, for an operationalized epistemology illuminates the role of belief in ordinary and scientific reasoning among *value-driven* human agents, namely serving as an instrument in the process of fulfilling practical goals. But not everyone agrees that the epistemic and the practical need to be so closely tied as we will see next.

2.2 A Non-Pragmatic Justification of Probabilism

Pragmatism is not the only road one can take in justifying probabilism. Since Joyce’s (1998) seminal paper “A Nonpragmatic Vindication of Probabilism,” there has been much thought given to the value of belief states independent of their role in practical reasoning. Specifically, the accuracy of belief has long been regarded as epistemically valuable since James (1896) forcefully demanded that we “believe truth!” A recent revival and lure towards *veritism*, or avoidance of inaccuracy in belief, has birthed a lively field of *accuracy-first epistemology*. Although my preferred justification for probabilism is the pragmatic one (as it will be made clear throughout), it is worth discussing the accuracy-based Bayesian movement provided its relative merits, in addition to it being a fascinating project overall.³

Accuracy-first epistemology draws heavily from decision theory, invoking measures of (epistemic) utility and dominance principles, which is why it has also

³See e.g. Joyce (1998), Greaves & Wallace (2006), Leitgeb & Pettigrew (2010a,b), Moss (2011), Pettigrew (2012), Easwaran & Fitelson (2015), Levinstein (2015), Pettigrew (2016), and Konek (forthcoming).

earned the title *epistemic decision theory*. Following Pettigrew's (2013) approach, there are a number of steps involved in vindicating probabilism in a non-pragmatic fashion. The first step is to accept that credences do indeed have epistemic value, and their accuracy is but one property making them valuable. I will not go through the arguments that attempt to support this assumption as it is highly contentious and would require more work than I can provide here, but since it is key to getting the project off the ground, we will take it for granted.

Next, the formal steps come. The first is to identify the 'vindicated' credence function for each world $w \in W$. We define the vindicated credence function at a world w as follows

$$v_w(A) = \begin{cases} 1 & \text{if } w \in A; \\ 0 & \text{otherwise.} \end{cases}$$

An individual is awarded maximal epistemic value for having credence 1 in A when $w \in A$ or credence 0 in A when $w \notin A$. The individual receives maximal epistemic disvalue if the reverse. What is taken to be epistemically valuable is an accurate belief, and one can see that an individual's belief is perfectly accurate when they have maximal credence in the event that obtains relative to a world w and perfectly inaccurate when they have maximal credence in the event that does not obtain relative to w .

On the assumption that credences vary by degree between 0 to 1, the next step involves constructing a distance measure that captures the proximity of a credence function from the vindicated or ideal credence function, which will ultimately get us closer to a *proper* scoring rule for credence functions. Let us define the following distance measure:

$$d(v_w, c) = \sum_{A \in \Theta} |v_w(A) - c(A)|^2 .$$

(To reiterate, we will only be concerned with finite sets of events, and this allows the distance measure to be well-defined.)

As Pettigrew states, once we put the above two formal steps together with the thesis that the epistemic utility of a credence function at a world is its proximity to the vindicated credence function at that world (pg. 900), we end up with a variant of the *Brier score*

$$B(c, w) = 1 - d(v_w, c) = 1 - \sum_{A \in \Theta} |v_w(A) - c(A)|^2$$

that was originally proposed by Glenn Brier (1950) and is well-known for its use in scoring weather forecasts. For our purposes, the presented version of the Brier score is a proper scoring rule that provides us with a measure of epistemic value.

The final step involves tying in components of decision theory where talk of epistemic value is replaced by talk of epistemic utility, and we introduce a utility function $U : \mathcal{O} \rightarrow \mathbb{R}$ that maps options from a set \mathcal{O} into the reals. Next, we state a general dominance principle:

DOMINANCE: For some options, $o, o' \in \mathcal{O}$, o is *strongly* dominated by o' relative to U if

- $U(o', w) > U(o, w)$ for all worlds $w \in W$,

OR o is *weakly* dominated by o' if

- $U(o', w) \geq U(o, w)$ for all $w \in W$,
- $U(o', w) > U(o, w)$ for at least one $w \in W$.

If o is dominated by o' and there is no o'' that dominates o' , then o is said to be an irrational option for an individual with utility function U .

Now, let us think of the set of options as a set of credence functions, \mathcal{C} , available for one to choose from, relative to some finite structure (W, \mathcal{F}) , and the measure of epistemic utility as B replacing U . Then, one can prove that for any credence function $c \in \mathcal{C}$, if c violates the axioms of finite probability, then there is a credence function c' satisfying the axioms and c' strongly Brier-dominates c . And if the credence function c does satisfy the axioms, then there is no credence function c' that weakly Brier-dominates c (Pettigrew 2014).

Another way to put it, a credence function that is a probability measure is strictly less inaccurate than a credence function that is not. Let \mathcal{I} be a measure of *inaccuracy* such that $\mathcal{I}(c, w) = 1 - B(c, w)$. The inaccuracy score has a ceiling of 1 (maximum inaccuracy) and a floor of 0 (minimum inaccuracy) where, like in golf, the lower the score the better. If c' is a probability measure and c is not, then accordingly $B(c', w) > B(c, w)$ for all $w \in W$, which entails that $\mathcal{I}(c', w) < \mathcal{I}(c, w)$ for all $w \in W$. Furthermore, there is no credence function c'' such $B(c'', w) > B(c', w)$ for some world $w \in W$. So c' is strictly less inaccurate than c . According to \mathcal{I} , any credence function c that is not a probability measure is

strictly worse in terms of epistemic utility than a credence function c' that is a probability measure, and thus c is irrational. These arguments suffice for establishing probabilism without appeal to the practical interests of an individual.⁴

In summary, there are at least two ways to justify a probabilistic account of rational credence. The justifications differ with each relying on contentious assumptions. On the pragmatic view, some worry that rational credence without further constraints is pure opinion based only on avoiding a sure loss. In the context of scientific reasoning, such view has not been entirely welcomed given a clash with the objective character of inquiry demanded within scientific methodology. While some who think practical values are indispensable from science may concede to a pragmatic view of belief, they are still likely to claim that pragmatic Bayesianism is insufficient without further rationality constraints imposed. As for the accuracy approach, a fetish toward truth makes for a more compelling story in selling a probabilistic view of rational credence, especially for those who think pragmatic Bayesianism's only place is in a gambling parlor. But the accuracy view faces a difficult conceptual challenge, namely justifying epistemic utility. From psychological and sociological viewpoints, it is difficult to see how any ordinary individual can separate practical interests and societal influences from human reasoning. Thus, it is unclear whether epistemic utility really exists or is a seductive fiction.

Regardless of the chosen justification, neither attempt is perfect. I do not wish to enter the debate in this project, just merely point out some issues already known to many. As I mentioned earlier, my preference is for the pragmatic view. While I do not give any meaningful defense beyond what has already be discussed, the reader may ultimately see some of its advantages in the chapters to come. With that said, we turn now to rational credences in less-than-optimal evidential situations.

2.3 Imprecision in Belief

Although orthodox Bayesianism has enjoyed much attention in addition to boasting a number of success stories scientifically and technologically, its inadequacy in certain situations has been brought to attention, which prevents it from serving as an all-encompassing formal method for modeling credences and admissible choices. The most illuminating instances of failure are those involving “Knightian

⁴While I have only described an epistemic utility-based justification for the basic probability axioms, others have proved that conditionalization, too, increases epistemic utility. See Greaves & Wallace (2006).

	<i>Red</i>	<i>Black</i>	<i>Yellow</i>		<i>Red</i>	<i>Black</i>	<i>Yellow</i>
Bet I	\$100	\$0	\$0	Bet III	\$100	\$0	\$100
Bet II	\$0	\$100	\$0	Bet IV	\$0	\$100	\$100

Table 2.1: Ellsberg Experiment: Three-Color

Uncertainty” or unmeasurable ‘risk’ (Knight 1921).⁵

In common parlance of the present day, some might call Knightian uncertainty *ambiguity*. Either way, one should quickly notice that ambiguity is not a property that can be modeled with precise Bayesian probabilities, for any classical (subjectivist) assessment of ambiguity will ultimately return *known* risks, but the Bayesian clearly has missed the aim of the exercise at that point. However, one might insist, like Savage did, that Knightian uncertainty or ambiguity may be manifested at times, but imprecise or vague probabilities have no role in a theory of rational choice. Followers of this line have a difficult time arguing the point, though, given particular empirical findings such as those found by Allais and Ellsberg.

Consider the two sets of bets in Table 2.1. Here is the relevant information you are given. There are 90 balls total in an urn. Of the 90, 30 of them are red and the remaining 60 are either black or yellow. The urn is well-mixed and you are offered bets on blindly drawing a ball of a specific color from the urn. In Bet I, you receive a \$100 reward if you draw a red ball and \$0 otherwise. In Bet II, you receive a \$100 reward if you draw a black ball and \$0 otherwise. In Bet III, you receive a \$100 reward if you draw either a red or yellow ball and \$0 if the ball drawn is black. In Bet IV, you receive a \$100 reward if you draw a black or yellow ball and \$0 if the ball drawn is red. There are two decision problems presented: the first consists in choosing between bets I and II and the second consists in choosing between bets III and IV. For the first problem, which bet do you prefer, or are you indifferent? For the second, which bet do you prefer, or are you indifferent?

Daniel Ellsberg (1961) developed the above test and surveyed his fellow decision theorists on the pair of problems. He reported the following: most surveyed had a preference for I to II, but in the second problem, most preferred IV to III. Interestingly, the reported preferences are inconsistent. This can be seen by decomposing the preference orderings. If I is (strictly) preferred to II, then the decision maker must consider a red ball being drawn more probable than a black ball being

⁵Frank Knight called for a distinction between risk and uncertainty where the former is a measurable quantity, or risk proper, while the latter he claimed is not able to be measured.

drawn. Yellow is irrelevant. In the second problem, the option of red or yellow *should* be strongly preferred to the option of black or yellow provided that red is considered to be more probable than black in the first problem, but this is not what has been observed. One conclusion we may draw from Ellsberg's experiment is that Bayes cannot always explain the preferences of real-world decision makers.⁶

Let us consider one more instance. You are given the following choices:

$$O_1 : \left\{ \begin{array}{l} 1.0 * \$1,000,000; \end{array} \right. \quad O_2 : \left\{ \begin{array}{l} 0.1 * \$5,000,000; \\ 0.89 * \$1,000,000; \\ 0.01 * \$0; \end{array} \right.$$

and then another two choices,

$$O_3 : \left\{ \begin{array}{l} 0.1 * \$5,000,000; \\ 0.9 * \$0; \end{array} \right. \quad O_4 : \left\{ \begin{array}{l} 0.11 * \$1,000,000; \\ 0.89 * \$0. \end{array} \right.$$

The test presented generates the so-called *Allais paradox* (Allais 1990). In the first problem, most subjects have reported a preference for O_1 to O_2 , but in the second, most prefer O_3 to O_4 . Given the reported preferences, it is clear that they are, in total, inconsistent with the expected utility hypothesis. But the preferences reveal another inconsistency, particularly in appetite for risk. The common preference of O_1 to O_2 is risk averse based on opting for a sure thing instead of an option with a higher expected reward, but a small probability of receiving nothing, while the common preference of O_3 to O_4 is risk seeking, for the probability of receiving nothing is greater in option 3. Since expected utility theory usually presumes risk neutrality, it does not account for the changes in an individual's risk attitudes. But these results indicate that risk is indeed a relevant factor, and so Bayesian decision theory comes up short once again.

Despite the empirical findings of Allais and Ellsberg pertaining to the psychology of decision makers that cast doubt on Bayesian decision theory, proponents of

⁶Ellsberg's experiment is not what many would consider as a scientific experiment. However, the effect of pairing known risks with ambiguous prospects, i.e. *ambiguity aversion*, has been observed in other empirical work. See Camerer & Weber (1992) who provide psychological evidence for the phenomenon.

Bayesianism might suggest that *normative* theories of credence are in focus rather than *descriptive*. However, economists have a tough time persuading firms, policy-makers, and even peers that the field of economics is a normative enterprise, especially since textbook definitions emphasize the descriptive character of the social science. Philosophers, on the other hand, have more leniency in regard to this matter. Since much of philosophy is centered on normative issues, the response is welcomed based on an *ideal* theory of credence and decision being delivered. But a shift toward normativity does not safeguard Bayesians, for some have notably disputed the optimality of Bayes.

One example is Isaac Levi (1974) who proposed that credences should be modeled with sets of probability measures instead of a single, point-valued probability measure. Of course, Levi was not the first to have such an idea, and he cites preceding proposals by I.J. Good (1952), C.A.B. Smith (1961), and Arthur Dempster (1967) to be similar to his own position he calls *revisionism*. But unlike the “intervalists” (Dempster, Good, and Smith), Levi claims that his view differs primarily with the introduction of *S-admissibility*—a decision rule that permits choosing an option O if and only if the minimum utility for possible outcomes w of O is maximum with respect to all options O' that maximize expected utility according to at least one credence function in an individual's set of credence functions (Levi 1974, 411). At first, one may judge Levi's view to be quite a distinct departure from orthodoxy, but the belief model and decision rule are consistent with the classical theory when a set of credence functions is a singleton, risk-averse otherwise.

After Levi made such a radical proposal, Bayesians not only were divided into objectivist and subjectivist camps, but they became divided on precision and imprecision. The precisionists did not seem to be impressed by his proposal, however, and one reason I suspect for why they were not moved is because Levi's initial work on the subject lacked a compelling justification for the epistemic rationality of *credal sets*. For those with an interest in inductive reasoning in science and the epistemic justification of credal attitudes, Levi's analysis left a gap to be filled and “obscure[d] the fact Bayesianism is...an epistemology” (Joyce 2005, 153), not just a theory of rational choice. Joyce, however, made an attempt at filling the gap through arguments in support of *imprecise probabilities* that rely on correctness in epistemic reaction, which I will briefly turn to.

The key for making imprecise probability an essential tool in the toolbox of methods is to illuminate the distinction between characteristics of evidence. In particular, evidence distinctively varies in *balance* and what Keynes (1921) called

weight. The former refers to how decisively data stands in favor or against some event. The latter refers to how much data is given. While balance and weight may already be familiar properties, there is yet another salient property of evidence, namely *specificity*, which is often overlooked. Information, for example, might be fully specific or unspecific. The Ellsberg example illustrates both. On one hand, the individual is supplied with fully specific statistical information about the chance of drawing a red ball from the urn. On the other hand, the individual has highly unspecific information about the chance of drawing a black ball or a yellow ball.⁷ We thus see that specificity of information may vary, and it is this particular dimension of evidence that does quite a bit of work in motivating imprecise probabilities.

To eliminate confusion induced by mixing precise and imprecise information à la Ellsberg, consider a simpler exercise.

BLACK & GREY COINS: A large number of coins have been painted black on one side and grey on the other. They are placed in an urn that is mixed well. The coins were made by a machine that may produce any bias β where $\beta \in (0, 1)$. You have no information about the proportion of coins of various biases that appear in the urn. How confident should you be that a coin drawn at random will come up black if tossed? (Joyce 2010, 283)

It does not seem at all correct to respond in an orthodox fashion for the reason that the unspecific information is consistent with a whole lot of unique, precise probability distributions. So one must tread carefully here in foregoing the pain of irrationality. The least risky option—risky in the sense of excluding a supported opinion—is to take the total set of probability distributions consistent with the evidence as one’s credence regarding the matter.⁸ While the reader may not be fully convinced yet that variation in specificity supplies a sufficient reason for considering the use of imprecise probability theory in modeling rational credences, the goal is to make the case in the subsequent chapters that draws on particular situations.

One final remark in regard to the formalism of imprecise probability. The reader may be curious as to why the mathematical notation of imprecise probability is not laid out in this section. The reason is that there is no one theory of imprecise probability. The term extends to a broad class of models with different properties,

⁷Joyce runs his own Ellsberg-style example to push the point. See pg. 168 in his (2005).

⁸One may see the risk as a kind of ‘inductive risk’ (Hempel 1965; Douglas 2000), but with respect to accepting or rejecting sharp credal attitudes modeled by probabilities rather than theories or hypotheses. Ruling out or rejecting a probability estimate from one’s credal state runs the risk of error. What exactly the risk is relativized to remains open.

some subtle but crucially important. Furthermore, their interpretations vary also, but there is a parallel account inspired by the pragmatic tradition as the reader will become intimately familiar with throughout the dissertation. Scoring rules, however, are much more controversial and appear to have little hope given a no-go result by Seidenfeld, Schervish, and Kadane (2012) showing that there is no proper scoring rule for imprecise probabilities. But as I am most lured toward the pragmatic view, the outcome is not a disappointment. As I tend to hover around a particular style throughout, differences are noted when appropriate and the chosen framework for each philosophical problem is explicated. With that said, we are now ready to move forward.

Chapter 3

Ambiguity Induced by Peer Disagreements

You and a colleague believe differently about the proposition that it will rain in London tomorrow. You are optimistic, i.e. you assign a probability greater than $\frac{1}{2}$ to the proposition, while your colleague is pessimistic, i.e. they assign a probability less than $\frac{1}{2}$ to the proposition. Neither you nor they are able to claim an epistemic advantage on the matter. You both have the same evidence, the same level of expertise, and the same cognitive skill. Needless to say, you are *epistemic peers*. You and your colleague thus find yourselves in a *peer disagreement*.¹

The issue of peer disagreement has received much attention in social epistemology over recent years. The question causing excitement: what is the rational response to a disagreement with an epistemic peer? In providing an answer, a now widely influential view in the literature asserts that you and your peer should respond to a disagreement on a particular matter by giving the same weight to each opinion (Elga 2007, 484). Upon following this recommendation, however, neither you nor your peer are permitted to maintain your originally held beliefs, but instead you both are required to revise by *splitting the difference* (Christensen 2007, 203).

EQUAL WEIGHT: If epistemic peers believe differently about a proposition, *A*, then, upon learning of their disagreement, the peers should give the same weight to each opinion and revise by splitting the difference.

There are at least three ways one can understand the equal weight response, which ultimately engender different assumptions about the nature of the evidence a

¹This chapter is largely based on a forthcoming article, “Resolving Peer Disagreements Through Imprecise Probabilities,” that is expected to appear in *Noûs*.

peer disagreement supplies, and how that evidence should guide a peer in changing their opinion. On one understanding, a peer disagreement provides evidence that either you or your peer is mistaken about the proposition in dispute. Such evidence is undermining in character. Your reaction to a peer disagreement ought to be the same as your reaction to receiving any other new but conflicting piece of evidence: you ought to *suspend judgment* on the proposition until more evidence becomes available (Feldman 2011). The suspension response is especially compelling if one accepts the following evidentialist thesis.

UNIQUENESS: An individual's evidence, on balance, supports at most a single, unique epistemic attitude towards a proposition A , for all propositions A .

Richard Feldman (2009, 2011) has defended the thesis at length insofar as a tripartite interpretation of belief is concerned—that is, you either *believe*, *disbelieve*, or *suspend judgment* on a proposition. The gist of his view is this. A body of evidence cannot reasonably support a categorical belief in both a proposition and its negation. In case of peer disagreement, shared evidence cannot reasonably support opposing views, and so one of the peers is mistaken. Without any indication of which peer made a mistake in their reasoning, the correct response is for both peers to suspend judgment until further notice as this response is uniquely determined by the evidence. Suspending judgment amounts to neither believing a proposition nor its negation. This move effectively *respects the evidence* (Feldman 2005).

If instead belief is viewed partially where an individual issues grades of credence to propositions that are represented mathematically by a unique, real-valued probability measure, p , then the suspension approach might entail each peer revising to a credence of $1/2$ for the disputed proposition. This way an individual is indifferent with respect to the truth of the proposition and its negation. Regardless of the interpretation of belief, however, the motivation for suspending judgment turns out to be the same on either account, namely that the evidence an individual receives from a peer disagreement undermines their current view. The suspension proposal is guided by the idea that one ought to respond to a peer disagreement by increasing one's uncertainty about the proposition(s) in dispute.

On another understanding of the view, a peer disagreement supplies you with a range of informed opinions, including your own, and so you ought to exploit this information. This assumption is widely held in the wisdom of crowds literature, for example (see Surowiecki 2005; Golub & Jackson 2010). By following such

line of thought, you might infer that the evidence obtained from a peer disagreement is *ameliorative* in character, and your judgment can be improved by adopting an *equally-weighted average* of the opposing opinions as your new credence rather than naïvely suspending judgment.² Evidence suggesting that judgments are improved, relative to how accurate they are, has recently been obtained through simulation results illustrating that a collective opinion improves in accuracy upon splitting the difference (Douven 2010).

A further way to understand the equal weight view is by forming a hybrid account that combines the first two understandings, entailing that a peer disagreement is undermining in character, but the judgments of peers ought to be improved by taking into account the new information and adopting some belief revision strategy, like equally-weighted averaging, i.e. split the difference. Peers effectively become conciliatory with one another by adopting the middle ground (literally) in opinion. This interpretation makes equal weight a compelling response to peer disagreement, for peers acknowledge the fact that each is equally likely to be mistaken and they respond in a way that improves their collective judgment about a proposition, thereby forcing *conciliation*. The intuitiveness and practicality of such strategy has attracted philosophers and social scientists alike toward an equal weight view.

3.1 Splitting the Difference

Much ink has already been spent on problems associated with an equal weight (or conciliatory) view. Let us set those matters to the side and focus our attention instead on the belief revision method that tends to be accepted in the literature as a representation of splitting the difference. One preliminary, though. Since philosophers are typically fixated on proposition talk rather than the conventional language of events in probability theory, I accommodate this habit by extending the standard probability framework for an arbitrary propositional language.

²Equally weighted averaging typically is assumed to be the belief revision method of the equal weight view, but whether it is representative of the philosophical view is up for debate (see Christensen 2011; Kelly 2013). Reasons for doubting the revisionary mechanism have surfaced from a list of problems brought against it in the literature (see e.g. Jehle & Fitelson 2009; Wilson 2010; Lasonen-Aarnio 2013). This list will be extended later on. Nevertheless, weighted averaging is the natural method for splitting the difference and has a long history in opinion aggregation tracing back to Stone (1961) who proposed weighted averaging to resolve group disagreements among Bayes agents with a common utility function. Stone used the phrase *opinion pool* to describe this general scenario, and *democratic opinion pool* for the special case when all opinions are equally weighted.

Let V be an interpretation for an arbitrary propositional language \mathcal{L} that associates all worlds $w \in W$ and primitive propositions of the language with a truth assignment such that for any world w and primitive proposition A , $V(w, A) = 1$ or $V(w, A) = 0$. With the introduction of an interpretation, V , we call a model $M = (W, \mathcal{F}, p, V)$ a *probability structure*, whereby $(M, w) \models A$ if and only if $V(w, A) = 1$, for all worlds w and primitive propositions A .

The correspondence between propositions and events is established through identifying $\llbracket A \rrbracket_M$ as the set of worlds of W in which the proposition A is true. The following illustrate the correspondences for connectives \vee , \wedge , and \neg .

$$\llbracket A \vee B \rrbracket_M = \llbracket A \rrbracket_M \cup \llbracket B \rrbracket_M \quad (\text{Disjunction})$$

$$\llbracket A \wedge B \rrbracket_M = \llbracket A \rrbracket_M \cap \llbracket B \rrbracket_M \quad (\text{Conjunction})$$

$$\llbracket \neg A \rrbracket_M = W \setminus \llbracket A \rrbracket_M \quad (\text{Negation})$$

With respect to a finite probability structure, an individual's credence towards either A or B being true, for example, is represented as $p(\llbracket A \vee B \rrbracket_M)$, but for simplicity, reference to M will be omitted and I will abuse notation by dropping $\llbracket \cdot \rrbracket$. So an individual's credence towards A or B will instead be denoted as $p(A \vee B)$. Throughout, we will generally be concerned with propositions, unless otherwise noted, and so an underlying probability structure will be assumed in the background.

Preliminaries aside, we turn our attention to a standard opinion pooling model representing the revisionary proposal for splitting the difference. For $i, j = 1, 2, \dots, n$, let $\{p_j\}$ be a set of probability distributions and $\{w_{ij}\}$ be a set of weighting assignments where ' w_{ij} ' is read as 'the weight that individual i assigns to probability distribution j '. In other words, weights are estimates for the reliability of the group members' opinions, and the estimates are determined by a function w that is (i) a mapping of opinions $\{p_j\}$ into the reals of the unit interval $[0,1]$, (ii) additive, and (iii) normalized to one relative to the set of opinions under consideration.

Now, for each member of the group, they revise by adopting a pooled opinion. A well-known model for opinion pooling due to Stone (1961) is the following

$$p_i^* = \sum_{j=1}^n w_{ij} p_j. \quad (3.1)$$

The model has generally been praised for its achievement of aggregating beliefs. However, it is difficult to see how the piece of mathematics represents the step-by-step procedure in resolving disagreement. A more natural pooling method in-

tuitively capturing the deliberation process, and yielding a consensus, has been shown by DeGroot (1974), Lehrer (1976), and Wagner (1978).

The road to achieving consensus described by DeGroot, Lehrer, and Wagner brings to use stochastic matrices. Let each member of a finite group of individuals form a *profile* of weight assignments, $\{w_{ij}\}$, for all opinions held by the group, including their own, that indicate the subjective estimate of reliability i assigns to j . The profiles are then bound together by row forming an $n \times n$ stochastic matrix.

In the simple case of two individuals, we are given the following matrix.

$$\mathbf{M} = \begin{bmatrix} w_{11} & w_{12} \\ w_{21} & w_{22} \end{bmatrix} \quad (3.2)$$

If members of the group happen to disagree in their forecasts, then they are able to resolve their differences through an iterative process.

Let \mathbf{P} be a column vector of opinions relative to individuals 1 and 2.

$$\mathbf{P} = \begin{bmatrix} p_1 \\ p_2 \end{bmatrix} \quad (3.3)$$

In the first round of deliberation, each member updates their opinion such that $\mathbf{P}^{(1)} = \mathbf{M}\mathbf{P}$, and then continuing on to $\mathbf{P}^{(2)} = \mathbf{M}\mathbf{P}^{(1)}$ until a consensus $\mathbf{P}^{(k)} = \mathbf{M}\mathbf{P}^{(k-1)}$ is reached at the k^{th} deliberation.

Using the collection of methods laid out, we are able to simply state that a group resolves a disagreement when individual members' opinions $p_i^* = \sum_{j=1}^n w_{ij}p_j$ agree, or collectively, $\mathbf{P}^{(k)} = \mathbf{M}\mathbf{P}^{(k-1)}$. As already mentioned, the latter appears to be a more natural representation of the procedure, for one can think of $^{(k)}$ as the number of deliberations needed for the group to resolve their disagreement once and for all. The success of the general model is driven by stationary probabilities and a one-step transition matrix \mathbf{M} of a Markov chain with k states (DeGroot 1974, 119).

Turning now to a group of epistemic peers, splitting the difference entails the following special case of the model (3.1)

$$p_i^* = \frac{1}{n} \sum_{j=1}^n p_j \quad (3.4)$$

where the weight distributed to the opinions is uniformly $1/n$, and we simplify as

written. The equal-weighting case is the most ideal scenario provided that the group resolves their disagreement in the first round of deliberation, which is not often the case when weighting assignments are not uniform among the groups' profiles.

Using the DeGroot, Lehrer, and Wagner setup, we have in the two-peer problem

$$\mathbf{M}_{1/n} = \begin{bmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{bmatrix} \quad (3.5)$$

and

$$\mathbf{P} = \begin{bmatrix} p_1 \\ p_2 \end{bmatrix}. \quad (3.6)$$

The result of $\mathbf{M}_{1/n}\mathbf{P}$ is the equally weighted opinions that peers now adopt, and consensus is obtained in a single deliberation. We therefore state that a set of epistemic peers split the difference just in case $\mathbf{P}^{(1)} = \mathbf{M}_{1/n}\mathbf{P}$. As long as the balance in respect is maintained by the group, members continue assigning equal weight to the opinions under varying conditions (Hartmann, Martini, & Sprenger 2009). So stably balanced respect will ultimately lead to an equally-weighted consensus.

Luckily for us, the peer disagreement problem typically focuses on the special case of two peers disputing a single proposition, and so splitting the difference becomes elementary without the need of generality (though, the general model, the Markov model in particular, is nice to have available for complex cases). To demonstrate, suppose in the example given at the beginning that you, p_y , and your colleague, p_c , have the following opinions regarding rain in London tomorrow: $p_y(\text{Rain}) = 0.4$ and $p_c(\text{Rain}) = 0.6$. The revised (collective) opinion according to the equal weight view is $p^*(\text{Rain}) = 1/2 (p_y(\text{Rain}) + p_c(\text{Rain})) = 0.5$, simply yielded by the model (3.4). The solution to a two-peer disagreement is easy to calculate with either method, but the general model powerfully captures the deliberation process in which peers establish a middle ground in opinion.

3.1.1 A Positive Consequence of Linear Pooling

One advantage opinion pooling strategies have over the naïve suspension of judgment approach is that pooling strategies in general, and equally-weighted averaging

in particular, yield a new credence that is guaranteed to fall within the *reasonable range* of informed opinions.

REASONABLE RANGE: For any group of peers, \mathbb{P} , whose credences in a proposition A range from x , the lowest credence in A , to y , the highest credence in A , a new credence is said to be within the *reasonable range* for members of \mathbb{P} if and only if its value is within the closed interval $[x, y]$.

To motivate why the Reasonable Range principle is indeed reasonable and a strict policy to suspend judgment is not, suppose that your credence for rain in London tomorrow is $8/10$ and your epistemic peer's is $9/10$. Upon learning of this disagreement, it would be unwise to advise either you or your peer to naïvely suspend judgment by revising to a credence of $1/2$. If the point is not immediately obvious, notice that you both judge it to be more likely to rain in London tomorrow than not, and so no strategy to resolve a disagreement among peers should mandate that each ought to suspend judgment on a proposition they both believe is overwhelmingly more likely to be true than false. Whatever uncertainty the peer disagreement may introduce, it should not destroy shared points of agreement.

Nevertheless, some may worry that certain disagreements do indeed support adopting a new opinion outside the range prescribed by the Reasonable Range principle. If you discover that you are party to a disagreement, which introduces you to variance where there was comparatively little or none before, then sometimes the reasonable response to a channel of information that increases your variance is to fault the channel rather than submit to constraints imposed by the information it delivers.³ For example, Christensen (2009, 759) devises an example of an individual who is confident that a treatment dosage for a patient is correct (0.97) and takes the opinion of a colleague who is slightly less confident in the same treatment dosage (0.96) as confirming evidence that warrants a confidence boost.

It is not at all obvious for why one ought to respond in such a way, however. While the individual in Christensen's example expresses low credence in the administered treatment being the incorrect dosage, the colleague has a slightly higher credence that the dosage is incorrect. Yet if the confidence-boost response were right, the individual would be licensed to infer from their colleague's judgment that the prospect of administering the incorrect dosage is even *lower* than one originally believed, which seems to be false unless the individual views their colleague's judgment to be biased away from the truth in a way that one's own judgment is not.

