London School of Economics and Political Science

Towards a Pluralistic View of Formal Methods

Ko-Hung Kuan

A thesis submitted to the Department of Philosophy, Logic, and Scientific Method
at the London School of Economics for the degree of Doctor of Philosophy
November 2020

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work.

I confirm that the literature review in the second chapter partly overlaps in general content with a chapter in my MSc thesis 'Coherence Preservation: A Threat to Probabilistic Measures of Coherence', where I provide an overview of the discussion. The positive proposal in this chapter is entirely new. I also confirm that the third chapter 'Beyond Linear Conciliation' is published in Synthese as (Kuan, 2020).

I declare that this thesis consists of 58,672 words.

Abstract

This thesis is a collection of three self-contained papers on related themes in the area of formal and social epistemology. The first paper explores the possibility of measuring the coherence of a set with multiplicative averaging. It has been pointed out that all the existing probabilistic measures of coherence are flawed for taking the relevance between a set of propositions as the primary factor which determines the coherence of the set. What I show in this paper is that a group of measures, namely the confirmation-based ones, can be saved from this problem if we adopt a nonlinear averaging function to measure the coherence of a set.

The second paper discusses how people should conciliate in disagreements. Some epistemologists take linear averaging as the only way of conciliating and claim that conciliating leads to fallacious results. In the paper, I show that the problem is not conciliating, but taking linear averaging as the only way to conciliate. Since there is no reason for us to insist on conciliating with linear averaging, we can adopt nonlinear averaging functions for conciliating and thereby avoid the formal deficiencies.

The third paper focuses on the pragmatic results of taking conciliating as a general strategy in disagreements. There is a potential dilemma about conciliating: if everyone always conciliates, it is likely for an epistemic bubble to arise. If everyone refuses to conciliate, an epistemic echo chamber may appear. A possible way of solving the dilemma is to develop a diachronic strategy which tells people how to both conciliate and update their estimate of their interlocutors' reliability. Although the three papers differ in the subject, they jointly offer some unifying reflections on the way we approach philosophical problems with formal tools. From

the first two papers, we see that a formal analysis of a philosophical position is incomplete if philosophers fail to consider a sufficiently wide range of formal tools. The third paper, on the contrary, shows that we should change the ordinary way of modelling a notion if required. This thesis concludes by proposing a pluralistic view of formal methods in formal epistemology.

Acknowledgements

Pursuing a PhD was a magical mystery tour. The experience would not have been so wonderful without many people. Among them, the most important ones are my supervisors, Anna Mahtani and Christian List. It would not be possible for me to finish this journey without their strong support throughout the years. They showed me how to do good philosophy and, more importantly, how to be an excellent teacher.

LSE is an amazing place for doing a PhD in Philosophy. There is a positive atmosphere within the department which encourages people to discuss all kinds of ideas. Thanks to Jason Alexander, Jonathan Birch, Richard Bradley, Liam Kofi Bright, Laurenz Hudetz, Lewis Humphreys, Becky Matthams, Andrea Pawley, Mike Otsuka, Miklos Redei, Brian Roberts, Ewan Rodgers, Johanna Thoma and Mary Wells for their help in all kinds of ways.

My fellow PhD students brought me a different type of learning experience. Joe Roussos, Nicolas Côté and Bastian Steuwer were my teammates in the Lunch Squad. I got to know lots of things at lunchtime every day, ranging from the definition of a sandwich to the current situation in Kashmir. I would also like to thank Fabian Beigang, Charles Beasley, Chloé de Canson, Paul Daniell, Christina Easton, Margherita Harris, Todd Karhu, Sophie Kikkert, David Kinney, Deren Olgun, Tom Rowe, Nicholas Makins, Chris Marshall, Silvia Milano, Aron Vallinder, Philippe van Basshuysen, Cecily Whiteley and James Wills for the fun time we had together.

Friends from Taiwan provided strong emotional support. Thanks to Hsuan-Chih Lin, Meng-Hsuan Yu, Tien-Chun Lo, Jay Jian, Dorothy Fang-En Chiang, Tony

H. Y. Cheng and Chung-Tang Cheng. The chance of speaking Taiwanese Hokkien in London is very precious. For a similar reason, I would like to thank Kevin Chun-Man Kwong, Matthew Man-Him Ip, Chun-Yin Salt Yeung and Howard Mok. It is really a waste of time hanging out with you guys, but I enjoyed it a lot.

My study in the UK was funded by the Ministry of Education of Taiwan. Hence, I would like to thank the Taipei Representative Office in the UK for taking care of everything. My guarantor and former supervisor, Syraya Chin-Mu Yang, supported me over the years and made this intellectual journey possible. It is hard to describe how grateful I am.

There are some other friends whom I chat with frequently, including Bing-Cheng Huang, Pei-Hua Huang, Yu-Ting Chou, Yao-Cheng Chang, Yu-Wei Liao, Ting-Hsuan Yu, Wei-Ting Lin, Simon Liu, Sylvia Tung Lin, Wei-Yi Chiang, Kais Wang, Wayne Chen, Cheng-De Lin, Xsinnski Ho and Ten-Herng Lai. Small talks with these people eased my pressure a lot.

Although my parents have no idea what I am working on, they have never doubted my decision. If I were them, I would probably dissuade my son from doing philosophy. I am glad that they are much wiser than I am. My sister Yen-Ting took care of family affairs during my days abroad. I truly appreciate their unconditional and unfounded trust.

Most importantly, thanks to my partner Tzu-Chi for her seemingly unlimited patience and tolerance. I believe that there are more challenges in the years to come. I also believe that we can overcome any difficulties together.

Contents

List of Figures

List of Tables

Chapter 1

Epistemology with numbers

A growing trend in contemporary epistemology is to reformulate classical issues with formal apparatuses. With the help of a variety of tools ranging from modal logic to decision theory, epistemologists not only reviewed the extensively discussed problems from a new perspective, but also discovered many new questions to work on. However, as novelty does not guarantee superiority, epistemologists need to justify this new approach by scrutinising its crucial features.

What are the advantages of doing philosophy with formal tools? The most apparent one is that these tools allow us to discuss problems with greater clarity and precision. To illustrate, imagine a case where two agents face a highly questionable claim and need to decide whether to act upon it. While both of them neither believe nor repudiate the claim, the first agent is a bit more confident in the claim than the second. Without some kind of formal tool, it is hard to correctly capture the exact nature of this case. Traditionally, epistemologists adopt a tripartite framework of beliefs, under which an agent can only be in three doxastic states concerning a proposition: believe, disbelieve or suspend judgement. Since both agents in the example neither believe nor repudiate the claim, they can only be taken as suspending judgement regarding the claim. The difference between them, hence, is beyond the expressive power of a tripartite framework. For the traditional framework to correctly represent the case, we might expand it with a new state of 'slightly more confident than suspending judgement'. Without a fourth doxastic state, one cannot correctly represent the difference between the two agents under a traditional framework.

Expanding a framework piecemeal in this fashion by adding more doxastic states is not really a satisfactory solution. If, in the example given, there is a third agent who is even more confident than the first agent yet still does not believe the questionable claim wholeheartedly, we would have to expand the framework with one more state to represent the doxastic state of this new agent. We can thus see that a traditional framework always runs the risk of being incomplete. No matter how many different doxastic states are included, it is always possible for a traditional framework to lack the proper doxastic state to represent a case correctly.

Compared to the traditional framework, a probabilistic framework, namely one which formalises the notion of belief with probability theory, provides an easier solution. Under a probabilistic framework, we can take the agents as having different degrees of belief and reformulate their doxastic states with numbers in between 0 and 1. When we say that the first agent, compared to second one, is more confident in the disputed claim, one may take her as having greater credence than the second agent. By formulating their doxastic states this way, we do not have to expand the framework every time a new doxastic state appears. As we can see, formal tools provide us with a more fine-grained framework and make it possible for us to carry out detailed discussions.

Apart from being good apparatuses for descriptive purposes, doing philosophical research with formal tools brings us another significant benefit. When we reformulate a notion with a formal theory, the constraints that can be derived from the theory could be taken as the constraints of the notion. In other words, from the formal theories we adopt to formalise philosophical notions, we can derive norms governing the notions we aim to capture.

We may thus see that the formal tools play two roles in philosophical research. On the one hand, they provide us with frameworks that are descriptively more accurate which facilitate and deepen our discussion concerning philosophical concepts. On the other hand, they may serve as additional sources generating norms about the notions we modelled. For these reasons, formal philosophy should be taken as a promising approach which generates significant results.

There are, however, some potential worries that accompany this approach. When one formulates a philosophical view with some formal tools and finds the view formally incorrect, it is not only possible that the theory is indeed wrong, but also well possible that the formal tools adopted are not the right tools for formulating the theory and hence bring up the problem. If we fail to spot where the problem really lies, we may attain an incorrect understanding of the theory modelled. Adopting formal tools to do philosophy, instead of bringing us forward in philosophical research, may lead us astray.

To avoid getting such an undesirable result, we should explore a sufficiently wide range of formal apparatuses when dealing with philosophical disputes. Then if we have formulated a theory with the ideal formal tool and obtain the consequence that the theory, formulated in the correct way, is mistaken, we may safely claim that the theory is indeed incorrect.

This thesis consists of three essays in formal epistemology, each of which individually contributes to a particular formal challenge in epistemology. Jointly, the essays highlight the fact that problems may arise when epistemologists neglect alternative ways of modelling a notion. In chapters two and three, I examine two formal approaches to issues in both traditional and social epistemology. With the cases I present, it can be seen that some philosophers overlook the deficiencies of the formal apparatuses they adopt and thus mistakenly take the incorrectly formalised philosophical position as problematic. Chapter four, compared to the other chapters, looks a bit like an outlier. It does not show that epistemologists mistakenly take the problem of a formal tool as the problem of a philosophical position. Instead, what it shows is that we should try out different ways of formalising the well known notions and explore possible ways to solve the problems surrounding them. Still, it is in accordance with the basic tenet of this thesis that we have to carefully reflect on the way we formulate philosophical notions, even for those that we are familiar with, and make substantial revisions when needed.

In conclusion, the overall upshot of the project, apart from making progress on three important problems in formal epistemology, is to call attention to the importance of exploring a variety of formal apparatuses. With a complete understanding of the formal tools we apply, we can correctly spot where the problem lies and tell whether it is the philosophical theory or the tool for formalising that is flawed. In the following sections, I will briefly introduce the three projects in this thesis.

1.1 Saving the confirmation-based measures of coherence

Philosophers have been trying to characterise the notion of coherence for decades. With this notion correctly defined, we may have a better understanding of both the notion of truth and the notion of justification. Some philosophers approach this issue from a comparative perspective and aim to establish a way of measuring the coherence of a set of propositions (Shogenji, 1999; Olsson, 2002; Fitelson, 2003; Douven and Meijs, 2007; Roche, 2013). If, according to an ideal measure of coherence, a greater degree of coherence of a set indicates a greater likelihood of truth, we may obtain a theory of truth based on the notion of coherence. Similarly, if a greater degree of coherence of a set indicates a greater degree of justification, we may account for the notion of justification in terms of coherence.

Most measures of coherence take the *relevance* between the propositions as the core factor determining the coherence of a set. If the contents of a set of propositions overlap to a great extent, we take the propositions as highly relevant and thus highly coherent. However, Koscholke and Schippers (2019) point out that the relevance-sensitive measures are flawed for failing to deal with cases where a common cause of a set appears. Given a set of relevant propositions, when a common cause of all these propositions appears and taken as background knowledge, the relevance between these propositions would be *screened-off* (Reichenbach, 1956). If we adopt the relevance-sensitive measures to calculate the coherence of such a set, we will get the counterintuitive result that the propositions are irrelevant to each other, and, since coherence is measured in terms of relevance, the set is neither coherent not incoherent. Moreover, if we, instead of taking a common cause as background knowledge, expand the set with the common cause, we would get another problematic result that a set may become less coherent when a common cause appears. Based on the two observations, Koscholke and Schippers conclude that epistemologists should give up the relevance-sensitive measures.

In this second chapter, I provide an overview of the search for an ideal probabilistic measure of coherence and raise some issues with the measures that have been established. It should be noted that, as coherence measures were the focus of my MSc dissertation, there is inevitably some overlap in the background discussion of existing coherence measures. With a thorough understanding of the literature, I will move on to reexamine the two problems Koscholke and Schippers (2019) raised. The crucial problem, as I will show, is not that the coherence of a set is measured in terms of relevance, but that the function we endorse lacks some important features. If we adopt a different averaging function which bears the required property to derive the result, the problem of common cause can be solved. As a consequence, it should still be allowed to measure the coherence of a set in terms of the relevance between the propositions in the set.

Koscholke and Schippers' criticism is a typical case in which epistemologists neglect the possibility of adopting a different formal apparatus and end up with a hasty conclusion. If we are aware of the blind spot in their reasoning and consider a sufficiently wide range of formal tools, we may correctly locate the problem and reach a more moderate, yet more accurate conclusion. The project of searching for an appropriate measure of coherence, thus, can be saved.

1.2 Beyond Linear Conciliation

An extensively discussed problem in the study of social epistemology is peer disagreement: when a person disagrees with her epistemic peer, how should she react? The epistemologists who endorse the Conciliatory View advise one to conciliate with the peer. Since one's interlocutor is one's epistemic peer who is equally likely to form a correct credence in the disputed claim, one should be epistemically modest and revise one's credence. On the contrary, some others claim that one should remain steadfast. When one's interlocutor forms a credence which differs from one's own credence, one may think that the interlocutor suffers from some cognitive defect. Hence, one should remain steadfast in the face of the disagreement. Since both views are supported by some strong arguments, the debate over an ideal solution to peer disagreement has not yet been settled.

In chapter three, I focus on a series of arguments against the Conciliatory View based on its formal features. Formal epistemologists criticise this view for a number of reasons. It is non-commutative with conditionalisation; it is path dependent, and it does not preserve the independence between propositions (Fitelson and Jehle, 2009; Gardiner, 2014; Elkin and Wheeler, 2018a). Failing to commute with conditionalisation means that one may switch the order between conciliating and conditionalising and obtain different outcomes. Failing to be path independent means that the outcome of conciliation varies with the order of the acquisition of new testimonies. Failing to preserve the independence between propositions means that one may suffer from a sure-loss and hence be deemed irrational. The three formal deficiencies urge people to abandon the Conciliatory View.

What I aim to show in this chapter is that the Conciliatory View can be saved if we conciliate with nonlinear averaging functions. Research in the study of opinion pooling shows that the three deficiencies are not problems of the Conciliatory View, but problems of linear averaging (Genest, 1984; Dietrich and List, 2016). Hence, one can get rid of these formal deficiencies by making conciliation with nonlinear averaging functions. After showing how the three deficiencies can be avoided, I will explore the features of nonlinear averaging functions and argue that they have properties that correctly capture people's intuitions concerning disagreement. The conclusion, therefore, is to suggest epistemologists develop a more fine-grained taxonomy for cases of disagreement. With a deliberate categorisation of different kinds of disagreement, epistemologists can pick the proper averaging rule to apply in each specific case, and eliminate possible formal deficiencies.

Chapter three again provides a case where epistemologists misfire. The real target of their argument is the proposition that the Conciliatory View is formally deficient if we conciliate with linear averaging, rather than the much stronger proposition that making conciliation, in general, leads to formal fallacies. By emphasising this fact, it can be shown that, for an argument concerning the formal features of a philosophical position to be valid, we must carefully consider sufficiently many ways of formalising the position.

1.3 Escaping an Echo Chamber

One way of evaluating different views concerning peer disagreement is to see the consequences of adopting each view. For example, if adopting the Conciliatory View leads to a defective epistemic community in which people are vulnerable to misinformation, we would have a reason to reject the Conciliatory View. A striking result is that when a group of people adopt the Conciliatory View, it would be quite likely for the group to form an epistemic bubble, namely a community with insufficient exposure to a diverse set of information sources. However, if one rejects the Conciliatory View and adopt the Steadfast View instead, it would be extremely likely for one to end up in an epistemic echo chamber, a community in which the members deem every external source of information unreliable. Hence, there seems to be a dilemma concerning whether to conciliate.

I will unpack this dilemma in chapter four. According to Christensen (2011), the mainstream views concerning peer disagreement can be categorised into two groups: the ones which suggest one to conciliate and the ones which do not. The crucial distinction between them is marked by the *Principle of Independence*: one needs to have a dispute-independent reason to deem one's interlocutor unreliable. The views conforming to this principle are the conciliatory ones, while the views violating it can be categorised as variants of the Steadfast View. By narrowing down the debate to the principle, we may, according to Nguyen's (2018) analyses of an epistemic bubble and an echo chamber, show that the Principle of Independence leads to a dilemma. After the dilemma is introduced, I will provide a possible solution based on an alternative understanding of reliability. If one follows the new strategy, the probability of one ending up in a defective community could be reduced. We may hence discuss the problem of disagreement from a new perspective and hopefully derive better solutions.

The core of my solution is an alternative formulation of the notion of reliability. The standard treatment of this notion takes it as the probability of one having the correct doxastic state. This formulation, however, leads to some strange results and thus cannot fully capture our ordinary understanding of reliability. To solve the problem, I propose a different understanding of the notion which, on the one hand, better fits the formal framework we adopt and, on the other hand, sheds light on a new response to the problem of peer disagreement. Again, the discussion about reliability shows that what we need is to consider different possible ways of formalising a philosophical notion. If we stick to the standard treatment, it would be hard to solve the dilemma. We may thus reaffirm the central tenet of this thesis that we must always carefully pick the formal apparatus when dealing with a philosophical problem.

Chapter 2

Saving the Confirmation-based Coherence Measures

2.1 Introduction

The notion of coherence has long played a central role in philosophy. On the one hand, some philosophers appeal to this notion to provide an account of truth, claiming that a true proposition must cohere with other true propositions. A highly coherent set, therefore, is very likely to be true. On the other hand, some take this notion to explain the nature of epistemic justification, arguing that a proposition is justified only if it is an element of a coherent set. Due to its philosophical significance, philosophers made various attempts to clarify the nature of the notion of coherence. This chapter focuses on one specific question: what is the proper way for us to compare the degree of coherence between different sets of propositions?

It should be noted that this question implicitly takes coherence as a graded notion. That is, the coherence of a set is not an all-or-nothing notion, but comes in different degrees. What we would like to obtain is a proper method to compare the degree of coherence between different sets of propositions. If we can find an ideal measure to correctly capture this notion, we may further explore the features of the notion of coherence and derive philosophically significant results from the method developed.