³Thanks to Richard Dawid for pushing this point.

Responding to a disagreement by adopting a judgment that falls outside the range of group opinion is reasonable only if your colleagues are not your epistemic peers. Otherwise, if every party to the disagreement is a peer and each peer's credence in A is between x and y , where $[x, y]$ is the smallest span covering the set of credences, then a response violating the Reasonable Range principle denies that the disagreement is in fact among epistemic peers or licenses one to deliberately move away, without reason, from the considered opinions of one's peers. In either case, individuals or factions of the group may be enticed to strengthen their view upon having a disagreement, leading to *belief polarization*. The evidence obtained through a peer disagreement, however, is not in any way suggestive of belief polarization, but rather a contraction in the group opinion if there is to be any movement at all. Even non-conciliationists, such as steadfasters, reject polarization.

On that note, equally weighted averaging is not the only response to a peer disagreement that satisfies the Reasonable Range principle. This is fortunate since there are instances when it is unreasonable to resolve a disagreement among peers by taking some or another *non-extreme* weighted average of peer opinions.⁴ If, for example, you are party to a peer disagreement in which nine out of ten agree yet one outlier does not, the rational response may be for the outlier to fall in line with the majority rather than for the majority to move partway to meet the outlier, especially once all of the evidence is considered, which includes the nine expert opinions that are in unison. Peerage does not confer infallibility after all.

In certain cases, what a peer learns in a disagreement with their equals is that they are in the wrong. For the time being, I only wish to point out that allowing a single peer to change their view to join a *steadfast* majority is an instance when the Reasonable Range principle is satisfied, but non-extreme weighted averaging is not. (I will consider issues with the steadfast view (Kelly 2011) in detail later on.) Moreover, any 'permissive' response to peer disagreement that allows a party to a disagreement to *stick to their guns* will trivially satisfy the Reasonable Range principle.⁵

⁴A weighted average is non-extreme just in case every peer's opinion takes values in the open interval $(0, 1)$, excluding 0 and 1.

⁵Permissive views suggest that a fixed body of evidence does not necessarily determine a uniquely rational judgment (Rosen 2001; Douven 2009; Kelly 2011; Schoenfield 2014; Kopec 2015), and thus Uniqueness is false. In a credal setting, where credences are represented by a probability measure, a trivial version of permissivism has been acknowledged since Savage's remark that theories of subjective probability "postulate that the individual concerned is in some ways 'reasonable,' but they do not deny the possibility that two reasonable individuals faced with the same evidence may have different degrees of confidence in the truth of the same proposition" (Savage 1954, 3). Non-trivial versions of permissivism arise when peers are presumed to share the same

Even though the Reasonable Range principle is satisfied by a variety of competing peer disagreement strategies—including Savage’s Minimax, calibrated maximum entropy, Maximax, and Levi’s E-admissibility—classical Bayesian methods that satisfy the Reasonable Range principle nevertheless appear to rule out an important insight from the suspension of judgment approach, namely that at least some peer disagreements increase one’s uncertainty. It is unlikely that evidence from every peer disagreement will turn out to be ameliorative in character. Sometimes the correct response to a peer disagreement is to become uncertain about the proposition in dispute. If true, how can one’s newfound uncertainty from a peer disagreement be reconciled with the Reasonable Range principle? That is one of the questions to be addressed in the positive proposal of this chapter.

Another issue to be addressed concerns a problem that conciliatory Bayesian views have in preserving some shared points of agreement among peers, which arises from the belief revision mechanism itself. It is this latter issue I turn to next, which ultimately gives way to the positive proposal in the subsequent sections.

3.1.2 Irrational Consequences of Linear Pooling

It has become common to discuss peer disagreement exclusively in terms of the special case of two peers disputing a single proposition,⁶ thereby neglecting other forms a peer disagreement may take and the different responses each form may warrant. For instance, a single outlier disagreeing with nine other peers illustrates how the distribution of group judgments may yield evidence warranting some members of the group to respond differently than others. One motivation for restricting attention to two-peer disagreements, however, is precisely to set aside disagreements like the one described that are easily defused by ‘swamping’ higher-order evidence (Kelly 2011). The restriction to two peers helps to bring the problem of peer disagreement into sharper focus by balancing the total evidence.⁷

The same, however, cannot be said for restricting attention to a single proposition. Any proposal for resolving a peer disagreement involving one proposition

values and same goals of inquiry, where it is a standard assumption in the judgment aggregation and opinion pooling literatures to fix such conditions by, for instance, stipulating a single, shared utility function. Because the plausibility of permissivism varies wildly depending both on how one models peer disagreement and how one formulates ‘permissivism’ in a particular model, a general discussion of permissivism goes beyond the scope of the current chapter.

⁶For example, see Christensen (2007, 2009), Elga (2007), Feldman (2011), Kelly (2011), Baltantyne & Coffman (2011), Schoenfield (2014), and Levinstein (2015).

⁷For discussions on high-order evidence in peer disagreements, see Christensen (2010), Kelly (2011), and Lasonen-Aarnio (2014).

should be able to handle a disagreement involving two or more. Yet, two peers disagreeing over two or more propositions puts pressure on non-extreme weighted averaging strategies. To see why, consider the following example.

HEADS AND RAIN: Forecaster One and Forecaster Two share the same data provided by the European Center for Medium Range Weather Forecasting and they each use this data to forecast rain in London for the following day (R). Forecaster One's credence in R is 0.4 while Forecaster Two's credence is 0.6. Included in their shared knowledge is information about a biased coin to be tossed today and the two forecasters disagree about that outcome, too. One's credence in the coin landing heads today (H) is 0.2 while Two's credence is 0.8. Despite their disagreements, both agree that rain in London tomorrow and the coin landing heads today are epistemically irrelevant to one another. So, while the forecasters disagree on rain tomorrow and they disagree on the coin landing heads today, both agree that there is no value in knowing the outcome of the coin toss for forecasting rain in London tomorrow.

In the given example, what Forecaster One and Forecaster Two particularly agree on is that rain in London tomorrow and the coin landing heads today are *stochastically independent*: that is, both $p_1(R \wedge H) = p_1(R)p_1(H)$ and $p_2(R \wedge H) = p_2(R)p_2(H)$, where p_1 and p_2 represent the credences of Forecaster One and Forecaster Two, respectively. So, however they decide to resolve their disagreements about today's coin toss and tomorrow's weather, their response should preserve the judgment that heads today yields *irrelevant evidence* for predicting rain tomorrow. Since the judgment of irrelevance given seems quite intuitive, strategies purporting to resolve peer disagreements ought to abide by the following principle.

PRESERVATION OF IRRELEVANCE IN EVIDENCE (PIE): If every member of a group of peers, \mathbb{P} , judges that their credence in a proposition A should remain unchanged whether or not another proposition B is true, and no member of the group changes their mind about the irrelevance of B to A after the disagreement becomes common knowledge to the group, then the resolution should preserve the judgment that B is irrelevant evidence to A .

A significant problem arising with *any* non-extreme weighted average of p_1 and p_2 that forecasters One and Two might propose to resolve their disagreement is that it will violate the PIE principle. Without loss of generality, consider the

		$p_1(\cdot \wedge \cdot)$	$p_2(\cdot \wedge \cdot)$	$p^*(\cdot \wedge \cdot)$	$p^*(\cdot)p^*(\cdot)$
H	R	0.08	0.48	0.28	0.25
H	$\neg R$	0.12	0.32	0.22	0.25
$\neg H$	R	0.32	0.12	0.22	0.25
$\neg H$	$\neg R$	0.48	0.08	0.28	0.25

Table 3.1: Forecasters p_1 and p_2 and Equally Weighted Average p^*

specific case of p^* in Table 3.1, which is the equally weighted average of p_1 and p_2 , i.e. $p^* = 1/2p_1 + 1/2p_2$. The ‘middle-ground’ determined by p^* fails to preserve independence between the coin toss today and the weather tomorrow as one can see from columns four and five,⁸ so resolving the forecasters’ disagreements by p^* does not satisfy the PIE Principle.

Although the PIE principle is seemingly compelling, not everyone agrees, and some nevertheless are unfazed by the violation. Lehrer and Wagner (1983), for instance, have argued that violations of the PIE principle are of “negligible epistemic significance.” Even critics of weighted averaging schemes, like Williamson (2015), would argue that the PIE principle should not constrain rational belief. The problem with flouting the PIE principle, though, is that a non-mandatory stance is not only unintuitive, but practically irrational. For according to p^* , heads today is epistemically relevant to forecasting rain tomorrow. Yet if Forecaster One and Forecaster Two adopt p^* while maintaining that heads today is irrelevant to rain tomorrow, they become vulnerable to suffering a sure loss.

To see the practical irrationality involved with violating the PIE principle, suppose that forecasters One and Two reconcile their disagreement by p^* , yet they continue believing that the coin landing heads today is irrelevant to rain in London tomorrow. A clever gambler may then compel them to accept a contract consisting of the following bets. The first stipulates that the gambler buys from the peers Ticket 1 for \$28 that pays \$100 if the coin lands heads today and it rains in London tomorrow, and pays nothing otherwise. The second stipulates that the gambler buys a second ticket, Ticket 2, for \$22 that pays \$100 if the coin lands tails today and it rains in London tomorrow, and pays nothing otherwise.⁹

⁸For example, to verify the first row of Table 3.1, $p_1(R \wedge H) = p_1(R)p_1(H) = (0.4)(0.2) = 0.08$, and $p_2(R \wedge H) = p_2(R)p_2(H) = (0.6)(0.8) = 0.48$, yet $p^*(R \wedge H) = \frac{p_1(R)p_1(H) + p_2(R)p_2(H)}{2} = 0.28 \neq 0.25 = \frac{p_1(R) + p_2(R)}{2} \times \frac{p_1(H) + p_2(H)}{2} = p^*(R)p^*(H)$.

⁹The betting argument is a variation of one that Henry Kyburg and Michael Pittarelli (1996) made against Levi’s E-admissibility decision rule, which, in Levi’s original form, presupposes non-extreme weighted averaging. Also, for the sake of the argument, it is assumed throughout that the utility of money for peers is linear.

	Ticket 3	Ticket 4	Net
$H \wedge R$	\$72	– \$25	\$47
$H \wedge \neg R$	– \$28	– \$25	– \$53
$\neg H \wedge R$	– \$28	\$75	\$47
$\neg H \wedge \neg R$	– \$28	– \$25	– \$53

Table 3.2: Forecaster One and Forecaster Two's Payoff

According to p^* , the bundle of tickets, Ticket 1 and Ticket 2, is fair according to the gambler—that is, $\$50(0.28) + \$50(0.22) - \$50(0.22) - \$50(0.28) = 0$. Now, suppose that the clever gambler gives the peers a seemingly advantageous opportunity to hedge by offering two more called-off bets similar to the first pair, only now the bets are arranged for the peers to judge as fair a pair of bets that swaps one of the values under the equally-weighted joint distribution with a value under the product of the equally-weighted marginal distributions.

For instance, suppose that the gambler sells to the peers Ticket 3 for \$28 that pays \$100 on heads and rain and \$0 otherwise, and he also sells to them Ticket 4 for \$25 that pays \$100 on tails and rain and \$0 otherwise. With this contract of four bets, Tickets 1-4, the peers are booked in an expected sure loss. (See Table 3.2 for the peers' payoffs and specifically the sum of the last column.) While the peers maintained that heads today is irrelevant to rain in London tomorrow after resolving their disagreement, the gambler cleverly chose to determine the payoff for Ticket 4 using the value of $p^*(\neg H)p^*(R)$, which exposes the peers to a sure loss.

One way around the problem is to double-down on linear pooling by adopting the new betting odds given by p^* as rational and come to accept that the coin and weather are in fact not independent in the reconciled judgment as originally thought. However, this response will enjoin the peers to place some value in the information provided by today's coin toss to further their epistemic goal of forecasting tomorrow's weather. So, according to this line, it would be rational for the peers to pay a fee, even if only a fraction of a cent, to learn the outcome of today's coin flip in order to better forecast tomorrow's weather. This is clearly absurd. While the move closes off the possibility of suffering a sure loss, it opens another for fortune tellers to sell to the peers epistemically useless information.

The result of the given argument is a challenge for conciliatory Bayesians. On the one hand, the measure p^* , which is the most intuitive credence for the Bayesian version of the equal weight view, cannot preserve epistemic irrelevance. Consequently, the Bayesian equal weight view does not abide by the PIE principle. This

argument applies generally to any conciliatory Bayesian who adopts a non-extreme weighted averaging of probabilities, and it extends to other conciliatory methods that fail to preserve independence.¹⁰ On the other hand, a conciliatory Bayesian who rejects the PIE Principle is committed to the view that a shared judgment of irrelevance among peers cannot, and should not, be preserved by any resolution strategy. But then the Bayesian without PIE becomes a mark for swindlers.

A way to escape the problem is simply to permit extreme weighted averaging. But this amounts to conciliation by ultimatum: you can hold any opinion you like so long as it is mine. A dictatorial response is hardly a conciliatory strategy. Without a principled reason for choosing one peer's judgment over another, there is little support for recommending the ultimatum strategy for resolving a disagreement among peers. An alternative response is simply to leave the set of peer judgments unchanged in which case everyone holds steadfast. Each peer in the set would satisfy the PIE principle by sticking to their guns and rejecting any change to their views. However, there are reasons for thinking that sticking to one's guns is not always the rational response either, namely because peers exposes themselves to a risk of regretting, a problem that will be discussed in detail in this chapter.

But before considering an argument against non-conciliatory responses to peer disagreements, a natural question to ask is whether there is some other view for reconciling the PIE principle with the demands of being conciliatory. In short, the answer is yes. It is straight-forward to formulate conciliatory responses to peer disagreements within the language of imprecise probability that satisfy both the Reasonable Range principle and the PIE principle. The first conclusion to draw from the proposed approach, which is introduced in the next section, is that one should question belief models that mandate a single, determinate subjective probability long before calling into question conciliatory responses that satisfy the PIE principle.

3.2 Set-Based Credences

What I will call *set-based credence* is a straightforward extension of numerically determinate credence pioneered by Ramsey (1926) and de Finetti (1931/1989) that was described in Chapter 2. A set-based credence, to be explained in this section, is

¹⁰See Dietrich & List (2016) and Stewart & Quintana (forthcoming) for thorough reviews of Bayesian linear pooling methods and their properties along with Wheeler (2012) for an objection to Williamson's objective Bayesian approach.

represented in terms of a (non-empty) set of probability measures, \mathbb{P} , each defined with respect to a finite probability structure, $(W, \mathcal{F}, \mathbb{P}, V)$. For the moment, one may think of \mathbb{P} as a set of Bayes agents, or set of peers, each with their own opinions about a fixed set of propositions.

For example, $\mathbb{P} = \{p_1, p_2\}$ represents the precise credal distributions of Forecaster One and Forecaster Two for all the relevant propositions in the HEADS AND RAIN example. Accordingly, the reasonable range in opinion for rain in London tomorrow is 0.4 to 0.6, and 0.2 to 0.8 for today's coin toss coming up heads, which are both given by \mathbb{P} . Observe that for any set of Bayes agents, \mathbb{P} , there is a probability measure $p \in \mathbb{P}$ whose value is the smallest, which is a *lower probability*, and a probability measure $p \in \mathbb{P}$ whose value is the largest, which is an *upper probability*, for all propositions A . Formally, a lower probability is represented by \underline{P} and an upper probability represented by \bar{P} such that for any proposition A , $\underline{P}(A) = \min\{p(A) : p \in \mathbb{P}\}$ and $\bar{P}(A) = \max\{p(A) : p \in \mathbb{P}\}$.

To accommodate conditional probability, for some proposition B and $\underline{P}(B) > 0$, there is a conditional lower probability, $\underline{P}(A|B) = \min\{p(A|B) : p \in \mathbb{P}\}$, and a conditional upper probability, $\bar{P}(A|B) = \max\{p(A|B) : p \in \mathbb{P}\}$, relative to a set $\mathbb{P}(A|B)$. If B is a logical truth, then conditional lower probability and conditional upper probability reduce to unconditional lower probability and unconditional upper probability, respectively. For the remainder of the chapter, assume that all lower and upper probabilities are defined with respect to a finite set of propositions, and I omit reference to the underlying probability structure when the context is clear.

In case the lower and upper probabilities (conditional or unconditional) are the same for all propositions in a set Θ , we say that the peers (so represented) are in *perfect agreement*. If a set of peers are in perfect agreement, the set \mathbb{P} is a singleton set consisting of a unique probability measure, p , realizing the upper and lower probabilities for every proposition.

PERFECT AGREEMENT: If $\underline{P} = \bar{P}$, then $\{p\} = \mathbb{P}$ and $p = \underline{P} = \bar{P}$.

If instead a peer disagreement occurs, then there is at least one proposition for which the upper and lower probabilities are not equal.

PEER DISAGREEMENT: A peer disagreement among \mathbb{P} occurs if and only if there is some $A \in \Theta$ such that $\underline{P}(A) \neq \bar{P}(A)$.

As one should observe from our newly extended theory, lower and upper probabilities provide a more informative approach to framing peer disagreements. However, they are not novel inventions. The idea traces back at least to Bernoulli (1713)

and Boole (1854), and developed further by Koopman (1940) and Halmos (1950). It was later observed that the language of events or propositions and *lower probabilities* is more limited in expressive capacity than the language of random variables and (lower) expectations or *lower previsions* (Smith 1961; Williams 1975; Walley 1991), an observation that has several far-ranging consequences. It is worth mentioning this point, but for the purposes of this chapter, I will set those developments to the side and restrict the discussion to a simple *lower probability model*, for lower probability is expressive enough to capture the central idea.

A lower probability model is commonly known as a type of imprecise probability model, which provides a general framework that we may use in representing and evaluating a variety of responses to peer disagreements. Every probabilistic account for peer disagreement satisfying the Reasonable Range principle can be represented and compared within this setting. As indicated above, a singleton set of probability measures is equivalent to a standard, numerically determinate probability model. In the imprecise probability setting, this model is the model of full agreement, and the Bayesian view of reconciling peer disagreement is simply one of specifying the method whereby a new model of full agreement is selected.

Although the proposal in general is not novel (see Walley 1981; Levi 1990), the basic approach detailed above ultimately pays particular attention to the subtle structural properties of the underlying set of probabilities that form the basis for lower and upper probability assessments. As it will be argued, this basis for upper and lower probability judgments plays a crucial role in modeling group opinions. Unlike a classical Bayes model, where all of the epistemically relevant information about an individual's credal commitments is allegedly captured by a single, numerically precise probability measure, lower and upper probabilities alone do not capture all epistemically relevant information about an individual's commitments. This subtly is a key difference from the likes of Levi (1980) and Walley (1991) who are committed to closed, convex sets of probabilities either as a consequence of rationality principles (Levi) or for mathematical expediency (Walley). The view given here is that convex bases ought to be permitted but not mandated.

Those remarks aside, a *set-based credence* for a proposition A induces a lower and upper probability on A relative to a base, \mathbb{P} . To see why the reasonable range as an interval determined by lower and upper probabilities, i.e. $[\min \mathbb{P}, \max \mathbb{P}]$, fails to capture all the information relevant to a peer disagreement, consider again the heads and rain example. The precise credences of Forecaster One and Forecaster Two are presented in the top two rows of Table 3.3 labeled (a). The bottom two rows, labeled

		H	R	$R H$	$R \wedge H$	$R \wedge H$
(a)	p_1	0.2	0.4	0.4	0.08	$\mathbb{P}_a[0.08, 0.48]$
	p_2	0.8	0.6	0.6	0.48	
(b)	p_3	0.2	0.4	0.6	0.12	$\mathbb{P}_b[0.12, 0.32]$
	p_4	0.8	0.6	0.4	0.32	

Table 3.3: Reasonable Ranges and Loss of Independence

(b), list the credences for a different pair of forecasters, Three and Four. Notice the symmetries between the two groups where One and Three hold identical views on heads today and on rain tomorrow and so do Two and Four. It appears that Three and Four are just duplicates of One and Two, respectively. However, group (a) differs from group (b) in the conditional judgments they endorse.

For group (a), the observation of heads today is irrelevant information to forecasting rain in London tomorrow. For group (b), though, heads today does provide relevant information to forecasting rain in London tomorrow, but Three and Four disagree with one another over how: Three believes that heads and rain are positively correlated, whereas Four believes they are negatively correlated. Despite this difference between group (a) and group (b), all four have the same reasonable range for the conditional judgment of rain given heads: $\mathbb{P}_{a \cup b}(R | H) = [0.4, 0.6]$. We see in the final column of Table 3.3, however, that the reasonable range of One and Two's set-based credence on the joint of heads and rain, $\mathbb{P}_a(R \wedge H) = [0.08, 0.48]$, and Three and Four's for the same joint, $\mathbb{P}_b(R \wedge H) = [0.12, 0.32]$, differ.¹¹

Although the reasonable ranges for each individual proposition regarding heads and rain and the reasonable range of the conditional judgment of rain given heads does not distinguish group (a) from group (b), the reasonable ranges for the joint of rain and heads do reveal a difference between the two groups—that is, $\mathbb{P}_a(R \wedge H) \neq \mathbb{P}_b(R \wedge H)$. So far, so good. However, if we were to pool (a) and (b) into a single group, the reasonable range for Three and Four on heads would be properly included in the reasonable range for One and Two. We then would be unable to distinguish between the merged group and the original pair by the reasonable range of opinions alone. The point generalizes such that R is *irrelevant* to H just in case both $\underline{P}(R | H) = \underline{P}(R | \neg H) = \underline{P}(R)$ and $\overline{P}(R | H) = \overline{P}(R | \neg H) = \overline{P}(R)$, where $\underline{P}(R)$ and $\underline{P}(H)$ are greater than 0. This means that H and R are *epistemically independent* when both H is irrelevant to R and R is irrelevant to H . In general, if

¹¹Thanks to Jennifer Carr for pointing this out.

H is epistemically independent of R under $\underline{\mathbb{P}}$, it does not follow that H and R are stochastically independent under every p in \mathbb{P} .¹²

Fortunately, the converse holds. That is, if R and H are stochastically independent under every p in \mathbb{P} , then R and H are epistemically independent under $\underline{\mathbb{P}}$. In the language of imprecise probability theory, $\underline{\mathbb{P}}$ defined in this way is an *independent lower envelope* (Walley 1991, 446). Notice that in the initial example where \mathbb{P} consists of just p_1 and p_2 , $\underline{\mathbb{P}}$ is an independent lower envelope, but adding either p_3 or p_4 to \mathbb{P}_a destroys this property. While the reasonable ranges for $\underline{\mathbb{P}}$ on $\mathbb{P}_{a \cup b}$ are the same as the reasonable ranges for $\underline{\mathbb{P}}$ on \mathbb{P}_a , not every p in $\mathbb{P}_{a \cup b}$ judges the two propositions independent. Intuitively, if One and Two agree that heads today and rain tomorrow are irrelevant to one another, adding someone else to the group who believes otherwise ends the consensus. So the notation, as laid out, allows for specifying a variety of commitments that a group of peers may have, and to work out the sometimes subtle consequences that follow from them.¹³

For instance, return to the initial heads and rain example. Forecaster One and Forecaster Two each judge that the coin landing heads today and rain in London tomorrow are independent, and their shared judgment of irrelevance becomes common knowledge to them upon learning of their disagreement. That is to say, since every p in \mathbb{P}_a —hereafter I will return to writing \mathbb{P} instead of \mathbb{P}_a —renders R independent of H , the basis set \mathbb{P} satisfies the conditions of an independent lower envelope. So the peers' individual *ex ante* judgments of epistemic irrelevance between H and R in \mathbb{P} ensure that their (shared) *ex post* set-based credences determined by $\underline{\mathbb{P}}$ and $\bar{\mathbb{P}}$ render H epistemically irrelevant to R and R epistemically irrelevant to H .

By contrast, if we replaced the two-element set \mathbb{P} by its convex hull, $Co(\mathbb{P})$,¹⁴ then $\underline{\mathbb{P}}$ based on $Co(\mathbb{P})$ would not be an independent lower envelope, even though the probability measures in $Co(\mathbb{P})$ realizing $\underline{\mathbb{P}}$ and $\bar{\mathbb{P}}$ satisfy epistemic independence. If this is not obvious, notice that $(1/2p_1 + 1/2p_2) \in Co(\mathbb{P})$, which does not preserve independence. The point here really is a familiar one given an earlier discussion, but just expressed in different terms, for the difference between the

¹²Even though irrelevance, epistemic independence, and stochastic independence (factorization) are logically equivalent for a single probability measure, assuming some regularization condition to avoid conditioning on propositions with probability 0, these three concepts are logically distinct for lower and upper probabilities. See Pedersen & Wheeler (2014) for discussion.

¹³See Walley (2000), de Cooman et al. (2011), Cozman (2012), and Pedersen & Wheeler (2014) for discussions on structural judgments in imprecise probability, and the differences between permutability and exchangeability (Walley 1991; de Cooman & Miranda 2007).

¹⁴In particular, replace \mathbb{P} with the set of probability measures constructed by all possible linear combinations of p_1 and p_2 —that is $Co(\mathbb{P}) = \{p' : p' = \lambda p_1 + (1 - \lambda)p_2, \text{ for all } 0 \leq \lambda \leq 1\}$.

original set \mathbb{P} and its convex hull $Co(\mathbb{P})$ is precisely the open set of all possible non-extreme weighted averages of p_1 and p_2 .

As a terminological aside, but one that helps connect together some of the work on imprecise probability, the convex hull of \mathbb{P} corresponds to Walley's *natural extension* of \mathbb{P} (1991), Levi's *credal set* (1980), and Joyce's *credal committee* (2010). From one point of view, the natural extension is the most naïve approach for representing credal judgments and conditional credal judgments because it ignores various structural judgments that may be in the original set \mathbb{P} . Walley discusses different extensions that incorporate different structural judgments yielding what Haenni et al. (2011) call different *parameterizations* of a set of probabilities. The independent lower envelope is but one.

With the technical remarks out of the way and an understanding of a probability model to be used throughout the chapter, we are now in a position to say what it means to have credence determined by $\underline{\mathbb{P}}$ and $\bar{\mathbb{P}}$ based on a set \mathbb{P} .

SET-BASED CREDENCE: A set-based credence in a proposition A is an individual's epistemic attitude determined by \mathbb{P} and the pair $\underline{\mathbb{P}}(A)$ and $\bar{\mathbb{P}}(A)$ —that is, \mathbb{P} is the credal basis for A determined by $\underline{\mathbb{P}}$ and $\bar{\mathbb{P}}$.

The point of describing set-based credences in the above way is the following. When assessing credences determined by $\underline{\mathbb{P}}$ and $\bar{\mathbb{P}}$ in the manner introduced, one must bear in mind the underlying probability structure, including the structure of \mathbb{P} itself.¹⁵ Fortunately, peer disagreements as described above supply the information necessary to specify each component of a probability structure, including the structure of \mathbb{P} , too. And these features allow one to work out subtle differences among a variety of judgments peers might have with respect to the considered propositions.

As further illustration of the value in identifying subtle features of set-based credences, suppose a group of peers disagree over judgments of evidential relevance. An instance of such disagreement occurred upon adding forecasters Three and Four to the original group of peers. In that case, a unanimous *ex ante* judgment of independence should not be preserved in the group's *ex post* judgments provided that the disagreement leads to uncertainty on the relevance of the disputed propositions. There are also cases where a group of peers is initially in agreement that two propositions are stochastically independent, but learning they are in disagreement over some probability judgment destroys this consensus and permits the peers to reject their initial judgments of independence and to affirm that one proposition is

¹⁵See Joyce (2010, 287).

relevant to the other. This possibility is the reason why the PIE principle includes the caveat that no member of the group changes their mind once the disagreement becomes common knowledge.

To see how common knowledge of a disagreement can undermine a prior judgment of irrelevance, consider the following example

FROM INDEPENDENCE TO DEPENDENCE: *There are two urns that both contain the same number of red and white balls. There are 99 balls of one color and a single ball of the other color in both urns, and this is common knowledge to two peers named Five and Six. Peer Five believes that both urns contain 99 red balls and 1 white ball, whereas peer Six believes that both urns contain 99 white and 1 red. Both Five and Six believe (falsely) that they are in agreement about the composition of the two urns; neither considers it ex ante to be a serious possibility that they may disagree. So each peer's ex ante belief about the urns is that a randomly drawn ball from the first urn is evidentially irrelevant for estimating the probability of drawing a red ball from the second urn. Now suppose the peers discover their disagreement with one another. Then, each peer will believe ex post that a randomly drawn ball from the first urn is highly relevant for estimating the probability of drawing a red ball from the second urn.*

In this example, the peers' *ex ante* judgments of independence should not be preserved in their *ex post* judgments.

The difference between the original heads and rain example and the two urns example is that in the former no member of the group changes their mind about any structural judgment of irrelevance upon discovering their disagreement, but in the latter everyone changes their mind about relevance upon discovering their disagreement. Notice, however, that the bases for the heads and rain example and for the two urns example both generate independent lower envelopes. What differentiates the original heads and rain example from the two urns example is that One and Two in the original example maintain the judgment that the marginal probabilities of heads and rain are independent, whereas this condition is not applicable to the two urns example and thus not binding on Five and Six.

In the language of imprecise probability theory, these two examples illustrate the difference between *strong independence* and independent lower envelopes (Miranda & de Cooman 2014). An independent lower envelope satisfies strong independence if the marginal distributions are stochastically independent. So, while the representations of the original heads and rain example and the two urns example




Figure 3.1: Reasonable Range

both satisfy the conditions for an independent lower envelope, only the representation of the heads and rain example satisfies the additional condition necessary for strong independence.

3.2.1 Why Set-Based Credences?

A set-based credence, as a representation of an individual's credal commitment towards a proposition X , induces a lower and upper probability representation of that commitment. The last section was meant to cash out the nuances of such representation and also to caution against the mistake of simply identifying an individual's credal commitment with the interval induced by \underline{P} and \bar{P} for some proposition, i.e. the reasonable range as an individual's credal commitment. One must also attend to the parameterization of \underline{P} , which is reflected both in the original topological structure of \mathbb{P} and by judgments made about properties of an extension that should or should not be preserved in light of a disagreement. Although the choice of extension for \underline{P} is unfamiliar to the classical Bayesian, the flexibility is merely a consequence of the increased expressive capacity of imprecise probability.

With that said, set-based credences typically do yield something resembling an interval of credal opinion. The lower probability and upper probability for rain tomorrow in London induced by p_1 and p_2 from our original example yields an interval constraint (of some kind) as depicted in Figure 3.1. The interval $[0.4, 0.6]$ is the reasonable range of opinion, but others have called the representation a *credal committee* (Joyce 2010; Bradley 2014) or *mental committee* (Moss 2015), which are simply alternative names for a *credal set* (Levi 1980). The idea is that the span between 0.4 and 0.6 captures some important features of indeterminacy in opinion, or imprecision in elicitation, that cannot be expressed by a determinate probability. In the peer disagreement problem, an indeterminate judgment for some proposition is imposed on each peer after they learn of equally credible estimates that nevertheless differ from their own. But again, one should not confuse set-based credences with interval-valued probabilities for the mentioned reasons.

Moreover, since the issue of peer disagreement is typically assumed to involve a group of Bayes agents, each peer's original credal judgment is a precise credence.¹⁶ To rehearse the de Finetti-Ramsey conception of credences, this means that each peer has a fair price for the proposition in question. In other words, if Forecaster One has a credence 0.4 for the proposition that it will rain tomorrow in London, then they are indifferent to engaging in two types of transactions. The first transaction calls on the forecaster to *buy* a gamble for \$.40 that pays \$1 if it rains in London and nothing otherwise. The second transaction calls on the forecaster to *sell* a gamble for the same price.

When an individual agrees to buy the gamble, they are agreeing to surrender a sure \$.40 to acquire the uncertain reward of \$1 on the condition that it rains in London tomorrow. Similarly, when the individual agrees to sell the gamble, they are agreeing to surrender an uncertain reward of \$1 on the condition that it rains to acquire a sure \$.40. In this tradition, an individual's credences can be identified with their commitment to a system of fair prices for buying and selling any finite number of gambles. The individual's commitment is rational if and only if the resolution of the bets behind such contracts does not result in a sure loss—that is, the prices committed satisfy the axioms of finitely additive probability.

The canonical subjective Bayesian view is rehearsed here to provide an illustration of the advantage that set-based credences have over the so-called *steadfast* response to peer disagreement, which proclaims that it is sometimes reasonable for you and your peer to maintain your original opinions (Kelly 2011). The argument to be given relies on what is learned by each peer in a disagreement from the perspectives of buyers and sellers. In particular, upon announcing a fair price of 0.4 for R , Forecaster One is unwilling to pay more than \$.40 for a gamble that returns \$1 if R is true, and Forecaster Two learns from this signal that Two's fair price for R may have been overestimated.

Think about how the forecasters One and Two would respond to gambles on R offered to them for less than \$.40. Both would see the gambles as bargains. So the span from 0 to 0.4 may be viewed as the range of agreement on buying prices for gambles on R . Each peer would respond to offers within this range in exactly the same way since each judges the expected value of $(R - \alpha)$ to be greater than 0

¹⁶The assumption can be relaxed permitting some or all agents having credal commitments that are indeterminate or to consider iterative peer disagreements that start with a group of standard Bayes agents but where indeterminacy is introduced by the resolution of a sequence of disagreements. We may even dispense with probabilities altogether and give a general qualitative account in terms of desirable gambles (Williams 1975; Walley 2000).

for buying prices $\alpha \leq \$0.40$. The two peers differ, however, in how they respond to offers for gambles on R that are priced between $\$0.40$ and $\$0.60$. Forecaster One will refuse to buy a gamble on R in this price range, but Forecaster Two will find any price up to $\$0.60$ acceptable. However, the disagreement should signal to Two that $\$0.60$ may be too high of a price to pay for a gamble on R . Therefore, Two's buying price for a gamble on R should change to agree with One's. This is simply what it means for Forecaster Two to change their original buying price expressed by the initial credence of 0.6 for R to a lower probability of 0.4 for R .

The roles are reversed upon turning to the selling price for R . In such case, Forecaster Two will not surrender a gamble on R that pays $\$1$ if R is true, $\$0$ otherwise, to acquire in exchange a sure reward of any amount less than $\$0.60$ while Forecaster One is willing to sell the gamble on R for as low as $\$0.40$. Forecaster One therefore is committed to unloading gambles on R for a price that Forecaster Two would never agree to match. The standpoints of the two peers are now reversed. Whereas both One and Two would agree to sell a gamble returning the uncertain reward of $\$1$ for a sure reward of $\$0.60$ or more since both judge the expected value of $(\beta - R)$ to be greater than 0 for selling prices $\beta \geq \$0.60$, Forecaster Two's refusal to sell for any price less than $\$0.60$ signals to Forecaster One that $\$0.40$ is too low of a price to sell a gamble on R to another. Therefore, Forecaster One's selling price should change to agree with Forecaster Two's. This is simply what it means for Forecaster One to change their original selling price expressed by the initial credence of 0.4 for R to an upper probability of 0.60 for R .¹⁷

The span between lower and upper probabilities for some proposition is determined by the range of credences expressed by a group of peers. As argued, there is no good reason to adopt an opinion outside of this range. Recall also that the discussion of the PIE principle knocked out conciliatory Bayes responses but left open the option of remaining steadfast in one's opinion. Kelly, for instance, ar-

¹⁷The assumption that the currency of trading is linear is important for pinning down an estimate of an individual's strength of belief in a proposition, and the operational details of the procedure for eliciting such credences are likewise important for making sense of such numbers. When those conditions are clearly specified and met, and strategic considerations are safe to leave aside, the talk of pricing the value of gambles translates directly to an individual's epistemic commitments.

What is novel about a theory of lower previsions, which the lower probability model belongs to, is that it allows an individual to commit to different buying and selling prices for a gamble. The theory of linear previsions, which standard precise Bayesian probability models belong to, does not allow an individual to commit to different buying and selling prices for a gamble, but instead takes for granted that there is a single number, the individual's fair price. There is nothing imprecise or indeterminate about the highest price you are willing to pay for a gamble or the lowest selling price you are willing to accept for it, regardless of whether those values are different or the same. So, the behavioral interpretation clarifies what set-based credences are meant to represent.

gues that sufficient ‘higher-order’ evidence is not always generated by two-person disagreements to warrant either peer to change their view (Kelly 2011). So peers in those situations are reasonable to remain steadfast. However, the steadfast response ignores the significance of the evidence supplied by a peer disagreement, essentially making it epistemically irrelevant to each peer’s credal commitment(s).

To remain steadfast in a peer disagreement is to ignore evidence that one should change their view. Suppose Forecaster One adopts a lower probability of 0.4 and an upper probability of 0.6 for reasons spelled out above, but Forecaster Two sticks to their guns and persists in viewing 0.6 as a fair price for R . Then, Forecaster Two would discover that Forecaster One refuses to pay more than \$.40 for a gamble on R but also refuses to sell the gamble to Two for less than \$.60. What Two learns from One is that One judges the expected value of $(R - \alpha)$ to be *negative* for prices α greater than \$.40, whereas Two judges their expected loss to remain at zero. Conversely, both Two and One judge One’s commitments to be non-negative in expectation. So the outcome of the disagreement is that Two receives evidence that they may be exposed to a loss whereas One receives no such evidence. This difference in judgment between One and Two may be defensible if Forecaster Two thought Forecaster One a fool or lacking information that Two had about rain tomorrow in London, but these differences are explicitly ruled out by the conditions of a peer disagreement.¹⁸ By remaining steadfast, then, Forecaster Two embraces an exposed risk of loss that Forecaster One does not without having a countervailing reason to persist in doing so.

The badness of such approach can be illuminated further by considering another intuitive principle.

MINIMAL RISK OF REGRETTING (MRR): For a set of peers, \mathbb{P} , each member minimizes the risk of regretting upon engaging in one of the two types of transactions described just in case each revises their lower and upper probabilities to the *group’s* min buying price and max selling price, respectively, for a gamble on A that pays \$1 if A is true, \$0 otherwise.¹⁹

¹⁸And the assumptions about a shared linear scale of value rule out differences in utility functions.

¹⁹The idea is partly inspired by Savage’s (1951) *minimax regret* criterion, Loomes & Sugden’s *regret theory* (1982), and Filiz-Ozbay & Ozbay’s (2007) *anticipated regret*. One thing to keep in mind here is that ‘regret’ is taken to be relative instead of absolute. If I agree to buy a gamble on a proposition A at any amount of money, I may experience a feeling of regret if A is false given my loss. The latter case is what I mean by absolute regret. I am instead considering a notion of regret that is relative to the steps taken prior to learning whether a proposition is true or not, and those steps are reflected upon after the fact, similar to the individual’s situation in the REGRET example.

	A	$\neg A$		A	$\neg A$
Steadfast	$(1 - p)$	$-p$	Steadfast	$(p - 1)$	p
SB Credence	$(1 - \underline{P})$	$-\underline{P}$	SB Credence	$(\bar{P} - 1)$	\bar{P}

Table 3.4: Buy-side Payoffs (left) and Sell-side Payoffs (right)

To see why this principle is indeed reasonable, consider the following example.

REGRET: You are looking to purchase a new TV. In your mind, your max price is \$500. You and a friend go to a local electronics shop to look at TVs and you come across a display with a nice, high-definition TV priced at \$500. You and your friend have similar knowledge about the model, but neither you nor your friend know the competitive market price. You say, “I think it’s worth \$500 and I’m going to buy it.” Your friend, however, disagrees thinking that the price is too high and that you could get it for a lower price elsewhere. They effectively signal that you may end up paying too much for the TV. Ignoring your friend, you buy the TV and walk out a happy customer. But your happiness is short-lived, for when you get home, you check Amazon and find that the same TV is selling for \$400 with free next day shipping.

Excluding post-purchase rationalization of your consumer decision, you are likely to experience a negative feeling of regret for buying the TV at the price you considered acceptable prior to and after learning your friend’s opinion. Taking your friend’s opinion under advisement instead may have saved you up to \$100, but the refusal to count their opinion as evidence against your own, and ultimately persisting in maintaining your initial opinion, is directly linked to your suffering. Analogously, the situation described corresponds to peers buying gambles on disputed propositions, and of course, we can just as easily concoct a scenario for the selling side, too, e.g. selling a used vehicle.

The point to be made is that remaining steadfast in the face of a peer disagreement exposes one to a risk of regretting. To what degree? The maximal risk of (relative) regret, r , one faces is the absolute difference of the group’s min buying and max selling prices, i.e. $|\underline{P} - \bar{P}|$. If one has a non-extreme opinion $p \in (\underline{P}, \bar{P})$, the non-maximal risk of regret on the buy-side is $|p - \underline{P}|$ and $|p - \bar{P}|$ on the sell-side.²⁰ Upon $|\underline{P} - \bar{P}| = 0$, the group is not exposed to a risk of regretting, *at least not to the best of their knowledge*. In case $\underline{P} \neq \bar{P}$ for at least one $A \in \Theta$, it is easy to

²⁰Notice that any opinion taken to be a non-extreme weighted average has some exposure.

see that members of the group may become exposed to a risk of regretting r , i.e. $|\underline{P}(A) - \bar{P}(A)| \geq r > 0$, when agreeing to buy or sell gambles on A . For what reason would one embrace a non-zero risk of regretting, especially upon being tipped off by a peer that one might be mistaken? I see no good reason at all.

Furthermore, the degree to which the steadfast might end up regretting (in a two-peer case) is made explicit in Table 3.4. Notice in the left table that p is an individual's initial fair price for a gamble on A that pays \$1 if A is true and \$0 otherwise. If $p = \underline{P}$, then there is no difference between the rows. But if $p \neq \underline{P}$, then $|p - \underline{P}| > 0$, and so one is exposed to a risk of regretting after consulting with the group given a willingness to buy a gamble on A at their fair price p . In particular, if one pays p for the gamble and A is true, then the gain $(1 - p)$ is smaller than $(1 - \underline{P})$, and if A is false, the loss $-p$ is bigger than $-\underline{P}$. In comparison to one's peer who engaged in a similar transaction but for a price \underline{P} , the buyer should experience a negative feeling of regret. So, revising to the lower probability appears to be the rational move here.

In addition, a similar argument can be made on behalf of the selling side. If an individual's selling price $p = \bar{P}$, then there is no difference between the rows in the right side of Table 3.4. But if the individual's selling price $p \neq \bar{P}$, then $|p - \bar{P}| > 0$, and so one is exposed to a risk of regretting after consulting with the group given a willingness to sell a gamble on A for as little as p . In particular, if one sells the gamble for a price p and A is true, then the loss $(p - 1)$ is bigger than $(\bar{P} - 1)$, and if A is false, the gain p is smaller than \bar{P} . Likewise, the seller should experience a negative feeling of regret, especially when comparing their behavior to their peers' behavior given that the peers would not sell the gamble for less than \bar{P} . So, revising to the upper probability appears to be the rational move here.

As it should be clear now, it is rational on two further dimensions for peers to adopt set-based credences in light of a peer disagreement: (1) to minimize a potential feeling of regret and (2) to maximize potential gains while minimizing potential losses. It is therefore irrational for peers to discount evidence indicating that they may have mis-priced gambles on the proposition(s) in dispute as those inclined persist in maintaining their originally held views.

3.3 Summary Thus Far

In Table 3.5, the results for each approach are given with respect to the desiderata put forth in this chapter. Of course, the set of principles is not exhaustive, but the

	Reasonable Range	Pres. Irrel. Evidence	Min. Risk of Reg.
Equal Weight	✓	X	X
Steadfastness	✓	✓	X
S-B Credences	✓	✓	✓

Table 3.5: Checklist

three constraints nevertheless are plausibly within the realm of what one should consider desirable properties expected to be obtained by any strategy purporting to resolve peer disagreements. Clearly, set-based credences fare well in comparison to its rivals, equal weight and steadfastness, for they satisfy all the principles.

In review, set-based credences trivially satisfy the Reasonable Range principle in which the lower \underline{P} and upper \bar{P} are the $\min \mathbb{P}$ and $\max \mathbb{P}$ relative to a set of propositions Θ . They abide by the Preservation of Irrelevant Evidence (PIE) principle when necessary. In particular, if two peers judge propositions A and B to be stochastically independent and they maintain this judgment after resolving disputes on each of these propositions, then \underline{P} is an independent lower envelope in which A and B are epistemically independent according to each $p \in \mathbb{P}$. However, the PIE principle's caveat that no peer changes their mind about the irrelevance once the disagreement becomes common knowledge does not require independence to be maintained after resolving a peer disagreement, and this is a good thing provided the instances when it is intuitive to break the consensus. So set-based credences adhere to the PIE principle on the right occasions, but do not force unanimity when there is genuine dispute over the epistemic relevance of propositions.

Set-based credences also obey the Minimal Risk of Regretting (MRR) principle given that a revision of each peer's lower and upper probabilities to the group's min buying and max selling prices for gambles on the relevant propositions in dispute minimizes the risk (at least to the best of the group's knowledge). Any opinion amounting to a proper subset $\mathbb{P} \subset [\underline{P}, \bar{P}]$ does not minimize the risk as much as set-based credences. And so with the MRR principle on the table now, a set-based credal approach is the dominating strategy covered in this chapter. Satisfying the latter principle, at least to me, is one of the biggest advantages of the proposed approach, especially given the significant role that epistemic attitudes play in guiding action. As Savage notes,

...it can be argued that all problems of statistics, including those of inference, are problems of action, for to utter or publish any statement is, after all, to take

a certain action. (1951, 55)

The point that Savage makes is that the epistemic realm is not cutoff or divorced from the practical realm: the two work in tandem. So an evaluation of peer disagreement-resolving strategies solely on epistemic merits is not well-grounded, for the practical consequences of reasoning from such newly found credences are non-negligible. Provided that a peer disagreement increases one's uncertainty with regard to the truth of a proposition, it is unwise to give a resolution that only takes into account the evidential import while writing off the practical risks of reasoning from increased uncertainty. The MRR principle is a safeguard against such neglect, illuminating those strategies with exposure to risk and those without.

An additional bonus yielded by the proposed view is a *unique* set-based credence for all parties to a peer disagreement. So the proposal may be viewed as embracing a central tenet of UNIQUENESS discussed earlier that reconciles a seemingly intractable conflict over the nature of the evidence that a peer disagreement generates. For those who prefer the tripartite view of *believe*, *disbelieve*, and *suspend judgment*, the conciliatory response to suspend judgment still has its appeal. But such response, given the limited options, forces an overestimation of the evidence from a peer disagreement, leading to maximal uncertainty. On the other hand, conciliatory Bayesians who restrict themselves to a single determinate probability measure interpret the evidence from a peer disagreement as being ameliorating. The set-based credence account embraces the insight from traditional suspension of judgment views that peer disagreements do not generate purely ameliorative evidence.²¹ But unlike the naïve suspension of judgment approach, the proposed view preserves ranges of agreement and comparative judgments that are lost by naïvely adopting a credence of $1/2$ for representing maximal uncertainty. For those who still have their doubts about numerical credences, whether precise or imprecise, a corresponding qualitative approach may be given, which I turn to next.

3.4 A Qualitative Account

An inclination towards the traditional, tripartite interpretation of belief remains quite common, and in some cases, there is good reason for not parting ways. In particular, talk of categorical or all-or-nothing belief plays a central role in ordi-

²¹Scott Sturgeon (2010) and Haenni et al. (2011) each consider interpreting the span between a lower and upper probability the degree to which an agent suspends judgment.

nary discourse since individuals very often admit only to either believing or not believing some proposition. The concept also plays an important part in heuristically describing practical reasoning in the following form: ‘S desires x ’, ‘S believes that φ -ing will satisfy S’s desire’, ‘So, S ought to φ ’ (Williams 1979). The familiarity of categorical belief talk in explanations of practical reasoning makes the convention difficult to eliminate from daily life. So why are we talking about credence, then?

Though I have no doubts about the conventionality of traditional belief talk, I also have little doubt that pressing one harder will reveal more specific epistemic attitudes. What I am getting at are comparative judgments. In serious discourse, ordinary individuals tend to reveal their uncertainty after being interrogated, and they make their position on some matter clear through comparisons of the relevant possibilities. Some think of the attitudinal state in terms of *comparative confidence* judgments (Hawthorne 2009; Fitelson ms.). As a representation of belief, an ordinal, comparative confidence model finds a place between categorical belief and numerical credence.²² Specifically, comparative confidence is more expressive than categorical belief, yet it maintains a qualitative character, unlike the unrealistic view that human agents have sharp numerical credences in their head.

On the issue of peer disagreement, some have returned to framing the problem in the language of categorical belief to circumvent various issues with conciliatory Bayesian views (Christensen 2011; Kelly 2013). A mere appeal to the fact that ordinary people talk in terms of categorical belief, and thus their disagreements are likely to be framed in the same way, may suffice as a reason for not acknowledging the failures of conciliatory Bayes. In other words, if ordinary individuals are not talking in terms of credences, why should we? But as Kelly (2011) argues, categorical belief’s simplistic structure only admits to there being *strong disagreements*—one believes X while another believes $\neg X$. However, disagreements are not always in fact extreme. So it would be unreasonable to evaluate the problem in such simple terms given the many disagreements that get excluded.

The back and forth oscillation in framing peer disagreements has mounted tension in choosing an appropriate interpretation of belief where categorical belief is held to be too simple of a model for appropriately addressing the philosophical problem, yet the alternative approach, i.e. numerical credences (precise or imprecise), expels a qualitative characterization of ordinary epistemic attitudes. Luckily,

²²The idea very much resembles comparative probability (Keynes 1921; Koopman 1940; Savage 1954; de Finetti 1974), but I will leave it loosely associated with the term ‘confidence’ to avoid any need to reduce the theory to additive probability.

however, the tension can be somewhat alleviated by using a qualitative comparative confidence model, which may also provide a viable response to Kelly's argument against a qualitative notion of belief.

In this section, I will adopt the language of *relative likelihood*, which yields a comparative confidence model that retains gradability, but in qualitative terms. The proposed framework will be nearly the same as described in Halpern (2003). In relating set-based credences to comparative confidence judgments, we will see that there are similar corresponding judgments. However, it is made clear that neither set-based credal judgments nor comparative confidence judgments are defined by the other due to the axioms of finitely additive probability conflicting with the relative likelihood constraints. Despite the misfortune, I briefly describe an alternative account of possibility theory that coheres with relative likelihood and resembles imprecise probability at least insofar as we are concerned.

To begin, let \succeq be a relation on a finite set of possible worlds W that is *reflexive* and *transitive*—that is, for all $w \in W$, $w \succeq w$, and for all $w, w', w'' \in W$, if $w \succeq w'$ and $w' \succeq w''$, then $w \succeq w''$. The relation \succeq is a *partial preorder* as that there may be at least two worlds, w and w' , that are not comparable: $w \not\succeq w'$ and $w' \not\succeq w$. In addition, \succeq is not a *partial order* given that the relation is not necessarily *anti-symmetric* (Halpern 2003, 45). Note, however, that \succeq can be a *total preorder* if all worlds are comparable, but it need not be total.²³ Given these details, we will regard a statement $\lceil w \succeq_S w' \ \& \ w' \not\succeq_S w \rceil$ as expressing that world w is (strictly) more likely than world w' from the perspective of an individual S . A statement $\lceil w \succeq_S w' \ \& \ w' \succeq_S w \rceil$ expresses that worlds w and w' are equally likely from the perspective of an individual S . And a statement $\lceil w \not\succeq_S w' \ \& \ w' \not\succeq_S w \rceil$ expresses that worlds w and w' are incomparable from the perspective of an individual S .²⁴

So far, we have defined a partial (possibly total) preorder on a finite set of possible worlds W , but for the moment we are unable to say anything about comparing subsets of W , i.e. propositions. To accommodate propositions, we introduce a relation \succeq^p on 2^W (where p stands for propositions). In comparing sets of possible worlds or propositions, an individual S regards a proposition X at least as likely as a proposition Y , i.e. $X \succeq_S^p Y$, just in case for all $w \in Y$, there is at least one $w' \in X$ such that $w' \succeq_S w$. If $\lceil w' \succeq_S w \ \& \ w \not\succeq_S w' \rceil$ for some $w' \in X$ and all

²³In such case, reflexivity is superseded by *completeness*, which says that all worlds $w, w' \in W$ are comparable—that is, either $w \succeq w'$ or $w' \succeq w$ (see Hawthorne 2009).

²⁴Usually, \succ is used for representing instances of strict confidence and \sim for instances of equivalence, but I will maintain use of the single relation \succeq throughout and write out the longhand expressions when called for.

$w \in Y$, it follows that $X \succeq_S^p Y$ but $Y \not\succeq_S^p X$. In that instance, S regards X to be (strictly) more likely than Y . On the other hand, if for all $w' \in X$ and all $w \in Y$ it is neither true that $\lceil w' \succeq_S w \ \& \ w \not\succeq_S w' \rceil$ nor $\lceil w \succeq w' \ \& \ w' \not\succeq w \rceil$, then $X \succeq_S^p Y$ and $Y \succeq_S^p X$. In that instance, S regards X and Y to be equally likely.

Moreover, the relation \succeq^p is an extension of \succeq and abides by the following constraints.

- If $X \subseteq Y$, then $Y \succeq^p X$ (*Respect for Subsets*).
- For all finite index sets I , if $X \succeq^p Y_i$ for all $i \in I$, then $X \succeq^p \cup_i Y_i$ (*Union Property*).
- If $X \succeq^p \{w'\}$, there is a $w \in X$ such that $\{w\} \succeq^p \{w'\}$ (*Determination by Singletons*).
- $\emptyset \not\succeq^p X$ where $X \neq \emptyset$ (*Conservativity*).

Some of the constraints should be intuitive. For example, Respect for Subsets says that a set is at least as likely as any of its subsets. Conservativity is fairly straightforward, too, as it implies that non-empty sets should be considered possible. The other two constraints may be unfamiliar and seem a bit strange. The Union Property states that if some set X is at least as likely as all other considered sets Y_i , then X is at least as likely as the union of all other considered sets Y_i . The Determination by Singletons constraint states that a set X is at least as likely as a singleton set $\{w'\}$ if a singleton subset of X is at least as likely as $\{w'\}$ (ibid. 46). As advertised, the likelihood of propositions is determined purely by singleton sets.²⁵

Furthermore, the Union Property and Determination by Singletons yield a *qualitative* property if \succeq^p is *total*. A relation \succeq^p is qualitative just in case for disjoint sets X_1 , X_2 , and X_3 , if $(X_1 \cup X_2) \succeq^p X_3$ and $(X_1 \cup X_3) \succeq^p X_2$, then $X_1 \succeq^p (X_2 \cup X_3)$. Essentially, the bulk of ‘likeliness’ is placed on X_1 if \succeq^p is qualitative, and the property happens to nicely accord with the intuitive notion of ‘qualitative’ belief in that a qualitative belief places most confidence in some

²⁵The Union Property is contentious for those who have probability on the mind since it does not necessarily hold in classical probability theory. For instance, suppose that $W = \{w_1, w_2, w_3\}$ and $p(\{w_1\}) = 0.45$, $p(\{w_2\}) = 0.25$, and $p(\{w_3\}) = 0.30$. It is clear, on a probabilistic interpretation of \succeq^p , that $\{w_1\} \succeq^p \{w_2\}$ and $\{w_1\} \succeq^p \{w_3\}$, but $\{w_1\} \not\succeq^p (\{w_2\} \cup \{w_3\})$. Additivity of probability measures prevents the singleton, $\{w_1\}$, from being at least as likely as the union of remaining singleton sets in the example, and so it turns out that the Union Property is not consistent with finite additivity. Thus, we encounter a problem in reducing \succeq^p to probability.

proposition while relatively less confidence in logically incompatible propositions. But note that even if \succeq^p is total, the relation is not necessarily qualitative (ibid. 47).

With the relevant details for modeling comparative confidence judgments laid out, we now turn to peer disagreement. In making things simple, consider a variation of the rain example where the set of propositions in which our peers, Forecaster One and Forecaster Two, are opinionated is $\mathcal{A} = \{\top, R, \neg R, \perp\}$. To model our peers' judgments, let \succeq_1^p and \succeq_2^p represent the relative likelihood relations of forecasters One and Two, respectively. Following the original example, One is pessimistic about rain in London tomorrow while Two is optimistic. It is obvious, then, that each hold at least the following judgments:

$$\text{One: } \neg R \succeq_1^p R \ \& \ R \not\succeq_1^p \neg R \quad || \quad R \succeq_1^p \perp \ \& \ \perp \not\succeq_1^p R \quad || \quad \top \succeq_1^p \neg R \ \& \ \neg R \not\succeq_1^p \top.$$

$$\text{Two: } R \succeq_2^p \neg R \ \& \ \neg R \not\succeq_2^p R \quad || \quad \neg R \succeq_2^p \perp \ \& \ \perp \not\succeq_2^p \neg R \quad || \quad \top \succeq_2^p R \ \& \ R \not\succeq_2^p \top.$$

A disagreement manifest simply by $\neg R \succeq_1^p R \ \& \ R \not\succeq_1^p \neg R$ and $R \succeq_2^p \neg R \ \& \ \neg R \not\succeq_2^p R$ —that is, One regards no rain in London to be more likely than rain, while Two regards rain in London tomorrow to be more likely than no rain. Both, however, are in agreement that neither of the propositions are logical falsehoods provided that One judges that $R \succeq_1^p \perp \ \& \ \perp \not\succeq_1^p R$ (and given $\neg R \succeq_1^p R \ \& \ R \not\succeq_1^p \neg R$, then $\neg R \succeq_1^p \perp \ \& \ \perp \not\succeq_1^p \neg R$ by transitivity) and Two judges that $\neg R \succeq_2^p \perp \ \& \ \perp \not\succeq_2^p \neg R$ (and given $R \succeq_2^p \neg R \ \& \ \neg R \not\succeq_2^p R$, then $R \succeq_2^p \perp \ \& \ \perp \not\succeq_2^p R$ by transitivity). Thus, the collections of judgments manage to capture the peers' "qualitative" uncertainty, yet they are expressive enough to capture strength in confidence in the propositions.

The upshot of the relative likelihood model is that it preserves points of agreement from a coarse-grained perspective even when there is an underlying fine-grained dispute. The Bayesian suspension model is not so fortunate. Specifically, imagine from a fine-grained perspective that a dispute occurs between One and Two over R , and we model the dispute with precise probabilities. For instance, suppose that $p_1(R) = 0.9$ and $p_2(R) = 0.85$. A suspension view entails a belief revision yielding $p_1(R) = p_2(R) = 0.5$. That response is absurd, however, given the expression of high confidence in R by both peers. From a coarse-grained perspective, the peers are in agreement. The agreement can be easily shown in the model for relative likelihood: $R \succeq_1^p \neg R \ \& \ \neg R \not\succeq_1^p R$ and $R \succeq_2^p \neg R \ \& \ \neg R \not\succeq_2^p R$. Or

we might define a collective judgment. Let \succeq_G^p be a collective, relative likelihood relation for a group G . In the above case, $R \succeq_G^p \neg R$ & $\neg R \not\succeq_G^p R$.

Although a categorical approach is not bound to run afoul like the Bayesian suspension approach provided that one might insist that both peers flat-out believe R , and therefore they are in agreement, the categorical model does not have enough structure to satisfactorily capture an individual's epistemic attitudes when the set of logically incompatible propositions has cardinality greater than two with respect to contingent propositions, e.g. $\{A, B, C, D\}$. In such case, the categorical model requires an individual to either believe one proposition in the set while disbelieving all the others or suspend judgment on every proposition. But this representation is not very informative. Using relative likelihood, on the other hand, we are able to at least provide a partial if not total preorder on the set of propositions, thereby leading to a more informative representation. So a comparative confidence model resolves the mounted tension with choosing a model of belief in the effort of framing peer disagreements, that is, if one considers there to be a tension at all.

Where do set-based credences fit in? Set-based credences appear to correspond with relative likelihood similar to precise credences. In the special case where $\mathbb{P} = \{p\}$, the model reduces to classical Bayes, and corresponding representations in relative likelihood are constructed accordingly.²⁶ If, however, \mathbb{P} is not a singleton set, then for all propositions $X, Y \in 2^W$, judgments like $\lceil X \succeq^p Y \text{ \& } Y \not\succeq^p X \rceil$ correspond to set-based credences if and only if $\underline{\mathbb{P}}(X) > \overline{\mathbb{P}}(Y)$. This should be intuitive, for regarding X to be (strictly) more likely than Y requires a minimum credence in X to be greater than the maximum credence in Y . As for indifference, judgments like $\lceil X \succeq^p Y \text{ \& } Y \succeq^p X \rceil$ correspond to set-based credal judgments if and only if $\mathbb{P}(X) = \mathbb{P}(Y)$ —that is, X and Y have the same lower and upper probabilities induced by \mathbb{P} .²⁷ Again, this should be rather intuitive.

Distinct from the Bayesian model, though, there are instances involving incomparable propositions in credence and relative likelihood alike. If set-based credences in X and Y overlap, i.e. neither strictly dominates or one properly contains the other, then the propositions X and Y are incomparable in which case $X \not\succeq^p Y$ and $Y \not\succeq^p X$. Fortunately for us, as defined above, a total preorder on sets of propositions is not mandated, only a partial preorder. So correspondence between imprecise probability and relative likelihood is not lost upon finding incomparable

²⁶Note that correspondence does not entail equivalence. See footnote 25 demonstrating that relative likelihood is not *reducible* to probability.

²⁷Note that $\mathbb{P}(X)$ and $\mathbb{P}(Y)$ might differ structurally, which may not be straightforwardly recognized in the relative likelihood model.

sets. With that said, I want to emphasize that I am not advocating a reduction of any kind here, which I will explain why in a moment, but merely pointing out that in many cases, there are seemingly corresponding judgments between models.