Many philosophers have tried to answer this question (Shogenji, 1999; Olsson, 2002; Fitelson, 2003; Douven and Meijs, 2007; Roche, 2013). However, all these measures are flawed in some aspects and thus fail to capture some of our intuitive understanding of coherence. Moreover, Koscholke and Schippers (2019) point out that all these measures yield an incorrect result if they take the *relevance* between the elements of a set as a factor determining the coherence of the set: given a set of propositions, when a common cause of the set of propositions appears, the relevance between the propositions would be *screened-off*. Hence, if we take the relevance of a set of propositions as a factor determining the coherence of the set, we would have to accept the counterintuitive result that, once the common cause of a set appears, the degree of coherence of the set becomes zero, which implies that the set is neither coherence nor incoherent. As most coherence measures take the relevance between propositions as the crucial factor determining the coherence of a set, Koscholke and Schippers' challenge seems to destroy the project of measuring coherence.

Although Koscholke and Schippers' argument looks convincing, it does not bring an end to the search for a proper measure of coherence. The crucial problem of the relevance-sensitive measures of coherence, as I will show, is that the number of mutual confirmations between propositions is not taken into account. If we can develop a measure which generates the result that the degree of coherence of a set increases with the number of mutual confirmation between elements, we may get rid of Koscholke and Schippers ' challenge.

In the following sections, I will first make a thorough review of the role of coherence in contemporary epistemology and introduce several traditional accounts concerning this notion.[1] With a complete survey of the defining features of this notion, I will move on to reexamine previous attempts at measuring coherence in

[1]Since the coherence measures were the focus on my MSc dissertation, part of the review of existing coherence measures in this chapter inevitably overlaps with what I wrote in my previous work. However, this section has all been developed, refined and extended. Also, the positive proposal in this chapter is entirely new.

terms of the probability of propositions. All these attempts, however, are unsatisfactory in some respects. After the relevant issues are clearly presented, I will propose a new way of measuring coherence which satisfies most of our intuitive requirements of an ideal coherence measure and, more importantly, is free from Koscholke and Schippers' criticism. By measuring coherence with this new measure, we can secure the project of measuring coherence and derive further results that are epistemically significant.

2.2 Two uses of coherence

2.2.1 The coherence theory of truth

The coherence theory of truth was originally proposed as an alternative to the more widely accepted correspondence theory of truth. To gain a thorough understanding of the coherence theory of truth and grasp the motivating idea behind it, we should begin with a comparison between the two theories.[2] The basic idea of the correspondence theory of truth is rather straightforward: for a proposition to be true, it must stand in a specific relation, namely correspondence, to some entities in reality. According to this theory, for the proposition 'F. Scott Fitzgerald is the author of *The Great Gatsby*' to be true, there must be a fact, namely the very fact that Fitzgerald wrote *The Great Gatsby*, which corresponds to the proposition and makes it true.[3] In other words, the truth-condition of a proposition is the obtaining of the corresponding fact in reality. Given this account, we know what it means for a proposition to be true.

One of the primary obscurities of the correspondence theory lies in the mysterious relation of correspondence. How does a proposition, an abstract entity, correspond to something in reality? For entities in the same category, it is compar-

[2]There are other theories of truth, such as pragmatic and deflationary. Since we do not have to know these theories to understand the underlying rationale of the coherence theory, there is no need to introduce these theories here.

[3]Recent discussions of the correspondence theory of truth take *truthmakers*, rather than facts, as the entity corresponded by the truth-bearer (Mulligan et al., 1984). For the sake of simplicity, I reformulate the view in its rudimentary form which takes facts as the truthmakers in reality.

atively easy to figure out the relation between them. Suppose there is a ball and a table. The ball may be on the table, beneath it, or stand in some other relations to the table. Since the table and the ball belong to the same category, people generally do not find the relation between them confusing. Similarly, it is not hard for people to sort out the relation between propositions. A proposition can bear several different relations to another. It can be the cause, the logical consequence or independent from another proposition. All these relations are complicated, but not as obscure as the cross-category correspondence relation between a proposition and a fact. Propositions are linguistic entities, while facts are not.[4] Since they belong to different ontological categories, the main challenge for the correspondence theorists is to explain the way propositions correspond to facts. Instead of trying to provide an account for the correspondence relation, some philosophers approach the notion of truth in a different way. They give up the idea that a proposition can bear a relation to entities in a different category and, as a result, embrace the view that the truth-condition of a proposition consists in its relation to other entities in the same category, namely other propositions. They claim that if a set of propositions 'hang together well', then they are true. By expanding this idea, philosophers develop the coherence theory of truth.

How does the notion of coherence account for the notion of truth? The coherentists argue that true propositions cannot contradict each other. For example, the true proposition that 'The English army won the Battle of Agincourt' is compatible with other true propositions about the Battle of Agincourt, such as 'The French were defeated in the Battle of Agincourt.' On the other hand, it is incompatible with false propositions like 'The English army did not win the Battle of Agincourt.' Following this line of reasoning, all the propositions that are true should form a set such that all the elements of it can be simultaneously true. Moreover, since all these true propositions describe the same reality, they would provide evidential

[4]Here I take propositions as the truth-bearer in the current discussion, which are linguistic entities. Some philosophers might criticise my formulation of the correspondence theory by claiming that propositions belong to the same category as facts. My response to this criticism is that whatever one takes the truth-bearers to be, they belong to a different category from the entities they correspond to.

support for each other. Consider a toy example:

(1) The English army won the Battle of Agincourt.

(2) King Henry V of England led the English troops into the Battle of Agincourt.

(3) King Henry V commanded well in the Battle of Agincourt.

None of these propositions entails another but they do confirm each other to some extent. Given (2) and (3), (1) becomes more likely to be true. We may, from this example, see that a set of true propositions that correctly depict reality should support each other and be very coherent. Hence, if a set of propositions cohere well with a true proposition, these propositions are likely to be true. The notion of coherence, hence, can be an useful indicator for the truth of some propositions. Once we find a highly coherent set which contains a true proposition, it is likely that all other propositions in the set, apart from the one known to be true, are also true.[5]

2.2.2 The coherence theory of justification

Epistemologists who give an account of epistemic justification in terms of coherence are motivated by a different debate. They aim to answer a fundamental question in epistemology: under what condition can we say that a proposition is justified? A straightforward answer to this question is that a proposition is justified when there is another proposition supporting it. In other words, there needs to be a *justifier* for that proposition. However, since we intuitively think that a justifier should itself also be justified, we need another justifier for the first justifier. Thus, for a proposition to be justified, there has to be a chain-like structure consisting of justifiers such that the justified proposition stands on one end of the chain. This structure of justification naturally gives rise to a question: how does such a chain of justifiers come to an end? There are, given our ordinary understanding of a

[5]Note that a proposition belonging to a highly coherence set could still be false, given that other propositions are not true. That is, the notion of coherence alone does not guarantee the truth of a proposition.

chain, three possibilities. First, a chain may extend infinitely and never reach an end. Second, it may stop at some point. Third, it may circle back to some previous propositions. The three ideas respectively evolve into three different theories about justification: infinitism, foundationalism and coherentism. Supporters of infinitism accept the result that epistemic justification is an infinite chain, while foundationalists claim that the chain stops at some particular propositions. If one accepts the former, one needs to explain how could an infinite regress be innocuous. If one accepts the latter, she needs to explain the nature of the end-points of a justification chain. Coherentists take the third route: they allow the chain to circle back and link to some proposition already in the chain. They claim that if the chain is long enough, it is acceptable that a proposition justifies a proposition in the same chain of justification.[6] Hence, if a proposition is involved in a very coherent set in which every proposition is justified by at least a proposition in the set, we may accept all the propositions on the chain as justified.

2.2.3 Characterising coherence

Serving as an explanation of truth and justification, the notion of coherence plays an crucial role in contemporary epistemology. However, an important question remain unanswered: what is coherence? Although we do seem to have a rough idea concerning a set of propositions 'hanging together', it is rather vague how this basic understanding allows us to derive a complete account for both the notion of truth and justification. To make the notion useful, we need a more accurate formulation.

The rudimentary versions of the coherence theory equate the notion of coherence with consistency and hold that a proposition is true if and only if it is a member of a consistent set. As long as a logically closed set of propositions does not

[6]Rorty (1979, p.178) explains the idea clearly by saying that

> '...nothing counts as justification unless by reference to what we already accept, and that there is no way to get outside our propositions and our language so as to find some test other than coherence.'

include any pair of contradictory propositions, it is coherent. Although this definition of coherence correctly captures a basic aspect of coherence, it is far from adequate. For any non-maximal consistent set, we may find a pair of contradictory propositions which are both consistent with the set.[7] To illustrate, consider a set S_1 which contains the following three propositions:

(p_1) Jay Gatsby owns a mansion.

(p_2) Jay Gatsby owns a yellow Rolls-Royce.

(p_3) Jay Gatsby inherited a large amount of money from Dan Cody.

The set S is consistent as its elements and their logical consequences do not contradict with each other. If we take coherence as consistency, the set is coherent. Now consider two further propositions:

(p_4) Tom Buchanan went to Yale.

(p_5) Tom Buchanan did not go to Yale.

Both (p_4) and (p_5) are consistent with S_1. If we simply define the notion of coherence as consistency, both $S_1 \cup \{p_4\}$ and $S_1 \cup \{p_5\}$ are coherent. If we adopt the coherence theory of truth, we should accept that all the members of a coherent set are true. Since both $S_1 \cup \{p_4\}$ and $S_1 \cup \{p_5\}$ are coherent, we may derive the result that both p_4 and p_5 are true. However, since p_4 contradicts p_5, it is impossible for both p_4 and p_5 to be true. We may thus see that the notion of coherence, if equated to consistency, cannot be taken as a proper account for the notion of truth. Coherentists must seek further conditions to define the notion of coherence in a more precise way.

A much stronger account, proposed by Ewing (1934), takes entailment as the defining feature of coherence: a set is coherent if every proposition in it *logically follows* from all other propositions in the set taken together. Consider a set composed of three propositions $p_1, p_2, p_1 \wedge p_2$. In this set, p_1 and p_2 follows from the

[7] A maximal set is one such that for every proposition p, either p or $\neg p$ is in the set.

rest of the set, namely $\{p_2, p_1 \wedge p_2\}$. Similarly, p_2 follows from the set $\{p_1, p_1 \wedge p_2\}$. The last proposition $p_1 \wedge p_2$ also follows from the conjunction of the rest of the set, namely the conjunction of p_1 and p_2. According to Ewing's definition, this is a coherent set. However, this account sets a very demanding standard for coherence which can only be met by a limited range of sets. Consider the set S_1 in the previous example. Intuitively, the set is highly coherent as all its elements show that Gatsby is a rich person. But since none of them entails the other propositions, the set does not satisfy Ewing's definition. It can thus be seen that this definition of coherence is overly narrow and fails to include many intuitively coherent sets.

Lewis (1946) provides a definition for coherence which can be regarded as a weaker version of Ewing's.[8] He claims that for a set of propositions S to be coherent, it should satisfy the condition that for any proposition p which is an element of S, if all other elements in S are assumed as true, the probability of p raises. That is, the probability of the proposition p conditioning on $S_1 \backslash \{p\}$ is greater than the unconditional probability of p.[9] This definition of coherence is preferred to Ewing's as it is less strict. The notion Lewis appeals to in order to define coherence is *probability raising*, rather than the much stronger logical entailment. We may consider the Gatsby example again to see this point. The proposition p_1 indicates that Jay Gatsby is very rich. If we assume this piece of information is true, the probability of Jay Gatsby owning an expensive car should increase. In other words, since all the propositions in S_1 indicate that Gatsby is rich, assuming the truth of each of them does make other propositions more probable.

Convincing as it seems, Lewis' definition of coherence is still far from adequate. He takes probability raising of a single proposition as the criterion for coherence,

[8]Lewis called the notion *congruence* in the original text. It has been generally agreed that congruence is identical to the notion of coherence.

[9]Chisholm (1966) provides a definition of coherence which is pretty similar to the one Lewis proposed:

> 'A set of propositions S is coherent just if S is a set of two or more propositions each of which is such that the conjunction of all the others tends to confirm it and is logically independent of it.'

The disadvantages of this definition are also pretty similar to problems of Lewis' definition.

but neglects the fact that coherence can also be a relation between different sets. Suppose that we have a set S which has two subsets S^* and S^{**}. According to Lewis' definition, we are in no position to tell whether S^* coheres with S^{**}. Lewis' definition only allows us to check if a set of propositions is coherent, but provides no way for us to check whether a set is coherent with another. Also, what Lewis' definition provides is a qualitative, rather than a quantitative notion. That is, it only tells us whether a set is coherent, but not whether one set is more coherent than another (Bovens and Olsson, 2000). Hence, Lewis' definition still fails to fully capture our understanding of coherence.

BonJour (1985, p.97-99) proposes a set of coherence criteria which provides a more complete and subtle characterisation of the notion of coherence:[10]

1. A system of propositions is coherent only if it is logically consistent.

2. A system of propositions is coherent in proportion to its degree of probabilistic consistency.[11]

3. The coherence of a system of propositions is increased by the presence of inferential connections between its component propositions and increased in proportion to the number and strength of such connections.

4. The coherence of a system of propositions is diminished to the extent to which it is divided into subsystems of propositions which are relatively unconnected to each other by inferential connections.

5. The coherence of a system of propositions is decreased in proportion to the presence of unexplained anomalies in the propositional content of the system.

[10] In the original text, BonJour considers the coherence between beliefs. Here I replace beliefs with propositions.

[11] BonJour characterises the notion of *probabilistic consistency* with two factors: (a) the number of conflicts between propositions in a system and (b) the degree of improbability involved in each case. If a set of propositions involves many conflicting pairs of propositions, it is probabilistically inconsistent. Also, if the members of a set are highly improbable, the set is probabilistically inconsistent.

These criteria highlight that the essence of coherence is the inferential connection between the elements of a set. If a set is closely connected, it is coherent. Also, the degree of coherence of a set is determined by the extent it is inferentially connected. A highly connected set, compared to a less connected one, is more coherent. With BonJour's criteria, we have a more sophisticated understanding of the notion of coherence.

2.3 Coherence and truth-conduciveness

The more complicated definitions of coherence, including Lewis and BonJour's, reveal an important point: the degree of coherence of a set can be understood in terms of the probabilities of the propositions included in the set. Put more precisely, we may measure the extent a set of propositions are connected with the probabilities of these propositions and thereby see how coherent a set is.

Surprisingly, based on the very idea that probability and coherence are correlated, Klein and Warfield (1994) claim that the degree of coherence of a set is negatively correlated to its likelihood of truth. In other words, the more coherent a set is, the less likely it is true. Their argument begins with two propositions:

1. Any set of propositions S_1 is more likely to be true than any other set of propositions $S_1 \cup S_2$, given that at least one element in S_2 is not entailed by S_1 and does not have an objective probability of 1.

2. To increase the coherence of a set of propositions S, one may expand the set with a proposition which is relevant to the propositions in the set. This proposition should not be entailed by S and does not have an objective probability of 1.

What the first claim says is that for any set of propositions, the more elements it has, the more likely it is false. However, to make a set of propositions more coherent, one has to expand it with a proposition which carries some information that provides inferential support to the propositions in the set. If a newly added proposition carries some information, it cannot be a proposition with objective probability of one since only tautologies are of maximal probability. The direct result that can be derived from the two observations is that a set can only be made more coherent when it is expanded with a proposition that is possibly false. In other words, when the coherence of a set increases, the probability for it to be false would also increase. When we see a highly coherent set which contains many highly specific propositions, we should infer that it is very likely to be false. Klein and Warfield thereby conclude that coherence, instead of being a truth-conducive notion, is negatively correlated with the likelihood of truth.

Their argument can be illustrated with a simple example. Consider the earlier set S_1 which includes three propositions about Jay Gatsby's wealth. Suppose there is another proposition describing how Gatsby earned his wealth:

(p_6) Jay Gatsby was a smuggler.

The original set S_1, if expanded with the new proposition p_6, would become more coherent since p_6 provides a reason supporting other propositions in S_1. As the fact that Gatsby was a smuggler can well explain where his wealth comes from, p_6 supports both p_1 and p_2 which describe Gatsby as a rich person. However, since p_6 is not tautologous, it is possible for p_6 to be false. As a direct result, the set $S_1 \cup \{p_6\}$, compared to S_1, is less likely to be true. We may thus see that the coherence of a set is negatively correlated to its likelihood of truth.

Klein and Warfield's observation seriously undermines the coherence theory of truth. If their argument is correct, we may see that that the notion of coherence is at best unrelated and at worst negatively correlated to truth. As a consequence, philosophers should give up the idea of taking coherence as an indicator of the truth of a set of propositions and abandon the coherence theory of truth, Similarly, if we adopt the coherence theory of justification, the result that follows from Klein and Warfield's argument is that a highly justified set of propositions, compared to a less justified one, may be less likely to be true. Such a result violates our ordinary understanding of the notion of justification. The coherentists, hence, have to either give up the idea of explaining justification in terms of coherence, or admit that epistemic justification is not truth-conducive. Since the latter option looks disastrous, giving up coherentism seems to be the only way out.

One way of responding to Klein and Warfield's challenge is to argue that the notion of coherence should be understood in an alternative way. If one can provide a more sophisticated way of measuring the degree of coherence of a set which shows that greater degree of coherence does guarantee greater likelihood of truth, one may save coherentism. To achieve this goal, the first step is to introduce the formal apparatus for establishing the desired formal definition of coherence.

2.3.1 Formal preliminaries

As we have seen in previous discussions, there are two primary requirements for a set of propositions to be coherent. First, the elements of a coherent set of propositions should be true or false together, or at least tend to be true or false together. That is, when a single proposition in a coherent set is assumed to be true, other propositions in the set should be more likely to be true. We can also understand this in terms of the content of these propositions. When a set of propositions tend to be true or false together, what is implied is that their contents overlap to a sufficiently large extent. If this condition is met, these propositions are likely to be true or false together. Second, the elements of a coherent set must bear strong *mutual support* with each other. In BonJour's words, there should exist strong *inferential connections* between the elements of a coherent set. A proper way of measuring the coherence of a set should capture the two factors.

The two aspects of a set could be represented in terms of probability. To show how to do this, we need some basic formal preliminaries. Since coherence is a property of a set of propositions, the first item we need is an algebra of propositions A, namely a set of propositions closed under negation and conjunction.[12] Secondly, we need a set of probability functions Pr which assign values within the

[12]For the sake of simplicity, we assume that the algebra is finite.

interval $[0,1]$ to every element in the algebra A.[13] Given the two basic entities A and Pr, a coherence measure $\mathcal{C}$ can be defined as a function which assigns a real number to a set of propositions given a specific probability function Pr, namely that $\mathcal{C} : A^n \rightarrow \mathbb{R}$ where n is the number of propositions in the set. Given a probability function Pr, a coherence measure takes a set of propositions as input, and generates the degree of coherence of that set.

2.3.2 Shogenji's coherence measure

In order to refute Klein and Warfield's criticism to coherentism, Shogenji (1999) provides a probabilistic coherence measure to show that coherence could be a truth-conducive notion. Given a set of propositions $S = \{p_1, \cdots, p_n\}$ and a probability function Pr, Shogenji measures the degree of coherence of S with the following formula:

Definition 2.3.1. Shogenji's coherence measure

$$\mathcal{C}_{Sh}(S) =: \frac{Pr(\bigwedge S)}{\prod_{i=1}^{n} Pr(p_i)}$$

Shogenji's measure divides the probability of the conjunction of all the propositions in S with the product of the probabilities of each proposition. The outcome, namely the quotient of the two probabilities, is taken to be the degree of coherence of S. Shogenji's original idea is quite elegant. In probability calculus, if a proposition p_i is independent from another proposition p_j, the probability of their conjunction would be equivalent to the product of their probabilities, namely $Pr(p_i)Pr(p_j)$. If we expand the idea, we may infer that if a set of propositions are all mutually independent, the probability of their conjunction would be equivalent to the product of the probabilities of each. If the propositions are not independent but relevant to a certain extent, the probability of their conjunction would be greater than the product of the probability of each. Shogenji takes the ratio between the two values as the degree of coherence of the set which indicates the

[13] A probability function is one which conforms to Kolmogorov's probability axioms.

degree of relevance between its elements. If the probability of the conjunction of all the elements in a set is high, the set is highly coherent. On the contrary, if the probability of the conjunction of all the elements is low or equivalent to to product of the probability of each proposition, the set is rather incoherent.

To see how Shogenji's measure works, we could compare two different sets of propositions. Suppose there are two sets $S_1 = \{p_1, p_2, p_3\}$, $S_2 = \{p_4, p_5, p_6\}$ and a probability function Pr such that $Pr(p_1) = Pr(p_4), Pr(p_2) = Pr(p_5)$ and $Pr(p_3) = Pr(p_6)$. Further assume that the main difference between the two sets is that the probability of the conjunction of all the propositions in S_1 is greater than S_2, namely that $Pr(p_1 \wedge p_2 \wedge p_3) > Pr(p_4 \wedge p_5 \wedge p_6)$. In other words, the elements of S_1 is more relevant with each other compared to S_2. Given Shogenji's measure $\mathcal{C}_{Sh}$, the degree of coherence of S_1 is

$$\mathcal{C}_{Sh}(S_1) = \frac{Pr(p_1 \wedge p_2 \wedge p_3)}{Pr(p_1)Pr(p_2)Pr(p_3)}$$

while the degree of coherence of S_2 is

$$\mathcal{C}_{Sh}(S_2) = \frac{Pr(p_4 \wedge p_5 \wedge p_6)}{Pr(p_4)Pr(p_5)Pr(p_6)}.$$

Since the product of the probabilities of p_1, p_2, p_3 equals the product of the probabilities of p_4, p_5, p_5, the denominator of $\mathcal{C}_{Sh}(S_1)$ is equivalent to the denominator of $\mathcal{C}_{Sh}(S_2)$. Given that $Pr(\bigwedge S_1)$ is greater than $Pr(\bigwedge S_2)$, Shogenji's measure gives the verdict that S_1 is more coherent than S_2. We can thus see that, given Shogenji's measure, a set is highly coherent if the probability of the conjunction of its elements is high. This measure correctly captures the intuitive idea that the degree of coherence of a set is determined by the extent the contents of the propositions overlap, namely the extent the propositions 'hang together'. According to this measure, a set if highly coherent only if the conjunction of its elements is highly likely to be true.

Shogenji's measure has an important merit that it is sensitive to the size of the proposition set being measured. Other things being equal, the more elements a set includes, the more coherent the set is.[14] This feature correctly reflects our intuitive idea that for any two sets of propositions, if the degree of relevance between elements of the two sets are the same, the one which has more elements should be considered as more coherent. The underlying rationale here is that it is harder, compared to a smaller set, for the elements of a bigger set to agree with each other. This point can be illustrated with an analogy. Imagine there are two groups of people. The first group has three members, while the second has thirty. It is much harder for members of the second group to reach a consensus, as it involves more members. Analogously, it is more difficult for the contents of propositions in a big set to agree with other propositions in the same set, compared to a smaller set. Therefore, when comparing two sets with the same degree of agreement, the one with greater size should be rendered with greater coherence.[15] This feature of coherence is well captured by Shogenji's measure, which can be illustrated by the following example:

Example 2.3.1. Given two sets of propositions $S_1 = \{p_1, \cdots, p_i\}$ and $S_2 = \{p_1, \cdots, p_i, p_{i+1}\}$. Suppose that $Pr(\bigwedge S_1)$ is equivalent to $Pr(\bigwedge S_2)$ and $Pr(p_{i+1})$ is smaller than 1. According to the given premises, the denominator of $\mathcal{C}_{Sh}(S_2)$ is smaller than the denominator of $\mathcal{C}_{Sh}(S_1)$. Hence, the degree of coherence of S_2 is greater than the degree of coherence of S_1 under Shogenji's measure, namely that

$$\mathcal{C}_{Sh}(S_1) = \frac{Pr(p_1 \wedge \cdots \wedge p_i)}{Pr(p_1)\cdots Pr(p_i)} < \frac{Pr(p_1 \wedge \cdots \wedge p_i \wedge p_{i+1})}{Pr(p_1)\cdots Pr(p_i)Pr(p_{i+1})} = \mathcal{C}_{Sh}(S_2)$$

With this case, we may see that Shogenji's measure yields the result that, other things being equal, a big set is more coherent than a small set.

Another factor which needs to be considered while measuring coherence is the *specificity* of elements of a proposition set. Two highly specific propositions, compared with two general ones, are less likely to agree with each other. Hence, a set of highly specific propositions, compared with a set of less specific ones, should be more coherent. This point can be illustrated by the following example:

[14]The most important factor here is the specificity of the propositions.

[15]'Having the same degree of agreement' here means that the probability of the conjunction of all the propositions are equal in both sets.

Example 2.3.2. Consider two pairs of propositions concerning the same subject matter but with different specificity:

(p_1) Gatsby lives in New York.

(p_2) Gatsby attended college.

(p_3) Gatsby lives on Long Island in New York.

(p_4) Gatsby attended Trinity College, Oxford.

In this example, (p_3) implies (p_1) and (p_4) implies (p_2). We can hence derive that $Pr(p_3) \leq Pr(p_1)$ and $Pr(p_4) \leq Pr(p_2)$ for any arbitrary probability function Pr. It can be further derived that $Pr(p_1)Pr(p_2)$ is greater than $Pr(p_3)Pr(p_4)$, which implies that the denominator of $\mathcal{C}_{Sh}(\{p_1, p_2\})$ is greater than the denominator of $\mathcal{C}_{Sh}(\{p_3, p_4\})$. On the condition that all other factors are equal, namely that $Pr(p_1 \wedge p_2)$ is equivalent to $Pr(p_3 \wedge p_4)$, we may get the result that $\mathcal{C}_{Sh}(\{p_3, p_4\})$ is greater than $\mathcal{C}_{Sh}(\{p_1, p_2\})$. That is, other thing being equal, Shogenji's measure generates the result that a set of highly specific propositions is more coherent than a less specific set, which is in accordance with our intuitive understanding of coherence.

Shogenji calls the size and specificity the *total individual strength* of a set. With this notion, he argues that given two sets with the same total individual strength, a coherent set of propositions, compared to a less coherent one, is more likely to be jointly true. We can again illustrate this point with an example. Suppose there are two equally specific sets $\{p_1, p_2\}$ and $\{p_3, p_4\}$ such that the propositions in each set are of the same degree of specificity. Assuming that equal specificity implies equal probability, we may derive that $Pr(p_1)Pr(p_2) = Pr(p_3)Pr(p_4)$. In this scenario, if $Pr(p_1 \wedge p_2)$ is greater than $Pr(p_3 \wedge p_4)$, the degree of coherence of the first set would be greater than the coherence of the second. Also, if $Pr(p_1 \wedge p_2)$ is greater than $Pr(p_3 \wedge p_4)$, namely that the contents of p_1 and p_2 overlap to a greater extent, then $\{p_1, p_2\}$ would be more likely to *jointly be true* than $\{p_3, p_4\}$. Suppose that both p_1 and p_3 are true. When we know that $\{p_1, p_2\}$ is more coherent than $\{p_3, p_4\}$, we can infer that the elements of $\{p_1, p_2\}$ are more likely to be jointly true. Given that p_1 is true, the whole set $\{p_1, p_2\}$ is more likely to be true than $\{p_3, p_4\}$. Shogenji thereby concludes that, given his probabilistic coherence measure, coherence *with truth* is truth-conducive. That is, coherence is truth-conducive on the condition that there are some true propositions in the set.

Given Shogenji's measure, Klein and Warfield's claim that coherence is negatively correlated with truth can be refuted. Suppose there is a set S which contains some true propositions. A proposition which coheres with the elements of S is highly likely to be true. Hence, expanding S with such a proposition would not make S more likely to be false.

In spite of its plausibility, some epistemologists are dissatisfied with Shogenji's measure is fallacious. Akiba (2000) points out that Shogenji's measure may actually be *falsity-conducive* and cannot properly measure the coherence of a set of propositions bearing the entailment relation. Consider the case in which the propositions p_1 entails p_2. The pairwise coherence between p_1 and p_2 would be:

$$\mathcal{C}_{Sh}(\{p_1, p_2\}) = \frac{Pr(p_1 \wedge p_2)}{Pr(p_1)Pr(p_2)} = \frac{Pr(p_1)}{Pr(p_1)Pr(p_2)} = \frac{1}{Pr(p_2)}$$

Since p_1 entails p_2, the probability of the conjunction of p_1 and p_2 is equivalent to the probability of p_1. The degree of coherence of the set $\{p_1, p_2\}$ would thus be the reciprocal of the probability of p_2. In other words, $Pr(p_2)$ is negatively correlated with $\mathcal{C}_{Sh}(\{p_1, p_2\})$. When $Pr(p_2)$ decreases, $\mathcal{C}_{Sh}(\{p_1, p_2\})$ increases. With such a result, Akiba claims that Shogenji's measure cannot show that coherence is truth-conducive.[16]

Apart from being falsity-conducive, Akiba points out another problem of Shogenji's measure which can be illustrated by the following example:

Example 2.3.3. When throwing a dice, one may believe in the following three propositions:

[16] A potential problem here is that the degree of coherence of $\{p_1, p_2\}$, according to Shogenji's measure, is solely determined by the probability of p_2. This result looks strange, as what we aim to measure is the coherence between two propositions. A possible response is to bite the bullet and claim that, since p_1 entails p_2, it is pointless to measure the coherence between the two propositions. As measuring the coherence of $\{p_1, p_2\}$ is pointless, it does not matter that the degree of coherence is determined only be the probability of p_2.

(p_1) The dice will come up two.

(p_2) The dice will come up an even number less than six.

(p_3) The dice will come up an even number.

The probability of the p_1 is one-sixth, while the probabilities of p_2 and p_3 are respectively two-sixth and one-half. Intuitively, a proposition is supposed to be extremely coherent with its consequence. Since both p_2 and p_3 are immediate results of p_1, the degree of coherence of $\{p_1, p_2\}$ should be the same as $\{p_1, p_3\}$. However, if we calculate the degree of coherence of the two sets with Shogenji's measure, we may obtain the result that the coherence of $\{p_1, p_2\}$ is 3, whereas the coherence of $\{p_1, p_3\}$ is 2. Such result indicates that Shogenji's measure fails to capture our intuitive understanding of coherence.

A further problem, as Akiba argues, is that the coherence of a singleton will always be 1. As a proposition is perfectly coherent with itself, 1 is supposed to be the maximal degree of coherence. However, if we arbitrarily combine two probabilistic independent propositions p_1 and p_2 into a proposition and measure the coherence of the singleton $\{p_1 \wedge p_2\}$, the degree of coherence of $\mathcal{C}_{Sh}(\{p_1 \wedge p_2\})$ would also be 1. This result, again, is quite counterintuitive. Since we have assumed that p_1 and p_2 are probabilistically independent, it should not be the case that the set $\{p_1 \wedge p_2\}$ is perfectly coherent. Because of these reasons, Akiba concludes that Shogenji's measure of coherence is inadequate.

Shogenji (2001) rejects all Akiba's criticisms. The claim that Shogenji's measure is falsity-conducive, according to Shogenji, does not really pose a threat to his measure. What Shogenji intends to show with his measure is exactly that a set of low probability, which can be interpreted as being highly specific, is more likely to be coherent compared to a less specific set. Akiba's criticism does not show that Shogenji's measure is falsity-conducive, but instead reveals the fact that the degree of coherence of a set increases with the specificity of the propositions in the set. Hence, in debating whether coherence is truth-conducive, the specificity of the propositions should be fixed. Akiba fails to see the underlying motivation of proposing a coherence measure and came up with a criticism that misses the point of measuring coherence..

As for the dice case, Shogenji provides an example to show that pairs of propositions bearing the entailment relation can differ in coherence.

Example 2.3.4. Consider the following propositions:

(p_1) The fossil was deposited 64-to-66 million years ago.

(p_2) The fossil was deposited 63-to-67 million years ago.

(p_3) The fossil was deposited more than 10 years ago.

In this case, p_1 entails both p_2 and p_3 but, intuitively, the set $\{p_1, p_2\}$ is more coherent than $\{p_1, p_3\}$ as the information provided by p_2 is far more specific than p_3. Hence, it should be acceptable that in Akiba's example, the degree of coherence of $\{p_1, p_2\}$ differs from the degree of coherence of $\{p_1, p_3\}$. It should be allowed that the coherence of a set containing a proposition p and its consequence differs from the coherence of another set containing p and a different consequence.

The last problem, namely the one concerning the coherence of the conjunction of two individual propositions, does not undermine Shogenji's measure either. Coherence is a relation between propositions, rather than a property of a single proposition. Measuring the coherence of a singleton set, hence, makes little sense. It does not tell us any information about the relation between the propositions in the set. Therefore, Akiba's arguments does not really show that Shogenji's measure is fallacious for failing to generate the correct degree of coherence for a singleton set.

There exist two other problems of Shogenji's measure The first one is the depth problem. Fitelson (2003) points out that Shogenji's measure does not take the coherence of the subsets of a set as a factor when measuring the overall coherence of the set. Given a set of propositions with n elements, Shogenji's measure can only calculate its n-wise coherence, but not its k-wise coherence for any $k < n$. A set might be very coherent when one considers only the pairs of propositions involved in the set, but not quite coherent when evaluated as a whole. Failing to capture the mixed nature of coherence, thus, is a serious shortcoming of Shogenji's measure. Consider the following example Schupbach (2011) provides:

Example 2.3.5. Police investigators caught eight robbery suspects, each of them are equally likely to have committed the crime. Three independent witnesses claimed that they have seen the criminal. In the first case, the witnesses provide the following set of testimonies respectively:

t_1: The criminal was either suspect 1, 2 or 3.

t_2: The criminal was either suspect 1, 3 or 4.

t_3: The criminal was either suspect 1, 2 or 4.

In the second case, the witnesses respectively provide three different testimonies:

t_1': The criminal was either suspect 1, 2 or 3.

t_2': The criminal was either suspect 1, 4 or 5.

t_3': The criminal was either suspect 1, 6 or 7.

The set of testimonies in the first case, intuitively, is more coherent than the testimonies in the second case. In the first case, the three testimonies indicates that it is very likely that suspects one to four are the real criminal. Compared to the first set, the information one may obtain from the set of testimonies in the second case is more ambiguous. Seven suspects were mentioned but, except suspect 1, all the suspects were mentioned only once. It is thus quite hard for us to make an inference about who the criminals are from the set of testimonies in the second case. We may thus see that the first set of propositions is intuitively more coherent. However, such difference cannot be reflected by Shogenji's measure which generates the result that the two sets are equally coherent. We may express this result formally:

$$\mathcal{C}_{Sh}(\{t_1, t_2, t_3\}) = \frac{Pr(t_1 \wedge t_2 \wedge t_3)}{Pr(t_1)Pr(t_2)Pr(t_3)} = \frac{Pr(t'_1 \wedge t'_2 \wedge t'_3)}{Pr(t'_1)Pr(t'_2)Pr(t'_3)} = \mathcal{C}_{Sh}(\{t'_1, t'_2, t'_3\})$$

All the testimonies are equiprobable, since they all point to three suspects. Hence, the denominator of $C_{Sh}(\{t_1, t_2, t_3\})$ equals to the denominator of $C_{Sh}(\{t'_1, t'_2, t'_3\})$. The conjunction of $\{t_1, t_2, t_3\}$ is that the criminal is suspect 1, which is the same as the conjunction of $\{t'_1, t'_2, t'_3\}$. Given Shogenji's measure, the two sets of testimonies are equally coherent. Such a result violates our intuition concerning the coherence of the two sets.

The problem, as we can see, stems from a feature of Shogenji's measure such that the *sub-coherence* of a set is not taken into account. In the given example, the overall coherence of $\{t'_1, t'_2, t'_3\}$ is not influenced by the fact that $\{t'_1, t'_2\}$ are less coherent than $\{t_1, t_2\}$. We may conclude that Shogenji's measure fails to generate the intuitive result that the first set of testimonies is more coherent. This is the the so-called *depth problem*.

The second problem is the problem of irrelevant addition. When a proposition which is totally irrelevant to a set S is added to S, the degree of coherence of that set, according to Shogenji's measure, remains the same. Again, this is a highly counterintuitive result. We may see this with an example.

Example 2.3.6. Recall the robbery example. Suppose a witness by accident provides another testimony:

(t_4) It is raining in Paris now.

This new testimony is totally irrelevant to t_1, t_2 and t_3. If we add this irrelevant testimony t_4 to the set $\{t_1, t_2, t_3\}$, the degree of coherence of the new set $\{t_1, t_2, t_3, t_4\}$ is:

$$\mathcal{C}_{Sh}(\{t_1, t_2, t_3, t_4\}) = \frac{Pr(t_1 \wedge t_2 \wedge t_3 \wedge t_4)}{Pr(t_1)Pr(t_2)Pr(t_3)Pr(t_4)}$$

Since t_4 is irrelevant to all other testimonies in the set, the probability of the conjunction of t_4 and other testimonies is equivalent to the product of their probabilities. Hence, this formula is equivalent to

$$\mathcal{C}_{Sh}(\{t_1, t_2, t_3, t_4\}) = \frac{Pr(t_1 \wedge t_2 \wedge t_3)Pr(t_4)}{Pr(t_1)Pr(t_2)Pr(t_3)Pr(t_4)}$$

As $Pr(t_4)$ appears in both the denominator and the numerator, we may remove it from the formula and derive the consequence that the degree of coherence of $\{t_1, t_2, t_3, t_4\}$ is equivalent to the degree of coherence of $\{t_!, t_2, t_3\}$.