A limitation on the theory of relative likelihood lies with comparative judgments, as represented by a partial preorder, not being reducible to credences (precise or imprecise). As already noted in footnote 25, the finite additivity constraint in classical probability theory is inconsistent with the Union Property. And given that any non-empty set, \mathbb{P} , is comprised of additive probability measures, then set-based credences do not necessarily satisfy the Union Property. The irreducibility of relative likelihood to probability is not so much a problem in simple cases of peer disagreement, though, for in many instances there are seemingly corresponding judgments. But if one is interested in the determination of credence rather than mere correspondence between different models, then one may find possibility theory (DuBois & Prade 1988) to be a better option over imprecise probability, which similarly captures the imprecision in belief expressed by set-based credences.

In possibility theory, we introduce an underlying, real-valued possibility distribution over a finite set of possible worlds W . The function $\pi : W \rightarrow [0, 1]$ is a mapping from W to the unit interval $[0, 1]$, and at least one $w \in W$ is assigned maximum possibility, i.e. $\pi(w) = 1$. With an underlying possibility distribution, we are able to construct a possibility measure $\Pi : 2^W \rightarrow \mathbb{R}$ for sets of possible worlds, i.e. propositions, that maps W to 1 and \emptyset to 0. For all other propositions $X \in 2^W$, their degree of possibility is $\Pi(X) = \sup_{w \in X} \pi(w)$. An epistemic interpretation given to Π yields a degree of possibility for a proposition, from an individual S 's perspective, ranging from 0, i.e. impossible, to 1, i.e. maximally possible.

Similar to imprecise probability, the credal model is bounded where Π is the upper bound or maximal degree of credence in a proposition X . Like upper probability $\bar{\mathbb{P}}$, Π defines a conjugate lower bound, $\mathcal{N} : 2^W \rightarrow \mathbb{R}$, the necessity measure. The necessity measure indicates, like its name suggests, how necessary or epistemically supported a proposition is where W is completely necessary or supported, i.e. 1, and \emptyset is completely unnecessary or impossible, i.e. 0. For all other propositions $X \in 2^W$, their degree of necessity is $\mathcal{N}(X) = \inf_{w \in X} \pi(w)$. Together with Π , a necessity measure provides a representation of imprecise belief when $\mathcal{N} \neq \Pi$.

A significant difference between possibility theory and probability theory resides in the definition for the possibility of the union of logically incompatible propositions. Specifically, $\Pi(X \cup Y) = \max\{\Pi(X), \Pi(Y)\}$ (*Maxitivity*). Notice that for a finite W , $\Pi(X) = \max_{w \in X} \pi(w)$ and $\Pi(Y) = \max_{w \in Y} \pi(w)$ for any

$X, Y \in 2^W$. For all finite index sets I , if $\Pi(X) \geq \Pi(Y_i)$ for all $i \in I$, then $\Pi(X) \geq \Pi(\cup_i Y_i)$. This is not hard to show. Suppose that $\Pi(X) \geq \Pi(Y_i)$ for all $i \in I$. Since $\Pi(\cup_i Y_i) = \max\{\Pi(Y_i)\}$ for all $i \in I$ and assuming there is no Y such that $\Pi(Y) > \Pi(X)$, then $\Pi(X) \geq \Pi(\cup_i Y_i)$. As a result of this, Π obeys an analogue to the Union Property. Relatedly, if $\Pi(X) \geq \Pi(\{w\})$, then it must be the case that there is a world $w' \in X$ such that $\pi(w') \geq \pi(w)$, otherwise $\Pi(X)$ would not be at least as great as $\Pi(\{w\})$. Thus, Π obeys an analogue to Determination by Singletons. As it turns out, relative likelihood reduces to possibility theory, at least when W is assumed to be finite.²⁸

In case one is interested in a unified theory bringing together fine-grained judgments and equivalent relative likelihood judgments, then possibility theory is a more suitable framework for modeling credence.²⁹ Furthermore, the close resemblance to imprecise probability allows for the same conceptual advantages of set-based credences to be had, namely the representation of increased uncertainty, at least with respect to disagreement. Possibility measures may also satisfy the proposed desiderata in the earlier sections if epistemic peers respond to a disagreement by adopting the min and max necessity and possibility measures, respectively, relative to a group of peers, \mathfrak{P} . Of course, there is much more work to be done, which would be quite interesting, but such work is beyond the scope of this discussion.

However, we might make a few naïve observations without going into too much detail. A possibility approach will satisfy the Reasonable Range principle. Possibility measures also satisfy the Minimal Risk of Regretting principle if the group's min and max are adopted and treated as betting rates for gambles. The PIE principle, however, is a little more controversial since independence is a complicated notion in possibility theory. Like imprecise probability, conditional independence is not necessarily symmetric, where symmetry is reserved for epistemic independence, which is not the same as stochastic independence. With some fancy footwork, though, like in the earlier discussion on parametrizing \mathbb{P} , the PIE principle may be said to be satisfied on the occasions when all members of a group initially agree that some propositions are independent and maintain the judgment throughout. Based on these naïve observations, it looks like possibility theory is a potential option for modeling a resolution to peer disagreements, too.

²⁸I leave it to the reader to uncover analogues for the other two constraints.

²⁹My own preference is toward set-based credences, at least insofar as peer disagreement is concerned, but for those who want to retain a qualitative image in addition to a fine-grained credal picture, relative likelihood and possibility theory is a viable approach, especially since the two frameworks can be unified, which leads to a comprehensive theory.

Chapter 4

Complete Ignorance in Imprecise Probability Theory

In a series of papers, John Norton (2007a, 2007b, 2008, 2010), a notable critic of the Bayesian program, has adamantly argued that the probability calculus is insufficient for modeling epistemic states, and as a result Bayesian epistemology turns out to be a failed attempt. One issue with a Bayesian theory in particular is that it is not fully encompassing of the possible epistemic states one might experience as it cannot capture the unique state of complete ignorance. According to Norton, this is because a probability measure generally fails to satisfy a desirable duality principle regarding invariance in belief and disbelief.

While an omission of ignorance from a Bayesian representation of an individual's epistemic attitudes poses a significant problem, those willing to deviate from the classical framework might insist instead that imprecise probability provides a better representation of epistemic states including states of ignorance (partial or complete).¹ Imprecise probability after all aims at modeling a wider range of belief states. However, Norton is unmoved by this alternative strategy of generalizing Bayes with sets of probability measures, for he claims there are many self-dual sets, but it remains unclear which set non-trivially captures the unique epistemic state.

Although I am sympathetic towards his contentions with classical Bayesian approaches, I beg to differ on the charge against imprecise probability. While it is not my intention to give a complete philosophical defense for a generalized Bayesian epistemology (see Joyce 2010), my aim in this chapter is to provide partial support for such a theory, at least in regard to having the ability to represent

¹See Benétreau-Dupin (2015) for a recent account.

complete ignorance. The strategy for addressing Norton's contentions begins with equating the unique epistemic state with a vacuous prior, i.e. $\{0, 1\}$. I then go on to systematically demonstrate that the representation trivially satisfies a central desideratum regarding duality, particularly through conjugacy relations held between lower and upper probability. However, these details alone are seemingly not enough to convince him, for Norton suggests that imprecise probability faces a problem of interpretation. After the discussion on the formal proposal, I show how one might interpret the representation, highlighting similarities with three-valued logics and Norton's envisaged logic of belief. I then provide what I find to be a more compelling interpretation of the $\{0, 1\}$ model through the theory of coherent lower previsions, which dispenses with belief and disbelief talk.

Of course, I am not the first to propose the $\{0, 1\}$ model as a representation of ignorance (see Walley 1991). Once the representation of belief is extended to sets of probability measures, it is quite natural for one to think about the state of complete ignorance in terms of vacuous priors. However, those familiar with the statistical and philosophical literature on imprecise probability know that the trivial model of uncertainty is not widely adopted, mainly because of a lurking problem associated with it. In particular, the adoption of vacuous priors exposes an individual to the problem of *belief inertia* (Walley 1991; Joyce 2010; Rinard 2013; Bradley 2014). To simply describe the problem, vacuous priors yield vacuous posteriors, and thus learning becomes impossible. If this challenge is left unresolved, then all of the effort put forth in representing complete ignorance with a vacuous prior is for nothing provided that the proposed account condemns an individual to an eternal state of complete ignorance.

To resolve the inductive learning problem, I propose an alternative updating method of *credal set replacement* towards the end of the chapter. But before revealing the magician's secret, let us first turn our attention to what started this whole investigation in the first place. In particular, let me motivate the challenge of representing complete ignorance probabilistically through a reconstruction of Norton's systematic arguments against a Bayesian epistemology, supplementing the arguments here and there.

4.1 Exiling Ignorance

Central to his complaint against Bayesian epistemology, Norton (2007a) claims that any formal theory of belief (and disbelief) should be self-dual, but a theory of

additive measures obeying Kolmogorov's axioms (1933/1956) is not self-dual.² In explicating the charge against a Bayesian theory of belief, we begin with the stipulation that an additive measure m representing belief has a dual additive measure M representing disbelief. Analogous to Boolean algebra, disbelief is the dual of belief similar to how False (0) is the dual of True (1). Unlike Boolean algebra, however, the additive measures are not exchangeable in a similar way the pair of Boolean operators \wedge and \vee and the pair of values 0 and 1 are exchangeable with one another, respectively, while still obeying the axioms of the algebra. In a theory of additive measures including the dual M , the dual measure maps back axioms that are entirely foreign to probability theory. To see this, let us first state some relational properties of these additive measures.

As Norton shows, there is a one-to-one correspondence between M and m with respect to a pair of contingent propositions A and $\neg A$,

$$m(A) \rightarrow M(A) = m(\neg A), \quad M(A) \rightarrow m(A) = M(\neg A). \quad (4.1)$$

Supposing that m is a probability measure, the axioms of finitely additive probability imply that $m(A) = 1 - m(\neg A)$ and $m(\neg A) = 1 - m(A)$, which induce dual measurements $M(A) = 1 - M(\neg A)$ and $M(\neg A) = 1 - M(A)$. The additive measure m , and subsequently the dual additive measure M , span the range from 0 to 1 with any high degree of belief in A yielded by m entailing a low degree of disbelief in A yielded by M . For example, if $m(A) = 1$, then $m(\neg A) = 0$. We derive from (4.1) the dual measurements, $M(A) = m(\neg A) = 0$ and $M(\neg A) = m(A) = 1$. In this particular instance, there is an *absence of disbelief* in A and *complete disbelief* in $\neg A$. The reciprocity of decreasing disbelief in A to increasing disbelief in $\neg A$ is a consequence of the additivity of the dual measure (2007a, 247).

The example just described suffices to show that the theory of additive measures is not *self-dual* for the reason that the axioms of finitely additive probability are not obeyed by the dual measure M , but instead, the dual measure obeys the following axioms:

(M1) $M(W) = 0$ for a finite set of worlds W ;

(M2) $M(A) \geq 0$ for all $A \in \mathcal{F}$ over W ;

(M3) $M(A \wedge B) = M(A) + M(B)$ if $A \wedge B = \perp$.³ (see 2007a, 234)

²For our purposes, we restrict ourselves to finite additivity.

³Like in the previous chapter, I will abuse notation by writing expressions like $A \wedge B$, which is

The unfamiliar axioms are sensible if M is interpreted as a measure of the *strength of disbelief* in propositions. Axiom M1 says that there should always be an absence of disbelief in logical truths and consequently there should always be full disbelief in logical falsehoods, i.e. $M(\perp) = 1$. Axiom M2 is straightforward and is equivalent to the non-negativity axiom of standard probability. Axiom M3 implies that disbelief does not decrease upon conjoining logically incompatible propositions.

By quick observation, we find that the peculiar axioms M1-M3 are jointly inconsistent with the finite probability axioms and thus the set of dual axioms lies beyond the axiomatic system of standard probability theory. As a consequence, a theory of additive measures representing belief and disbelief is not self-dual. The lack of self-duality poses a problem for Bayesian epistemology since a plausible representation of belief ought to respect interchangeability with its dual of disbelief, at least on Norton's view, just the same as True (1) respects its interchangeability with its dual False (0) in Boolean algebra.

Moreover, the Bayesian faces a further problem arising from a lack of self-duality, namely the inability to represent complete ignorance provided that measures m and M only yield degrees of belief and disbelief by their additivity, which subsequently sends ignorance into exile. Although not entirely explicit, Norton's claim is true if complete ignorance is reasonably assumed to be represented by a low value $i \in (0, 1)$. I say 'reasonably assumed' here since a mid or high value would conflict with the very notion of ignorance if m is an increasing function representing the *strength of belief*. At this juncture, though, it is hard to identify a value that i could take on for $i \approx 0$ implies nearly complete disbelief in a proposition. The only way to strike a balance between belief and disbelief in propositions A and $\neg A$ is by assigning each proposition a value $i = 1/2$, but again, i cannot be $1/2$ as such value indicates a high grade of belief in m 's range. We thus encounter a technical difficulty in defining i .

The problem is actually more troubling than it may seem given the formal assumptions held with respect to m and M . In particular, it turns out to be impossible to give a representation of complete ignorance relative to some contingent propositions A and $\neg A$ if an ignorance value $i < 1/2$ is uniformly assigned. We would end up with $m(A \vee \neg A) = 1 > m(A) + m(\neg A)$ and $M(A \wedge \neg A) = 1 < M(A) + M(\neg A)$, where the former statement is inconsis-

short hand for $\llbracket A \wedge B \rrbracket_M := \llbracket A \rrbracket_M \cap \llbracket B \rrbracket_M$ where M is a *probability structure*. As a reminder to the reader, a probability structure is a model $M = (W, \mathcal{F}, m, V)$ where V is an interpretation for a propositional language \mathcal{L} that associates all worlds w and primitive propositions X with a truth assignment such that for any world w and primitive proposition X , $V(w, X) = 1$ or $V(w, X) = 0$.

tent with the finite probability axioms and the latter statement is inconsistent with the mirroring axioms. In order for the measures to satisfy the respective set of axioms, some $i < 1/2$ may be assigned to A (or $\neg A$), but then so must $1 - i$ be assigned to $\neg A$ (or A), which entails that a more favorable degree of belief is given to $\neg A$ (or A) and a more favorable degree of disbelief is given to A (or $\neg A$). While the latter is necessarily required by the simple mathematical systems, the additivity property of both m and M precludes a representation of complete ignorance (or in the confirmational sense, neutral support).

Elsewhere, Norton (2008) has stipulated that a proper representation of complete ignorance should preserve what he calls *invariance under negation*—that is, there is no disproportional support (or belief) for a contingent proposition A over $\neg A$ or *vice versa*. But as we saw above, probability measures cannot satisfy the criterion if the numerical value representing ignorance is assumed to be a low value in $(0, 1)$. So the representation of complete ignorance with additive measures does not satisfy the following desideratum.

SELF-DUALITY FOR COMPLETE IGNORANCE (SD): An epistemic state of complete ignorance is invariant in its contingent propositions under the dual map given by [(4.1)]—that is, the epistemic state is self-dual in its contingent propositions, so that $m(A) = M(A) = m(\neg A)$ for all contingent A . (2007a, 247)

If (SD) is a required constraint on a formal representation of complete ignorance, the immediate thought that comes to mind is to abandon the idea that there is some low, fixed ignorance value i and instead adopt the *principle of indifference* (Keynes 1921), for in the simple case of contingent propositions A and $\neg A$, the principle entails equivalence: $m(A) = M(A) = m(\neg A)$. Intuitively, the principle of indifference is the seemingly natural way of representing a lack of belief by m , relative to contingent propositions A and $\neg A$. And so it appears that a Bayesian theory is able to satisfy (SD) after all upon supplementing with the principle of indifference.

Not so fast. Norton has stirred up trouble for this maneuver and suggests that the principle of indifference faces difficulty in providing a viable representation of the epistemic state sought. For one thing, he points out that the additivity of probability measures gives rise to the long-standing principle of indifference paradoxes (e.g. Keynes' countryman, Bertrand's paradox, von Mises' wine-water paradox). He insists that representing belief with a non-additive measure instead may resolve the paradoxes of indifference (see Norton (2008) for further details on his analysis

of the paradoxes, which I will leave to the reader).

A further issue with the principle of indifference is that a uniform prior fails to distinguish between disbelief and ignorance upon considering a fine-grained set of logically incompatible propositions with cardinality n and $n > 2$ (Norton 2010, 504). To see this, consider a set of propositions $\{A_1, A_2, \dots, A_n\}$ with each proposition A_i denoting a singleton set $\{w\}$ for all corresponding worlds w in a finite W . Additivity normalizes the marginal probabilities to unity: $\sum_{\{w\} \subset W} p(\{w\}) = p(W) = 1$. But additivity implies that the equal probabilities assigned to singleton sets $\{w_1\}, \{w_2\}, \dots, \{w_n\}$ approach 0 as n increases, which, by the argument from above, entails that the attitude toward each proposition is on the side of disbelief given the dual additive measure—that is, $M(A_i) = 1 - 1/n > 1/n = m(A_i)$ if $n > 2$. So the indifference device fails to distinguish between ignorance and disbelief.

Let me demonstrate the point through an example. In a simple case where an individual considers only contingent propositions A and $\neg A$, ‘ $\neg A$ ’ is typically regarded as the “catch-all” term, which is the set of worlds not in $\llbracket A \rrbracket$. Now, consider throwing a normal die. Only one of the six sides of the die will show face-up, and so we can reasonably say that there are six possible worlds. If we consider the proposition that an even number will come up on a random throw, call it E , assign to it probability $p(E) = p(\{w_2, w_4, w_6\}) = 1/2$. The opposite in this case is $\neg E = \{w_1, w_3, w_5\}$ and the proposition is also assigned probability $1/2$. It appears that we have a case of complete ignorance regarding the propositions E and $\neg E$ based on there being equal support for and against each proposition.

Suppose instead that we consider a different set of propositions $\{A, B, C\}$ relative to the same set of worlds such that $A = \{w_1, w_2\}$, $B = \{w_3, w_4\}$, and $C = \{w_5, w_6\}$. We abide by the principle of indifference and so each proposition is assigned probability $1/3$. With a focus on any single proposition in the set, the the dual measurement is $(1 - 1/3) > 1/2$. Specifically, $M(A) = 2/3$, $M(B) = 2/3$, and $M(C) = 2/3$. The difference between the two cases described is that the equal probabilities assigned by m in the latter case yield fairly high degrees of disbelief by the dual measure M , whereas in the former instance there is equal belief and disbelief provided that $m(E) = M(E) = m(\neg E) = 1/2$. The point to be made is that the principle of indifference is not always a principle of indifference, but sometimes it is a principle of (uniform) disbelief, at least in a theory that invokes an additive measure m and its dual additive measure M .

Since there can be any number of logically incompatible propositions or hypotheses (possibly infinite) to be considered when performing an experiment, the

number of propositions in a set under consideration may often exceed two. By entertaining a set of logically incompatible propositions whose number of elements is greater than two, the principle of indifference implies that $p(A) = 1/n$ for all A in the set and $1/n < 1/2$. It follows from a theory of belief and dual disbelief measures that each proposition is disbelieved to a specific degree rather than being epistemically neutral, which is illustrated through the dual measurements.

The problems discussed in this section, I think, pose a serious challenge for Bayesians in attempting to capture a state of complete ignorance. It is quite unfortunate that belief and dual disbelief measures banish ignorance from the kingdom of probability, but apparently, there is not much one can do given the additivity requirement, which is unable to be straightforwardly abandoned. An alternative strategy to circumvent the issue is to look beyond classical probability and turn to a more expressive framework for modeling belief, which we turn to next.

4.2 Vacuous Priors

Retreating from a Bayesian approach to representing complete ignorance, we consider a less-than-conventional method of imprecise probability here. The idea is to capture one's ignorance through a set of probabilities. But this attempt is also unsatisfying from Norton's perspective. Generally speaking, sets of probability measures violate (SD). However, there are some sets that satisfy the desideratum, and one in particular that I think is a viable candidate. In what is to follow, I propose a specific, though obvious, approach and demonstrate that imprecise probability provides a proper representation of complete ignorance after all. But before we get to the positive account, let us consider Norton's contention with the proposed method. Here is one passage illuminating a central concern.

Let the set of measures $\{m_i\}$, where i varies over some index set, be a candidate representation of complete ignorance. Under the dual map given by [(4.1)] this set is not mapped back to itself. Instead it is mapped to the corresponding set of additive dual measures $\{M_i\}$. That is, a set of additive measures fails to be self-dual, whether the set is convex or not. (2007a, 248)

From this passage, perhaps we should be skeptical as to whether imprecise probability can adequately represent ignorance given the violation of the duality principle. Let us make Norton's point vivid with a non-convex set of probability measures $\{p_1, p_2\}$ whose numerical values for some contingent proposition A

we suppose are $\{0.33, 0.55\}$. The formula for determining the dual measure from above is applied to each individual probability measure yielding a set of dual measures $\{M_1, M_2\}$ with values $\{0.67, 0.45\}$, i.e. $M_1 = 1 - p_1$ and $M_2 = 1 - p_2$. Notice in this example that the set of additive probability measures is not self-dual.

Despite the many sets of additive measures failing to be self-dual, implying that imprecise probability theory is not self-dual, we do get the following acknowledgment from Norton of there being sets that are indeed self-dual.

While sets of additive measures are not self-dual, we can readily define sets of measures that are self-dual. The simplest is just the set consisting of some additive measure m and its dual M , that is $\{m, M\}$...Clearly many such sets are possible. (2007a, 248)

Norton, however, is not convinced that a subclass of probability sets obtaining the property of self-duality is good enough for modeling ignorance with imprecise probability, mainly because self-duality alone does not determine the correct set (and dual set) that uniquely represents the epistemic state of complete ignorance.

However, it seems that he has overlooked what is seemingly the most plausible set capturing the unique state, namely a vacuous prior, $\{0, 1\}$, where the upper bound is the dual of the lower bound and the lower the dual of the upper. Although this suggestion may be intuitive to some, the model has very little appeal from Norton's perspective. One reason that will be touched on later is that he finds there to be a difficulty in interpreting sets of probabilities, but more on this contention in the coming sections. For the moment, I only wish to claim that the duality principle is satisfied by the $\{0, 1\}$ model.

How exactly does $\{0, 1\}$ as a representation of the epistemic state satisfy the (SD) desideratum? To illustrate, let us define a finite lower probability structure $(W, \mathcal{F}, \mathbb{P}, \underline{P}, V)$, in relation to a propositional language \mathcal{L} , where W is a finite set of worlds, \mathcal{F} is an algebra over W , \mathbb{P} is a non-empty set of probability measures (not necessarily convex) with each measure defined on \mathcal{F} , and \underline{P} is a special functional representing *lower probability* that is $\inf\{p(A) : p \in \mathbb{P}\}$. For propositions A and $\neg A$ that partition W such that neither proposition is equal to W nor \perp , suppose that the lower probability, relative to \mathbb{P} , for each proposition is the smallest admissible value, i.e. $\underline{P}(A) = 0$ and $\underline{P}(\neg A) = 0$. Accordingly, these values imply complete disbelief in each of the propositions based on the earlier discussion, but one need not to worry as we only have a partial picture at this point.

Uncovering the rest, we look to the conjugacy relations of lower and upper

probabilities, i.e. $\underline{\mathbb{P}}(A) = 1 - \overline{\mathbb{P}}(\neg A)$ and $\overline{\mathbb{P}}(A) = 1 - \underline{\mathbb{P}}(\neg A)$, where $\overline{\mathbb{P}}$ represents *upper probability*, which is $\sup\{p(A) : p \in \mathbb{P}\}$. Through conjugacy, we automatically obtain upper bounds, $\overline{\mathbb{P}}(A) = 1$ and $\overline{\mathbb{P}}(\neg A) = 1$, which these values imply complete belief in each of the propositions given the earlier discussion. The set of probabilities, \mathbb{P} , assigned to A is thus $\mathbb{P}(A) = \{\underline{\mathbb{P}}(A), \overline{\mathbb{P}}(A)\} = \{0, 1\}$.⁴ Let $\mathcal{M}(A)$ be the dual set of $\mathbb{P}(A)$, which, if decomposed, is the set $\{M(A), M(\neg A)\}$ containing each individual dual measure of each probability measure of $\mathbb{P}(A)$. Accordingly, $\mathcal{M}(A) = \mathbb{P}(\neg A) = \{\underline{\mathbb{P}}(\neg A), \overline{\mathbb{P}}(\neg A)\} = \{0, 1\}$. Thus, the values contained in $\mathbb{P}(A)$ are equivalent to those of $\mathcal{M}(A)$. We now have the full picture of an individual's state of complete ignorance regarding the contingent propositions A and $\neg A$, clearly indicating that the individual is entirely *un-opinionated* on the matter, which is consistent with the very idea of complete ignorance.

Provided the analysis so far, a lower probability $\underline{\mathbb{P}}$ and its conjugate upper probability $\overline{\mathbb{P}}$ are important features to the demonstration, and imprecise probability theory in general. The special functionals $\underline{\mathbb{P}}$ and $\overline{\mathbb{P}}$ are non-additive in the way that additivity is understood in classical probability—the property that Norton takes serious issue with. Rather, $\underline{\mathbb{P}}$ is super-additive and $\overline{\mathbb{P}}$ is sub-additive. This is made obvious upon summing the lower and upper probabilities for A and $\neg A$, yielding $\underline{\mathbb{P}}(A) + \underline{\mathbb{P}}(\neg A) = 0$ and $\overline{\mathbb{P}}(A) + \overline{\mathbb{P}}(\neg A) = 2$, whereas $\underline{\mathbb{P}}(A \vee \neg A) = 1$ and $\overline{\mathbb{P}}(A \vee \neg A) = 1$, for it is certain that one of the propositions is true. The upshot of non-additivity with respect to $\underline{\mathbb{P}}$ is that assigning probability zero to both A and $\neg A$ is possible. The rest of the epistemic picture is then easily uncovered. Specifically, if the infimum probability for both A and $\neg A$ is 0, then the supremum probability for each proposition is *necessarily* 1 as determined by the conjugacy relations.

To summarize the formal argument for the $\{0, 1\}$ ignorance model thus far, the first condition to be met for ensuring duality is that each contingent proposition in $\{A, \neg A\}$ obtains a lower probability 0. Then, we stipulate that \mathcal{M} is the dual set or mirror image of \mathbb{P} . Through the conjugacy relations of lower and upper probability for contingent propositions A and $\neg A$, \mathbb{P} satisfies (SD): $\mathbb{P}(A) = \mathcal{M}(A) = \mathbb{P}(\neg A) = \{0, 1\}$. That is to say that \mathbb{P} here is self-dual as it fulfills the invariance under negation condition. Without mandating convexity, it is actually $\underline{\mathbb{P}} = 0$ for both A and $\neg A$ that (trivially) satisfies (SD) given that the lower probability automatically defines a conjugate upper $\overline{\mathbb{P}} = 1$. Thus, the base, \mathbb{P} , is set

⁴The set is not convex. If it were, then for any probability $p \in [0, 1]$, $1 - p$ also lies in the linear span—that is, for each real-valued function in the convex hull, its dual is also in the space. It follows that the dual of every measure p is in the closed, convex set $[0, 1]$. I will not assume convexity as a necessary requirement, however, since it can lead to unintuitive consequences.

of probabilities equivalent to its dual set for propositions A and $\neg A$.

The point generalizes to n logically incompatible propositions where $2 < n < \infty$. Let $\Theta_n = \{A_1, A_2, \dots, A_n\}$ be a fine-grained partition of a finite set of worlds W . Suppose that the supremum probability measure for all propositions in Θ_n is equal to 1.⁵ We automatically obtain an infimum probability of 0 through conjugacy for each proposition in Θ_n . So, $\mathbb{P}(A_1) = \mathbb{P}(A_2) = \dots = \mathbb{P}(A_n) = \mathcal{M}(A_1) = \mathcal{M}(A_2) = \dots = \mathcal{M}(A_n) = \mathbb{P}(\neg A_1) = \mathbb{P}(\neg A_2) = \dots = \mathbb{P}(\neg A_n) = \{0, 1\}$. To show that this is the case, let there be a set of probability distributions $\mathcal{P}_n = \{p_1, p_2, \dots, p_n\}$ such that p_1 assigns probability 1 to $A_1 \in \Theta_n$, p_2 assigns probability 1 to $A_2 \in \Theta_n, \dots, p_n$ assigns probability 1 to $A_n \in \Theta_n$. Notice that $p_1(A_1) = p_2(A_2) = \dots = p_n(A_n) = 1$ implying that $M_1(A_1) = M_2(A_2) = \dots = M_n(A_n) = 0$. Next, observe that for $i, j, k = 1, 2, 3, \dots, n$ and epistemic state \mathbb{P} determined by $\underline{\mathbb{P}}$ and $\bar{\mathbb{P}}$, we have $\underline{\mathbb{P}}(A_i) = \inf \mathcal{P}_n(A_i) = p_j(A_i)$ and $\bar{\mathbb{P}}(A_i) = \sup \mathcal{P}_n(A_i) = p_k(A_i)$ and $j \neq k$, relative to the set of probability distributions \mathcal{P}_n . With no other distributions specified in the set \mathcal{P}_n , then $\mathbb{P}(A_i) = \{\underline{\mathbb{P}}(A_i), \bar{\mathbb{P}}(A_i)\} = \{0, 1\}$ for all $A_i \in \Theta_n$. Accordingly, the dual set $\mathcal{M}(A_i) = \{\underline{\mathcal{M}}(A_i), \bar{\mathcal{M}}(A_i)\} = \{0, 1\}$ for all $A_i \in \Theta_n$. Therefore, \mathbb{P} here satisfies (SD) with respect to Θ_n . In this case, it is actually $\bar{\mathbb{P}} = 1$ for all $A_i \in \Theta_n$ that (trivially) satisfies (SD) with respect to the set of n logically incompatible propositions under consideration.

From the discussion just had, we see that the imprecise probability approach maintains a distinction between ignorance and disbelief in addition to fulfilling the duality principle, unlike Bayesian approaches. So there are two conclusions we may draw: (i) imprecise probability has more expressive capabilities with respect to epistemology than Bayes and (ii) Norton's proposed desideratum is satisfied by the selected model. The latter does not seal the deal by Norton's lights, however, since he claims that the representation gives rise to a muddled interpretation.

4.2.1 Interpreting Vacuous Priors

Although a vacuous prior (trivially) satisfies the duality principle, the representation is subject to a problem of interpretation. In particular, Norton claims that a set of additive measures cannot totally represent an epistemic state once the dual

⁵For this demonstration, I go the opposite way starting with the upper probability instead of the lower probability. But since each defines the other through conjugacy, then it does not matter which one we initially specify. I do it this way in the general case, though, for it seems slightly easier, I think, to walk through the steps in proving the point.

measures are included. This is because probability measures and their duals behave differently where an additive measure m is a non-decreasing monotonic function while the dual measure M is a non-increasing monotonic function. The total representation, then, combines opposing epistemic notions, which becomes clear once one thinks about how m and M behave as belief and disbelief measures, respectively, over propositions and their logical consequences. As a reminder to the reader, the measures act in opposition to one another.

One way to sidestep the conceptual confusion pointed out is through devising some sort of classification scheme that separates the additive probability measures and the dual additive measures, similar to the way \mathbb{P} and \mathcal{M} were described. While the additive probability measures in the set \mathbb{P} are not individually self-dual, I have shown that the set $\mathbb{P} = \{0, 1\}$ is self-dual—that is, the set is invariant under negation with respect to contingent propositions. But even so, we still cannot escape the problem of interpretation provided that $M(\cdot) = p(\neg\cdot)$, and thus \mathcal{M} contains the dual of every $p \in \mathbb{P}$. There is no clear way of interpreting the sets of additive measures if they behave in opposing ways provided that one has a distinctive mark of belief while the other a distinctive mark of disbelief (2007a, 248). At this point, it seems tempting to give up on a probabilistic representation of belief altogether considering the trouble encountered. We might instead search for a formal framework that does not rely on combining belief and disbelief measures.

A plausible alternative is Norton's own envisaged non-additive schema for a logic of belief or degrees of support, which seemingly does the trick in accommodating complete ignorance (or neutral support in a confirmational sense).

$[T|B] = 1$, for all propositions T deductively entailed by B ;

$[A|B] = I$, for all contingent propositions A ;

$[F|B] = 0$, for all propositions F that logically contradict B . (2010, 505)

The values 1 and 0 yielded by the measure are meant to represent maximal and minimal support, respectively, and I a neutral or ignorance value. The conditional measure, $[\cdot|\cdot]$, presented above has its purpose in developing a non-probabilistic confirmation theory, but for our purposes, I will simply write the unconditional measure, $[\cdot]$, and omit B . This move, however, should not be problematic since we are primarily focused on complete ignorance as opposed to the evidential support for contingent propositions.⁶

⁶The conditional form $[A|B]$ is thought of as the degree that B confirms A with the exception

Moreover, in the instance that a proposition $A = \top$, then $[A] = 1$, and consequently $[\neg A] = 0$. Or, if $A = \perp$, then $[A] = 0$, and consequently $[\neg A] = 1$. On the assumption that an individual is completely ignorant toward A and $\neg A$, then we have the special case of $[A] = [\neg A] = I$. Given the latter, it should be easy to see now that $[\cdot]$ is self-dual with respect to ignorance where contingent propositions that are not determined to be better supported than their negations (or alternatives) are assigned the value I . In general, for a set of n logically incompatible propositions $\{A_1, A_2, \dots, A_n\}$, none of which are better supported than any other, $[A_1] = [A_2] = \dots = [A_n] = I$. So far, so good.

In addition, the proposed representation does not face the interpretational problem brought against sets of probability measures provided that each contingent proposition is just assigned the neutral or ignorance value I by $[\cdot]$, thereby avoiding combined belief and disbelief when both belief and dual disbelief measures are invoked in the logic. To see the difference, let $[A] = I$, which indicates that one is *categorically* ignorant towards A (and $\neg A$), while $\mathbb{P}(A) = \{0, 1\}$ indicates that one holds both complete belief and complete disbelief in A (and $\neg A$) given that each precise value is the dual of the other. Comparatively, the former proposal is actually quite elegant, and seemingly more so than the latter. For $[\cdot]$, as a type of epistemic valuation, tends to resemble a type of valuation familiar within non-classical logics for vagueness and indeterminacy. One may regard this new logic of belief, then, as sufficient for representing epistemic indeterminacy broadly construed.

To see the family resemblance with non-classical logics, let the propositions of an arbitrary propositional language \mathcal{L} either be true (1), false (0), or indeterminate (#) like in Łukasiewicz's (\mathbb{L}_3) and Kleene's (\mathbb{K}_3) three-valued logics (see Gottwald 2015). The logical connectives \neg , \wedge , and \vee are defined in \mathbb{L}_3 by the truth-tables in Table 4.1 (Note: the truth-tables for \neg , \wedge , and \vee are the same in \mathbb{K}_3 , but the systems differ on \rightarrow , which is omitted here). Now, let us introduce a Łukasiewicz valuation

that if $B = \perp$, then $[A|B]$ is undefined. Furthermore, Norton extends the model with a qualitative support relation \leq that is reflexive, antisymmetric, and transitive, yielding a partial order. For propositions A , B , C , and D , if $[A|B] \leq [C|D]$, then D confirms C at least as strongly as B confirms A , and if $[A|B] < [C|D]$, then D confirms C more strongly than B confirms A where $<$ is the strict partial order. Accordingly, the relation \leq is said to obey the following axioms:

$$\begin{aligned} &[\perp|W] \leq [A|B] \leq [W|W]; \\ &[\perp|W] < [W|W]; \\ &[A|A] = [W|W] \text{ and } [\perp|A] = [\perp|W]; \\ &[A|B] \leq [C|D] \text{ or } [C|D] \leq [A|B] \text{ (universal comparability);} \\ &[A|B] \leq [B|C] \text{ if } A \rightarrow B \text{ and } B \rightarrow C \text{ (monotonicity). (2008, 68)} \end{aligned}$$

(Note that for all contingent propositions, universal comparability entails *completeness*.)

\neg		\wedge	1	#	0	\vee	1	#	0
1	0	1	1	#	0	1	1	1	1
#	#	#	#	#	0	#	1	#	#
0	1	0	0	0	0	0	1	#	0

Table 4.1: Truth-tables for NOT, AND, and OR in \mathbb{L}_3

on \mathcal{L} . For all atomic propositions $A \in \mathcal{L}$,

$$\mathbb{L}V(A) = \begin{cases} 1 & \text{if } A \text{ is true;} \\ \# & \text{if } A \text{ is indeterminate;} \\ 0 & \text{otherwise.} \end{cases}$$

Upon considering an atomic proposition A of \mathcal{L} , suppose that $\mathbb{L}V(A) = \#$. Then, $\mathbb{L}V(\neg A) = \#$ according to the truth-table for negation. What this means is that A is neither true nor false, and likewise for $\neg A$, or that the truth-values for A and $\neg A$ are yet to be determined. Such valuation is seemingly accurate in some situations, particularly when a sentence is semantically indeterminate. Notably, \mathbb{L}_3 has in the past been applied to semantic vagueness for borderline cases where a proposition is seemingly neither determinately true nor determinately false. Now, if we were to relate the two frameworks, then notice an obvious parallel here to propositions that are indistinguishable with respect to their levels of support. For example, it is classically undetermined whether the sentence ‘John is tall’ is true or false (absent of some contextual factor). An individual, then, is not permitted to believe or disbelieve the proposition expressed by the sentence given that neither the proposition nor its negation is better supported. Observe that similar valuations are given in this particular case within each theory described. The motivation for each approach seems to be quite clear from this example.

In comparing the formal systems, there is a striking resemblance structurally between Norton’s $[\cdot]$ and $\mathbb{L}V$ for atomic propositions of a propositional language \mathcal{L} closed under \neg . In particular, for all atoms $A \in \mathcal{L}$, both functions are similar in their formal semantics with the only difference lying in the interpretations where $[\cdot]$ is epistemic and $\mathbb{L}V$ is truth-oriented. Although it was not described above, Norton’s logic of belief also extends to \wedge and \vee (see 2008, 67-69), with very similar structures to that of \mathbb{L}_3 , which makes his proposed framework equally expressive. (I leave it to the reader to check.) As it turns out, Norton’s proposed logic of belief that accommodates complete ignorance is a three-valued logic assigning 1 to all truths, 0 to all falsehoods, and I to all neutrally supported contingent propositions

whose classical truth values are unknown. While simple and elegant, the logic, however, comes at a price that is especially expensive in comparison to the price of the alternative imprecise probability model.

One cost associated with a three-valued logic in epistemology is a loss of generality. Although there is a seamless connection between $[\cdot]$ and a traditional interpretation of belief where belief is assigned 1, disbelief 0, and suspended judgment I , the *gradability* of belief, disbelief, and suspension of judgment is removed from the picture, at least in the special case outlined. While categorical belief plays an important role in everyday reasoning, so too does partial belief, especially in situations involving calculative reasoning. In order to maintain partial belief, a model more general than the three-valued logic is needed. Imprecise probability is such a model, which allows for the gradability of judgments. An imprecise probability model also preserves a type of categorical model, for the three-valued approach is a limiting case. Specifically, all propositions that are deductively true have a lower probability 1, all propositions that are necessarily false have an upper probability 0, and all contingent propositions that are neutrally supported have a value of $\{0, 1\}$, which replaces I . Given the choice between the three-valued logic of belief and imprecise probability, the latter should be preferred on the grounds of expressiveness. But Norton might point to the partial order \leq as an extension (see footnote 6), which differentiates the strength of belief.

Even so, there is a further cost with the simplistic model, and one that is common among three-valued logics. It is identified in the truth-tables above. In particular, \mathbb{L}_3 generates absurd consequences. For one, suppose that $\mathbb{L}V(A) = \#$ for some contingent A . The proposition $A \wedge \neg A$ is assigned the value $\#$ according to the truth-tables for negation and conjunction. Whether the formal logic is meant to supply laws of correct inference or epistemic norms of belief, on neither view does it seem correct to say that $A \wedge \neg A$ is indeterminate. One may be ignorant of which proposition is true or at least better supported, but surely one knows *a priori* that $A \wedge \neg A = \perp$. A similar absurdity arises when \wedge is replaced by \vee . Holding the same values fixed for contingent A and $\neg A$, the truth-table for disjunction tells us that $A \vee \neg A = \#$. Again, one might be ignorant as to whether A or $\neg A$, but surely one knows *a priori* that $A \vee \neg A = \top$. These absurdities arise from the logic itself, for we have just seen that the laws of excluded middle and non-contradiction are neither guaranteed by $\mathbb{L}V(\cdot)$ nor by $[\cdot]$, laws that most, I think, find desirable. It is only when all atomic propositions in the language have classical truth-values that the laws are upheld. Classicality in the epistemic sense, however, fails to ac-

commodate ignorance, so a three-valued system is needed. But then a three-valued logic for belief gives rise to the above absurdities.

While it is reasonable to think that Norton would attempt to escape these absurdities in describing a logic for belief, he instead doubles down on giving up maximal belief or support for the logical truth $A \vee \neg A$. Specifically, if $[A] = [\neg A] = I$, then $[A \vee \neg A] = I$ on his approach. Norton suggest that the representation is feasible since it avoids what he calls the *inductive disjunction fallacy*. In short, the inductive disjunction fallacy is where a (large) disjunction of neutrally supported propositions yields a strongly believed or supported proposition (2010, 309). We can easily gather that his target is the probability calculus. To see this, let $\Theta_n = \{A_1, A_2, \dots, A_n\}$ be a set of logically incompatible propositions each having equal belief or support $1/n$. Using a probability measure p to represent belief, disjunctions get more support: $p(A_1) = 1/n < p(A_1 \vee A_2) < p(A_1 \vee A_2 \vee A_3) < p(A_1 \vee A_2 \vee \dots \vee A_n) = 1$. Additivity is the culprit once again, and it leads to the so-called inductive disjunction fallacy.

It is not clear, however, why Norton considers an increase in belief or an accumulation of support for disjunctions of logically incompatible propositions to be fallacious. In the simple case of A and $\neg A$, an individual knows *a priori* that $w \in A \vee w \in \neg A$, where w is the actual world, or in the general case that $w \in A_1 \vee w \in A_2 \vee \dots \vee w \in A_n$. So it seems obvious that the proposition $A \vee \neg A$ (or $A_1 \vee A_2 \vee \dots \vee A_n$) should be more strongly believed (or supported) by an individual than any single contingent proposition, and thus the disjunction fallacy does not seem to be a fallacy at all.

Things continue only to get worse for Norton's three-valued approach when assigning contingent propositions A and $\neg A$ a value I and consequently assigning $A \vee \neg A$ a value I as this is inconsistent in the proposed framework. Propositions that are known to be logical truths should be assigned maximum value 1, i.e. $[T] = 1$, on Norton's account, but it turns out that there are logical truths that end up being assigned an alternative value, i.e. I . So the system itself produces inconsistency. Despite the attractiveness and simplicity of the proposed logic together with the measure $[\cdot]$ stably maintaining self-duality, the endorsement of the inductive disjunction fallacy over the preservation of logical truths and rejection of logical falsehoods costs the approach reasonably desirable properties, and ultimately leads to an inconsistent system, while the purported fallacy is not a very convincing reason for throwing out logical truths to begin with. Therefore, Norton's three-valued approach to representing complete ignorance is implausible.

Imprecise probability, on the other hand, avoids the described absurdities provided that the conjunction of contingent propositions A and $\neg A$ is contradictory (or the intersection is empty). Therefore, $\mathbb{P}(A \wedge \neg A) = \{p(A \wedge \neg A)\} = \bar{\mathbb{P}}(A \wedge \neg A) = 0$. (The dual set is opposite $\mathcal{M}(A \wedge \neg A) = \{M(A \wedge \neg A)\} = \underline{\mathbb{M}}(A \wedge \neg A) = 1$, indicating complete disbelief.) We also find that $\mathbb{P}(A \vee \neg A) = \{p(A \vee \neg A)\} = \underline{\mathbb{P}}(A \vee \neg A) = 1$. (The dual set is opposite $\mathcal{M}(A \vee \neg A) = \{M(A \vee \neg A)\} = \bar{\mathbb{M}}(A \vee \neg A) = 0$, indicating complete lack of disbelief.) This ought to be clear, for the mathematical statements are implied by the axioms of finitely additive probability. Specifically, $A \wedge \neg A = \perp$ and $A \vee \neg A = W$ where W is finite. Since normalization requires that $p(W) \neq 1$ and that $p(W) \neq 1$ for any p , then $\underline{\mathbb{P}}(W) = 1$. Consequently, $p(A \vee \neg A) = \underline{\mathbb{P}}(A \vee \neg A) = 1$. Moreover, $p(\perp) = 1 - p(W)$, and so $p(A \wedge \neg A) = \bar{\mathbb{P}}(A \wedge \neg A) = 1 - p(W) = 0$. Thus, we have shown that the absurdity of being less than certain in logical truths is rendered unacceptable in this framework. It seems to me, then, that imprecise probability is still our best representation, at least with respect to a viable, non-classical system.

Finally, I would like to return to an earlier point about the categorical approach to representing belief as a special case in imprecise probability. As a reminder, $[\cdot] = I$ (and $\mathbb{L}V(\cdot) = \#$) is somewhat analogous to $\{0, 1\}$ in imprecise probability, at least insofar as our interest in a formal representation of ignorance is concerned. Actually, Norton's basic schema is closely related to a special case of imprecise probability such that for all proposition A , either $\mathbb{P}(A) = \{1\}$, $\mathbb{P}(A) = \{0\}$, or $\mathbb{P}(A) = \{0, 1\}$, where $\{0, 1\}$ replaces I . Imprecise probability generally, not just the special case, conforms to almost all of the axioms of the system in footnote 6 (in an unconditional form): for all propositions A, B , $\mathbb{P}(\perp) \leq \mathbb{P}(A) \leq \mathbb{P}(W)$, $\mathbb{P}(\perp) < \mathbb{P}(W)$, $\mathbb{P}(A) \leq \mathbb{P}(B)$ if $A \rightarrow B$. The only axiom that is not necessarily satisfied, however, is universal comparability, i.e. $\mathbb{P}(A) \leq \mathbb{P}(B)$ or $\mathbb{P}(B) \leq \mathbb{P}(A)$, namely because there are instances where imprecise probabilities for contingent propositions are incomparable, and so \leq is not necessarily total. But it is unclear why universal comparability is proposed as a necessary requirement anyway since Norton states that \leq partial. Nevertheless, by replacing I with the self-dual set $\{0, 1\}$, we end up with a logic of belief that conforms to an axiomatic system similar to Norton's, but imprecise probability does not throw out logical truths.

While imprecise probability models are quite powerful, the problem of interpretation introduced at the beginning of this section still lingers as I have yet to provide a feasible answer, and it is not one that can simply be shrugged off. Those who are familiar with the philosophical literature on imprecise probability should

not be surprised that Norton has taken issue given the common parlance. Starting with van Fraassen (1990), sets of probability measures compatible with one's opinion or judgment received the name 'representor'. More recently, Joyce (2010) has continued with the metaphor by dubbing an imprecise probability as an individual's 'credal committee' that personifies probability measures as individual committee members holding definite opinions on a proposition. And in decision situations, Moss (2015) uses the term 'mental committee' referring to a set of precise opinions about how one ought to act. When you hear these kind of metaphors, it is hard not to immediately strip imprecise probabilities down to their parts—that is, decompose them into their individual measures.

The "reductive" interpretation of imprecise probabilities opens the door to the conceptual concern, I think, namely because the attention is turned toward classifying each additive measure either as a belief or disbelief measure. However, decomposing the representation into individual probability measures and dual measures *à la* mental or credal committees is not the right way to think about imprecise probabilities, at least not in the case of ignorance. Possibly a better way to understand the representation is by taking the set itself, not the measures of which it is composed, as a representation of *suspended judgment* on a proposition (see Sturgeon 2008; Haenni et al. 2011). In case suspension of judgment is maximal, an individual has a vacuous prior, otherwise they suspend judgment to a degree—that is, $\underline{P}, \bar{P} \in (0, 1)$. The ability to represent grades of suspended judgment is quite fortunate, for such capability is not one available to alternative three-valued approaches or any traditional conception of belief for that matter.

Furthermore, the suspension interpretation is consistent with having both belief and dual disbelief measures in a set of measures since the notion of suspended judgment after all implies an equal balance between belief and disbelief—that is, the scale is not tipped one way or the other. But this explanation is still on the reductive track. One should realize, however, that the sets themselves represent generic epistemic attitudes in contingent propositions A and $\neg A$. In case of complete ignorance, the support for A is absolutely minimal, likewise for $\neg A$, but alongside the lack of evidence against A is absolutely maximal, likewise for $\neg A$, which implies that there is no favoring of A or $\neg A$, and hence it is a state of complete ignorance.

By putting the focus on the set itself rather than individual measures, we "blackbox" the representation (Good 1962), thereby closing off the view except for the outer dimensions that provide a general indication of the support and lack of evidence against. Intuitively, complete ignorance is just like a blackbox, and

so the metaphor is rather fitting. Furthermore, the latter reading of vacuous priors is more compelling, I think, than a metaphorical committee since a conflicted committee betters accords with conflicting evidence rather than no evidence at all. Although some may remain unconvinced by the given interpretation, I turn now to a different way of thinking about the representation, one that emphasizes the importance of the dimensions of the blackbox rather than its contents.

4.2.2 A Behavioral Interpretation

Following de Finetti (1974), Williams (1975), and Walley (1991), let us move forward with the language of *previsions*. I have already introduced the idea earlier in Chapter 2, and briefly touched on it again in Chapter 3. Recall that the type of prevision or fair price we considered before is the amount of money that an individual is willing to spend or accept for a special kind of gamble $I_A(w)$ with an uncertain reward where I is an indicator of $A \subseteq W$ such that

$$I_A(w) = \begin{cases} 1 & \text{if } w \in A; \\ 0 & \text{otherwise.} \end{cases}$$

In earlier discussions, a probability was considered to be an individual's prevision for a gamble $I_A(w)$, which depended on the truth of a proposition A . The price one gives effectively represents how likely the individual considers the proposition in question to be true.

Up until this point, however, we have started with an epistemic notion of probability and moved in the direction of *linear* previsions, followed by lower and upper previsions, essentially giving precise probabilities, along with lower and upper probabilities, behavioral interpretations. Starting with a theory of previsions instead, *lower previsions* in particular, is simpler for the reason that a lower probability model is subsumed by it. The purpose of introducing a more expressive language in this subsection is to (a) emphasize that rationality is grounded in something that each one of us is familiar with, namely acting under uncertainty, and (b) subsequently making clear what a rational individual's behavioral commitments are in a state of ignorance formally represented by a vacuous prior, $\{0, 1\}$, which, in short, is *abstention*. Ultimately, I aim at providing an intuitive interpretation of $\{0, 1\}$, or at least extending its interpretation from the epistemic to the practical.

In adopting a theory of lower previsions as described by Miranda & de Cooman (2014, 28-34), we begin by defining a gamble f , which is a function from a finite

set of worlds W to the reals \mathbb{R} , i.e. $f : W \rightarrow \mathbb{R}$. The gamble f , with value $f(w)$ when world $w \in W$ is actual, represents an uncertain reward expressed in units of a linear utility scale, and serves as the underlying asset involved in two types of transactions. First, one may decide to buy f for a price x , indicating that the individual finds the uncertain reward $(f - x)$ desirable. Second, one may decide to sell f for a price y , indicating that the individual finds the uncertain reward $(y - f)$ desirable. An individual's lower prevision $\underline{P}(f) = \sup\{x \in \mathbb{R} : f - x \in \mathcal{D}\}$ ⁷ is the supremum buying price for f , implying that they also find acceptable $\underline{P}(f) - \epsilon$ for all $\epsilon > 0$. On the other hand, an individual's upper prevision $\overline{P}(f) = \inf\{y \in \mathbb{R} : y - f \in \mathcal{D}\}$ is the infimum selling price for f , implying that they also find acceptable $\overline{P}(f) + \epsilon$ for all $\epsilon > 0$. In simple terms, an individual finds a gamble f acceptable to buy for any price up to $\underline{P}(f)$, while the individual also finds the gamble f acceptable to sell for any price as low as $\overline{P}(f)$.

Equivalently, selling f for a price y amounts to buying $-f$ for a price $-y$, and therefore if an individual is willing to accept one of the transactions, then they should also be willing to accept the other under the same conditions. This fact establishes conjugacy relations between lower and upper previsions such that $\overline{P}(-f) = -\underline{P}(f)$. It is fortunate that we are able to define lower and upper previsions in terms of one another, for we only need to primitively specify one of the functionals in a model, which a lower prevision is often chosen, hence the theory of lower previsions. Much of this should ring a bell from previous discussions, but notice that the language introduced is quite general where probabilities are not the only things we can model within the framework. For instance, \underline{P} and \overline{P} may represent the bid and ask prices, respectively, for a contract in a futures or options market. Considering the latter example, notice the general target a prevision aims at modeling, namely a behavioral disposition.

Moreover, now that we have a handle on lower and upper previsions, let us turn to a coherence property in order to identify (minimally) rational behavior involving the two types of transactions. To begin, suppose that an individual has a lower prevision on a subset \mathcal{X} of the set of all gambles $\mathcal{L}(W)$ on a finite set of worlds W —that is, the individual has a lower prevision $\underline{P} : \mathcal{X} \rightarrow \mathbb{R}$. Provided an obvious desire to avoid sure losses in either of the two types of transactions described, one's lower previsions for a subset of gambles $\mathcal{X} \subseteq \mathcal{L}(W)$ should never result in the loss of utiles. Formally,

⁷ \mathcal{D} is the set of gambles found desirable by an individual.

$$\sup_{w \in W} \sum_{i=1}^n [f_i(w) - \underline{P}(f_i)] \geq 0 \text{ for all natural } n \geq 0 \text{ and all } f_1, \dots, f_n \in \mathcal{X}. \quad (4.2)$$

If the condition is not satisfied, then there are gambles $f_1, \dots, f_n \in \mathcal{X}$ and $\epsilon > 0$ such that $\sum_{i=1}^n (f_i - (\underline{P}(f_i) - \epsilon)) \leq -\epsilon$. This means that the sum of desirable transactions leads to a loss of at least ϵ no matter the outcome. This is assuming that a positive linear combination of acceptable transactions is considered acceptable.

More strongly, though, rational lower previsions are constrained by *coherence*, whereby an individual's supremum acceptable buying price for f is not raised upon considering a positive linear combination of other acceptable (finite number of) gambles. For all $n, m \geq 0$ and $f_0, f_1, \dots, f_n \in \mathcal{X}$,

$$\sup_{w \in W} \sum_{i=1}^n [f_i(w) - \underline{P}(f_i)] - m[f_0(w) - \underline{P}(f_0)] \geq 0. \quad (4.3)$$

If $m = 0$, then it follows that a lower prevision avoids a sure loss—that is, (4.3) reduces to (4.2) when $m = 0$.

If the domain of a lower prevision \underline{P} is a linear space, i.e. \mathcal{X} is closed under linear combinations, then the coherence property may be described in another way. In particular, \underline{P} is coherent upon satisfying:

- $\underline{P}(f) \geq \inf f$ for all $f \in \mathcal{X}$ (accepting sure gains);
- $\underline{P}(f + g) \geq \underline{P}(f) + \underline{P}(g)$ for all $f, g \in \mathcal{X}$ (super-linearity);
- $\underline{P}(\lambda f) = \lambda \underline{P}(f)$ for all $f \in \mathcal{X}$ and real $\lambda > 0$ (positive homogeneity).

As Miranda and de Cooman state, a lower prevision on an arbitrary domain is coherent only if it is extended to a lower prevision on a linear space, and so a coherent lower prevision will consequently satisfy the above conditions.

Now, consider a special case where the domain \mathcal{X} of a lower prevision is a set of indicators of propositions, relative to some structure (W, \mathcal{F}) , each being bounded by 0 and 1. A lower prevision \underline{P} on \mathcal{X} is a lower probability. In particular, $\underline{P}(I_A)$ is an individual's supremum buying price for a gamble I on a proposition A that pays \$1 if $w \in A$ and \$0 otherwise. Alongside, a lower prevision $\underline{P}(I_A)$ defines a conjugate upper prevision $-\overline{P}(-I_A)$ or the infimum selling price for the gamble $-I_A$. And lastly, the lower prevision \underline{P} obeys the above conditions and therefore

avoids sure losses. Observe that we have started with a theory of lower previsions and deduced lower probabilities rather than going in the other direction using some expectation operator (Miranda 2008), which makes for a smooth transition on the technical front. More important to the story, coherent lower previsions is a theory of rational behavior under uncertainty. Thus, the special case of lower probability is a model of an individual's behavioral dispositions for specific transactions. The common parlance of epistemology is absent from the theory, and so the earlier confusion about what the imprecise probability model represents is eliminated.

Although we need not rely on the formal notation for lower previsions since the thing that we care most about is the pragmatic interpretation of lower probabilities, a detailed discussion of the theory of lower previsions helps make clear what lower probabilities actually are, at least according to those inclined toward a subjectivist view. So we can stay the course with the original notation for imprecise probability now that we are provided with a better understanding of what the mathematical model is supposed to represent. Fixing our attention on the pragmatic interpretation of imprecise probabilities results in a loss of interest in the individually precise measures of a set of probability measures, which were of prior concern. Our interest is now shifted to an individual's *determinate* behavioral dispositions and indecision with respect to accepting gambles, which may be said to reflect their epistemic state. As a consequence, the (self-dual) sets of probabilities are no longer in focus. Let us consider an example to demonstrate the point.

Suppose that a cutting machine at a coin making factory is not well-calibrated. It tends to cut heads and tails evenly sometimes, producing a fair coin, but it also sporadically shaves too much off the side to be heads on some coins and tails on others. The unbalanced coins have at most a 10% bias towards heads or tails. We enter the factory just as the machine has got done cutting, and a mixed lot is put into a box. I randomly select a coin from the box. Before, I would have asked you what your credence is that the coin will land heads up if I toss it. Your answer should be $\mathbb{P}(H) = \{0.4, 0.6\}$. If rational belief is determined by evidential support, then this is the correct belief model, for your evidence supports a credence in H up to 0.4, but the evidence against H is no less than a credence of 0.6. If evidentialism (Conee & Feldman 2004) is true, then it suggests imprecise probability in this case.

But we have given up the purely epistemic talk in this section in favor of the pragmatic interpretation of the lower probability model. To better understand your epistemic state, I ought to ask you a different question. If I were to offer you a gamble on H prior to the toss that will pay \$1 if heads lands face up and \$0

otherwise, what is the maximum amount of money you are willing to pay for the gamble? Assuming that rational decisions are influenced by one's beliefs, whatever they may be, then your supremum buying price should not be any more than \$.40 for the gamble. Let us go to the other side now. If I ask you to sell the gamble to me, what is the minimum amount of money that you are willing to accept? From a recollection of information provided to you about the coin, your infimum selling price should not be any less than \$.60.

Let us alter the example to accommodate a state of complete ignorance. Suppose we walk into the coin factory and the floor man tells us that the coin cutting machine is exhibiting extremely erratic behavior. Because there is no accurate estimation of potential bias due to the machine's behavior, the factory will have to throw out the batch. I quickly snatch a random coin out of the box before the batch is taken to the dump. I ask you how much you would pay for the opportunity to win \$1 if the coin lands head and \$0 otherwise? Since you have no idea what the actual chance of the coin landing heads is, you should not pay anything for the gamble. Otherwise, you put yourself at risk of paying too much for it. More precisely, you cannot rule out the possibility that the coin is completely biased toward tails based on what you have learned, and so any positive amount exchanged for the gamble may result in an unnecessary loss for you.

If, instead, I wanted you to sell me the gamble, what is the smallest amount you would let it go for? Without further information, you should not let it go for less than \$1, which just amounts to an even swap whatever the outcome is. Why be so risk-averse in this instance? Because you could easily sell the gamble for too small of a price. More precisely, you cannot rule out the possibility that the coin is completely biased toward heads based on what you have learned, and so any amount less than \$1 exchanged for the gamble may result in a loss for you. If this story is consistent with your inclinations, then your behavior is indicative of your epistemic state, which in a lower probability model implies $\mathbb{P}(H) = \{0, 1\}$ and $\mathbb{P}(T) = \{0, 1\}$. There is nothing unclear nor unintuitive about one refraining from taking any action as a result of being in a state of ignorance. The behavioral theory described in this section suggests just that, and so we end up with a compelling interpretation of the $\{0, 1\}$ complete ignorance model without reference to beliefs, representors, or credal committees.

4.3 Belief Inertia

In response to Norton, Benétreau-Dupin (2015) has recently developed a similar defense of imprecise probabilities for representing ignorance. What divides Benétreau-Dupin and I on the matter, however, is the permissible sets of probability measures we consider for representing ignorance. On his view, he excludes the closed, convex set $[0, 1]$ (and ultimately the subset $\{0, 1\} \subset [0, 1]$) in order to avoid the inductive learning problem or *belief inertia* (Rinard 2013). He says,

There is however a good reason not to be content with such an extreme representation of ignorance. Indeed, in that set \mathfrak{J} of all possible probability distributions will be extremely sharp probability distributions that require an unreasonably large—or even infinite—number of updatings before they can yield posteriors distributions that are significantly different...Such distributions in \mathfrak{J} are said to be dogmatic, and consequently the whole set \mathfrak{J} is dogmatic. A representation of complete ignorance \mathfrak{J} , and generally any vacuous prior, entails a vacuous posterior. This should prevent such a set from being used in an inferential process in which we may hope to move away from a state of ignorance after a certain number of iterations of Bayesian updating. This representation of ignorance by means of a family of credal functions, although it satisfies Nortons criteria for ignorance, is incompatible with learning. (2015, 1534)

On the view I have proposed, however, $\{0, 1\}$ is said to be the proper representation of ignorance, and I would go as far as saying that the vacuous prior is the only intuitive representation, which should be clear from the discussion on how to bet when in a state of complete ignorance. Although convexity is not mandated on the account I have laid out, $\{0, 1\}$ is still a dogmatic prior. As Rinard (2013, 4) writes:

On the set of [probability] functions model, updating proceeds by individually conditionalizing each function in your representor on your new evidence. Each function in your representor will have its posterior probability for H1 determined by Bayes rule as follows: $\Pr(H1|E) = \Pr(E|H1)\Pr(H1) / [\Pr(E|H1)\Pr(H1) + \Pr(E|H2)\Pr(H2)]$. In this case, all the functions in your representor agree on the likelihoods, as they are fixed by objective chances in accordance with the Principal Principle: $\Pr(E|H1) = 1$ and $\Pr(E|H2) = 1/10$. Substituting 1 –

$P(H1)$ for $P(H2)$ and simplifying yields $\Pr(H1|E)=\Pr(H1)/[1/10+9/10\Pr(H1)]$.
 $\Pr(H1|E)=1$ when $\Pr(H1) = 1$, and $\Pr(H1|E) = 0$ when $\Pr(H1) = 0$.