Given Shogenji's measure, no matter how many irrelevant propositions are added to a set, as long as they are independent from the other propositions in the set, the degree of coherence of the set remains the same. This result is highly counterintuitive. When a set is extended with irrelevant propositions, people tend to consider the new set as less coherent then the original set since the newly added propositions do not provide any support to any proposition in the original set. Again, Shogenji's coherence measure fails to capture what we think about the notion of coherence.

Because of the two problems, Shogenji's measure cannot be adopted as an ideal coherence measure. Coherentists need to invent a different measure which is free from the two problems to show that coherence is truth-conducive.

2.3.3 Shogenji's measure generalised

Upon realising that Shogenji's measure is flawed, Schupbach (2011) provides a revised measure which is free from the two problems. The common root of the depth problem and the problem of irrelevant addition is that Shogenji's original measure does not take the coherence of the subsets into account. Hence, Schupbach comes up with the idea to measure the coherence of a set at different levels and take the weighted average as the overall degree of coherence of the set. The first notion that needs to be introduced, hence, is the r-wise coherence of a set:

Definition 2.3.2. r-wise Shogenji coherence

For a set of propositions $S = \{p_1, \cdots, p_k\}$, $[S]^r$ represents the set of all subsets of S with r elements. Given an ordering $\langle S_1, \cdots, S_m \rangle$ of the members of $[S]^r$, the degree of *r-wise* coherence of S is measured as:

$$\mathcal{C}_{Sc}{}^r(S) =: \frac{\sum_{i=1}^{m} s(S_i)}{m}$$

in which m is the number of elements in $[S]^r$ and $s(S)$ is the logarithm of Shogenji's original coherence measure:[17]

$$s(S) =: log\left(\frac{Pr(\bigwedge S)}{\prod_{i=1}^{n} Pr(p_k)}\right)$$

Suppose one wants to measure the 3-wise coherence of a set S. She should collect all subsets of S with three elements, calculate the coherence of all these subsets respectively with the original Shogenji measure, and average the logarithm of the degrees of coherence of these subsets. The outcome is the desired 3-wise coherence of S.

With the notion of r-wise coherence, we can calculate the weighted coherence of a set by giving a weigh vector to each r-wise coherence:

Definition 2.3.3. The generalised Shogenji measure

Given a set of propositions with k elements $S = \{p_1, \cdots, p_k\}$ and a weight vector $\langle \mu_1, \cdots, \mu_{k-1} \rangle$ which assigns different weights to r-wise coherence for every r such that $\sum_{i=1}^{k-1} \mu_i = 1$, the degree of coherence of S is measured as:

$$\mathcal{C}_{Sc}(S) =: \sum_{i=1}^{k-1} \mu_i \mathcal{C}_{Sc}{}^{i+1}(S)$$

What this measure generates is the weighted average of the coherence of the set at all levels.[18] Here μ_i is the weight for the $i+1$-wise coherence of S. For example, μ_2 is the weight of the 3-wise coherence.

[17] Here Schupbach used logarithm to simplify the numbers to be calculated.

[18] Note that the coherence of singleton sets are intentionally neglected. Given Shogenji's idea the coherence is a relation between propositions, this consequence should be acceptable.

This scheme allows us to define different coherence measures by changing the value of the weight vector. The simplest one is generated by assigning equal weight to all r-wise coherence:

Definition 2.3.4. Straight average

$$\mathcal{C}_{SA}(S) =: \frac{\sum_{r=2}^{k} C^r(S)}{k-1}$$

This formula sums up the coherence of S at all levels and divide it by the number of levels. Hence, this measure assigns equal weight to the degree of coherence of each level.

Given the scheme of coherence measures, we can define another measure which assigns greater weight to the higher level coherence of a set.

Definition 2.3.5. Deeper Decreasing

Let the scheme assign decreasing weights to decreasing k

$$\mu_i = \frac{i}{(k-1)+(k-2)+\cdots+1} = \frac{2i}{k(k-1)}$$

The degree of coherence of $S = \{b_1, \cdots, b_k\}$ is measured as:

$$C_{DD}(S) =: \sum_{i=1}^{k-1} \frac{2i}{k(k-1)} C^{i+1}(S) = \frac{\sum_{i=1}^{k-1} iC^{i+1}(S)}{k(k-1)/2}$$

With this measure, the pairwise coherence of a set is assigned with the lowest weight, while the k-wise coherence the greatest.

On the contrary, we may also define a measure which assigns greater weight to the r-wise coherence of a set when r is distant from k:

Definition 2.3.6. Deeper Increasing

Let the weight of each i-wise coherence be

$$\mu_i = \frac{k-i}{(k-1)+(k-2)+\cdots+1} = \frac{2(k-i)}{k(k-1)}$$

The degree of coherence is thus measured as

$$C_{DI}(S) =: \sum_{i=1}^{k-1} \frac{2(k-1)}{k(k-1)} C^{i+1}(S) = \frac{\sum_{i=1}^{k-1} (k-i)C^{i+1}(S)}{2k(k-1)}$$

This measure assigns the greatest weight to the pairwise coherence and the lowest weight to the k-wise coherence of a set.

All the measures generated this way are free from the depth problem, for they all take the coherence of subsets of a set into account while measuring the coherence of the set. Let us reconsider the two sets of testimonies in the given example, namely $\{t_1, t_2, t_3\}$ and $\{t'_1, t'_2, t'_3\}$. The depth problem stems from the fact that Shogenji's measure fails to reflect the difference in the pairwise coherence of the two sets. That is, although we know that $\mathcal{C}_{Sh}(\{t_1, t_2\}) > \mathcal{C}_{Sh}(\{t'_1, t'_2\})$, $\mathcal{C}_{Sh}(\{t_1, t_3\}) > \mathcal{C}_{Sh}(\{t'_1, t'_3\})$ and $\mathcal{C}_{Sh}(\{t_2, t_3\}) > \mathcal{C}_{Sh}(\{t'_2, t'_3\})$, the overall coherence of the two sets $\{t_1, t_2, t_3\}$ and $\{t'_1, t'_2, t'_3\}$ are equal. By taking the pairwise coherence of a set as a factor determining the overall coherence of the set, the problem can be solved.

Moreover, C_{SA} and C_{DI} are free from the problem of irrelevant addition.[19] This point can be seen by considering the pairwise coherence of each set. If an irrelevant proposition is added to a set, the pairwise coherence of that set decreases, as the newly added proposition is not coherent with any proposition in the original set. As a result, the overall coherence of the set decreases when expanded with an irrelevant proposition. Revising this way, Shogenji's measure may still be taken as a good measure for coherence.

2.3.4 Olsson's coherence measure

Although Shogenji claims that being specificity-sensitive is an advantage of his coherence measure, Olsson (2002) criticises Shogenji's measure for having this feature. He points out that if a coherence measure is specificity sensitive, the upper bound of the degree of coherence of a set, given that measure, would be determined by the specificity of its elements. We may illustrate this deficiency with a simple example.

Example 2.3.7. Suppose there are four propositions p_1, p_2, p'_1 and p'_2 such that

[19] C_{DD} is similar to the original $\mathcal{C}_{Sh}$ that it assigns less weight to subsets that are small. Although not as serious as Shogenji's original measure, C_{DD} is still vulnerable to the problem of irrelevant addition.

$Pr(p_1) = Pr(p_2) = 0.5$, $Pr(p'_1) = Pr(p'_2) = 0.3$. The degree of coherence of $\{p_1, p_2\}$, according to Shogenji's measure, is

$$\mathcal{C}_{Sh}(\{p_1, p_2\}) = \frac{Pr(p_1 \wedge p_2)}{Pr(p_1)Pr(p_2)} = \frac{Pr(p_1 \wedge p_2)}{0.25}$$

Since $P(p_1) = P(p_2) = 0.5$, when p_1 and p_2 coincide perfectly, $\{p_1, p_2\}$ is maximally coherent. The degree of coherence of $\{p_1, p_2\}$, according to Shogenji's measure, is

$$\mathcal{C}_{Sh}(\{p_1, p_2\}) = \frac{Pr(p_1 \wedge p_2)}{Pr(p_1)Pr(p_2)} = \frac{0.5}{0.25} = 2$$

.

On the other hand, if p'_1 and p'_2 coincide perfectly, the set $\{p'_1, p'_2\}$ would also be maximally coherent. The degree of coherence of the set would then be

$$\mathcal{C}_{Sh}(\{p'_1, p'_2\}) = \frac{Pr(p'_1 \wedge p'_2)}{Pr(p'_1)Pr(p'_2)} = \frac{0.3}{0.09} = 3.3\bar{3}$$

If we suppose that both $\{p_1, p_2\}$ and $\{p'_1, p'_2\}$ are maximally coherent, $\{p_1, p_2\}$ will be rendered a degree of coherence lower than $\{p'_1, p'_2\}$ simply because p_1 and p_2 are more probable than p'_1 and p'_2. Such a result, as Olsson sees, is undesirable. We can imagine cases in which the set $\{p'_1, p'_2\}$ is intuitively moderately coherent, yet still measured as more coherent than a perfectly coherent but less specific set of propositions $\{p_1, p_2\}$. We may further expand this idea and see that a set with perfectly coherent propositions may still not be rendered maximally coherent if the measure has no upper bound. In sum, Olsson claims that the degree of coherence of a set should not be bounded by the probability of its elements. That is, the maximal degree of coherence of a set should not be determined by the specificity of that set.

The underlying problem revealed by this case is that Shogenji's measure does not have a *maximal value*. That is, adopting Shogenji's measure may lead to the consequence that no matter how coherent a set is, one can always arbitrarily create another set which is more coherent. Since there is no maximal degree of coherence, a set of logically equivalent propositions, which is supposed to be the most coherent set that can possibly be perceived, do not have a maximal degree of coherence. Such a result does look quite problematic.

Being aware of the defect of Shogenji's measure, Olsson develops another coherence measure which is free from these problems:

Definition 2.3.7. Olsson's coherence measure

Given a set $S = \{p_1, \cdots, p_n\}$, the degree of coherence of S is:

$$\mathcal{C}_O(S) =: \frac{Pr(\bigwedge S)}{Pr(\bigvee S)}$$

With Olsson's measure, the degree of coherence of a proposition set is no longer bounded by the probability of elements in the set but takes $[0,1]$ as range. For a set of propositions which do not agree on anything, the set has minimal degree of coherence. On the other hand, a set of propositions $\{p_1, \cdots, p_n\}$ is maximally coherent when $Pr(p_1 \wedge \cdots \wedge p_n)$ equals $Pr(p_1 \vee \cdots \vee p_n)$.

Olsson's measure is free from the problem of irrelevant addition. Suppose there are two sets of propositions $S = \{p_1, p_2\}$ and $S' = \{p_1, p_2, p_3\}$. If p_3 is irrelevant to p_1 and p_2, the denominator of $\mathcal{C}_O(S')$ would be greater than the denominator of $\mathcal{C}_O(S)$ and the numerator of $\mathcal{C}_O(S')$ would be smaller than the numerator of $\mathcal{C}_O(S)$. We may derive the result that

$$\mathcal{C}_O(S) = \frac{Pr(p_1 \wedge p_2)}{Pr(p_1 \vee p_2)} > \frac{Pr(p_1 \wedge p_2 \wedge p_3)}{Pr(p_1 \vee p_2 \vee p_3)} = \mathcal{C}_O(S')$$

With Olsson's measure, expanding a set with irrelevant propositions leads to a decrease in its coherence. Thus, Olsson's measure fares better than Shogenji's in capturing our ordinary idea in this aspect.

Although Olsson's measure is free from some problems of Shogenji's, it is not impeccable. Siebel (2005) points out that with Olsson's measure, adding necessary truths to a set makes the set less coherent. Consider a set of propositions $\{p_1, p_2\}$ such that both p_1 and p_2 are not necessary truths. If one expands the set with a necessary truth, say p_t, the denominator would become one. As a result, the overall coherence of $\{p_1, p_2, p_t\}$ would be lower than the coherence of $\{p_1, p_2\}$.

$$\mathcal{C}_O(\{p_1, p_2\}) = \frac{Pr(p_1 \wedge p_2)}{Pr(p_1 \vee p_2)} > \frac{Pr(p_1 \wedge p_2 \wedge p_t)}{Pr(p_1 \vee p_2 \vee p_t)} = \frac{Pr(p_1 \wedge p_2)}{Pr(p_1 \vee p_2 \vee p_t)} = \mathcal{C}_O(\{p_1, p_2, p_t\})$$

When expanded with a necessary truth p_t which is irrelevant to p_1 and p_2, $Pr(p_1 \wedge p_2)$ remains the same, while $Pr(p_1 \vee p_2 \vee p_t)$ raises.[20] Therefore, adding p_t lowers the degree of coherence of the original set.

Although Siebel's observation is correct, it may not cause substantial harm to Olsson's measure. Given a set of propositions $\{p_1, \cdots, p_n\}$, if one adds a necessary truth which is irrelevant to all elements of that set, it does not seem wrong to consider the new set as less coherent than the original one. Take the robbery case in the last section for example. Suppose that a witness provides the testimony

t_4 : Five plus seven equals twelve.

Since this testimony is totally irrelevant to the robbery, it should not be regarded as coherent with the original set of testimonies. According to Olsson's measure, the degree of coherence of $\{t_1, t_2, t_3, t_4\}$ is lower than the degree of coherence of $\{t_1, t_2, t_3\}$, which correctly captures this idea. Hence, the point Siebel criticises should be taken as an advantage, rather than a shortcoming.

The real problem of Olsson's measure lies in it being *size-uninformative*. The degree of coherence of a set, according to Olsson's measure, does not increase with its size. Consider two sets $B = \{p_1, p_2\}$ and $B' = \{p'_1, \cdots, p'_{100}\}$. The size of the latter, as we can see, is much bigger than the first. If $Pr(p_1 \wedge p_2) = Pr(p'_1 \wedge \cdots \wedge p'_{100})$ and $Pr(p_1 \vee p_2) = Pr(p'_1 \vee \cdots \vee p'_{100})$, then, according to Olsson's measure, the degree of coherence of B is equivalent to B'. Such a consequence is quite dubious. As previously discussed, other things being equal, people tend to take sets with greater size as more coherent. We can illustrate this with a revised version of the robbery example:

Example 2.3.8. Police investigators caught eight suspects for a robbery. Each of them are equally likely to have committed the crime. In the first scenario, two

[20] The idea of irrelevance here differs from the notion of probabilistic independence. If what we meant by the term irrelevance is actually probabilistic independence, then there would not be any necessary truth that is irrelevant to other propositions. Siebel uses this term in an ordinary way and refer to our everyday understanding of irrelevance.

independent witnesses claimed that they have seen the suspect and provided the following set of testimonies:

t_1: The criminal was either suspect 1, 2 or 3.

t_2: The criminal was either suspect 1, 3 or 4.

In the second scenario, there are one hundred independent witnesses who claimed that they have seen the suspect and provided the following set of testimonies:

t_{1-50}: The criminal was either suspect 1, 2 or 3.

t_{51-100}: The criminal was either suspect 1, 3 or 4.

Intuitively, the set of testimonies in the second scenario is more coherent than in the first scenario as the size of the set of testimonies is much larger than the set of testimonies in the first scenario. More precisely, if we do want to take coherence as a notion which accounts for justification, the second set must be more coherent as it provides stronger justification for suspect 1 and 3 being the real criminal. Olsson's measure fails to capture this feature of coherence but takes the two sets as equally coherent instead. Hence, it does not generate the intuitive verdict. Compared to Olsson's, Shogenji's measure fares better in this aspect. When the size of a set increases, the denominator of Shogenji's measure decreases. Consequently, the degree of coherence of the set increases.[21]

We have seen that both Shogenji and Olsson's measures are flawed. Although Shogenji's measure is size-informative, it has the undesirable feature that the maximal degree of coherence is determined by the specificity of its elements. On the other hand, although Olsson's measure has a fixed upper bound, it is not size-informative. If we agree that a coherence measure should be size-informative but

[21]Recall that the denominator in Shogenji's measure is the product of the probabilities of all the propositions in the set that is being measured. Since all the values are in the interval $[0,1]$, the denominator can only decrease when the size increases when all the values remain unchanged. Thus, the overall coherence of a set can only increase when the size increases.

not bounded by the specificity of its elements, both Shogenji and Olsson's measures are fallacious. Thus, coherentists need to find better measures to capture this notion.

2.3.5 Coherence as the average strength of confirmation

Both Shogenji and Olsson's measures are based on the idea that coherence is about the overlapping of the contents between several propositions. The more the contents of a proposition set overlap, the more coherent it is. As both measures are unsatisfactory, coherentists need to measure the notion of coherence in a more sophisticated way. Since it is generally accepted that coherence is the *mutual support* between the elements of a set, a possible approach is to take the degree of coherence of a set as the average degree of confirmation between all the pairs of elements in that set.

Based on the idea that coherence should be measured by the degree of mutual support between elements, Fitelson (2003) proposes a coherence measure based on the notion of *mutual confirmation*. His measure generates the degree of coherence of a set in two steps. First, we calculate the degree of confirmation between every pair of combinations of propositions in the set. Second, we calculate the average degree of confirmation between all such pairs in the set. With this two-step process, we may generate the average degree of mutual support between all possible combinations of propositions in a set and take this value as the degree of coherence of that set. To formally construct the desired measure, Fitelson first introduced a two-place function $F(p, p')$.[22] This function measures the degree a proposition p' confirms another proposition p,[23] which is defined as the following:

Definition 2.3.8. Fitelson's measure for support

Given any pair of propositions p and p' and a probability function Pr, the degree that p' confirms p, denoted by $F(p, p')$, is defined as:

[22] This function is a modification of the measure of factual support proposed by Kemeny and Oppenheim (1952).

[23] Here p and p' can also be sets of propositions. We can just take the conjunction of all elements of a set as a single proposition, and measure it in the way suggested.

$$F(p,p') =: \begin{cases} \dfrac{Pr(p'|p) - Pr(p'|\neg p)}{Pr(p'|p) + Pr(p'|\neg p)} & \text{if } p' \text{ does not entail } p \text{ and } p \text{ does not entail } p' \\ 1 & \text{if } p \text{ entailS } p' \text{ and } p \text{ is not inconsistent} \\ -1 & \text{if } p' \text{ entails } \neg p \end{cases}$$

With this function, Fitelson defines his coherence measure as follows:

Definition 2.3.9. Fitelson's coherence measure

Suppose S is a set of propositions $\{p_1, \cdots, p_n\}$. The degree of coherence of S is defined as:

$$\mathcal{C}_F(S) =: \frac{1}{[\![M]\!]} \sum_{\langle X,Y\rangle \in M} F(\bigwedge X, \bigwedge Y)$$

where M is the set of all pairs of non-empty non-overlapping subsets of S defined as $\{\langle X, Y\rangle | X, Y \in (\wp(S)/\varnothing) \wedge X \cap Y = \varnothing\}$ and $[\![M]\!]$ is the cardinality of M.