Rinard has thus shown that a vacuous prior is dogmatic and an unsuitable model for ignorance. Or maybe it is the case that canonical Bayesian updating has a problem? As Benétreau-Dupin and Rinard seem to suggest, the vacuous prior has to go, not conditionalization.

However, we need not buy into the false dilemma, and instead we might look for a different solution in which an alternative belief updating rule is made available just in case an individual is in a state of complete ignorance while still maintaining conditionalization for inductive inference when priors are imprecise yet non-extreme. A proposal has been systematically detailed by Nic Wilson (2001), which is driven by a notion of *implausibility*. The idea is intuitive: reject all (precise) probability distributions that are implausible in light of new information. I might offer a simple updating procedure here motivated by the same idea.

Let us consider a simple set of propositions, $\Theta = \{A, \neg A\}$, relative to a set of worlds W . Suppose an individual is completely ignorant as to whether A or $\neg A$. So the individual adopts vacuous priors, $\mathbb{P}(A) = \mathbb{P}(\neg A) = \{0, 1\}$. Then, they learn some new evidence E and no longer are completely ignorant toward A and $\neg A$. To escape what is seemingly an eternal state of ignorance, the individual successfully updates their beliefs through *credal set replacement*:

$$\mathbb{P}(A \parallel E) = (\mathbb{P} \setminus \{p_i\}) \cup \{p_j\} \text{ such that } p_j \in (0, 1), \text{ for all } i \text{ and } j. \quad (4.4)$$

In plain terms, the updated set of probability measures is the union of the *remaining set* $(\mathbb{P} \setminus \{p_i\})$ and a set of plausible probability measures $\{p_j\}$ such that no p_j is extreme.⁸ The set $\{p_i\}$ contains all probability measures in $\mathbb{P}(A)$ that are rendered *implausible* by the recently learned information. As a result, the set $\mathbb{P}(A \parallel E)$ contains the remaining probability measures not rendered *implausible* along with any new probability measure(s) that the information makes plausible. If there are no new additions, then the joining set is empty and replacement just becomes *reduction*. In the case of complete ignorance, however, the original set $\mathbb{P}(A) = \{0, 1\}$ is effectively replaced by a non-empty $\mathbb{P}(A \parallel E) = \{p_j\}$. If $\{p_j\} = \{0, 1\}$, then $\mathbb{P}(A \parallel E) = \mathbb{P}(A)$, otherwise, $\mathbb{P}(A \parallel E) \neq \mathbb{P}(A)$, which the latter suggests that

⁸The idea for the belief update method is inspired by the AGM belief revision function (*) (Alchourrón, Gärdenfors, & Makinson 1985). In the future, it would be fruitful to explore a stronger connection, but I leave that for another day.

the information learned does indeed bear evidential relevance to the propositions under consideration. (I am being cautious here since credal set replacement can preserve independence as we will see later.)

How is it that one can non-arbitrarily justify replacing $\{0, 1\}$? This will depend on the information that an individual receives. Initially, the individual has no clue whether A or $\neg A$ is true, hence the vacuous priors $\{0, 1\}$. But upon acquiring E , if they learn that neither A nor $\neg A$ are logically implied by the evidence E , then any probability p warranted by the evidence is non-extreme such that $p(A) < 1$ and $p(\neg A) < 1$ provided that the evidence reduces the degree of ignorance, but they now are aware that neither A nor $\neg A$ are deductively true given E . As a result, probability 1 assigned to both A and $\neg A$ ought to be considered implausible, at least at the current time, and removed from one's set of probabilities. The conjugate lower probabilities, i.e. 0, also get removed, so the set $(\mathbb{P}(A) \setminus \{p_i\})$ is empty. The union of the empty set with a non-empty set of probability measures $\{p_j\}$ given the evidence E is just the set of plausible probability measures given the evidence E .

But how do we determine what the set of plausible probability measures is relative to E ? There is no precise rule that I am aware of for determining such set at this moment, but I might say that the set can sometimes be determined fairly easily. With that said, the evidence itself might uniquely determine the set of plausible probability measures. Consider, for example, a coin toss. Let $W = \{w_1, w_2\}$, $H = \{w_1\}$, and $\neg H = (W \setminus \{w_1\}) = \{w_2\}$, where H stands for the proposition that a coin lands heads and $\neg H$ stands for the proposition that a coin does not land heads. Suppose that our individual has no clue whether the given coin is fair or biased and consequently has vacuous priors, $\mathbb{P}(H) = \mathbb{P}(\neg H) = \{0, 1\}$. Then, they learn that the coin is fair. How should the individual revise their beliefs in light of learning the new information? For one thing, the information rules out that H is deductively true given the evidence and the same goes for $\neg H$. So the vacuous priors get extinguished, but what replaces them? Since the expectation for a fair coin landing heads is $1/2$, the individual should consider $p(H) = 1/2$ to be a plausible probability. Thus, $1/2 \in \mathbb{P}(H \parallel E)$ and $0, 1 \notin \mathbb{P}(H \parallel E)$. This example illustrates a successful credal set replacement and so learning is possible after all with vacuous priors. Of course, a more systematic rule for determining $\{p_j\}$ is desirable, which may come in time. For the moment, though, I only wish to propose a rule that avoids inertia while capturing intuitive belief change.

Is this alternative approach meant to displace conditionalization? Not at all. In fact, the two might work nicely together, especially in assessing whether two

propositions, X and Y , are independent. Suppose that $\mathbb{P}(X)$ contains a finite number of probabilities including the extremes 0 and 1. Conditioning on Y , however, will not affect the dogmatic probability measures assigning either 0 or 1 come what may, but the interior probabilities between the lower and upper probabilities might have better luck. Nevertheless, whether X and Y are actually stochastically independent cannot be determined in the lower and upper conditional probabilities in this case, for they will always come out to be the same as the prior lower and upper probabilities, which is just a forced consequence of conditional probability.

In order to assess whether X and Y are epistemically independent (though, not necessarily stochastically independent), X may be conditioned on Y together with a credal set replacement that rids the dogmatic (conditional) probability measures. If X and Y are probabilistically correlated for at least one p in the remaining set $(\mathbb{P} \setminus \{0, 1\})$, then X and Y are not said to be epistemically irrelevant. In this case, the joining set $\{p_j\}$ may be non-empty. On the other hand, if $p(X|Y) = p(X)$ for all $p \in (\mathbb{P} \setminus \{0, 1\})$, then X and Y are at least epistemically irrelevant, whereas X and Y are *epistemically independent* just in case $p(X|Y) = p(X)$ and $p(Y|X) = p(Y)$ for all $p \in (\mathbb{P} \setminus \{0, 1\})$. The subsets X and Y are said to be stochastically independent if and only if $p(X \wedge Y) = p(X)p(Y)$ for all $p \in (\mathbb{P} \setminus \{0, 1\})$.⁹ If X and Y are either epistemically independent or stochastically independent, then the joining set in credal set replacement is $\{ \}$ provided that no new probability distributions are rendered plausible by learning an irrelevant proposition Y .

We see now that credal set replacement need not be in competition with conditionalization, for replacement (or reduction) and Bayesian conditioning work well together in providing an individual with an informative belief state in light of obtaining some new piece of information. Credal set replacement eliminates dogmatism while Bayesian conditionalization illuminates more precise credal commitments in the process of learning.

4.4 Summary

In summarizing this chapter, if (SD) is a necessary condition for a proper representation of complete ignorance as Norton proclaims, then by the formal analysis provided above, the imprecise probability method of assigning the set $\{0, 1\}$ to contingent propositions is a sufficient approach to representing an epistemic state of

⁹For discussions on independence in imprecise probability, see Cozman (2012) and Pedersen & Wheeler (2014).

complete ignorance regarding those propositions. The crucial properties of imprecise probability that allow the argument to go through are the conjugacy relations of lower and upper probability since these properties trivially guarantee vacuous priors $\{0, 1\}$ for a contingent proposition and its negation if the smallest admissible probability is set as the lower bound on each proposition.

Furthermore, I have compared the approach with Norton's seemingly plausible account of a non-additive schema, but showed that the former maintains desirable properties for a logic of belief that the latter forfeits. The imprecise probability model has also been given a plausible behavioral interpretation in which common epistemic parlance is made irrelevant, and so the conceptual confusion is put to rest. Finally, I offered a novel response to the inductive learning problem where an individual revises vacuous priors through credal set replacement and more extensively credal set replacement with conditioning.

Chapter 5

Probabilistic Confirmation Theory with Imprecise Probabilities

Reasoning correctly about propositions whose truth values are uncertain can be challenging, especially if the evidence available distracts one from the truth by inducing imprecision in belief. This may happen when a body of evidence contains conflicting peer opinions (Chapter 3; Elkin & Wheeler forthcoming), unspecific statistical information (Ellsberg 1961; Joyce 2005), or indeterminate chance hypotheses (Fine 1988; Hájek & Smithson 2012). If an individual is unfortunate enough to find oneself in such a scenario, how should they react?

To reiterate the ongoing line throughout the dissertation, an evidentialist, like Joyce (2005), claims that one should adjust their attitude accordingly, which amounts to adopting an opinion matching the character of the evidence. However, in representing a less-than-precise epistemic state formally, the canonical tool, i.e. Bayesian probability, falls short as we already know, for Bayes is capable of representing known uncertainty, but it cannot model ranges of uncertainty. Overcoming this limitation can be attained upon adopting a formal theory of imprecise probability once again, and one in which beliefs are represented with *convex* sets or interval-valued probabilities instead of single, point-valued probabilities.¹

In the previous chapters, we have seen how imprecise probability may be of service in modeling a plausible solution to peer disagreement and accommodating the epistemic state of complete ignorance. But to be a successful epistemology matching up to Bayes, imprecise probability is in need of a story about confirma-

¹See Levi (1974), Gilboa & Schmeidler (1989), van Fraassen (1990), Walley (1991), Sturgeon (2008), Joyce (2010), Bradley (2014), Augustin et al. (2014), and Benétreau-Dupin (2015) who assume or discuss imprecise probabilities as closed, convex sets of probability measures.

tion or inductive support. In an attempt to extend the epistemology, proponents might naïvely propose a theory of confirmation resembling ordinal Bayesian confirmation. However, I illustrate in this chapter that the task is not that simple, for a number of ordinal confirmation theories in imprecise probability theory can be stated, each having merits but problems, too. As we will come to see, the problems are non-trivial, and so we should be skeptical as to whether there is such a confirmation theory sufficient for scientific reasoning and epistemology in general. The conclusion may come as a surprise to the reader, but like any responsible philosopher, scientist, or engineer, one should admit to any methodological shortcomings, and it appears that confirmation is but one of imprecise probability's limitations.

This chapter will proceed in the following way. Section 5.1 introduces what has become the canonical theory of imprecise probability in which the beliefs of an individual are represented by convex sets of probability measures bounded by lower and upper probabilities. The subsequent subsection 5.1.1 explicates a particular advantage that the model enjoys over orthodox Bayes. In section 5.2, I briefly rehearse the specifics of ordinal Bayesian confirmation theory along with accompanying background details that provide a template for constructing a confirmation theory in imprecise probability. In section 5.3, I return to imprecise probability and consider ways of defining confirmational relations within the framework through a number of candidates: *extremity*, *previsions-based*, *sensitivity*, and *interval dominance*. Following each theory's description, I highlight its unique merits and limitations. Finally, I conclude briefly in section 5.4 on the prospects of a probabilistic confirmation theory with imprecise probabilities.

5.1 Imprecise Probability Revisited

In adhering to the apparent tradition of the philosophical literature, let \mathcal{P} be a non-empty set of probability measures with each measure $p \in \mathcal{P}$ defined on an algebra \mathcal{F} over a finite set of worlds W .² For all elements $X \in \mathcal{F}$, there is a probability measure with the smallest point-value and a probability measure with the largest point-value relative to $\{p(X) : p \in \mathcal{P}\}$. The probability measures with the smallest and largest point-values for every $X \in \mathcal{F}$ realize a lower and an upper probability, $\underline{P}(X)$ and $\bar{P}(X)$, respectively. Provided that a lower probability automatically induces a conjugate upper probability (and *vice versa*), we may primitively spec-

²See Levi (1974), Sturgeon (2008), and Joyce (2010).

ify a lower probability structure $(W, \mathcal{F}, \mathcal{P}, \underline{\mathbf{P}}, V)$ or an upper probability structure $(W, \mathcal{F}, \mathcal{P}, \bar{\mathbf{P}}, V)$ in relation to a propositional language \mathcal{L} . By convention, lower probability is commonly taken to be primitive,³ but nothing in our discussion significantly depends on which side is to be taken.

The formalism laid out is enough to provide us with a framework for representing static imprecise credences throughout this chapter. In particular, an individual's credence in a proposition A , relative to a lower probability structure $(W, \mathcal{F}, \mathcal{P}, \underline{\mathbf{P}}, V)$, is represented by the non-empty set of probability measures $\mathcal{P}(A)$ called a *credal set* (Levi 1980). Credal sets are assumed to be closed under convex combinations—that is, for any $\lambda \in [0, 1]$ and $p_1, p_2 \in \mathcal{P}$, $\lambda p_1 + (1 - \lambda)p_2 \in \mathcal{P}$. Closure under convex combinations induces intervals, $\mathcal{P} = [\underline{\mathbf{P}}, \bar{\mathbf{P}}]$, and so as things stand, the representation of credence is interval-valued.⁴ Everything described here should be familiar with the exception of convexity. In earlier chapters, we considered non-convex sets, \mathbb{P} , but now we will look at a generalization, \mathcal{P} , for the purpose of examining the most general probabilistic confirmation theory later on.

At this point, we have established an imprecise probability representation for a static belief state, but how does the theory accommodate learning? The common belief updating proposal of imprecise probability is generalized Bayesian conditionalization. Since the probability measures in a credal set, \mathcal{P} , are individually precise, then a proposition A is conditioned on some newly acquired evidence E where $\underline{\mathbf{P}}(E) > 0$ for every $p \in \mathcal{P}(A)$, which yields a closed, convex set of conditional probabilities, $\mathcal{P}(A|E)$, bounded by lower and upper conditional probabilities, $\underline{\mathbf{P}}(A|E)$ and $\bar{\mathbf{P}}(A|E)$. The individual then adopts a new credal set $\mathcal{P}'(A) = \mathcal{P}(A|E) = [\underline{\mathbf{P}}(A|E), \bar{\mathbf{P}}(A|E)]$ until $\underline{\mathbf{P}}'(A) = 1$ or $\bar{\mathbf{P}}'(A) = 0$.

The basic theory is now complete, but how should the belief model be interpreted? To jog the reader's memory, imprecise probabilities may be regarded as guides to one's previsions and elicited via a betting scheme the same as before. The key difference, however, is that we no longer require fair prices since one may have one-sided *lower* and *upper* previsions, or in the language of British sports books, 'back' and 'lay' rates. To demonstrate the point, let us consider a set of hypotheses $\Theta = \{H, \neg H\}$, which partitions W with respect to a lower probability structure $(W, \mathcal{F}, \mathcal{P}, \underline{\mathbf{P}}, V)$. We may now determine an individual's imprecise credences regarding Θ through identifying their betting rates for a gamble I_H that rewards them \$1 if hypothesis H is true, \$0 otherwise, and likewise for $I_{\neg H}$.

³See Williams' (1975) theory of previsions that takes the upper to be primitive.

⁴Convexity is a requirement held by Levi (1974), Walley (1991), and Joyce (2010).

Suppose that I_H is a gamble that pays \$1 if ocean heat is a cause of increasing global temperatures. Imagine now that a scientific observer is inclined to exchange up to \$.45 for the gamble. The maximum buying price is one's lower prevision or the highest amount that they will back I_H . On the other side, though, it turns out that they refuse to sell the gamble for less than \$.60. The minimum price is one's upper prevision or the lowest amount that they will lay I_H . Like before, previsions reflect an individual's credences, but in this case, the observer has an imprecise credence in H where the back price corresponds to the lower probability of their credal set while the lay price corresponds to the upper probability, and we find that $\underline{P}(H) \neq \bar{P}(H)$. Similarly, if the observer were offered the gamble $I_{\neg H}$ instead (where $\neg H$ denotes 'ocean heat is not a cause of increasing global temperatures'), then they should have a back price of \$.40 and a lay price of \$.55 (determined by conjugacy), which correspond to the coherent lower and upper probabilities for $\neg H$, respectively, and again, $\underline{P}(\neg H) \neq \bar{P}(\neg H)$. If any price in (\$.45, \$.60) is asked or offered for I_H and any price in (\$.40, \$.55) is asked or offered for $I_{\neg H}$, the observer rejects the gambles, thereby revealing their range of uncertainty.

Because an individual might have different reactions toward contingent propositions due to epistemic imprecision or maybe cognitive susceptibility to loss aversion⁵, imprecise probability provides a more suitable framework for modeling credence than Bayes, and there ought to be a general inclination toward it since, after all, there is little if anything to be lost methodologically by adopting the model and much more to be gained.

⁵See Kahneman, Knetsch, & Thaler (1990) and Tversky & Kahneman (1991) on differences in 'willingness to pay' (WTP) and 'willingness to accept' (WTA). Their empirical findings suggest that there is often a gap between WTP and WTA, which plausibly stems from an *endowment effect* (Thaler 1980)—that is, some good increases in value once it is owned or added to one's endowment. Kahneman & Tversky (1979) explained the phenomenon on the basis that many individuals are *loss averse* when an expected loss looms large in comparison to an equal-sized expected gain relative to a reference point e.g. current state of wealth. This idea might provide a plausible explanation for why an individual has non-equal buying (WTP) and selling (WTA) prices for risky contracts.

For our purposes, we can capture a similar idea in the language of lower previsions. On the buy side, suppose that one has a maximum buying price of $x < \$1$ for an asset I_H and surrenders x for the chance to increase their wealth w by $\$1 - x$. On the sell side, the individual's selling price for the asset is y where $x \leq y \leq \$1$. The selling price is strictly greater just in case the individual values the \$1 to be surrendered with the contract I_H more than before it became part of her endowment. Of course, the difference here between the buying and selling prices will depend on the individual's degree of loss aversion determined by the ratio $\frac{WTA}{WTP}$. The explanation may breakdown when introducing other factors, however, but loss aversion is one way to account for differences in lower and upper previsions in certain instances.

5.1.1 Rejection of Indifference From an Aversion to Regret

Before turning to Bayesian confirmation, I would like to first discuss an advantage of imprecise probability. We will begin by imagining that you are almost fully ignorant with respect to some contingent propositions A and $\neg A$. How should you respond to any minimal and equally balanced background information relating to the propositions? A typical solution involves constraining your credences in A and $\neg A$ by the *principle of indifference* (POI) (see Keynes 1921; White 2009; Pettigrew 2014). The principle is motivated by the intuition that an individual should have uniform credences if, in the general case, no proposition in a set of $n < \infty$ logically incompatible propositions is favored more than any other.

While we have already considered technical arguments against the principle in the previous chapter, there is a pragmatic argument to be given against it. The purpose of introducing the pragmatic argument is to illustrate why an imprecise prior in reasoning, ordinary or scientific, is reasonable at times and a uniform prior motivated by the principle of indifference or maximum entropy is not, which may be shown through the following example. Suppose that an experimenter presents you with two urns, A and B, containing 100 balls each. Here is your information:

- Urn A contains a mixture of 50 red and 50 black balls.
- Urn B contains some unknown mixture of red and black balls.

You are asked the following questions. What is your expectation for blindly drawing a red ball from Urn A? It is obvious that your expectation should be $1/2$ based on the statistical information provided. Next question, what is your expectation for blindly drawing a red ball from Urn B?⁶ The answer is not immediately clear.

Let *Red* denote ‘a red ball is drawn from Urn B’ and *Black* denote ‘a black ball is drawn from Urn B’. As an (objective) Bayesian, you adhere to POI with respect to the set of propositions $\{Red, Black\}$ and arrive at credences $p(Red) = 0.5$ and $p(Black) = 0.5$. This is the usual story. Here is one explanation for having these credences toward the propositions. You have no evidence indicating that there are more red balls than black or more black than red in Urn B. Thus, you are only in a position to believe that a randomly drawn ball from Urn B will be red or a randomly drawn ball will be black. Representing your belief state by a probability measure p , the above probability distribution (allegedly) accommodates your epistemic ignorance on the matter correctly.

⁶This example was constructed by Ellsberg (1961).

Supposing that you take your epistemic position to be properly accounted for in this instance, you might reconsider when facing risks associated with acting on such credences. In particular, your fair prices for gambles on *Red* and *Black* increase a risk of regretting, for they might be too high or too low. Your fair prices therefore leave you exposed to being swindled by a clever party. To illustrate, let I_{Red} and I_{Black} be regarded as the same kind of gambles described previously with one of them paying out \$1 and the other paying out \$0, depending on which proposition is true. By adhering to POI, your prices turn out to be coherent given that you do not expect a sure loss come what may. So far, so good. Now suppose that a gambler is willing to sell you one of the above gambles at your fair price, either I_{Red} or I_{Black} , but not both. Since you are indifferent toward *Red* and *Black*, it does not matter to you which gamble you choose, so you agree to I_{Red} at a price of \$.50. Here is where things could go wrong. Even though you are not booked in sure loss, \$.50 would be an unfavorable price to pay for the gamble if it turns out that there are less than 50 red balls in the urn, which is unknown to you.

Let us develop the situation to make things clear. After agreeing to the arrangement, the seller tells you that they will announce the number of red balls in the urn, which is known to them, after the draw. The seller then gives you the option to void the contract and get your money back or stay the course. Since the new information has no bearing on your beliefs at this moment in time, you might as well stay the course—that is, the new information does not improve your epistemic situation with respect to the color of the ball to be blindly drawn. So on we go. A red ball is drawn! Your wealth is now guaranteed to increase by \$.50. The seller, although disappointed in the outcome, announces that there were only a total of 5 red balls in the urn. Once the excitement of winning settles, you realize that you should have paid a lot less, especially because you could have very easily been \$.50 poorer given an objective probability 0.95 that a black ball would be drawn. Despite the investment turning out to be profitable for you, it is certain that you would have preferred to pay less now that you know the objective probability of a red ball being drawn, which ultimately reveals how fortunate you were.

As the agreement was initially setup, the seller conned you by offering the gamble at your fair price. They benefited from your ignorance and made you overpay for the gamble—things just did not go their way. Of course, though, the gambler would always sell the bet for \$.50 to anyone willing since they would profit in the long run. *Black* is significantly favored to *Red* on every random draw from the mixed urn. Realizing this, you may be left in a state of *regret* when reflecting on

the agreement at the fair price determined by POI.⁷ The point is that adopting a uniform prior in the face of ignorance can inadvertently expose one to practical risks of regretting. Even though adhering to the principle of indifference seemed optimal in the above scenario at first, it turns out that it was not. The possibility that the proportion of red balls in the urn is less than $1/2$ was not one that you could have ruled out given your information. So, a uniform prior fails at times in appropriately accounting for one's epistemic ignorance.

Consider an alternative to the current case that puts pressure on precise credences in general, not just uniform priors. Suppose that you agree to I_{Red} at the price of \$.50 under the original conditions. Then, the seller generously and truthfully tells you that your estimate is off—black is favored to red. The seller gives you the option to void the contract and get your money back or revise your fair price and take a rebate. You now have good reason to abandon the principle of indifference and revise your fair price, but how are you to adjust your credences in light of the new information? You might try conditioning on the new information, but it does not seem like your epistemic situation will improve a whole lot. This is because the seller has supplied you with fairly unspecific information, and any precise probability estimate favoring *Black* will ultimately amount to a guess.

Just like in the first scenario, it is possible that you do not overpay for the gamble since if, for example, the undisclosed number of red balls in the urn is 30 and you revise your price to $$.x < $.30$, then you are getting a bargain. But of course, the clever gambler did not set things up in your favor. Instead, they offer you I_{Red} at the price of \$.40 and will return to you a \$.10 rebate. The price that

⁷See Loomes & Sugden (1982) for a theory of regret in decision making. More recently, some economists (e.g. Filiz-Ozbay & Ozbay 2007) have studied the anticipation of regret, particularly in first price sealed bid auctions, which bears relevance to the discussion here. In first price sealed bid auctions, bidders hold private values v for an object at auction. The bidders' values are independently and identically drawn from $[\underline{v}, \bar{v}]$ with bidder i having a bidding strategy $b_i : [\underline{v}, \bar{v}] \rightarrow [0, \infty)$ and $b_i \leq v_i$. The winning bid is $b_w > b_j$ for all $j \neq w$ and the winner pays b_w for the item auctioned. If $b_i = b_j = b_w$ for any i and j , then the object is randomly awarded to either i or j . At the conclusion of the auction, *ex post* winner regret may be felt by i if i wins the object with a bid b_i , but could have won with a lower bid $b^* < b_i$. For instance, if i values an item at \$100, bids value v_i , wins the item, but learns afterward that the second highest bid was only \$5, then i will experience a feeling of regret since if bids increase incrementally by \$1, then $b_w = \$6$ yielding a profit of $v_i - b_w = \$94$. If i anticipates winner regret prior to bidding, then i may lower the bid b_i in order to minimize the risk of regretting. Similarly, an individual placed in our scenario above lacks information regarding Urn B's composition, which may induce an anticipation of regret in exchanging \$.50 for I_{Red} . Again, the proportion of red balls in the urn may very well be less than $1/2$ and so one may anticipate such possibility leading to a refusal of paying the asked price. Unfortunately, you, the objective Bayesian, neglected such possibility upon adopting a uniform prior and paid the price (literally).

the seller offers is consistent with the announcement that black is favored to red, and so the gamble might appear attractive. From your perspective, things seem to be going all right at this stage and so you agree to take the rebate, but once again, you pay too much for the gamble given that there are only 30 red balls in the urn. As a result, you will regret agreeing to the contract, win or lose, after the final announcement is made even though you thought you reasoned well on the matter.

The take home message is that ignorance, full or partial, subjects an individual to unwanted risk associated with acting on precise judgments. An individual would be rational to minimize the risk by better accommodating their epistemic position just in case they are in a state of ignorance.⁸ Considerations of regret may nudge one along by inducing an aversion to willingly take the risk, and the individual might then adopt imprecise credences that match the available evidence if any. In the first case, for example, one's evidence at the very best warrants $\mathcal{P}(Red) = (0, 1)$ and $\mathcal{P}(Black) = (0, 1)$. In the second case, $\mathcal{P}(Red) = (0, 0.5)$ and $\mathcal{P}(Black) = (0.5, 1)$. Although these imprecise credences could require refraining from taking action, one is better off refraining rather than embracing an opportunity of being swindled by another party, or by nature while it refuses to fully reveal itself.

5.2 Bayesian Confirmation Theory

In determining whether evidence stands in favor of a scientific theory or hypothesis, Bayesian confirmation theory provides an answer, and that answer has had much influence in the philosophy of science for several decades.⁹ There is very little mystery surrounding its prominence based on how well Bayesian confirmation captures the relation between theory and evidence and also its ability to accommodate surprising new evidence. The ordinal version of the theory clearly illustrates this and is simply written in the following way.

BAYESIAN CONFIRMATION THEORY:

⁸Pettigrew (2014) has developed an argument from risk (of epistemic disutility) but in defense of POI. The difference between the conclusions we draw is forced by presuppositions on both sides, as I see it, where he is interested in minimizing the risk of epistemic disutility resulting from a non-uniform prior with greater expected inaccuracy than a uniform prior, while I am interested in minimizing the risk of unnecessary loss in utility resulting from acting on one's precise judgments.

⁹Some notable figures who endorse the method include Bovens & Hartmann (2003), Howson & Urbach (2005), and Fitelson & Hawthorne (2010) just to name a few. However, Bayesian confirmation has its notable critics also including Glymour (1980), Mayo (1996), and Norton (2011).

- $p(H|E, K) > p(H|K)$ (Confirmation)
- $p(H|E, K) < p(H|K)$ (Disconfirmation)
- $p(H|E, K) = p(H|K)$ (Irrelevance).¹⁰

In plain terms, Bayesian confirmation theory states that a theory or hypothesis H is (i) confirmed if evidence E increases the probability of H conditional on background knowledge K , (ii) disconfirmed if E decreases the probability of H conditional on K , or (iii) neither confirmed nor disconfirmed if E makes no difference to H conditional on K . On the face of it, the theory is minimal and elegant, but there are some important details associated with the theory, which I will now discuss.

The irrelevance condition is a fundamental concept in classical probability theory, and it implies that H and E are stochastically independent. Leaving K in the background, if H and E are stochastically independent, then the following equality obtains

$$p(H \wedge E) = p(H)p(E), \quad (5.1)$$

which implies

$$p(E|H) = p(E). \quad (5.2)$$

In case a theory or hypothesis H and evidence E are stochastically independent, by symmetry, E is not evidentially relevant to H and H is not evidentially relevant to E . It is thus intuitive that no confirmation or disconfirmation occurs upon learning some irrelevant piece of information.¹¹

Moreover, if H and E are not independent, then H may either be confirmed or disconfirmed by E , depending on whether Bayesian confirmation is given a strong or weak interpretation. On a strong interpretation, a theory or hypothesis H is (dis)confirmed just in case $p(H|E, K) > t$ (or $p(H|E, K) < 1 - t$) where t is some threshold typically above $1/2$. On a weak interpretation (the version presented above), a theory or hypothesis H is (dis)confirmed as long as E increases (or decreases) the probability of H given K . The former is referred to as *absolute confirmation* and the latter *incremental confirmation*. Of the two, Bayesians typically prefer the incremental version, for it was I.J. Good who remarked that if

¹⁰Assuming *regularity*—that is, $p(E) > 0$ for all propositions E .

¹¹It is important to keep this property in mind since confirmational irrelevance in imprecise probability does not necessarily imply stochastic independence.