The set M contains all non-empty non-overlapping subsets of S which stands for all the pairs of combinations of propositions in a set. Take a set of propositions $S = \{p_1, p_2, p_3\}$ for example. According to the definition given, M is the set

$$\{\langle\{p_1\},\{p_2\}\rangle, \langle\{p_1\},\{p_3\}\rangle, \langle\{p_1\},\{p_2.p_3\}\rangle, \langle\{p_2\},\{p_1\}\rangle, \langle\{p_2\},\{p_3\}\rangle, \langle\{p_2\}, \{p_1,p_3\}\rangle, \langle\{p_3\},\{p_1\}\rangle, \langle\{p_3\},\{p_2\}\rangle, \langle\{p_3\},\{p_1,p_2\}\rangle, \langle\{p_1,p_2\},\{p_3\}\rangle, \langle\{p_1,p_3\},\{p_2\}\rangle, \langle\{p_2,p_3\},\{p_1\}\rangle\}.$$

Given any set of propositions, every non-empty subset $X \in \wp(S\backslash\varnothing)$ is confirmed or disconfirmed by another subset $Y \in \wp(S\backslash\varnothing)$ to a certain degree. We may derive the degree of confirmation by calculating the degree the conjunction of all elements in X confirms the conjunction of all elements in Y. By averaging the degree each X is confirmed or disconfirmed by every other non-empty element of $\wp(S\backslash\varnothing)$, one may measure the strength of mutual confirmation between all the subsets in S and take this value as the degree of coherence of S.

Fitelson's measure is free from the depth problem. Given any set, the degrees of coherence of all its subsets are taken into account. We may take the robbery case introduced earlier for example again. Recall that we have a set of testimonies $\{t_1, t_2, t_3\}$:

t_1: The criminal was either suspect 1, 2 or 3.

t_2: The criminal was either suspect 1, 3 or 4.

t_3: The criminal was either suspect 1, 2 or 4.

Suppose every suspect is equally susceptible, the probability of every testimony from t_1 to t_3 is $\frac{3}{8}$ as each testimony points out three suspects. The degree of coherence is the average value of the set $\{F(t_1,t_2)$, $F(t_1,t_3)$, $F(t_2,t_1)$, $F(t_2,t_3)$, $F(t_3,t_1)$, $F(t_3,t_2)$, $F(t_1,t_2 \wedge t_3)$, $F(t_2,t_1 \wedge t_3)$, $F(t_3,t_1 \wedge t_2)$, $F(t_1 \wedge t_2,t_3)$, $F(t_1 \wedge t_3,t_2)$, $F(t_2 \wedge t_3,t_1)\}$. With the function $F(X,Y)$ defined above, we can derive that

$$F(t_1,t_2) = F(t_1,t_3) = F(t_2,t_1) = F(t_2,t_3) = F(t_3,t_1) = F(t_3,t_2) = \frac{Pr(t_1|t_2) - Pr(t_1|\neg t_2)}{Pr(t_1|t_2) + Pr(t_1|\neg t_2)}$$

Since $Pr(t_1|t_2)$ is $\frac{2}{3}$ and $Pr(t_1|\neg t_2)$ is $\frac{1}{5}$, the value of $F(t_1,t_2)$ is $\frac{7}{13}$. We can further derive that

$$F(t_1,t_2 \wedge t_3) = F(t_2,t_3 \wedge t_1) = F(t_1,t_3 \wedge t_2) = \frac{1}{4}$$

$$F(t_1 \wedge t_2,t_3) = F(t_2 \wedge t_3,t_1) = F(t_1 \wedge t_3,t_2) = \frac{1}{5}$$

With these values, we may calculate the average degree of mutual support between all the combinations of propositions in the set, which is roughly 0.38. On the other hand, we have another set of testimonies $\{t'_1, t'_2, t'_3\}$:

t'_1: The criminal was either suspect 1, 2 or 3.

t'_2: The criminal was either suspect 1, 4 or 5.

t'_3: The criminal was either suspect 1, 6 or 7.

Similarly, since $Pr(t'_1) = Pr(t'_2) = Pr(t'_3) = \frac{3}{8}$. However, the degree of mutual support between these propositions are much weaker. We may derive that

$$F(t'_1,t'_2) = F(t'_1,t'_3) = F(t'_2,t'_1) = F(t'_2,t'_3) = F(t'_3,t'_1) = F(t'_3,t'_2) = \frac{-1}{11}$$

$$F(t'_1,t'_2 \wedge t'_3) = F(t'_2,t'_3 \wedge t'_1) = F(t'_1,t'_3 \wedge t'_2) = 1$$

$$F(t_1' \wedge t_2', t_3') = F(t_2' \wedge t_3', t_!') = F(t_1' \wedge t_3', t_2') = \frac{1}{5}$$

The average degree of mutual confirmation of the set $\{t_1', t_2', t_3'\}$ is roughly 0.254, which is lower than $C_F(\{t_1, t_2, t_3\})$. It can thus be seen that Fitelson's measure correctly captures our intuitive idea that $\{t_1, t_2, t_3\}$ is more coherent than $\{t_1', t_2', t_3'\}$.

Fitelson's measure is also immune to the problem of irrelevant additions. A proposition, if irrelevant to a set, does not confirm any proposition in that set. Hence, expanding a set with an irrelevant proposition would reduce the degree of confirmation between its subsets, and further reduce the degree of coherence of the whole set.

Similar to Olsson's measure, Fitelson's measure does have a maximal value for perfectly coherent proposition sets. For two perfectly coherent sets of propositions which differ in their specificity, Fitelson's measure renders them with equal maximal coherence. Again, this result is in accordance with our common understanding of coherence. For these reasons, Fitelson's coherence measure seems like an ideal way of measuring coherence.

Although Fitelson's measure looks quite promising, Bovens and Hartmann (2003) provide an example to cast doubt on its validity:

Example 2.3.9. Imagine two criminal scenarios: In the first one, there are 100 suspects, 6 of them play chess, 6 of them are from the Trobriand island but only one of the suspects is a Trobriand chess player. Let p_1 stand for 'the culprit is a chess player', p_2 for 'the culprit is a Trobriand', we may measure how strongly p_1 confirms p_2 and the other way round:

$$F(p_1, p_2) = F(p_2, p_1) = \frac{Pr(p_1|p_2) - Pr(p_1|\neg p_2)}{Pr(p_1|p_2) + Pr(p_1|\neg p_2)} = \frac{\frac{1}{6} - \frac{5}{94}}{\frac{1}{6} + \frac{5}{94}} = \frac{16}{31}.$$

The degree of coherence of $\{p_1, p_2\}$ is thus

$$C_F(\{p_1, p_2\}) = \frac{\frac{16}{31} + \frac{16}{31}}{2} = \frac{16}{31} \approx 0.52$$

The second case involves 100 suspects. 85 of them are rugby players. 85 of them are from Uganda. 80 rugby players are from Uganda. Let p_3 be 'The culprit is a rugby player', p_4 be 'The culprit is from Uganda',we may derive the following result:

$$F(p_4, p_3) = F(p_3, p_4) = \frac{Pr(p_3|p_4) - Pr(p_3|\neg p_4)}{Pr(p_3|p_4) + Pr(p_3|\neg p_4)} = \frac{\frac{80}{85} - \frac{1}{3}}{\frac{80}{85} + \frac{1}{3}} = \frac{31}{65}$$

Therefore, the coherence of the set $\{p_3, p_4\}$ is

$$C_F(\{p_3, p_4\}) = \frac{\frac{31}{65} + \frac{31}{65}}{2} \approx 0.48$$

Since the overlapping part between elements of $\{p_3, p_4\}$ is greater than the overlapping part between elements of $\{p_1, p_2\}$, $\{p_3, p_4\}$ is intuitively more coherent than $\{p_1, p_2\}$.[24] However, Fitelson's measure gives us the counterintuitive verdict that $\{p_1, p_2\}$ is more coherent than $\{p_3, p_4\}$.[25]

Fitelson's measure, like other measures, fails to correctly capture our conception of coherence. Thus, coherentists need to find out if this new measure can be refined to get rid of problematic cases.

2.4 Douven and Meijs' scheme of coherence measures

Douven and Meijs (2007) develop a scheme for coherence measures which, similar to Fitelson's measure, takes the degree of coherence of a set as the average degree of mutual confirmation between all its subsets. We may plug different Bayesian measures of confirmation in the scheme and generate a variety of coherence measures. Hence, it can be seen as a generalisation of Fitelson's measure.

[24]It is generally accepted by epistemologists that the degree of coherence of a set is determined by the extent its elements overlap.

[25]One may claim that this criticism is based on the idea that the coherence of a set is determined by the extent its elements overlap. If one gives up this idea, Fitelson's measure would no longer be problematic. One who aims to save Fitelson's measure this way must explain why $\{p_1, p_2\}$ is more coherent than $\{p_3, p_4\}$, which is rather counterintuitive.

They begin with an introduction to three major types of confirmation measures: the *difference measure*, *ratio measure* and *likelihood measure*.

Definition 2.4.1. Confirmation measures

Given a probability function Pr, the degree a proposition p' confirms p can be measured in the following ways:

Difference measure: $d(p,p') =: P(p|p') - P(p)$

Ratio measure: $r(p,p') =: \dfrac{P(p|p')}{P(p)}$

Likelihood measure: $l(p,p') =: \dfrac{P(p|p')}{P(p|\neg p')}$

These confirmation measures can be generalised to measure the degree of confirmation between sets by taking a set as the conjunction of all its elements:

Definition 2.4.2. Measure for the confirmation between sets

Based on the three measures, the degree a set S' confirms another set S can be measured as:

Difference measure: $d(S,S') =: Pr(\bigwedge S | \bigwedge S') - Pr(\bigwedge S)$

Ratio measure: $r(S,S') =: \dfrac{Pr(\bigwedge S | \bigwedge S')}{Pr(\bigwedge S)}$

Likelihood measure: $l(S,S') =: \dfrac{Pr(\bigwedge S | \bigwedge S')}{Pr(\bigwedge S | \overline{\bigwedge S'})}$

Based on these confirmation measures, we may define a scheme of coherence measures. Let d,r,l stand respectively for the three measures and m be the variable for different confirmation measures. Further define $[S]$ as the set of ordered pairs of non-empty, non-overlapping subsets of S. Put formally, $[S] = \{\langle S', S'' \rangle | S', S'' \subset S \backslash \{\varnothing\} \land S' \cap S'' = \varnothing\}$. With this definition, we can establish the following scheme of coherence measures:

Definition 2.4.3. The scheme for confirmation-based coherence measures

Given a set $S = \{p_1, \cdots, p_n\}$. With an ordering $\langle \hat{S}_1, \cdots, \hat{S}_{[\![S]\!]} \rangle$ of the members of $[S]$, the degree of coherence of S is given by the function

Table 2.1: The degree of coherence in the murder case

	Case 1.	Case 2.
$\mathcal{C}_{Sh}$	80.3	9
$\mathcal{C}_O$	0.0043	0.818
$\mathcal{C}_F$	0.97559	0.97561
$\mathcal{C}_d$	0.0084	0.8
$\mathcal{C}_r$	80.3	9
$\mathcal{C}_l$	80.9	81

$$\mathcal{C}_m(S) =: \frac{\sum_{i=1}^{[\![S]\!]} m(\hat{S}_i)}{[\![S]\!]}$$

for $m \in \{d, r, l\}$. Here $[\![S]\!]$ stands for the number of elements in the ordering.

For example, given a set $S^* = \{p_1, p_2\}$, the degree of coherence of S^* under the difference measure is

$$\mathcal{C}_d(S^*) = \frac{d(p_1, p_2) + d(p_2, p_1)}{[\![S]\!]} = \frac{P(p_1|p_2) - P(p_1) + P(p_2|p_1) - P(p_2)}{2}$$

Douven and Meijs (2007) claim that $\mathcal{C}_d$ is the least problematic coherence measure among all that have been proposed. To show this, they provide several test cases:

Example 2.4.1. Two murder cases

Case 1. A murder happened in a city with 10,000,000 inhabitants. 1,059 among them are Japanese, 1059 among them own Samurai swords while only 9 of them are Japanese owning Samurai swords.

Case 2. A murder happened on a street with 100 inhabitants. 10 of them are Japanese, 10 of them own Samurai swords, and 9 of them are Japanese who own Samurai swords.

Let p_1 stand for the proposition 'The murderer is Japanese' and p_2 for the proposition 'The murderer owns a Samurai sword.' Degrees of coherence of $S = \{p_1, p_2\}$ under different coherence measures in two cases are listed in Table 2.1.

In the first case, the two groups, namely the Japanese people and the Samurai sword owners, overlap to a very small extent. On the contrary, the two propositions overlap with each other very well in the second case. The intuition, hence, should be that the coherence of S in case 2 is much greater than the coherence of S in case 1. $\mathcal{C}_{Sh}$,$\mathcal{C}_F$,$\mathcal{C}_r$,$\mathcal{C}_l$ all fail to capture this intuition. Fitelson's measure $\mathcal{C}_F$ and the likelihood measure $\mathcal{C}_l$ render the set S with similar degree of coherence in both cases. $\mathcal{C}_r$ and $\mathcal{C}_{Sh}$ renders S with greater coherence in case 1 than in case 2, which is even worse than $\mathcal{C}_F$ and $\mathcal{C}_l$. Only $\mathcal{C}_d$ and $\mathcal{C}_O$ correctly represent the expected huge difference between the coherence of S in the two cases.

Another example, originally provided by Bovens and Hartmann (2003), shows that Olsson's measure leads to an unacceptable result:

Example 2.4.2. Consider two sets $S = \{p_1, p_2\}$ and $S' = \{p_1, p_2, p_3\}$ such that

(p_1) Our pet is a bird.

(p_2) Our pet is a ground dweller.

(p_3) Our pet is a penguin.

The probability distribution of these propositions is shown in Figure 2-1. Intuitively, S' is more coherent than S. Since only a few kinds of birds are ground dwellers, the set $\{p_1, p_2\}$ is not very coherent. Expanded with p_3, the set becomes more coherent. However, $\mathcal{C}_O$ generates the problematic result that the degree of coherence of S is equivalent to the degree of coherence of S'. Olsson's measure, thus, should be abandoned.

With these examples, Douven and Meijs (2007) show that $\mathcal{C}_d$ is the only coherence measure which does not generate unacceptable outcomes. The difference coherence measure $\mathcal{C}_d$, hence, should be taken as the correct way of measuring coherence.

Roche (2013) provides a variant to Douven and Meijs's coherence measure. He claims that although $\mathcal{C}_d$ is free from problems of other coherence measures, it generates unacceptable results for some other cases. Consider the following example:

Example 2.4.3. Suppose there are 10 suspects over a murder. Each of the suspects has equal probability of 0.1 of being the murderer. 6 of them have committed

Figure 2-1: The probability distribution of the penguin case.

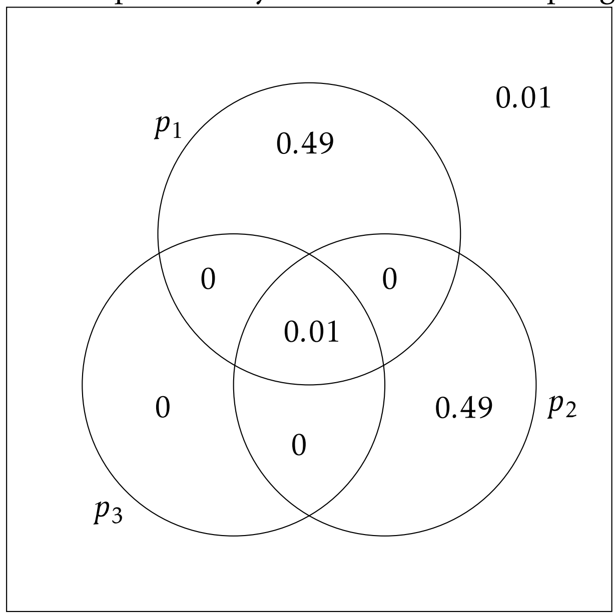

both pickpocketing and robbery, 2 of them have only committed pickpocketing and another 2 committed only robbery. Let the set S^* be $\{p_1, p_2\}$ and

(p_1) The murderer has committed robbery.

(p_2) The murderer has committed pickpocketing.

Since there are six suspects who committed both robbery and pickpocketing, the set S^* is intuitively quite coherent. However, according to $\mathcal{C}_d$, the coherence of S^* equals

$$\mathcal{C}_d(S^*) = \frac{d(p_1, p_2) + d(p_2, p_1)}{2} = -0.05$$

Since the outcome is negative, $\mathcal{C}_d$ indicates that S^* is incoherent. This violates our intuition that S^* is fairly coherent.

To get rid of this problem, Roche suggested we measure coherence with a new confirmation measure which differs from $\mathcal{C}_d, \mathcal{C}_r, \mathcal{C}_l$:

$$R(p, p') =: \begin{cases} Pr(p|p') & \text{if } p \text{ does not entail } p' \text{ and } p \text{ does not entail } \neg p'. \\ 1 & \text{if } p \text{ entails } p' \text{ and } p \text{ is consistent.} \\ 0 & \text{if } p \text{ entails } \neg p'. \end{cases}$$

By plugging the confirmation measure R in Douven and Meijs' scheme, we may obtain Roche's coherence measure $\mathcal{C}_R$ which is:

$$\mathcal{C}_R(S) =: \frac{\sum_{i=1}^{[\![S]\!]} R(\hat{S}_i)}{[\![S]\!]}$$

This measure is invulnerable to all the problematic cases for other confirmation-based coherence measures. Thus, Roche claimed that $\mathcal{C}_R$ is an ideal way for measuring coherence.

We can see Roche's coherence measure as just another variant of Fitelson's measure. The only difference between the two is that Roche does not measure the degree of mutual support between two propositions in terms of confirmation, but takes it as the probability of one proposition conditioned on another. Hence, Roche's measure is, like many others, a *confirmation-based* measure. The degree of coherence of a set of propositions, given Roche's measure, is positively correlated to the degree of confirmation between the propositions.

The approach Douven and Meijs propose can be further expanded by considering other confirmation measures. We may plug in a set of different confirmation measures into Douven and Meijs's scheme and generate the corresponding confirmation-based coherence measures.[26] By doing so, we would have a set of coherence measures that have different features. If we examine each confirmation-based measure with different examples, we can expect to find one which is impeccable and could be adopted as the correct coherence measure.

2.5 The problem of common cause

All these coherence measures introduced can be categorised as *relevance-sensitive*. That is, the core factor determining the coherence of a set, according to these measures, is the *relevance* between the propositions. Shogenji and Olsson's measures take the probability of the conjunction of a set of propositions as the primary factor. By measuring this quantity, we may know how coherent a set is. Here the

[26] For a list of significant Bayesian confirmation measures, see Table 2.2.