$p(H|E, K)$ is close to unity (and above t) but less than $p(H|K)$, one should not claim that the evidence is confirming (1968, 134).¹²

Furthermore, notice that Bayesian confirmation theory as described above is *qualitative* provided that there is no mention of numerical degrees of confirmatory support that some evidence lends. Instead, confirmation, disconfirmation, and irrelevance are defined above in terms of ordinal relations over posterior and prior credences. In the literature, however, one can find a variety of quantitative measures of support, and there has been much discussion devoted to figuring out which measure is sufficient for determining the degree to which some evidence E (dis)confirms a theory or hypothesis H .¹³ For the purposes of this chapter, I will leave numerical degrees of support out of the picture and focus strictly on ordinal Bayesian confirmation theory to keep things simple.

Now that we have a handle on the basic details of Bayesian confirmation, interpretations and caveats aside, what benefits are enjoyed by adopting the probabilistic approach in scientific reasoning? One benefit of Bayesianism is a capability of ranking theories and hypotheses from the best to least supported, which is absent from deductive accounts of confirmation. As a result, the assessment of scientific theories and hypotheses is no longer an all-or-nothing or “yes/no” matter. Instead, the status of theories and hypotheses may vary by incremental support yielded through the available evidence. Secondly, confirmational relations in the Bayesian framework are not determined purely by syntax, which is a good thing as some of the philosophical puzzles associated with deductive accounts of confirmation are avoided. Meanings of terms turn out to be relevant in the probabilistic confirmation theory (Hájek & Joyce 2008).

Beyond its theoretical advantages over competing approaches, Bayesian confirmation has purported success stories, and proponents continue finding novel applications in science. A recent example is the Bayesian *No Alternatives Argument* put forth by Dawid, Hartmann, and Sprenger (2014). Using the formal machinery, they have shown that a failure to find an alternative theory T' in scientific inquiry produces (non-empirical) support for the default theory T under consideration, which ultimately is said to confirm T . The NAA argument has positive implications for scientific theories that are unable to be empirically tested in theoretical physics, for example, such as string theory and the multiverse (Dawid 2013).

¹²Philosophically, I will remain neutral on this point, primarily because an exploration of both interpretations in imprecise probability will be fruitful later on.

¹³See Fitelson (1999) for a nice overview of confirmation measures.

While a Bayesian approach to scientific reasoning offers various benefits, Bayesian confirmation theory by no means comes out fully unscathed in philosophical analysis. Despite the problems that have been raised in the literature, though, one would be hard-pressed in denying that Bayesianism in general has had significant influence on the philosophy of science.

5.3 Ways of Describing Confirmation in Imprecise Probability Theory

With an understanding of the most influential theory of confirmation in the philosophy of science, let us attempt to make sense of probabilistic confirmation for when prior credences are imprecise. In moving forward, it may seem natural to generalize the Bayesian definitions already given, resulting in generalized Bayesian confirmation. But an immediate problem surfaces with such a naïve maneuver. Specifically, the required class of probability measures satisfying each ordinal condition is undefined.

The hope of easily constructing a suitable ordinal confirmation theory with imprecise probabilities in the image of Bayes turns out to be short lived, and I will demonstrate just how difficult it is to adequately define confirmation in imprecise probability theory through an examination of four candidate theories. I do not claim, however, that the compiled list is complete, for there may be some plausible theory overlooked. But the explication will suffice to show that it is far from clear whether there is a unique theory of confirmation that is sufficient for inductive reasoning with imprecise probabilities.

5.3.1 Confirmational Extremity

I will begin by introducing the strongest account of (incremental) confirmation in imprecise probability, which we might call *extremity*. It is simply written as follows.

EXTREMITY:

- $\mathcal{P}(H|E, K) > \mathcal{P}(H|K)$ **(Confirmation)**
- $\mathcal{P}(H|E, K) < \mathcal{P}(H|K)$ **(Disconfirmation)**

- $\mathcal{P}(H|E, K) = \mathcal{P}(H|K)$ **(Irrelevance)**.¹⁴

In plain terms, a theory or hypothesis H is said to be (i) confirmed if E increases the probability of H conditional on K for every probability measure $p \in \mathcal{P}$, (ii) disconfirmed if E decreases the probability of H conditional on K for every probability measure $p \in \mathcal{P}$, or (iii) neither confirmed nor disconfirmed if E makes no difference to H conditional on K for any one probability measure $p \in \mathcal{P}$.

Without much consideration, one should quickly recognize that these conditions are awfully demanding based on the initial assumption that credal sets include the convex hull between the minimum and maximum probabilities. The span between lower and upper probabilities contains an incredibly vast set of points (infinite), which makes any one of the confirmational conditions extremely difficult to satisfy upon theoretically calculating each conditional point-valued probability measure relative to H given some new evidence E and background knowledge K .

In an attempt to skirt the problem, we could relax the convexity assumption and instead adopt the following.

WEAK EXTREMITY:

- $\mathbb{P}(H|E, K) > \mathbb{P}(H|K)$ **(Confirmation)**
- $\mathbb{P}(H|E, K) < \mathbb{P}(H|K)$ **(Disconfirmation)**
- $\mathbb{P}(H|E, K) = \mathbb{P}(H|K)$ **(Irrelevance)**.¹⁵

The *weak extremity* theory employs a bounded set of probability measures, \mathbb{P} , that is not necessarily convex. On conceptual grounds, one may regard \mathbb{P} as the set of *plausible* probability measures from an individual's perspective.

Opting for a non-convex set, \mathbb{P} , is reasonable in certain cases like the Ellsberg experiment discussed earlier. Considering such task, a probability $p(\text{Red}) = 0.25816$ is in the set $\mathcal{P}(\text{Red}) = (0, 1)$, but an individual may rule out that probability given the possible compositions of the urn and corresponding probability distributions. Specifically, one might reason that it is possible for 25 balls in the urn to be red or 26, but surely it is not the case that the number of red balls is 25.816. A degree of belief 0.25816 in *Red* does not track any objective chance related to the problem (assuming that chances are mapped to the unit interval). Thus, it appears

¹⁴ Assuming $\bar{\mathbb{P}}(E) \geq \underline{\mathbb{P}}(E) > 0$.

¹⁵ Assuming $\bar{\mathbb{P}}(E) \geq \underline{\mathbb{P}}(E) > 0$.

unnecessary from the individual's perspective to include the above probability in their set, which breaks convexity.

There is another reason, too, for relaxing the convexity assumption, namely to obtain flexibility in *parameterizing* \mathbb{P} in such a way that structural judgments, like independence, are preserved, which vanish with closed, convex credal sets (Haenni et al. 2011).¹⁶ Interval-valued credal sets seem intuitive when characterizing imprecision in belief by spanning the gap from the minimum to maximum level of confidence, but they tend to mask more subtle credal commitments one might hold as discussed in the third chapter.

Putting the differences between belief models aside, either approach to the issue at hand enjoys the advantage of engendering a novel perspective on probabilistic confirmation, distinct from the classical Bayesian theory. This is done by working an old and familiar idea into the picture to make a brand new one. In particular, we may call upon supervaluationism to enhance our understanding of confirmation from these theories and add new notions to the repertoire: *superconfirm*, *superdisconfirm*, and *superirrelevant*.¹⁷

Imprecise probability and supervaluationism seem to relate nicely, and the overlap has recently been highlighted by Rinard (2015), describing the connection as follows.

We can apply this supervaluationist strategy to doxastic imprecision by seeing each function in your [credal] set as one admissible precisification of your doxastic state. Functions excluded from your set are inadmissible precisifications. Whatever is true according to all functions in your set is *determinately* true; if something is true on some, but not all functions in your set, then it's *indeterminate* whether it's true. For example, if all functions in your set have $\Pr(A) > \Pr(B)$, then it's determinate that you're more confident in A than B. If different functions in your set assign different values to some proposition P, then for each such value, it's indeterminate whether that value is your credence

¹⁶As a reminder, irrelevance does not necessarily entail stochastic independence in imprecise probability as it does in classical probability. There is a vast literature dedicated to the issue (e.g. Walley 1991; Cozman 2012; Augustin et al. 2014; Pederson & Wheeler 2014). I will not dwell on the point, but it is important to know that propositions X and Y are completely stochastically independent when $\forall p \in \mathcal{P}, p(X \wedge Y) = p(X)p(Y)$ (Pederson & Wheeler 2014, 1325). There will be times, however, when the joint is not factorizable with respect to at least one probability measure in the set. As a consequence, E might be confirmationally irrelevant to H in imprecise probability theory without the propositions being stochastically independent. If the independence property is desired for certain judgments, it is optimal to adopt \mathbb{P} instead of \mathcal{P} .

¹⁷Thanks to Branden Fitelson and Matt Kotzen for suggesting similar ideas in discussion.

in P. (Rinard 2015, 2)

Tailoring the idea to confirmation, a supervalutionist interpretation of the extremity theories supplies us with the following definitions: E **superconfirms** H given K just in case $\mathcal{P}(H|E, K) > \mathcal{P}(H|K)$ (or $\mathbb{P}(H|E, K) > \mathbb{P}(H|K)$), E **superdisconfirms** H given K just in case $\mathcal{P}(H|E, K) < \mathcal{P}(H|K)$ (or $\mathbb{P}(H|E, K) < \mathbb{P}(H|K)$), or E is **superirrelevant** to H given K just in case $\mathcal{P}(H|E, K) = \mathcal{P}(H|K)$ (or $\mathbb{P}(H|E, K) = \mathbb{P}(H|K)$).

A benefit of the extremity theory of confirmation with a supervalutionist interpretation is in confirming that confirmation, disconfirmation, or confirmational irrelevance *determinately* obtains. To demonstrate the difference between classical and supervalutionist probabilistic confirmation, suppose that one has a precise subjective probability distribution p regarding a finite set of logically incompatible theories or hypotheses. The individual acquires new evidence and updates their credences such that they have a new credence p^* in some H where $p^*(H) = p(H|E, K) > p(H|K)$. The new evidence confirms H according classical Bayesian confirmation theory. But since precise numbers are sensitive to Bayesian conditioning (even more so as numbers are sharpened), the outcome could have been different had the individual's prior credence been different.

For instance, suppose that one's prior credence is instead $p'(H|K)$ and the difference is $|p'(H|K) - p(H|K)| = \epsilon$ where ϵ is small. Nothing guarantees that the same confirmational verdict is given if likelihoods are very sharp. The posterior probability $p'(H|E, K)$ may come out less than or equal to the prior probability $p'(H|K)$, entailing either disconfirmation or irrelevance according to Bayesian confirmation theory. In contrast to orthodoxy, if a credence is represented by a set of probability measures like in the extremity approaches, the confirmational relation is better confirmed the coarser the set becomes, i.e. confirmational relations hold over a *neighborhood* (or a finite set) of probability measures. Assuming that any one of the confirmational conditions of *extremity* (or *weak extremity*) is satisfied with non-singleton sets, a supervalutionist interpretation suggests that the new evidence is *determinately* confirming, disconfirming, or irrelevant, whereas the Bayesian has no such reassurances. In some sense, we have added a counterfactual element for a Bayes agent where evidence is said to confirm H if the positive condition were to be satisfied had the precise Bayes agent's prior been different, relative to a particular region of the unit interval.

While this is all well and good, the extremity accounts face substantial prob-

lems. The most glaring issue concerns the stipulation that either positive, negative, or zero probabilistic correlation uniformly obtains over credal sets. To illustrate the difficulty, here is just one example where we end up with mixed results, assuming fixed, almost equal, and extremely precise likelihoods.

A counterexample to extremity was computed using SageMathCloud. First, a “Bayes” function is defined for computing posterior probabilities. Next, a simple algorithm is written that computes posterior probabilities for 1000 random prior probabilities and determines whether each posterior is strictly smaller than the corresponding random prior (note: the $<$ confirmation relation for comparison was arbitrarily chosen, which can be changed unproblematically). All posteriors are calculated with arbitrary fixed and almost equal likelihoods l_1 and l_2 . If a posterior probability is smaller than the corresponding random prior probability, then “Yes” is printed along with the prior probability, 1 minus the prior probability, and the posterior probability, otherwise “No” and the respective data. Below is the code written in Python.

```
def Bayes(a, b, c):
    return (a * b) / ((a * b) + ((1 - a) * c))

for i in range(0,1000):
    l1 = 0.4999999999999998
    l2 = 0.4999999999999999
    x = RealField().random_element(0, 0.99)
    if Bayes(x, l1, l2) < x:
        print ‘‘Yes’’, x, (1 - x), Bayes(x, l1, l2)
    else: print ‘‘No’’, x, (1 - x), Bayes(x, l1, l2)
```

$$\begin{aligned}
 p_1(H|E, K) &= \\
 &= \frac{(0.4999999999999998)(0.0741474379509381)}{(0.4999999999999998)(0.0741474379509381) + (0.4999999999999999)(0.925852562049062)} \\
 &= 0.0741474379509379
 \end{aligned}
 \tag{5.3}$$

$$\begin{aligned}
p_2(H|E, K) &= \\
&= \frac{(0.4999999999999998)(0.987310939847310)}{(0.4999999999999998)(0.987310939847310) + (0.4999999999999999)(0.0126890601526900)} \\
&= 0.987310939847310
\end{aligned}
\tag{5.4}$$

In an extreme case of having a near uninformative credal set $\mathcal{P}(H|K) = (0, 1)$, conditioning on new evidence leads to the above opposing confirmational verdicts at the very least. Equation 5.3 shows that E lowers the probability of H given K according to $p_1 \in \mathcal{P}$, precisely disconfirming H . Equation 5.4, however, shows that E is precisely irrelevant to H given K according to $p_2 \in \mathcal{P}$. As a consequence, no extremity confirmational condition is satisfied in this instance. Although this is a trivial example with high precision, it does point out a significant problem for extremity.

Moreover, the above pair of confirmationally-opposing outcomes is one example from several generated by updating 1000 random priors falling within $(0, 1)$. Imagine how often we obtain mixed results by conditioning each point-valued measure in just a small convex region of a credal set. With a very high chance of obtaining mixed results, we come to realize that extremity is silent when there is no single comparative relation for all precise posterior probabilities of credal sets. One might be tempted to say that when none of the confirmational conditions are satisfied, then it is *confirmationally indeterminate* as to whether the new evidence confirms a theory or hypothesis. Reflecting on the supervaluationist discussion, the suggestion seems correct as we might demand that a single confirmation relation holds for every precisification. By adopting this line, however, we permit even a single probability measure to tyrannically undermine a strong majority, thus giving the first to revolt the keys to the city. The weak extremity theory is not immune to the problem either, and the vulnerability increases as the cardinality of a non-convex set, \mathbb{P} , increases.

In the event that there is a lack of unanimity on the confirmational verdict of a theory or hypothesis H , an inquirer may be kept waiting for an eternity by demand of the extremity theories prior to making a judgment. Ultimately, scientific progress will come to a halt if inquirers are constrained in such a way, for many will be long dead before there is a verdict on which theory or hypothesis is best supported by the evidence. This consequence makes extremity in imprecise probability implausible.

5.3.2 A Behavioral Confirmation Theory

Instead of holding such strict standards for confirmation in imprecise probability, we might follow the subjectivist tradition and leave the confirmational status of a theory or hypothesis to be determined by the behavioral reactions of an individual in response to new evidence. In particular, we might say that evidence confirms or disconfirms a theory or hypothesis just in case an individual is disposed to act differently in response to the evidence, otherwise the evidence is irrelevant if one is disposed to act in the same way as they would have prior to learning the evidence.

Some might be alarmed by a highly subjective interpretation of confirmation like this one, but it is doubtful that one can successfully separate intellectual interests from practical interests for *value-driven* individuals. So let us at least entertain the possibility of such kind of confirmation theory, which is written in the following way.

PREVISIONS-BASED CONFIRMATION:

- $\underline{P}(H|E, K) > \underline{P}(H|K) \ \& \ \overline{P}(H|E, K) \geq \overline{P}(H|K)$ **(Confirmation)**
- $\underline{P}(H|E, K) \leq \underline{P}(H|K) \ \& \ \overline{P}(H|E, K) < \overline{P}(H|K)$ **(Disconfirmation)**
- $\underline{P}(H|E, K) = \underline{P}(H|K) \ \& \ \overline{P}(H|E, K) = \overline{P}(H|K)$ **(Irrelevance)**.¹⁸

In plain terms, a theory or hypothesis H is said to be (i) confirmed by E if an individual would now pay more for the gamble H but would not sell the same gamble for any less than before, (ii) disconfirmed by E if an individual would not pay any more for H than before but would sell the gamble now for less, or (iii) neither confirmed nor disconfirmed by E if an individual would maximally buy and minimally sell H at the same prices as before.

How is it that adjustments in previsions correlate with judgments of confirmation and disconfirmation? Suppose that an individual's lower prevision for a gamble H is marked up from \$.50 to \$.60 upon learning E —that is, one is now willing to pay a maximum of \$.60, compared to a previous maximum price of \$.50,

¹⁸I use the italic \underline{P} and \overline{P} to denote lower and upper previsions. A linear prevision P is defined as a function on a set of gambles \mathcal{G} to the reals \mathbb{R} . Lower and upper previsions are one-sided supremum buying and infimum selling prices, respectively. Given the focus on a special kind of gamble, I_X , coherent previsions are bounded by 0 and 1—that is, an individual's supremum buying rate and infimum selling rate for a special gamble I_X are neither negative nor exceed 1. Also, I abuse notation here, using the variable H to represent the special gamble instead of I_H . It should be assumed that H is shorthand for I_H throughout this subsection.

to win \$1 if H is true, \$0 otherwise. Assume in this instance that there is not a decrease in the individual's upper prevision. Moreover, an increase in one's supremum buying price implies an increase in credence in the theory or hypothesis under consideration since the individual's lower prevision for the gamble corresponds to their lower probability for the theory or hypothesis. Thus, we say that a theory or hypothesis is confirmed by E if the highest price one is willing to buy the gamble H increases (and the lowest selling price does not decrease). A positive adjustment in the lower prevision (and non-decreased adjustment in upper prevision) signals that the individual finds the theory or hypothesis to be more plausible, hence why they would pay more for the gamble but not sell it for any less than before.

Imagine instead that an individual's upper prevision for a gamble H is discounted from \$.70 to \$.60 upon learning E , implying that one would now sell the gamble H for as low as \$.60. Assume in this instance that there is not an increase in the individual's lower prevision. Moreover, the individual is willing to exchange the gamble H for a smaller fee than previously. A decrease in the infimum selling rate implies a decrease in credence in the theory or hypothesis under consideration since the individual's upper prevision for the gamble corresponds to their upper probability for the theory or hypothesis. Thus, we say that a theory or hypothesis is disconfirmed by E if the lowest price that one is willing to part ways with the gamble H decreases (and the highest buying price does not increase). A downward change in the upper prevision (and non-increased adjustment in lower prevision) signals that the individual finds the theory or hypothesis to be less plausible, hence why they would sell the gamble to another for a smaller fee but not buy it for any more than before.

With respect to irrelevance, suppose that neither the individual's lower prevision, i.e. \$.50, nor their upper prevision, i.e. \$.70, for the gamble H changes after learning E . Then, the individual's lower and upper probabilities for the theory or hypothesis under consideration are not at all affected by learning evidence E and so it is irrelevant. From the individual's point of view, the theory or hypothesis is neither confirmed nor disconfirmed by E . The irrelevance condition is fairly straightforward in the previsions-based theory.

Now that an explanation of confirmation in terms of adjustments to previsions has been provided, there are some advantages that come with the proposed theory. First, it is more relaxed than the extremity theories and consequently makes for a palatable theory of confirmation in imprecise probability where confirmation and disconfirmation become realistically obtainable. Second, with the theory being

motivated by behavior, it is easy to see how it can be adopted in practice. Of course, many will likely be averse to adopting such a theory because of the subjectivity involved. But methodologically, the behaviorist approach can be quite useful for eliciting judgments of confirmation and disconfirmation held by inquirers. The betting scheme is merely a device for identifying the opinions of experts.

Despite the theory's promise, there are specific confirmationally-relevant instances that are not accounted for. Two instances that come to mind are *attraction* and *repelling* of lower and upper previsions. Let us call the first *precision* since an individual's credence becomes more precise as the lower and upper previsions converge. In regard to the second, it has been named *dilation* (Seidenfeld & Wasserman 1993; White 2009; Bradley & Steele 2014; Pederson & Wheeler 2014). Check these instances against the confirmational criteria for the previsions-based account. You will find that such instances are not subsumed by any of the confirmational categories given by the theory. Again, we might admit that deviant instances fall under a category of *confirmational indeterminacy*. In regard to dilation, the categorization seems apt, for dilation maintains imprecision with respect to credence in a theory or hypothesis. Precision, however, is not very fitting. As an individual becomes more precise on some matter, one dimension of their uncertainty is lessened. In such case, it is unintuitive to describe the epistemic state as 'confirmationally indeterminate'.

A way around the mis-categorization of precision is to place the instances under confirmation. However, doing so calls for a revision to the theory where the condition $\bar{P}(H|E, K) \geq \underline{P}(H|K)$ in the confirmation definition is dropped, but the revision exposes it to an unusual problem. For example, suppose that $\mathcal{P}(H|K) = [a, b]$ and $\mathcal{P}(\neg H|K) = [1 - b, 1 - a]$ where $a, b \in (0, 1)$. Upon conditioning on E , suppose that $\mathcal{P}(H|E, K) = [a + \epsilon, b - \epsilon']$ and $\mathcal{P}(\neg H|E, K) = [1 - (b - \epsilon'), 1 - (a + \epsilon)]$ where ϵ and ϵ' are greater than 0, but $(\epsilon + \epsilon') \leq |a - b|$. Conditioning both H and $\neg H$ on E and K leads to instances of precision. Notice, though, that H and $\neg H$ **both** are confirmed simply by the lower conditional previsions being larger than the respective lower unconditional previsions. A theory of confirmation that simultaneously confirms logically incompatible theories or hypotheses spells disaster.

5.3.3 Confirmational Sensitivity

To follow up on the deficiencies of the previsions-based theory, I introduce what we might call the *sensitivity* theory. This account is by far the most complex given its

extension to the various instances described earlier, many of which are unfamiliar to classical Bayesians.

Remaining neutral on the interpretation of lower and upper probabilities in this section, the theory is written in the following way.

SENSITIVITY:

- $\underline{P}(H|E, K) > \underline{P}(H|K) \ \& \ \bar{P}(H|E, K) \geq \bar{P}(H|K)$ **(Confirmation)**
- $\underline{P}(H|E, K) = \underline{P}(H|K) \ \& \ \bar{P}(H|E, K) > \bar{P}(H|K)$ **(U-Confirm)**
- $\underline{P}(H|E, K) \leq \underline{P}(H|K) \ \& \ \bar{P}(H|E, K) < \bar{P}(H|K)$ **(Disconfirmation)**
- $\underline{P}(H|E, K) < \underline{P}(H|K) \ \& \ \bar{P}(H|E, K) = \bar{P}(H|K)$ **(L-Disconfirm)**
- $\underline{P}(H|E, K) = \underline{P}(H|K) \ \& \ \bar{P}(H|E, K) = \bar{P}(H|K)$ **(Irrelevance)**
- $\underline{P}(H|E, K) > \underline{P}(H|K) \ \& \ \bar{P}(H|E, K) < \bar{P}(H|K)$ **(Precision)**
- $\underline{P}(H|E, K) < \underline{P}(H|K) \ \& \ \bar{P}(H|E, K) > \bar{P}(H|K)$ **(Dilation).¹⁹**

There is little surprise here provided the previous discussion. And it should be clear that the sensitivity theory inherits all of the benefits of the previsions-based theory along with accounting for fine-tuning and coarse-graining of credences by new evidence, i.e. precision and dilation.

Additionally, two new definitions have been introduced, U-Confirm and L-Disconfirm. The former is satisfied just in case the evidence increases the upper side of a credal set while the lower side remains unfazed. The latter is satisfied just in case the evidence decreases the lower side of a credal set while the upper side remains unfazed. In a confirmational sense, U-Confirm signals that evidence partially confirms a theory or hypothesis H and L-Disconfirm signals that evidence partially disconfirms a theory or hypothesis H , where credal sets are anchored by the opposite end in either instance. These additions to a theory of probabilistic confirmation capture partiality and allow it to “cover all of the bases.”

Clearly, though, the sensitivity theory does not inherit the simplicity of previsions-based confirmation. Of course, it is desirable to have something said about the confirmational status of a theory or hypothesis when an unfamiliar situation arises, i.e. precision, dilation, U-confirm, L-disconfirm, but satisfying this desire comes at a price. Although some might not be so concerned with the lack

¹⁹Assuming $\bar{P}(E) \geq \underline{P}(E) > 0$.

of simplicity, the trade-off can be made apparent by comparing sensitivity with ordinal Bayesian confirmation theory. Bayesian confirmation theory is quite simple and elegant with its tripartite division accounting for every possible posterior state. Sensitivity on the other hand surely is neither simple nor elegant. By Occam's razor, why not Bayes, then?

Simplicity might not be a compelling reason to reject sensitivity, but there are further problems with pieces of the theory. Specifically, U-Confirm and L-Disconfirm do not accord with the conceptual notions of confirmation and disconfirmation. To see this, let us treat a credal set as an individual's metaphorical *credal committee* (Joyce 2010). If upon obtaining new evidence, the pessimistic voting committee members remain steadfast while the optimistic voting committee members further their optimism, it is incorrect to say that some kind of confirmation took place, for the disagreement has been exacerbated. L-Disconfirm faces a similar worry.

To make the point clear, suppose that $\mathcal{P}(H|K) = [a, b]$, where $a, b \in (0, 1)$, and upon obtaining new evidence E , $\mathcal{P}(H|E, K) = [a, b + \epsilon]$, where $\epsilon > 0$. The optimists of the credal committee are positively swayed by the new evidence, but the pessimists refuse to budge. Notice that epistemic imprecision increases such that $|a - b| < |a - (b + \epsilon)|$. Since the inquirer's epistemic imprecision increases upon learning E , it is counterintuitive to claim that E is partially confirming evidence. Similarly, if instead $\mathcal{P}(H|E, K) = [a - \epsilon, b]$, then again epistemic imprecision increases such that $|a - b| < |(a - \epsilon) - b|$. One should feel reluctant in calling E partially disconfirming evidence. The notions of (subjective) probabilistic confirmation and disconfirmation typically imply an increase and decrease in belief in a theory or hypothesis H , not an increase in epistemic imprecision. However, I have just shown that U-Confirm and L-Disconfirm uniquely increase an individual's epistemic imprecision towards a theory or hypothesis. Thus, U-Confirm and L-Disconfirm fail to accord with the conceptual notions of probabilistic confirmation and disconfirmation. Unlike dilation, though, the movement in credal committees is lopsided. So what do instances fitting the U-Confirm and L-Disconfirm patterns express? It is not so clear. Despite confirmational sensitivity in imprecise probability theory covering all of the bases, there are trade-offs and unintuitive notions that one must accept if adopted.

5.3.4 Interval Dominance Confirmation

The final candidate I will consider is *interval dominance* confirmation. Unlike the others, interval dominance entails an absolute confirmation theory. To see this, let us begin by defining a strict interval dominance criterion.

Strict Interval Dominance (SID): For any propositions A, B and closed convex credal set \mathcal{P} , A strictly interval dominates B if and only if $\underline{P}(A) > \bar{P}(B)$. Otherwise, B is *interval undominated*.

In the conditional case, A given E will be strictly more probable than B given E or *vice versa* if SID is satisfied. Otherwise, neither proposition interval dominates the other upon learning E .

We will use SID as the criterion for defining confirmation in imprecise probability and propose the following.

INTERVAL DOMINANCE:

- $\underline{P}(H|E, K) > \bar{P}(\neg H|E, K)$ (Confirmation)
- $\bar{P}(H|E, K) < \underline{P}(\neg H|E, K)$ (Disconfirmation)
- $\underline{P}(H|E, K) = \underline{P}(H|K) \ \& \ \bar{P}(H|E, K) = \bar{P}(H|K)$ (Irrelevance).²⁰

In plain terms, a theory or hypothesis H is said to be (i) confirmed if E increases the lower probability of H conditional on K above the upper probability of $\neg H$ conditional on K or (ii) disconfirmed if E decreases the upper probability of H conditional on K below the lower probability of $\neg H$ conditional on K . Irrelevance is held to be the same as in the previsions-based and sensitivity theories.

The requirement for confirmation is stronger than in the previous confirmation theories we have seen and entails an absolute confirmation condition where the lower probability for H conditional on E and K satisfies a threshold t such that $t = \inf(\bar{P}(\neg H|E, K), 1)$. Notice also that under certain conditions, absolute Bayesian confirmation theory is just a special case of the interval dominance theory when $\mathcal{P} = \{p\}$, for $p(H|E, K) \geq t$ and $t > \frac{1}{2}$ provided that $p(H|E, K) = \underline{P}(H|E, K) > p(\neg H|E, K) = \bar{P}(\neg H|E, K)$ if and only if $p(H|E, K)$ is at least $t = \inf(\bar{P}(\neg H|E, K), 1)$, which must be a value greater than $1/2$.

²⁰Assuming $\bar{P}(E) \geq \underline{P}(E) > 0$.

What is significantly different about the interval dominance theory compared to the other confirmation theories that turn out to be incremental is that the latter have no stipulation on a minimum lower probability required for confirmation and disconfirmation whereas the former does. If, for example, $\mathcal{P}(H|E, K) = [a, b]$, $\mathcal{P}(H|K) = [c, d]$, $b \geq d$, $a > c$, and $a, b, c, d \in (0, 1)$, then E confirms H given K according to the previsions-based and sensitivity theories (possibly the extremity theories as well). On the other side, if $\mathcal{P}(\neg H|E, K) = [1 - b, 1 - a]$, $\mathcal{P}(\neg H|K) = [1 - d, 1 - c]$, $(1 - b) \leq (1 - d)$, and $(1 - a) < (1 - c)$, then $\neg H$ given K is disconfirmed by E on the previsions-based and sensitivity theories (possibly the extremity theories as well). In the interval dominance theory, however, E does not necessarily confirm H given K (or disconfirm $\neg H$) unless if $a > (1 - a)$. Thus, the interval dominance confirmation theory may deliver a different confirmational verdict than the incremental confirmation theories in imprecise probability.

From one perspective, the interval dominance theory seems the most plausible. If we consider three disjoint theories, A , B , and C , and are able to determine a ranking of theories conditional on E and K , then we obtain a total preorder of theories by how well they are confirmed. For example, suppose that A is strictly more probable than B and B is strictly more probable than C conditional on E and K . Then, $\underline{\mathcal{P}}(A|E, K) > \overline{\mathcal{P}}(B|E, K)$ and $\underline{\mathcal{P}}(B|E, K) > \overline{\mathcal{P}}(C|E, K)$. In this case, it is intuitive that A is better confirmed than B and B is better confirmed than C . An incremental theory, on the other hand, will supply confirmational verdicts even when imprecise probabilities “overlap” while theories or hypotheses in actuality turn out to be incomparable in support. It is hard to say which candidate is the most likely to be true when none dominate. The interval dominance theory gives confirmational verdicts only when at least one theory or hypothesis is dominating. It therefore avoids confirmational ambiguity in assessment.

However, the interval dominance confirmation theory runs afoul in certain instances, namely when

$$\begin{aligned} \mathcal{P}(H|E, K) &= [a + \epsilon, b + \epsilon'] \\ &> \\ \mathcal{P}(H|K) &= [a, b], \\ \\ \mathcal{P}(\neg H|E, K) &= [1 - (b + \epsilon'), 1 - (a + \epsilon)] \\ &< \\ \mathcal{P}(\neg H|K) &= [1 - b, 1 - a], \end{aligned}$$

and $(1 - (b + \epsilon')) > (b + \epsilon')$, where ϵ and ϵ' are both greater than 0. In this example, $\neg H$ given E and K strictly interval dominates H given E and K . But since $b + \epsilon' > b$, E decreases the lower probability of $\neg H$ given K , yet the interval dominance theory classifies the instance as confirmation of $\neg H$ given K . The outcome is quite unusual since E incrementally reduces the support for $\neg H$, but the verdict is that E confirms $\neg H$. Here is where the absolute theory seemingly goes wrong, and this point against the interval dominance theory is a generalization of the point made by Good against absolute Bayesian confirmation theory.

One way out might be to say that such instance leads to *confirmational indeterminacy*. The addition of a confirmational indeterminacy category might also subsume non-domination by any candidate (full incomparability) noted above. But in adopting the go-to confirmational indeterminacy category, however, the theory is in need of revision since it accommodates too few confirmational notions. Furthermore, formal amendments to the theory, like confirmational indeterminacy, are not obvious on conceptual grounds, for in the above example, $\neg H$ reigns supreme as the most well confirmed theory, despite a slight evidential setback. Without there being a quick fix in sight, we find that an interval dominance theory of confirmation falls short of being sufficient.

5.4 Summary

To rehearse the main findings in this chapter, we have learned of ways in which Bayesian confirmation theory may be extended by imprecise probability theory. However, we lose sight of confirmation, for it remains entirely unclear on how confirmation should be defined when imprecise probabilities are deployed in an epistemology. Although each candidate theory examined above has some benefit, we have also observed that each faces serious challenges.

So where do we go from here? As I mentioned earlier, the four candidate theories do not exhaust the space of possibilities, and so there may be some plausible definition of confirmation in imprecise probability theory that has not been considered here. However, at this time, the challenges associated with each theory suffice to show that there is no obvious plausible confirmation theory with imprecise probabilities for scientific reasoning or epistemology in general. As a result, imprecise probability theory has so far failed to accommodate a central epistemic notion that is familiar in Bayesian circles. Interestingly, a non-singleton set of probability measures can make all the difference with respect to probabilistic confirmation theory.

Of course, my judgment may be premature since an entirely plausible theory may be discovered. But even if there is no other plausible theory left to consider, a positive stance on the matter could be embraced. In particular, one of the mentioned theories may be adopted, despite there being limitations, for each theory does have unique advantages after all. Or one might instead adopt a pluralistic attitude and employ different definitions of confirmation in different contexts. Negatively, however, one may admit that (ordinal) confirmational relations are unlikely to be defined in a satisfactory way in imprecise probability theory anytime soon and eschew the idea altogether. Regardless of the route taken, we have discovered that making sense of confirmation in imprecise probability is not at all an easy task.

Conclusion

Throughout the dissertation, I suggested ways that imprecise probability might be applied in epistemology and philosophy of science. On the issues of peer disagreement and ignorance, the use of imprecise probability is justified. Confirmation, on the other hand, turned out to be troublesome. But I recommend not giving up hope and furthermore encourage the pursuit of a probabilistic confirmation theory with imprecise probabilities. If no plausible theory should exist, that, however, does not mean we should be quick to abandon imprecise probability. Such a discovery would only indicate that the model(s) has limitations like any other. That is not very troubling, at least insofar as I am concerned. Anyhow, it would be nice to end on a positive note about imprecise probabilities, and this final discussion provides a prime opportunity to do so by describing future work.

There are a number of issues that would be interesting to analyze with sets of probabilities beyond the present dissertation. I will describe just two here. The first can be seen as a relative to the problem of peer disagreement. In particular, the coherence of multiple witness reports has become an important issue in social epistemology, especially in gauging the epistemic contribution that eyewitness reports make in legal cases, for example. Bovens & Hartmann (2003) studied the issue from a purely traditional Bayesian perspective, and their account is a powerful one indeed. Furthermore, the proposal simplifies the issue by encoding reliability in likelihoods and precludes collusion or tainted testimonies through an independence assumption regarding reports R_i where $i = 1, 2, 3, \dots, n$. While the suggested approach is normative, psychologists have recently found evidence in support of sensitivity towards Bayesian aggregation of information (Harris & Hahn 2009).

The successful Bayesian analysis of testimony provides legal experts with a unique way of aggregating information. However, in considering the chapter on peer disagreement, it is entirely possible for a group of experts to arrive at a collective estimate that is imprecise, i.e. $\mathbb{P}(H) = \{x_1, \dots, x_n\}$ such that $x_1 \neq x_n$. Forget peerhood, this is quite common in legal cases given different individual information sets, bias, and ulterior motives. Unfortunately, the latter cannot be extinguished in the real world. So to safeguard against adopting a misleading opinion, the assessors of H should adopt the whole lot of opinions. In doing so, the conditional

judgments (witness reports factored in) are likely to come out imprecise assuming an imprecise prior. The story becomes more complicated at this point beyond the simple assessment yielded by Bovens and Hartmann. Studying the complications might result in novel insights on the mess trial by jury leads to.

The second issue I will mention is very recent and has attracted significant attention. It is the idea of *transformative experience* (Paul 2014). L. A. Paul describes the notion through two key factors.

EPISTEMIC TRANSFORMATIVE EXPERIENCE: An experience is epistemically transformative if the only way to know what it is like to have the experience is to actually have it yourself.

PERSONAL TRANSFORMATIVE EXPERIENCE: An experience is personally transformative if it changes your point of view, including your core preferences.

A **transformative experience** is the conjunction of the above two types of experiences.¹ Drawing on this interesting notion, Paul argues that it would be extremely difficult for an individual to make an informed choice if one lacks the relevant experience, epistemic or personal, that is salient to a decision problem. Her two popular examples are choosing between becoming a vampire or remaining human, and more realistically, choosing between having a child or not.

One would think that classical decision theory would have something to say on this matter, but classical decision theory is exactly what she targets as a failure in these kinds of situations. Without having the appropriate experience, she says, one would be challenged in assigning cardinal utilities to outcomes. To complicate matters even further, Paul illuminates a concern about personal identity, namely whether one's current self will be identical to the future self. Based on the definition of transformative experience, it seems that there will be differences between the future and current self. So how could one possibly know in the present moment what utilities the future person would assign to possible outcomes?

Metaphysics aside, there is a way to sensibly analyze the choice problems raised by Paul. On the epistemic issue, one ought to admit to being completely ignorant about what a particular experience is like if they never had such an experience. From Chapter 4, we saw how the state can be modeled. As for the personal

¹These definitions are taken directly from Paul's "Teaching guide for transformative experience" <<http://www.lapaul.org/papers/teaching-guide-for-transformative-experience.pdf>>.

issue, utilities could, in fact, be imprecise the same way credences can be imprecise (see Levi 1974). Combining the two, an individual ends up with indeterminate preferences. So in taking a step back from the classical notions that are targeted by Paul, we may broaden the framework to accommodate choice problems where one lacks the necessary experiences for making a rational decision. Now, I do not suggest that the proposal eliminates the philosophical problem. Instead, the proposal is meant to clarify what the problem is. The real problem is in justifying a decision rule. On one extreme, a rule for choosing based on indeterminate preferences might make all options permissible while on the other extreme, all options might be made impermissible. It is not clear what the right decision rule is in cases pertaining to transformative experience, and thus it would be a very fruitful avenue of research to explore, which may even be *transformative*.

Again, there are quite a number of philosophical questions that may be helpfully addressed by using imprecise probability as a formal tool. Those mentioned above are nowhere near exhaustive, and so further investigation of a generalized Bayesian epistemology should be warmly welcomed in future research.

Bibliography

- [1] Alchourrón, C.E., Gärdenfors, P., & Makinson, D. (1985). On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *The Journal of Symbolic Logic*, 50(02), 510-530.
- [2] Allais, M. (1990). Allais paradox. In *Utility and Probability*. Palgrave Macmillan.
- [3] Augustin, T., Coolen, F.P.A., de Cooman, G., & Troffaes, M.C.M. (2014). *Introduction to Imprecise Probabilities*. Wiley.
- [4] Ballantyne, N. & Coffman, E.J. (2011). Uniqueness, Evidence, and Rationality. *Philosophers' Imprint*, 11(18).
- [5] Bayes, T. (1764). An Essay Towards Solving a Problem in the Doctrine of Chances. *Philosophical Transactions of the Royal Society of London*, 53, 37?418. Reprinted in *Studies in the History of Statistics and Probability*. E.S. Pearson and M.G. Kendall (Eds.). Charles Griffin (1970).
- [6] Benétreau-Dupin, Y. (2015). The Bayesian Who Knew Too Much. *Synthese*, 192(5), 1527-1542.
- [7] Bernoulli, J. (1713/2006). *The Art of Conjecturing, Together with Letter to a Friend on Sets in Court Tennis*. Translated by Edith Sylla. Johns Hopkins University Press.
- [8] Bonjour, L. (1985). *The Structure of Empirical Knowledge*. Harvard University Press.
- [9] Boole, G. (1854). *An Investigation of the Laws of Thought: On Which are Founded the Mathematical Theories of Logic and Probabilities*. Dover Publications.

-
- [10] Bovens, L. & Hartmann, S. (2003). *Bayesian Epistemology*. Oxford University Press.
- [11] Bradley, S. (2014). Imprecise Probabilities. *Stanford Encyclopedia of Philosophy*. URL <https://plato.stanford.edu/entries/imprecise-probabilities>.
- [12] Bradley, S. & Steele, K. (2014). Uncertainty, Learning, and the “Problem” of Dilation. *Erkenntnis*, 79(6), 1287-1303.
- [13] Brier, G.W. (1950). Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review*, 78(1), 1-3.
- [14] Camerer, C. & Weber, M. (1992). Recent Developments in Modeling Preferences: Uncertainty and Ambiguity. *Journal of Risk and Uncertainty*, 5(4), 325-370.
- [15] Carnap, R. (1962). *Logical Foundations of Probability*. University of Chicago Press.
- [16] Christensen, D. (2007). Epistemology of Disagreement: The Good News. *The Philosophical Review*, 116(2), 187-217.
- [17] Christensen, D. (2009). Disagreement as Evidence: The Epistemology of Controversy. *Philosophy Compass*, 4(5), 756-767.
- [18] Christensen, D. (2010). Higher-Order Evidence. *Philosophy and Phenomenological Research*, 81(1), 185-215.
- [19] Christensen, D. (2011). Disagreement, Question-Begging, and Epistemic Self-Criticism. *Philosophers' Imprint*, 11(6).
- [20] Conee, E. & Feldman, R. (2004). *Evidentialism: Essays in Epistemology*. Clarendon Press.
- [21] Cozman, F. (2012). Sets of Probability Distributions, Independence, and Convexity. *Synthese*, 186(2), 577-600.
- [22] Crupi, V. (2013). Confirmation. *Stanford Encyclopedia of Philosophy*. URL <https://plato.stanford.edu/archives/win2016/entries/confirmation/>.
- [23] Dawid, R. (2013). *String Theory and the Scientific Method*. Cambridge University Press.

- [24] Dawid, R., Hartmann, S., & Sprenger, J. (2014). The No Alternatives Argument. *The British Journal for the Philosophy of Science*, 66(1), 213-234.
- [25] de Cooman, G. & Miranda, E. (2007). Symmetry of Models Versus Models of Symmetry. In *Probability and Inference: Essays in Honor of Henry E. Kyburg, Jr.* W. Harper and G. Wheeler (Eds.). Kings College Publications.
- [26] de Cooman, G., Miranda, E., & Zaffalon, M. (2011). Independent Natural Extension. *Artificial Intelligence*, 175(12-13), 1911-1950.
- [27] de Finetti, B. (1974). *Theory of Probability (Vols. I & II)*. Wiley.
- [28] de Finetti, B. (1989). Probabilism. *Erkenntnis*, 31(2-3), 169-223.
- [29] DeGroot, M.H. (1974). Reaching a Consensus. *Journal of the American Statistical Association*, 69(345), 118-121.
- [30] Dempster, A.P. (1967). Upper and Lower Probabilities Induced by a Multi-valued Mapping. *The Annals of Mathematical Statistics*, 38(2), 325-339.
- [31] Dietrich, F. & List, C. (2016). Probabilistic Opinion Pooling. In *Oxford Handbook of Probability and Philosophy*. A. Hájek and C. Hitchcock (Eds.). Oxford University Press.
- [32] Douglas, H. (2000). Inductive Risk and Values in Science. *Philosophy of Science*, 67(4), 559-579.
- [33] Douven, I. (2009). Uniqueness Revisited. *American Philosophical Quarterly*, 46(4), 347-361.
- [34] Douven, I. (2010). Simulating Peer Disagreements. *Studies in History and Philosophy of Science Part A*, 41(2), 148-157.
- [35] Dubois, D. & Prade, H. (1988). *Possibility Theory: An Approach to Computerized Processing of Uncertainty*. Plenum.
- [36] Easwaran, K. & Fitelson, B. (2015). Accuracy, Coherence, and Evidence. *Oxford Studies in Epistemology*, 5, 61-96.
- [37] Edwards, W. Lindman, H. & Savage, L.J. (1963). Bayesian Statistical Inference for Psychological Research. *Psychological Review*, 70(3), 193-242.

- [38] Elga, A. (2007). Reflection and Disagreement. *Noûs*, 41(3), 478-502.
- [39] Elkin, L. & Wheeler, G. (forthcoming). Resolving Peer Disagreements Through Imprecise Probabilities. *Noûs*.
- [40] Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms. *The Quarterly Journal of Economics*, 75(4), 643-669.
- [41] Feldman, R. (2005). Respecting the Evidence. *Philosophical Perspectives*, 19(1), 95-119.
- [42] Feldman, R. (2009). Evidentialism, Higher-Order Evidence, and Disagreement. *Episteme*, 6(3), 294-312.
- [43] Feldman, R. (2011). Reasonable Religious Disagreements. In *Social Epistemology: Essential Readings*. A. Goldman and D. Whitcomb (Eds.). Oxford University Press.
- [44] Filiz-Ozbay, E. & Ozbay, E.Y. (2007). Auctions with Anticipated Regret: Theory and Experiment. *The American Economic Review*, 97(4), 1407-1418.
- [45] Fine, T.L. (1988). Lower Probability Models for Uncertainty and Nondeterministic Processes. *Journal of Statistical Planning and Inference*, 20(3), 389-411.
- [46] Fitelson, B. (1999). The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity. *Philosophy of Science*, 66, S362-378.
- [47] Fitelson, B. (ms.). *Coherence*.
- [48] Fitelson, B. & Hawthorne, J. (2010). How Bayesian Confirmation Theory Handles the Paradox of the Ravens. In *The Place of Probability in Science. Boston Studies in the Philosophy of Science*, 284. Springer.
- [49] Gilboa, I. & Schmeidler, D. (1989). Maxmin Expected Utility with Non-unique Prior. *Journal of Mathematical Economics*, 18(2), 141-153.
- [50] Glymour, C. (1980). *Theory and Evidence*. Princeton University Press.
- [51] Goldberg, S.C. (2001). Testimonially Based Knowledge from False Testimony. *The Philosophical Quarterly*, 51(205), 512-526.

- [52] Goldman, A.I. (1979). What is justified belief?. In *Justification and Knowledge*. Springer.
- [53] Goldman, A. I. (1999). *Knowledge in a Social World*. Clarendon Press.
- [54] Golub, B. & Jackson, M.O. (2010). Naive Learning in Social Networks and the Wisdom of Crowds. *American Economic Journal: Microeconomics*, 2(1), 112-149.
- [55] Good, I.J. (1962). Subjective Probability as the Measure of a Non-Measurable Set. In *Logic, Methodology, and Philosophy of Science*. E. Nagel, P. Suppes, and A. Tarski (Eds.). Stanford University Press.
- [56] Good, I.J. (1968). Corroboration, Explanation, Evolving Probability, Simplicity, and a Sharpened Razor. *The British Journal for the Philosophy of Science*, 19(2), 123-143.
- [57] Gottwald, S. (2015). Many-Valued Logic. *Stanford Encyclopedia of Philosophy*. URL <https://plato.stanford.edu/entries/logic-manyvalued>.
- [58] Graham, P.J. (1997). What is testimony? *The Philosophical Quarterly*, 47(187), 227-232.
- [59] Greaves, H. & Wallace, D. (2006). Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, 115(459), 607-632.
- [60] Haenni, R., Romeijn, J.W., Wheeler, G., & Williamson, J. (2011). *Probabilistic Logics and Probabilistic Networks*. Springer.
- [61] Hájek, A. (2011). Interpretations of Probability. *Stanford Encyclopedia of Philosophy*. URL <http://plato.stanford.edu/entries/probability-interpret>.
- [62] Hájek, A. & Joyce, J. (2008). Confirmation. In *Routledge Companion to the Philosophy of Science*. S. Psillos and M. Curd (Eds.). Routledge.
- [63] Hájek, A. & Smithson, M. (2012). Rationality and Indeterminate Probabilities. *Synthese*, 187(1), 33-48.
- [64] Halmos, P.R. (1950). *Measure Theory*. Van Nostrand Reinhold Company.
- [65] Halpern, J.Y. (2003). *Reasoning About Uncertainty*. MIT Press.

- [66] Harris, A.J. & Hahn, U. (2009). Bayesian Rationality in Evaluating Multiple Testimonies: Incorporating the Role of Coherence. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(5), 1366-1373.
- [67] Hartmann, S. & Eva, B. (ms.) Bayesian Argumentation and the Value of Logical Validity.
- [68] Hartmann, S., Martini, C., & Sprenger, J. (2009). Consensual Decision-Making Among Epistemic Peers. *Episteme*, 6(02), 110-129.
- [69] Hartmann, S. & Suppes, P. (2010). Entanglement, Upper Probabilities and Decoherence in Quantum Mechanics. In *EPSA Philosophical Issues in the Sciences*. Springer.
- [70] Hawthorne, J. (2009). The Lockean Thesis and the Logic of Belief. In *Degrees of Belief*. F. Huber and C. Schmidt-Petri (Eds.). Springer.
- [71] Hempel, C.G. (1945). Studies in the Logic of Confirmation. *Mind*, 54(213), 1-26, 97-121.
- [72] Horsten, L. & Douven, I. (2008). Formal Methods in the Philosophy of Science. *Studia Logica*, 89(2), 151-162.
- [73] Howson, C. & Urbach, P. (2005). *Scientific Reasoning: The Bayesian Approach*. 3rd edition. Open Court.
- [74] Hume, D. (1888). *A Treatise of Human Nature*. L.A. Selby-Bigge (Ed.). Clarendon Press.
- [75] James, W. (1896). *The Will to Believe: And Other Essays in Popular Philosophy*. Longmans, Green, and Company.
- [76] Jaynes, E.T. (1957). Information Theory and Statistical Mechanics. *Physical Review*, 106(4), 620-630.
- [77] Jeffrey, R. (1965). *The Logic of Decision*. University of Chicago Press.
- [78] Jeffrey, R. (1983). Bayesianism with a Human Face. In *Testing Scientific Theories, Minnesota Studies in the Philosophy of Science*, 10, 133-156.
- [79] Jehle, D. & Fitelson, B. (2009). What is the "Equal Weight View"? *Episteme*, 6(3), 280-293.

- [80] Joyce, J.M. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science*, 65(4), 575-603.
- [81] Joyce, J.M. (2005). How Probabilities Reflect Evidence. *Philosophical Perspectives*, 19(1), 153-178.
- [82] Joyce, J.M. (2010). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives*, 24(1), 281-323.
- [83] Kahneman, D., Knetsch, J., & Thaler, R. (1990). Experimental Tests of the Endowment Effect and the Coase Theorem. *Journal of Political Economy*, 98(6), 1325-1348.
- [84] Kahneman, D. & Tversky, A. (1979). Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, 47(2), 263-292.
- [85] Kelly, T. (2011). Peer Disagreement and Higher Order Evidence. In *Social Epistemology: Essential Readings*. A. Goldman and D. Whitcomb (Eds.). Oxford University Press.
- [86] Kelly, T. (2013). Disagreement and the Burdens of Judgment. In *The Epistemology of Disagreement: New Essays*. D. Christensen and J. Lackey (Eds.). Oxford University Press.
- [87] Keynes, J.M. (1921). *A Treatise on Probability*. Macmillan & Co.
- [88] Knight, F.H. (1921). *Risk, Uncertainty, and Profit*. Hart, Schaffner and Marx.
- [89] Kolmogorov, A.N. (1933/1956). *Foundations of the Theory of Probability*. Translated by N. Morrison. Chelsea Publishing Company.
- [90] Konek, J. (forthcoming). Epistemic Conservativity and Imprecise Credence. *Philosophy and Phenomenological Research*.
- [91] Koopman, B.O. (1940). The Axioms and Algebra of Intuitive Probability. *Annals of Mathematics*, 41(2), 269-292.
- [92] Kopec, M. (2015). A Counterexample to the Uniqueness Thesis. *Philosophia*, 43(2), 403-409.
- [93] Kyburg, H.E. & Pittarelli, M. (1996). Set-based Bayesianism. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 26(3), 324-339.

- [94] Lackey, J. (1999). Testimonial Knowledge and Transmission. *The Philosophical Quarterly*, 49(197), 471-490.
- [95] Lasonen-Aarnio, M. (2013). Disagreement and Evidential Attenuation. *Noûs*, 47(4), 767-794.
- [96] Lasonen-Aarnio, M. (2014). Higher-Order Evidence and the Limits of Defeat. *Philosophy and Phenomenological Research*, 88(2), 314-345.
- [97] Lehrer, K. (1976). When Rational Disagreement is Impossible. *Noûs*, 10(3), 327-332.
- [98] Lehrer, K. & Wagner, C. (1983). Probability Amalgamation and the Independence Issue: A Reply to Laddaga. *Synthese*, 55(3), 339-346.
- [99] Leitgeb, H. & Pettigrew, R. (2010a). An Objective Justification of Bayesianism I: Measuring Inaccuracy. *Philosophy of Science*, 77(2), 201-235.
- [100] Leitgeb, H. & Pettigrew, R. (2010b). An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy. *Philosophy of Science*, 77(2), 236-272.
- [101] Levi, I. (1961). Decision Theory and Confirmation. *Journal of Philosophy*, 58(21), 614-625.
- [102] Levi, I. (1974). On Indeterminate Probabilities. *The Journal of Philosophy*, 71(13), 391-418.
- [103] Levi, I. (1980). *The Enterprise of Knowledge*. MIT Press.
- [104] Levi, I. (1990). *Hard Choices: Decision Making Under Unresolved Conflict*. Cambridge University Press.
- [105] Levinstein, B.A. (2015). With All Due Respect: The Macro-Epistemology of Disagreement. *Philosophers' Imprint*, 15(13), 1-20.
- [106] Lewis, D. (1986). A Subjectivist's Guide to Objective Chance. In *Philosophical Papers, Vol. 2*. Oxford University Press.
- [107] Loomes, G. & Sugden, R. (1982). Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty. *The Economic Journal*, 92(368), 805-824.

- [108] Mayo, D. G. (1996). *Error and the Growth of Experimental Knowledge*. University of Chicago Press.
- [109] Miranda, E. (2008). A Survey of the Theory of Coherent Lower Previsions. *International Journal of Approximate Reasoning*, 48(2), 628-658.
- [110] Miranda, E. & Cooman, G.D. (2014). Lower Previsions. In *Introduction to Imprecise Probabilities*. T. Augustin, F.P.A. Coolen, G. de Cooman, and M.C.M. Troffaes (Eds.). Wiley.
- [111] Miranda, E. & de Cooman, G. (2014). Structural Judgements. In *Introduction to Imprecise Probabilities*. T. Augustin, F.P.A. Coolen, G. de Cooman, and M.C.M. Troffaes (Eds.). Wiley.
- [112] Moore, G.E. (1939). *Proof of an External World*. H. Milford.
- [113] Moss, S. (2011). Scoring Rules and Epistemic Compromise. *Mind*, 120(480), 1053-1069.
- [114] Moss, S. (2015). Time-slice Epistemology and Action Under Indeterminacy. In *Oxford Studies in Epistemology*, 5. T.S. Gendler and J. Hawthorne (Eds.). Oxford University Press.
- [115] Norton, J.D. (2007a). Disbelief as the Dual of Belief. *International Studies in the Philosophy of Science*, 21(3), 231-252.
- [116] Norton, J.D. (2007b). Probability Disassembled. *The British Journal for the Philosophy of Science*, 58(2), 141-171.
- [117] Norton, J.D. (2008). Ignorance and Indifference. *Philosophy of Science*, 75(1), 45-68.
- [118] Norton, J.D. (2010). Cosmic Confusions: Not Supporting Versus Supporting Not. *Philosophy of Science*, 77(4), 501-523.
- [119] Norton, J.D. (2011). Challenges to Bayesian Confirmation Theory. *Philosophy of Science*, Vol. 7: *Handbook of the Philosophy of Science*. P.S. Bandyopadhyay and M.R. Forster (Eds.). Elsevier.
- [120] Paul, L.A. (2014). *Transformative Experience*. Oxford University Press.

- [121] Pedersen, A.P. & Wheeler, G. (2014). Demystifying Dilation. *Erkenntnis*, 79(6), 1305-1342.
- [122] Pettigrew, R. (2012). Accuracy, Chance, and the Principal Principle. *The Philosophical Review*, 121(2), 241-275.
- [123] Pettigrew, R. (2013). Epistemic Utility and Norms for Credences. *Philosophy Compass*, 8(10), 897-908.
- [124] Pettigrew, R. (2014). Accuracy, Risk, and the Principle of Indifference. *Philosophy and Phenomenological Research*, 92(1), 35-59.
- [125] Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford University Press.
- [126] Ramsey, F.P. (1926). Truth and Probability. In *The Foundations of Mathematics and other Logical Essays*. R.B. Braithwaite (Ed.). Kegan, Paul, Trench, Trubner, and Co.
- [127] Rinard, S. (2013). Against Radical Credal Imprecision. *Thought: A Journal of Philosophy*, 2(2), 157-165.
- [128] Rinard, S. (2015). A Decision Theory for Imprecise Probabilities. *Philosophers' Imprint*, 15(7).
- [129] Rosen, G. (2001). Nominalism, Naturalism, Epistemic Relativism. *Noûs*, 35(15), 69-91.
- [130] Rosenkrantz, R.D. (1977). *Inference, Method and Decision: Towards a Bayesian Philosophy of Science (Vol. 115)*. Springer.
- [131] Savage, L.J. (1951). The Theory of Statistical Decision. *Journal of the American Statistical Association*, 46(253), 55-67.
- [132] Savage, L.J. (1954). *The Foundations of Statistics*. Wiley.
- [133] Schoenfield, M. (2014). Permission to Believe: Why Permissivism is True and What It Tells Us About Irrelevant Influences on Belief. *Noûs*, 48(2), 193-218.
- [134] Seidenfeld, T. & Wasserman, L. (1993). Dilation for Sets of Probabilities. *The Annals of Statistics*, 21(3), 1139-1154.

- [135] Seidenfeld, T., Schervish, M.J., & Kadane, J.B. (2012). Forecasting with Imprecise Probabilities. *International Journal of Approximate Reasoning*, 53(8), 1248-1261.
- [136] Smith, C.A.B. (1961). Consistency in Statistical Inference and Decision. *Journal of the Royal Statistical Society. Series B (Methodological)*, 23(1), 1-37.
- [137] Stewart, R. & Quintana, I. (forthcoming) Probabilistic Opinion Pooling with Imprecise Probabilities. *Journal of Philosophical Logic*.
- [138] Stone, M. (1961). The Opinion Pool. *The Annals of Mathematical Statistics*, 32(4), 1339-1342.
- [139] Sturgeon, S. (2008). Reason and the Grain of Belief. *Noûs*, 42(1), 139-165.
- [140] Sturgeon, S. (2010). Confidence and Coarse-Grained Attitudes. In *Oxford Studies in Epistemology*, 3. T.S. Gendler and J. Hawthorne (Eds.). Oxford University Press.
- [141] Surowiecki, J. (2005). *The Wisdom of Crowds*. Anchor.
- [142] Teller, P. (1973). Conditionalization and Observation. *Synthese*, 26(2), 218-258.
- [143] Thaler, R. (1980). Toward a Positive Theory of Consumer Choice. *Journal of Economic Behavior and Organization*, 1(1), 39-60.
- [144] Tversky, A. & Kahneman, D. (1991). Loss Aversion in Riskless Choice: A Reference-Dependent Model. *The Quarterly Journal of Economics*, 106(4), 1039-1061.
- [145] van Fraassen, B.C. (1984). Belief and the Will. *The Journal of Philosophy*, 81(5), 235-256.
- [146] van Fraassen, B.C. (1990). Figures in a Probability Landscape. In *Truth or Consequences*. Springer.
- [147] von Neumann, J. & Morgenstern, O. (2007). *Theory of Games and Economic Behavior*. Princeton University Press.
- [148] Wagner, C. (1978). Consensus Through Respect: A Model of Rational Group Decision-Making. *Philosophical Studies*, 34(4), 335-349.

- [149] Walley, P. (1981). Coherent Lower (and Upper) Probabilities. Statistics Research Report 22, University of Warwick, Coventry.
- [150] Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall.
- [151] Walley, P. (2000). Towards a Unified Theory of Imprecise Probability. *International Journal of Approximate Reasoning*, 24(2-3), 125-148.
- [152] Wheeler, G. (2012). Objective Bayesian, Calibration, and the Problem of Non-Convex Evidence. *The British Journal for the Philosophy of Science*, 63(4), 841-850.
- [153] White, R. (2009). Evidential Symmetry and Mushy Credence. In *Oxford Studies in Epistemology*. T. Szabo Gendler and J. Hawthorne (Eds.). Oxford University Press.
- [154] Williams, B. (1979). Internal and External Reasons. In *Rational Action*. R. Harrison (Ed.). Cambridge University Press.
- [155] Williams, P.M. (1975). Notes on Conditional Previsions. School of Mathematical and Physical Sciences, University of Sussex. Reprinted in *International Journal of Approximate Reasoning*, 44(3), 366-383, 2007.
- [156] Williamson, J. (2010). *In Defence of Objective Bayesianism (Vol. 24)*. Oxford University Press.
- [157] Williamson, J. (2015). Deliberation, Judgement and the Nature of Evidence. *Economics and Philosophy*, 31(1), 27-65.
- [158] Wilson, A. (2010). Disagreement, Equal Weight, and Commutativity. *Philosophical Studies*, 149(3), 321-326.
- [159] Wilson, N. (2001). Modified Upper and Lower Probabilities Based on Imprecise Likelihoods. In *Proceedings of the 2nd International Symposium on Imprecise Probabilities and Their Applications*.