Table 2.2: Bayesian Confirmation measures

Measure	Proponent
$Pr(e\|h)/Pr(e)$	Keynes (1921)
$Pr(e \mid h) - Pr(e)$	Carnap$_1$(1950)
$Pr(e \wedge h) - Pr(e)P(h)$	Carnap$_2$(1950)
$[Pr(h \mid e) - Pr(h \mid \neg e)]/[Pr(h \mid e) + Pr(h \mid \neg e)]$	Kemeny and Oppenheim (1952)
$[Pr(e\|h) - Pr(e)]/[(1 - Pr(e))Pr(h)]$	Rescher (1958)
$Pr(h \mid e) - Pr(h \mid \neg e)$	Nozick (1981)
$Pr(h \mid e) - Pr(h)$	Mortimer (1988)
$Pr(e \mid h) - Pr(e \mid \neg h)$	Christensen (1999)
$min[Pr(e\|h), Pr(e)]/Pr(e) - min[Pr(\neg E\|H), P(\neg e)]/P(\neg e)$	Crupi et al. (2007)
$[log_2 Pr(e\|h) - log_2 Pr(e)]/ - log_2 Pr(e)$	Shogenji (2012)

probability of the conjunction can be understood as the relevance between these propositions. One may see that both Shogenji and Olsson's measures are relevance-sensitive. The confirmation-based measures go another route by taking the average degree of mutual confirmation between a set of propositions as the coherence of that set. Since the degree of confirmation is determined in part by the probability of the conjunction between propositions, what we are measuring is still the relevance between a set of propositions. We may thus conclude that all these measures are relevance-sensitive.

Although the idea of measuring coherence in term of relevance looks promising, Koscholke and Schippers (2019) point out that all the *relevance-sensitive* measures of coherence generate a counterintuitive outcome when we consider cases involving a common cause. Suppose that a proposition p_3 is the common cause of two relevant propositions p_1 and p_2. Given these assumptions, the relations between p_1, p_2 and p_3 satisfy the following conditions:

(1)] $Pr_{p_3}(p_1) > Pr(p_1)$ and $Pr_{p_3}(p_2) > Pr(p_2)$

(2)] $Pr(p_1 \wedge p_2) > Pr(p_1) \cdot Pr(p_2)$

(3)] $Pr_{p_3}(p_1 \wedge p_2) = Pr_{p_3}(p_1) \cdot Pr_{p_3}(p_2)$

where Pr_{p_3} is the probability function P conditioned on p_3. Put formally, it is just the proposition that $Pr_{p_3}(\cdot) = Pr(\cdot|p_3)$. Condition (1) states that p_1 and p_2 are both confirmed by their cause p_3. This condition is a natural consequence of the fact that p_3 is the cause of both p_1 and p_2. Condition (2) state that p_1 and p_2 are not probabilistically independent. As we have assumed, p_1 and p_2 are relevant to a certain extent. Hence, the probability of their conjunction is greater than the product of the probabilities of each. As p_3 is the cause of both p_1 and p_2, the relation between them must satisfy the third condition that, given p_3, p_1 and p_2 are no longer probabilistically relevant. That is, the relevance between p_1 and p_2 is *screened-off* by the presence of the common cause p_3. Before we know that p_1 and p_2 have a common cause, we take them as relevant as they are, in some occasions, jointly true. Given a common cause, we know that both p_1 and p_2 are consequences of this common cause and would no longer see them as mutually relevant. They are, given the common cause, two independent consequences of the common cause. It is not the case that p_2 is true by virtue of p_1 being true or the other way round. What makes p_1 and p_2 jointly true in some occasions is their common cause p_3. Hence, when we know that p_3, p_1 and p_2 would no longer be seen as relevant. This result can be presented formally. The probability of the conjunction of p_1 and p_3, conditioned on p_3, can be expanded as

$$Pr_{p_3}(p_1 \wedge p_2) = Pr(p_1 \wedge p_2|p_3) = \frac{Pr(p_1 \wedge p_2 \wedge p_3)}{Pr(p_3)}$$

Since p_3 is the common cause of p_1 and p_2, $Pr(p_1 \wedge p_2 \wedge p_3)$ is equivalent to $Pr(p_3)$. We may derive the result that $Pr_{p_3}(p_1 \wedge p_2)$ is 1. On the other hand, since p_3 causes both p_1 and p_2, $Pr_{p_3}(p_1)$ and $Pr_{p_3}(p_2)$ are both 1. Hence, $Pr_{p_3}(p_1 \wedge p_2)$ is equivalent to $Pr_{p_3}(p_1) \cdot Pr_{p_3}(p_2)$, which implies that p_1 is irrelevant to p_2.[27]

This kind of case reveals a problem of the relevance-sensitive coherence measures. Consider a set of two propositions $\{p_1, p_2\}$. The degree of coherence of the set, according to the relevance-sensitive measures, is the average of the degree p_1 confirms p_2 and the degree p_2 confirms p_1. When a common cause p_3 appears,

[27] Here is it implicitly assumed that p_3 is a deterministic, rather than probabilistic, cause of both p_1 and p_2.

one should take it as background knowledge and update the probability function with p_3. The coherence of the set $\{p_1, p_2\}$, thus, would be the average degree of mutual confirmation between p_1 and p_2 conditioned on p_3. Since the relevance between p_1 and p_2 is screened-off by the appearance of p_3, p_1 and p_2 do not confirm each other after p_3 is taken as background knowledge. The relevance-sensitive measures thereby generate the outcome that the set $\{p_1, p_2\}$ is neither coherent nor incoherent when p_3 appears, which is absurd. This problem can be illustrated with a concrete case. Suppose we adopt Mortimer's confirmation measure, the degree of coherence of $\{p_1, p_2\}$ given p_3 is the average of $Pr_{p_3}(p_2|p_1) - Pr_{p_3}(p_2)$ and $Pr_{p_3}(p_1|p_2) - Pr_{p_3}(p_1)$. From the fact that p_3 is the common cause of p_1 and p_2, p_1 and p_2 are independent under Pr_{p_3}, namely that $Pr_{p_3}(p_2|p_1)$ is equivalent to $Pr_{p_3}(p_2)$ and $Pr_{p_3}(p_1|p_2)$ is equivalent to $Pr_{p_3}(p_1)$. The degree of coherence of $\{p_1, p_2\}$, hence, is 0. If we do have the intuition that the set $\{p_1, p_2\}$ remains fairly coherent when a common cause p_3 appears, we would have to accept Koscholke and Schippers' claim that relevance-sensitive coherence measures are incorrect.

One may argue that the problem stems from an incorrect way of treating the new proposition p_3 and claim that, when a common cause p_3 appears, we should expand the set $\{p_1, p_2\}$ with p_3 instead of conditioning on it. That is, we should not update our probability function by taking p_3 as background knowledge, but should simply expand the original set $\{p_1, p_2\}$ with their common cause. By doing so, it would not be the case that p_1 is no longer relevant to p_2 in the presence of p_3 and, consequently, we may still measure the coherence of the set $\{p_1, p_2, , p_3\}$ in terms of the relevance between them. We can thus get rid of the problem Koscholke and Schippers point out.

To argue against this response, Koscholke and Schippers provide two reasons to strengthen their argument. First, compared to expanding the set, it is more natural to update the probability function with conditionalisation when a new proposition appears. Second, even if we choose not to condition on the common cause but expand the set with it, the relevance-sensitive measures may still yield counterintuitive results. Koscholke and Schippers show that when a set is ex-

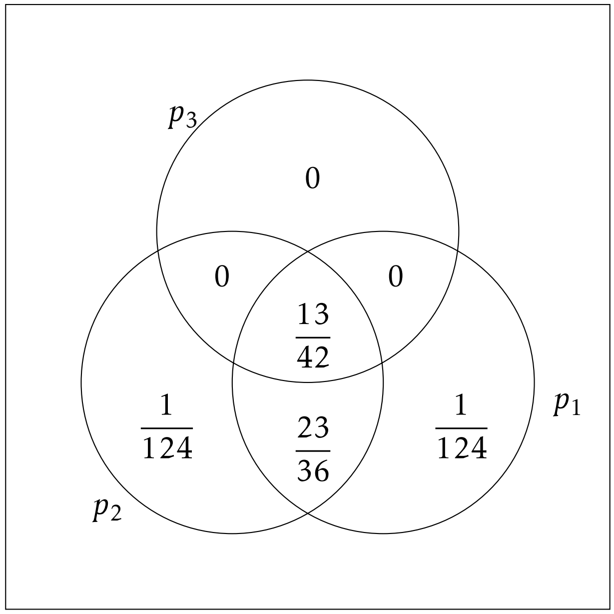

Figure 2-2: A common cause scenario where $\{p_1, p_2, p_3\}$ is less coherent than $\{p_!, p_2\}$.

panded with a common cause of its members, its degree of coherence, given the relevance-sensitive measures, may decrease. Since a cause is supposed to be very coherent with its effects, this result again shows that the relevance-sensitive measures generate mistaken results. This problem can be illustrated with a concrete case. Consider the probability distribution in Figure 2-2.[28] Given this distribution, the degree of coherence of $\{p_1, p_2, p_3\}$ is lower than the coherence of $\{p_1, p_2\}$ given most confirmation measures. For instance, if we adopt Mortimer's confirmation measure, the degree of coherence of $\{p_1, p_2, p_3\}$ is approximately 0.30, while the degree of coherence of $\{p_1, p_2\}$ is 0.35. Since a cause is supposed to be coherent with its effect, $\{p_1, p_2, p_3\}$ should not be less coherent than $\{p_1, p_2\}$. Koscholke and Schippers therefore conclude that the relevance-sensitive measures fail to correctly capture the notion of coherence.

[28] This is the same case presented in Koscholke and Schippers (2019).

2.6 The requirements for an ideal coherence measure

After a survey of the attempts of measuring coherence, we may list some requirements a proper coherence measure should meet.

Definition 2.6.1. Size-informativeness

A coherence measure $\mathcal{C}$ is size-informative if, for any two sets of propositions S and S' and any probability function Pr, if S has more elements than S' then, other things being equal, $\mathcal{C}(S) > \mathcal{C}(S')$.[29]

Given any two sets of propositions of difference sizes, if they are equal in all other aspects, the one with more elements should be taken as more coherent. Suppose there are two sets of propositions. One contains two propositions, another contains two hundred. If the extent the propositions overlap in the first set is equal to the extent they overlap in the second set, we would consider the second set as more coherent as it contains more elements. The underlying idea is fairly straightforward. Compared with a small set, it is less likely for the propositions of a bigger set to agree with each other. When it happens that the contents of a big set agree with each other to the same extent as the agreement of the contents of a smaller set, the bigger set should be taken as more coherent.

Definition 2.6.2. Specificity-informativeness

A coherence measure $\mathcal{C}$ is size-informative if, for any two sets of propositions S and S' and any probability function Pr, if the elements of S are more specific than the elements of S', then, other things being equal, $\mathcal{C}(S) > \mathcal{C}(S')$.

This requirement states that if the information conveyed by a set of propositions is very specific, namely of low probability, then it should be taken as more coherent than another set which provides less specific information. The underlying idea of this requirement is the same as the one for size-informativeness. As it is less likely for a set of highly specific propositions to agree with each other, when it

[29]The crucial factors that need to be held fixed here are the extent of the overlapping part between propositions and the specificity of these propositions.

happens to be so, such a set should be regarded as highly coherent. However, this requirement is not equally convincing as size-informativeness. One might consider a set of equivalent but unspecific propositions as more coherent than a set of specific but not completely equivalent propositions. Whether this requirement should be adopted, thus, is still an open question.

Definition 2.6.3. Maximal coherence

A coherence measure $\mathcal{C}$ satisfies maximal coherence if it renders any set of equivalent propositions maximally coherent.

This requirement may be more doubtful than the former ones. At first glance, it does not seem problematic that a coherence measure has no maximal value. However, if a coherence measure $\mathcal{C}$ fails to satisfy this condition, then, according to $\mathcal{C}$, it is always possible to make a set more coherent by expanding it with a new proposition. We may consider a simple case in which one measures the degree of coherence of a set S with Shogenji's measure $\mathcal{C}_{Sh}$ which has no upper bound. Suppose one expands a set S with a non-tautologous proposition p equivalent to some element in S. The numerator of the measure, namely the probability of the conjunction of all propositions in S, remains the same. But since p is not a tautology, the denominator of the measure decreases. Thus, the overall degree of coherence increases.

If the primary purpose of measuring the coherence of a set is to tell whether the propositions in the set are justified, adopting a coherence measure having no upper bound implies that the propositions can never be fully justified. That is, for any set of propositions, no matter how justified they are, it is always possible to make them more justified by expanding the set with some new propositions.[30] Similarly, if the purpose of measuring coherence is to measure the likelihood of truth of a set of propositions, adopting a measure with no upper bound implies

[30] It is possible to get rid of this problem by setting a threshold for justification. That is, one may take a value x to be the threshold and claim that if the degree of coherence of a set S is greater than x, then S is justified. However, if the coherence measure one adopts has no upper bound, it would be rather difficult to set the threshold since the range is too broad.

that the propositions can never be absolutely true. For any set of true propositions, no matter how likely to be true they are, there exists an arbitrary set which does not differ significantly from the set, but is more likely to be true. Hence, an ideal coherence measure should have an upper bound. If there is no upper bound, we would have to accept the undesirable consequence that the notion of coherence cannot correctly inform us about the degree of justification of a set, nor its likelihood of truth.

Definition 2.6.4. Irrelevant additions

A coherence measure $\mathcal{C}$ is sensitive to irrelevant additions if, given a set of propositions S and a proposition p which is irrelevant to all the propositions in S, $\mathcal{C}(S) > \mathcal{C}(S \cup \{p\})$.

Here what is meant by an 'irrelevant proposition' is one which is independent from every proposition in the original set. This requirement is in accordance with BonJour's fifth criterion of coherence. Given a set of propositions S, an irrelevant proposition p does not confirm any element in that set. Hence, when the irrelevant proposition p is added to the set S, the overall coherence of the set S should decrease.

All these requirements reflects some of our intuitive understanding of the notion of coherence. For a measure to correctly capture the notion of coherence, it should meet all the requirements listed. However, some of these requirements are incompatible. For example, the requirement of maximal coherence contradicts the requirement of size-informativeness. Suppose that a coherence measure $\mathcal{C}$ has an upper bound. We may arbitrarily construct a set S such that S is maximally coherent. For any set S' bigger than S, it is at best equally coherent to S since S is maximally coherent. Hence, any coherence measure which has an upper bound would fail to be size-informative. Similarly, the requirement is incompatible with the requirement of specificity-informativeness. Since there is no measure which satisfies all these requirement, we could give up the idea of finding one which perfectly satisfies all requirements but instead aim at finding a measure which satisfies some weaker version of these requirements.

2.7 Saving the relevance-sensitive coherence measures

We have seen that, given all these requirements and test cases, the measures that are the most likely to meet most of them are the confirmation-based measures, namely the variants of Fitelson's measure. However, Koscholke and Schippers' criticism has shown that all relevance-sensitive measures, including the confirmation-based ones, fail to generate the correct result when a common cause appears. Their argument is significant for the following reasons. First, Koscholke and Schippers' criticism is based on the appearance of a common cause of a set, which is a very general phenomenon. One cannot get rid of this problem simply by arbitrarily excluding the sets involving propositions with a common cause. Second, relevance is a crucial, if not the most crucial, factor when measuring coherence. Their criticism does not only imply that we have to abandon most coherence measures that have been developed, but also that, when trying to develop new coherence measures, the relevance between propositions should not be taken as a factor. Formally speaking, we can no longer appeal to the probability of the conjunction between propositions as a factor determining the coherence of a set. We would then have to find another notion which characterises the notion of coherence, which is apparently very difficult. For these reasons, Koscholke and Schippers' argument seems to bring an end to the project of searching for an ideal coherence measure.

Although Koscholke and Schippers' argument is convincing, there is still hope for a relevance-sensitive coherence measure. The strategy I am going to take here is to show that the second problem they point out, namely the problem of expanding a set with its common cause, can be avoided by adopting a different average function. After introducing this function and show how to avoid the problem, I will further argue that the first problem does not really pose a threat to the relevance-sensitive measures. With both arguments rejected, the confirmation-based coherence measures can be saved.

To solve the problem of expanding a set with a common cause, a thorough analysis of the problem is required. What the problem shows is that expanding a set with a common cause of its elements may make the set less coherent. This is actually not a surprising result. Let us, for sake of simplicity, call the relation between a proposition p and another proposition which p confirms a *confirmation relation*. Recall that the confirmation-based measures are motivated by the idea that the degree of coherence of a set is the average degree of mutual confirmation between its elements. When one expands a set with a new proposition, the new proposition may confirm many propositions or combinations of propositions in the original set. Thus, the number of confirmation relations that we need to consider increases significantly. For a set of two elements, there are only two relations that we need to take into account. But for a set of three elements, there are twelve confirmation relations to consider. If one expands a set of two elements with a new proposition, the number of the confirmation relations to consider increases significantly. For the new set to be of greater coherence, the average strength of all these new confirmation relations, brought in by the new proposition, needs to be higher than the average degree of confirmation of the set prior to the expansion. To see this, consider a set S_1 containing two elements p_1 and p_2. Suppose that the degree of coherence of S_1, according to a confirmation-based measure, is k. What this means is that the linear average of the degree p_1 confirms p_2 and the degree p_2 confirms p_1 is k. When the set is expanded with a common cause p_3, there are ten more confirmation relations in the set. Thus, we need to consider many more combinations of propositions.[31] For the new set, call it S_2, to be more coherent than S_1, the average degree of all ten new confirmation relations need to be greater than k. However, this requirement cannot always be satisfied as nothing guarantees that the average of these new confirmation relations would be greater than k. Hence, we may derive the strange result that expanding a set with a common cause of all the propositions in the set may lead to a drop in its coherence.

[31] Check the example in Section 2 to see all the confirmation relations that need to be taken into account.

This problem does not only occur when a set is expanded with a common cause, but also happens with a proposition confirming every element. That is, given a set of propositions, it is possible to find a new proposition which confirms every element of the set but, when the set is expanded with this proposition, the overall coherence of the set decreases. We may call this the problem of set expansion. Consider the probability distribution in Figure 2-3. The degree to which p_3 confirms both p_1 and p_2 is stronger than the degree of mutual confirmation of p_1 and p_2. Based on the idea that coherence is the strength of mutual confirmation between the elements, $\{p_1, p_2, p_3\}$ is supposed to be more coherent than $\{p_1, p_2\}$. However, the degree of coherence of $\{p_1, p_2, p_3\}$, given all the confirmation measures in Table 2.2, is lower than the degree of coherence of $\{p_1, p_2\}$. For example, adopting Kemeny and Oppenheim's confirmation measure leads to the result that the degree of coherence of $\{p_1, p_2\}$ is approximately 0.2, while the coherence of $\{p_1, p_2, p_3\}$ is approximately 0.12. Given this case, we can see the core of the problem of expansion more clearly. Although the degree to which p_3 confirms both p_1 and p_2 is stronger than the mutual confirmation between p_1 and p_2, expanding the set $\{p_1, p_2\}$ with p_3 brings in some much weaker confirmation relations and results in a decrease of the overall degree of coherence.

The root of the problem of set expansion is that the confirmation-based measures take only the average strength, but not the number of confirmation relations into account. When p_3 is added to the set $\{p_1, p_2\}$, it brings in ten more positive confirmation relations. According to BonJour's third coherence criterion, the degree of coherence of this set should increase. When a set is expanded with a new proposition which confirms many elements of a set, the set should, in ordinary cases, become more coherent as there are now many more confirmation relations.[32] Since the confirmation-based measures of coherence take the average degree of confirmation as the coherence of a set, they fail to capture this intuition and generate unacceptable results. The solution to the problem of set expansion,

[32]It is possible for there to be a case such that a proposition p_3 confirms both p_1 and p_2 but not $p_1 \wedge p_2$. In such a case, the degree of coherence of $\{p_1, p_2, p_3\}$ could be lower than $\{p_1, p_2\}$. However, the problem I point out remains true for other cases.

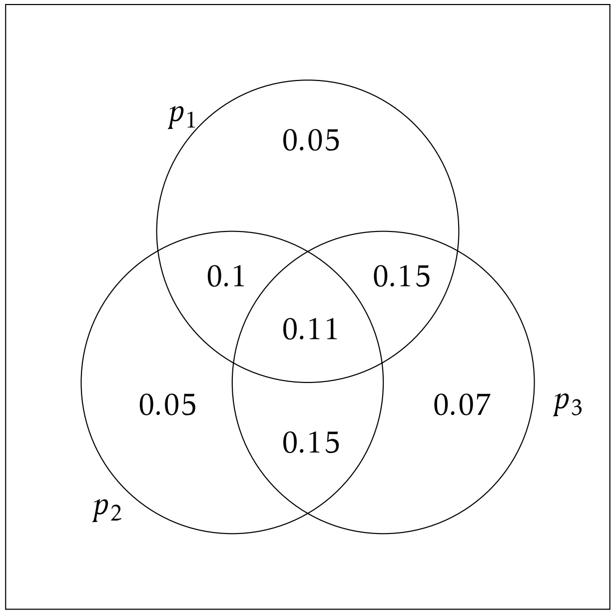

Figure 2-3: A probability distribution where p_3 confirms both p_1 and p_2

hence, is to construct a measure such that the degree of coherence increases not only with the average degree of confirmation, but also with the number of confirmation relations.

2.8 Changing the way of averaging

A widely discussed function in the literature of opinion pooling is *multiplicative averaging*.[33] This function takes the normalised product of a set of values as its outcome and has the property of *synergy*: when the input values are above a certain threshold, the outcome generated by this function would be higher than the linear average of the inputs. Moreover, *synergy* gets stronger when the number of inputs increases. This feature allows us to capture the idea that the degree of coherence of a set should increase with the number of confirmation relations in the set.[34]

To measure the degree of coherence with multiplicative averaging, what we

[33]See Dietrich (2010); Easwaran et al. (2016) and Dietrich and List (2016).

[34]The result is correct in some cases. If p_3 confirms both p_1 and p_2 but does not confirm $p_1 \wedge p_2$, it should be correct that the set $\{p_1, p_2, p_3\}$ is less coherent than $\{p_1, p_2\}$. Nevertheless, the problem remains significant. We may think of a case in which a proposition p_i confirms a set of propositions $p_1, ..., p_n$ and their conjunction $p_1 \wedge ... \wedge p_n$. Expanding the set $\{p_1, ...p_n\}$ with p_i may still lead to a decrease of the overall coherence of the set because of the reason mentioned here.

need is a simple modification of the scheme of confirmation-based measures. Instead of summing up the strength of confirmation relations and divide it by the number of confirmation relations, we take the product of the degrees of them and normalise it with a factor ζ. Suppose S is a set of propositions, the scheme of multiplicative confirmation-based measures $\mathcal{C}_{\mathcal{M}}$ can be formulated as:

$$\mathcal{C}_{\mathcal{M}}(S) =: \zeta \prod_{\langle S_i, S_j \rangle \in M} c(\bigwedge S_i, \bigwedge S_j)$$

where M is the set of all pairs of non-empty and non-overlapping subsets of S defined as before and c is a Bayesian confirmation measure which generates the degree $\wedge S_j$ confirms $\wedge S_i$. ζ is a normalisation factor which guarantees that the outcome satisfies the probability axioms:

$$\zeta = \frac{1}{\prod_{\langle S_i, S_j \rangle \in M} c(\bigwedge S_i, \bigwedge S_j) + \prod_{\langle S_i, S_j \rangle \in M} (1 - c(\bigwedge S_i, \bigwedge S_j))}$$

This normalisation factor guarantees that the value generated by $\mathcal{C}_{\mathcal{M}}$ falls in the interval $[0, 1]$. There are several reasons for us to normalise the outcome. First, since most confirmation measures takes $[0, 1]$ as range, the product of the degrees of confirmation decreases with the number of input values. That is, the more inputs there are, the lower the average of them. Such a feature contradicts with the ordinary understanding that the degree of coherence of a set increases with the number of confirmation relations in the set. If we normalise the result, this problem can be eliminated.

An even stronger reason for normalising the result is that normalisation brings us with a significant feature that the values *synergise* with each other. If, among all the input values, there are multiple values greater than 0.5, the outcome would be greater than the linear average of these values. A toy example may illustrate the difference between the multiplicative confirmation-based measures and the linear confirmation-based measures. For the sake of simplicity, we only calculate the pairwise coherence, namely the coherence between singletons here. Consider a set

$\{p_1, p_2\}$ such that for some confirmation measure c, $c(p_1, p_2) = c(p_2, p_1) = 0.7$. According to the standard confirmation-based measure, the degree of coherence of $\{p_1, p_2\}$ is 0.7. Suppose we add a new proposition p_3 to the set such that $c(p_1, p_3) = c(p_2, p_3) = c(p_3, p_2) = c(p_3, p_1) = 0.6$, the degree of coherence of $\{p_1, p_2, p_3\}$ would be 0.64, which is lower than the degree of coherence of $\{p_1, p_2\}$. As we have seen, this is the primary problem of confirmation-based measures of coherence. Although p_3 does confirm both p_1 and p_2, the overall coherence of the set, after expanding with p_3, is lower than the original set $\{p_1, p_2\}$. If we calculate the degree of coherence with multiplicative averaging, the coherence of $\{p_1, p_2\}$ would be approximately 0.84, whereas the coherence of $\{p_1, p_2, p_3\}$ is roughly 0.96. We can thus see that the positive confirmation relations synergize with each other and generate the more intuitive result that the coherence of $\{p_1, p_2, p_3\}$ is greater than the coherence of $\{p_1, p_2\}$. The number of confirmation relations in a set, given multiplicative averaging, is positively correlated with its overall degree of coherence. With the property of synergy, we may get rid of the problem of set expansion.

There are, however, several technical worries concerning the multiplicative confirmation-based measures that need to be addressed. The first problem is that synergy goes in both directions. Without any modification, the *break-even* point of the multiplicative average function is 0.5. When the input values are greater than the threshold 0.5, their multiplicative average would be greater than their linear average. But if the values are below 0.5, the outcome would be lower than the linear average. Hence, if the degrees of the confirmation relations involved in a set are all below 0.5, the number of confirmation relations would be negatively correlated to the degree of coherence of a set. This is an undesirable result, as the degree of confirmation between two propositions is, in many cases, below 0.5. Adopting the multiplicative coherence measures would thus lead to the absurd consequence that the more confirmation relations there are, the less coherent the set is. Take the probability distribution in Figure 2-3 for example. If we plug in Kemeny and Oppenheim's confirmation measure and calculate the degree of coherence of the set $\{p_1, p_2, p_3\}$, the outcome would be $7.93347917 \times 10^{-14}$, while the degree of coherence of $\{p_1, p_2\}$ is 0.061. In this case, the set $\{p_1, p_2, p_3\}$ involves far more confirmation relations than the set $\{p_1, p_2\}$. However, as the degree of these confirmation relations are all below 0.5, the effect of synergy goes downwards and generates the undesirable result that $\{p_1, p_2, p_3\}$ is much less coherent than $\{p_1, p_2\}$.

Another problem of this new coherence measure is that there may be propositions which disconfirm other propositions. In such cases, the degree of confirmation would be lower than 0. If we adopt multiplicative averaging, the overall degree of coherence would be negative when the number of input values below 0 is odd. This is obviously an undesirable result. If one highly coherent set happens to contain a proposition which slightly disconfirms another proposition, the overall degree of coherence of the set becomes negative.

To solve the two problems, we should move the input values from $[-1, 1]$ to $[0, 1]$ with the following function:

$$f(x) = 0.5x + 0.5$$

The use of this function can be illustrated with a simple example. Suppose that the degree p_1 confirms p_2 is -0.2. Given $f(x)$, $c(p_1, p_2)$ becomes 0.4. With $f(x)$, we may move every possible value in the range $[-1, 1]$ to the range $[0, 1]$. The break-even point, as previously explained, is 0.5. After calculating the multiplicative average of all these values, we may use the inverse of $f(x)$ to move the outcome back to the original range. If we calculate the degree of coherence this way, the problem of negative degree of confirmation can be avoided. Also, we do not need to worry that the break-even point is too high.[35]

If one is unhappy with the natural break-even point of 0.5, there is also a way to shift it. Easwaran et al. (2016) show that this could be done by adding another value to calibrate the function. Suppose we want to shift the break-even point to k, we can multiply every input value with $(1-k)$. The multiplicative confirmation-based measures should thus be reformulated as the following:

[35]Thanks to Catrin Campbell-Moore, one of my examiners, for suggesting me to revise the measure this way.

$$\mathcal{C}_{\mathcal{M}}(S) =: \zeta \prod_{\langle S_i, S_j\rangle \in M} (1-k)\cdot c(\bigwedge S_i, \bigwedge S_j)$$

The normalisation factor should be revised accordingly as:

$$\zeta = \frac{1}{\prod_{\langle S_i,S_j\rangle \in M}(1-k)\cdot c(\bigwedge S_i, \bigwedge S_j) + \prod_{\langle S_i,S_j\rangle \in M} k(1-c(\bigwedge S_i, \bigwedge S_j))}$$

With such a modification, we may set the break-even point at k. Any value greater than k, when averaged with some other values also greater than k, would synergise and generates an outcome greater than their linear average. Suppose we plug in Mortimer's confirmation measure for c and set the break-even point to 0.01, the degree of coherence of $\{p_1, p_2, p_3\}$ in Figure 2-2 would be 0.9999, while the degree of coherence of $\{p_1, p_2\}$ is 0.9289. The multiplicative confirmation-based coherence measures generates the correct result that $\{p_1, p_2, p_3\}$ is more coherent than $\{p_1, p_2\}$.

There is another technical issue concerning the choice of confirmation measures. It should be noted that multiplicative averaging is originally used to calculate the average of a set of credences, which are values within the interval $(0, 1)$.[36] If the input value is beyond this scope, there would be undesirable formal consequences. Since Keynes (1921) and Rescher's (1958) confirmation measures generate values greater than 1 for some cases, they cannot be adopted to measure the degree of coherence of a set.[37] Also, Kemeny and Oppenheim (1952), Crupi et al. (2007) and Shogenji's (2012) measures take the degree a cause confirms its consequence as 1. This feature leads to the consequence that the normalisation factor equals to one for every set which contains a single cause-consequence pair.[38] We would then be unable to normalise the result. For this reason, we should also

[36]Here I assume the requirement of *regularity*, namely that one's credences over propositions should not take 1 and 0 as its value except logical truths and falsehoods.

[37]It is possible, of course, to move the values to the interval $[0, 1]$ with a linear function. However, since these confirmation measures have different ranges, there is no unified function to move the values. It would thus be more complicated to measure coherence with these confirmation measures.

[38]If the degree of confirmation between a cause and its consequence is one, when we calculate the multiplicative average of a set containing a cause-consequence pair, the denominator of the normalisation factor of the multiplicative average would also be one. The value of the normalisation factor would thus be one.

Table 2: Results of adopting the multiplicative confirmation-based measures

Measure	$\{p_1, p_2\}$	$\{p_1, p_2, p_3\}$
Carnap_1	0.2022	0.6151
Carnap_2	0.0836	0.2977
Nozick	0.3363	0.8572
Mortimer	0.2022	0.6151
Christensen	0.3363	0.8572

abandon these measures. The options we are left with, thus, are Carnap, Nozick, Mortimer,and Christensen's confirmation measures.

Apart from all these, there remains a kind of case that may lead to technical problems: independent propositions. For any pair of independent propositions, the degree of confirmation between them is 0. Given the multiplicative confirmation-based measures, the degree of coherence of any set containing a pair of independent propositions would be 0. This result is somewhat strange. A set including a pair of independent propositions can still be coherent, as long as other elements bear strong mutual confirmation. To avoid erroneous outcomes, this kind of case needs to be treated separately.

To avoid this potential problem, the whole process of measuring the coherence of a set S should be carried out in several steps. First, we collect all the possible combinations of propositions in S. Second, we calculate the degree of mutual confirmation between all these combinations and sort the results into two groups: the non-independent ones and the independent ones. For any two propositions p and p', if the degree p confirms p' is not zero, we categorise the value to the first group. On the contrary, if the degree p confirms p' is zero, we put the value into the group of independent propositions. After sorting all the outcomes into the two groups, the next step is to move all the values in the non-independent group from the interval $[-1, 1]$ to the interval $[0, 1]$. Once we calculate the multiplicative average of all these values, we move the result back to the interval $[-1, 1]$. By doing so, we may derive the overall coherence of the set.

This measure gets us out of the problem of set expansion. Consider the probability distribution in Figure 2-3.[39] The results of adopting the multiplicative confirmation-based measure is listed in Table 2. As we can see from the table, all the confirmation measures generate the more intuitive result that the set $\{p_1, p_2, p_3\}$ is more coherent than the set $\{p_1, p_2\}$. The problem of set expansion, hence, can be solved by adopting the multiplicative confirmation-based measures of coherence.

There is a possible doubt concerning this process of measuring coherence. In the method of measuring coherence I proposed in this section, the confirmation relations between independent propositions are ignored. If we do so, we may get the counterintuitive result that expanding a set with an irrelevant proposition does not make the set less coherent. In other words, the problem of irrelevant addition rises again. To solve this problem, a possible strategy is to count the number of such relations, and calibrate the final result with this value. For example, suppose that there are x null confirmation relations in a set S. We may, after deriving the degree of coherence of S with the confirmation-based measures, calibrate the final result with x. By doing so, these null confirmation relations may still change the overall degree of coherence of S.

One may claim that it is *ad hoc* to save the confirmation-based coherence measures this way. Apart from solving the problem of set expansion, generating the degree of coherence of a set with multiplicative averaging does not seem to provide us with any additional feature which better captures the notion of coherence. This criticism seems to overlook the crucial advantage of this approach. Measuring coherence with multiplicative averaging does not only solves the problem of set expansion, but also successfully captures the intuition that the coherence of a set increases with the number of confirmation relations. As it is a better tool for us to measure coherence, adopting such a tool should not be deemed *ad hoc*.

[39] Here I do not consider Koscholke and Schippers's example. In their original example, as presented in 2-2, the common cause p_3 disconfirms the conjunction of p_1 and p_2. Hence, it is natural in their case that $\{p_1, p_2\}$ is more coherent than $\{p_1, p_2, p_3\}$.

2.9 Back to the problems of common causes

Since the problem of set expansion can be solved by adopting the multiplicative confirmation-based measures, we may now come back to reexamine Koscholke and Schippers' first criticism that the confirmation-based measures generate counter-intuitive results when a common cause is taken as background knowledge.

A brief review of their argument is required before I defend my solution here. Suppose there is a set of propositions $\{p_1,...,p_n\}$, a probability function Pr which assigns a value to each proposition and a proposition p which is the common cause of $p_1,...,p_n$. When p is given as background knowledge, we should update Pr by conditioning it on p. By doing so, we may obtain a new function $Pr_p(\cdot)$ which is equivalent to $Pr(\cdot|p)$. However, according to the updated probability function Pr_p, the propositions $p_1,...,p_n$ are not relevant since their relevance is screened-off by the appearance of the common cause. Based on such an observation, Koscholke and Schippers claim that the relevance-sensitive measures of coherence are fallacious. If we take the relevance between a set of propositions as a factor determining its coherence, we would have to accept the result that, once the common cause of a set appears, the set becomes neither coherent nor incoherent. Since a set, intuitively, does not become less coherent when a common cause appears, Koscholke and Schippers conclude that the relevance-sensitive measures are incorrect.

Koscholke and Schippers' criticism misses the point of measuring coherence and hence fails to undermine the relevance-sensitive measures of coherence. As stated in the very beginning, when one measures the coherence of a set, say $\{p_1,...,p_n\}$, one either aims to know whether the propositions in the set are justified or whether they are true. Suppose that one aims to know whether these propositions are true by measuring the coherence of $\{p_1,...,p_n\}$. When their common cause p appears, the set $\{p_1,...,p_n\}$ becomes neither coherent nor incoherent. Consequently, we do not know whether these propositions are true. However, it is pointless to measure the coherence of $\{p_1,...,p_n\}$ when p appears. If one knows that a common cause p is true, its consequences $p_1,...,p_2$ would follow. Hence, there is no need to measure the coherence of $\{p_1, ..., p_n\}$. Similarly, if what one wants to know is whether all the elements of $\{p_1, ..., p_n\}$ are justified, one does not have to measure the coherence between them either. Since p is given as background knowledge, $\{p_1, ..., p_n\}$ are all well justified.[40]

There are cases in which p does not necessitate p_1 and p_2. In this kind of case, the consequences do not follow from the cause and we may still have a reason to measure the coherence between the propositions in question. My response here, given the existence of such cases, would be incomplete. However, if the common cause we consider here is a probabilistic one, Koscholke and Schippers's argument would not work. The relevance between the consequences and the common cause, given standard probability calculus, would not be screened-off. In Koscholke and Schippers's original argument, when p_3, a common cause of p_1 and p_2, is taken as background knowledge, the relevance between p_1 and p_2 would be screened-off since both $Pr_{p_3}(p_1 \wedge p_2)$ and $Pr_{p_3}(p_1)Pr_{p_3}(p_2)$ equal to one. In a case where p_3 does not necessitate p_1 and p_2, it is possible that $Pr_{p_3}(p_1 \wedge p_2)$ differs from $Pr_{p_3}(p_1)Pr_{p_3}(p_2)$. Hence, Koscholke and Schippers's argument would fail in such cases.

In sum, although the appearance of a common cause does screen-off the relevance between the elements in a set, it does not imply that the confirmation-based measures are fallacious. As long as coherence is taken as a notion which grounds other notions and is not valuable *per se*, the confirmation-based measures are free from Koscholke and Schippers (2019)' criticism. If one agrees that people do not measure the coherence of a set for the sake of coherence, Koscholke and Schippers' criticism does not hold.

2.10 Conclusion

Koscholke and Schippers provide two arguments to show that when the common

[40]One may argue that a cause may not justify its effect. Here I take it as a separate issue topic is beyond the scope of this paper.

causes of a set appears, the relevance-sensitive measures generate counterintuitive results. I have shown in this paper that, first, their first criticism misses the point of measuring coherence. In a situation where the common cause of a set is present, there is no need for one to measure the coherence of that set. Second, although the second problem they point out does pose threat to the relevance-sensitive measures, we may get rid of it by measuring the degree of coherence of a set with the multiplicative averaging function. Both their criticisms, hence, can be dissolved. What remains to be explored are the other features of this new coherence measure. If one can show that the multiplicative confirmation-based measures are conducive to other desirable properties, we may have a proper coherence measure that is useful in some aspects. If they do not, we should move on to search for other possible coherence measures that correctly captures all our intuitions. In any case, the feature of synergy, given by the multiplicative average function, is valuable and should be kept while measuring the coherence of a set.

Chapter 3

Beyond Linear Conciliation

3.1 Introduction

The Conciliatory View of peer disagreement holds that when one disagrees with their epistemic peers, one should compromise with their peers by revising their credence in the proposition at issue (Christensen, 2007; Elga, 2007a; Feldman, 2006).[1] Despite its intuitive plausibility, many epistemologists find this view untenable. Some claim that conciliating is a self-abasing act (Pettit, 2006; van Inwagen, 1996), while others argue that it is not truth-conducive (Kelly, 2010). Among the arguments against the Conciliatory View, the ones that focus on its formal deficiencies deserve special attention. It has been pointed out that there are three deficiencies in the Conciliatory View. First, it does not commute with the Bayesian rule of *conditionalisation* since the outcome of conciliation is partially determined by whether one updates before making conciliation (Fitelson and Jehle, 2009; Wilson, 2010). Second, the Conciliatory View is path dependent (Gardiner, 2014). In a case where one makes multiple conciliations with their peers at different times, the final result is determined by the temporal order in which one makes the conciliations with each peer. Third, the Conciliatory View does not preserve one's judgement concerning the relevance between propositions (Elkin and Wheeler, 2018b).

[1] The formulation here involves the notion of *credence*. I will provide a reason for formulating this way in section 3.

That is, one's judgement that two propositions are irrelevant may not be well preserved after they conciliates with others. Since all three features bring about some unacceptable consequences, the Conciliatory View seems seriously flawed.

The criticisms concerning the formal deficiencies of the Conciliatory View, however, are misguided. Studies in opinion pooling have shown that it is linear averaging, the function that is generally adopted to make conciliation, that has the three formal deficiencies. Since linear averaging is by no means the only legitimate way to make conciliation, one can save the Conciliatory View by adopting nonlinear average functions. To show this, I will first introduce the Conciliatory View and its most prominent form, the Equal Weight view. After reformulating the two views in a formal framework, I will demonstrate the three formal deficiencies of the Conciliatory View and explain how can they be solved by making conciliation with *geometric* and *multiplicative* average functions. To further justify the approach of nonlinear conciliation, I will point out that some features of the nonlinear average functions better reflect our intuitions about disagreement. As a result, some misconceptions in the study of peer disagreement may be clarified. The conclusion, hence, is that we should embrace a pluralistic view concerning conciliation and give up the idea that there is a single conciliating rule which can be applied in every case of peer disagreement. Conciliationists should develop a taxonomy of different kinds of disagreement and find out the proper average function to apply for each kind.

3.2 Disagreement between peers

Consider the following scenario: Albert, a brilliant historian who specialises in the Victorian era, wants to solve the mystery of Jack the Ripper. Having spent years reviewing all the evidence relevant to the Whitechapel murders, he becomes very confident that it was the Polish barber Aaron Kosminski who committed the atrocity. However, his colleague Bridget firmly believes the opposite. Like Albert, Bridget is also an expert in Victorian Britain who has reviewed all the evidence related to the murders. Unlike Albert, Bridget considers it extremely unlikely that Kosminski is Jack the Ripper. Knowing that Bridget, as a historian, is as good as himself, how should Albert respond to their disagreement concerning this controversy?

Albert and Bridget's case is a typical example of *peer disagreement*. Two symmetric assumptions need to be highlighted for one to see why these kinds of circumstances constitute a real problem for epistemologists. First, since it is assumed that Bridget is Albert's *epistemic peer*, they are symmetric with their reliability concerning this issue. We may unpack this assumption a bit further by assuming that Albert and Bridget were both educated in prestigious universities, trained in similar ways and had equally outstanding track records. Given these conditions, they are equally likely to have the correct credence concerning the historical fact in question with the same body of evidence. Precisely because of such peerhood, any reason which allows Albert to cast doubt on Bridget's credences concerning specific historical facts should also allow Bridget to cast doubt on Albert's credences. Hence, Albert cannot dismiss Bridget's disagreement but must take it seriously. Second, they are symmetric with the evidence they possess respectively. The evidence Albert has is to a great extent, if not exactly, identical to the evidence Bridget has.[2] Because of the parity of evidence between them, one cannot expect either of them to change their mind after reviewing the evidence available to themselves. It is the two symmetries that make peer disagreement a thorny problem for social epistemologists.

There are many variants of the standard peer disagreement case that can be generated by revising the two symmetric assumptions. Regarding the symmetry of reliability, one may specify the reliability of the interlocutors involved. When all the interlocutors are highly reliable, we have the case of expert disagreement. On the contrary, when all the interlocutors are unreliable, we have the case of layperson disagreement. Regarding the symmetry of evidence, one may assume

[2]Some epistemologists consider cases in which the individuals have pieces of evidence that are only accessible to themselves. Here it is assumed that this kind of evidence does not play a crucial role in their judgement.

different degrees of information they have concerning the evidence their peers possess. In some cases, all the interlocutors are fully aware that others have the same body of evidence as they do. Each of them not only knows that the other interlocutors do possess some evidence but also knows the content of the evidence. In a slightly different case, the interlocutors do not share the full content of the evidence. What they know is that all the interlocutors possess evidence of the same strength.[3] If we relax the notion of peerhood further, we may derive cases in which the interlocutors know that the other interlocutors have some pieces of evidence, but have no information concerning the strength of the evidence others possess. The most radical case would be one where each interlocutor has no idea whether their peers have any piece of evidence. Due to the highly varied nature of all these different kinds of cases, we should consider the possibility of dealing with each case in different ways.

One may doubt whether these variants still count as peer disagreement, especially cases of the latter kind in which the symmetry of evidence is weakened. Indeed, without assuming the symmetry of evidence, the problem of peer disagreement might be thought to be overly easy and loses its philosophical significance. However, as King (2012) points out, a perfect case of peer disagreement, namely one which satisfies both symmetric assumptions, is rather rare. The ultimate goal of the study of peer disagreement should not be finding a solution that is only applicable to the hardest cases. If we take it to be the final goal, the study of peer disagreement may be somewhat trivial as the result is extremely limited. What we should aim for, instead, is finding out a solution that applies to a wider set of cases.[4] Following this line of thought, then, cases without the symmetry of evidence, though deviant from the standard cases of peer disagreement, are still worth to be discussed.

[3]This definition of peer disagreement can be found in Matheson (2014).

[4]Matheson (2014) also holds the view that the purpose of studying peer disagreement is to find a solution applicable to other cases.

3.3 The Conciliatory View and the Equal Weight View

One of the most widely accepted solutions to peer disagreement is the *Conciliatory View*: Whenever a disagreement occurs among a group of epistemic peers, each one involved should make a compromise with the others. The motivation is reasonably straightforward: Since no one is epistemically impeccable, it is always possible for an individual to have incorrect credence in a proposition. Thus, when one disagrees with their epistemic peers, one should realise that they might have made a mistake and revise their credence in the proposition in doubt. We can perhaps see the plausibility of this view from another perspective. In the face of peer disagreement, a person who refuses to change their credence can be criticised for ignoring their own fallibility. To be epistemically modest, one should choose to conciliate when involved in a disagreement.

A question immediately follows: How, in practice, should one make a conciliation? Conciliation can be made in many different ways. One can conciliate either by giving up their credence entirely and accept whatever their peers say, or by making a minimal revision of their original credence concerning the proposition in question. Although the two ways lead to remarkably different outcomes, they both count as conciliating. If the Conciliatory View suggests a variety of ways of dealing with disagreement, it would be overly general and hence lack significance. Thus, conciliationists cannot merely claim that conciliation is the proper solution to peer disagreement, but have to provide precise instructions concerning how people should revise their credence in the face of disagreement.

One way to establish a more elaborate formulation of the Conciliatory View is to reconsider the core assumption of peer disagreement. Recall that all the individuals involved in disagreement are assumed as equally reliable. All their credences concerning the proposition in dispute, therefore, are equally likely to be correct and should be treated in the same way. Thus, the most promising form of the Conciliatory View is to assign equal weight to all the disagreeing peers' credences and take the average as the outcome of conciliation. Call this the Equal Weight View (henceforth the EWV).[5] Take the Jack the Ripper controversy for example. In the given scenario, Albert is nearly sure that Aaron Kosminski is the one who committed the Whitechapel murders, while Bridget is almost certain that Kosminski is not. If Albert adopts the EWV and assigns equal weight to both his credence and Bridget's, he would come to have moderate credence in the proposition under dispute. In the following sections, I will take the EWV as a view representing other Conciliatory Views since they share all the important formal properties which we are going to discuss.

The core idea of the EWV, as stated, is to respect the fact that all the individuals involved in a genuine peer disagreement are equally reliable. Hence, everyone's opinion should be given equal weight. From this core claim, one may infer that the EWV should be formulated within a Bayesian framework which represents an individual's doxastic state as credence.[6] To see this, consider a case where three individuals A, B and C disagree over the truth of a proposition. Both A and B believe that the proposition is true, while C disbelieves. Suppose that they all adopt the EWV and intend to revise their credence in the proposition, what would the outcome be? There is no answer if we adopt the traditional tripartite conception of belief which says that one either believes, disbelieves or suspends judgement concerning a proposition. The individuals should not jointly believe or disbelieve the proposition since there is no consensus among them. The remaining option, namely suspend judgement, is also incorrect. A joint suspension of judgement concerning the issue implies that C's disbelief outweighs A and B's beliefs, which leads to a violation of the EWV. If instead of taking this overly coarse-grained framework, we choose to represent their doxastic states in terms of credences, this problem can be solved easily. In brief, since credences can be properly split and represent all the possible outcomes of conciliation, the EWV should be formulated formally. Following the same line of reasoning, any non-trivial form of the Con-

[5]For a detailed introduction to the EWV, see Feldman (2006); Elga (2007a) and Christensen (2007).

[6]The EWV can also be formulated in terms of imprecise probability. See Elkin and Wheeler (2018b).

ciliatory View is also fundamentally formal.[7]

Since the EWV is essentially formal, we need to formulate it within a formal framework. The first item required is the set of all the propositions that can be subject to disagreement. Take a non-empty finite set of possible worlds Ω as primitive, a single proposition can be defined as a subset of Ω, which is the set of all the worlds where the proposition is true. The set of all propositions, following this definition, should be defined as 2^{Ω}, namely the power set of Ω. An individual's credence over each proposition can hence be defined as a function. Let $Pr_i(\cdot)$ represent the credence function of individual i which assigns a value in the interval $[0,1]$ to every possible world ω in the set Ω, where the sum total of values across the worlds in Ω is 1. As a direct result, every credence function also assigns a value to every proposition in 2^{Ω}, namely the sum of its values for the constituent worlds. Every credence function is formally a probability function.

The most widely accepted version of the EWV takes the linear average of the disagreeing individuals' credences as the outcome of conciliation, which can be formulated as the following:

Definition 3.3.1. The Linear Equal Weight View

Given a case in which n individuals $1, ..., n$ disagree over the proposition $P \in 2^{\Omega}$, the outcome of conciliation should be

$$\sum_{i=1}^{n} \frac{1}{n} Pr_i(P).$$

That is, one may divide the sum of individual credences with the number of individuals involved and take the outcome as the result of conciliation.

An example may illustrate how the Linear EWV works. Suppose that, in the Jack the Ripper case, Albert's credence in Kosminski being Jack the Ripper (P) is

[7] Kelly (2010, p.187) has another argument supporting the claim that the EWV should be presented within a formal framework.

0.9 while Bridget's is 0.1. Taking their credence functions respectively as $Pr_1(\cdot)$ and $Pr_2(\cdot)$, the result of assigning equal weight to both their credences in P is:

$$\frac{1}{2}Pr_1(P) + \frac{1}{2}Pr_2(P) = \frac{1}{2} \cdot 0.1 + \frac{1}{2} \cdot 0.9 = 0.5.$$

According to the Linear EWV, they should have 0.5 credence in Kosminski being the real murderer. This result correctly captures our intuition that they should both have moderate credence over P after conciliating with each other.

3.4 Three formal deficiencies of the Linear Conciliatory View

3.4.1 Non-commutativity with conditionalisation

Although the Linear EWV seems to be an ideal rule for making conciliation, it has three major formal deficiencies. First of all, it fails to commute with the Bayesian rule of *conditionalisation*. As one of the defining features of Bayesianism, conditionalisation suggests any individual who acquires a piece of evidence E update their credence by conditioning their credence in any proposition on E. Apart from being a handy and plausible rule for updating credence, conditionalisation is also the Bayesian norm for the diachronic coherence of one's credences. If one does not update their credence with the rule of conditionalisation upon receiving new evidence, one takes the risk of having diachronically incoherent credences over a set of propositions.[8]

It has been pointed out by many epistemologists that the Linear EWV fails to commute with conditionalisation as switching the order between conciliating and updating leads to different outcomes (Fitelson and Jehle, 2009; Wilson, 2010). To illustrate, consider the Jack the Ripper example. Let $Pr_1(\cdot)$ stand for Albert's credence function, $Pr_2(\cdot)$ for Bridget's and $Pr_{1+2}(\cdot)$ for their joint credence function

[8]Although the rule of conditionalisation is a core Bayesian norm, it is not undoubtedly true. Some philosophers argue that one may reject conditionalisation yet still be rational.(Bacchus et al., 1990; Hild, 1998; Arntzenius, 2003) For the projects aiming at vindicating conditionalisation, see Greaves and Wallace (2005); Briggs and Pettigrew (2020) and Pettigrew (2019).

obtained by conciliating with the Linear EWV. Assume that they have the following credences over the proposition P and a piece of evidence E:

P	E	$Pr_1(\cdot)$	$Pr_2(\cdot)$	$Pr_{1+2}(\cdot)$
T	T	0.285	0.05	0.1675
T	F	0.615	0.05	0.3325
F	T	0.015	0.15	0.0825
F	F	0.085	0.75	0.4175

Table 3.1: Albert and Bridget's credences over P and E

The first row stands for the possible world in which P and E are both true. The value given by $Pr_1(\cdot)$, hence, is Albert's credence in $P\&E$. Given this credence distribution, if Albert and Bridget decide to first make a conciliation concerning their credences in P and E, their joint credence in $P\&E$ and E would respectively be the following:

$$Pr_{1+2}(P\&E) = \frac{1}{2}Pr_1(P\&E) + \frac{1}{2}Pr_2(P\&E) = \frac{1}{2}(0.285 + 0.05) = 0.1675,$$

$$Pr_{1+2}(E) = \frac{1}{2}Pr_1(E) + \frac{1}{2}Pr_2(E) = \frac{1}{2}((0.285 + 0.015) + (0.05 + 0.15)) = 0.25.$$

By applying the rule of conditionalisation, we may derive their joint credence of P conditional on E:

$$Pr_{1+2}(P|E) = \frac{Pr_{1+2}(P\&E)}{Pr_{1+2}(E)} = \frac{\frac{1}{2}(Pr_1(P\&E) + Pr_2(P\&E))}{\frac{1}{2}(Pr_1(E) + Pr_2(E))} = \frac{0.1675}{0.25} = 0.67.$$

On the other hand, if Albert and Bridget choose to first update their credences on the evidence E respectively, they would have the following credences in P conditional on E:

$$Pr_1(P|E) = \frac{Pr_1(P\&E)}{Pr_1(E)} = \frac{0.285}{0.3} = 0.95$$

$$Pr_2(P|E) = \frac{Pr_2(P\&E)}{Pr_2(E)} = \frac{0.05}{0.2} = 0.25.$$

If they, after updating with E respectively, make a conciliation, their joint credence would be the linear average of the two values:

$$Pr_{1+2}(P|E) = \frac{1}{2}Pr_1(P|E) + \frac{1}{2}Pr_2(P|E)) = \frac{1}{2}(0.95 + 0.25) = 0.6.$$

This case shows that if the individuals make conciliation with the Linear EWV, the order of conciliating and updating determines the outcome of their conciliation.

What is wrong with the Linear EWV failing to commute with conditionalisation? First, the order of updating and conciliating is, in most cases, irrelevant to the disagreement itself. It is unacceptable to let an irrelevant factor determine the outcome of conciliation. Suppose that, in the Jack the Ripper case, Albert and Bridget decide to conciliate with each other and have moderate credence in Kosminski being the murderer. After the conciliation, they find a ledger containing the names of suspects that they have never seen. Both Albert and Bridget update with this new piece of evidence and come to have a new credence in Kosminski being the murderer. In a different case, they find the ledger and each update their credence before they conciliate. Since conciliation does not commute with conditionalisation, the result of conciliation in the second case may differ from the result in the first case. This is rather problematic since the time they receive the ledger is irrelevant to whether Kosminski is Jack the Ripper. If the purpose of conciliating is to make disagreeing peers come to have credences that are as close to the truth as possible, we should not take an irrelevant factor into account. Second, if the outcome of conciliation is sensitive to the time of updating, the individuals involved in disagreement would be vulnerable to manipulation. Suppose that, in the Jack the Ripper case, another historian Claire possesses a new piece of evidence E' which is unknown to both Albert and Bridget. Further suppose that Claire intends the outcome of Albert and Bridget's conciliation to be as close to 1 as possible. Knowing that Albert and Bridget are about to conciliate, Claire would choose to conceal E' until the conciliation is made. By revealing E' to Albert and

Bridget after they make a conciliation, she can make their joint credence closer to 1, which is the result she intends. This feature is again undesirable, as an ideal rule should leave no space for manipulation. In other words, conciliating with the Linear EWV puts people under the risk of being manipulated. The Linear EWV, therefore, is a problematic way of resolving peer disagreement.

3.4.2 Path dependence

Apart from failing to commute with conditionalisation, the Linear EWV also fails to be path independent (Gardiner, 2014). That is, if one makes multiple conciliations according to the Linear EWV, the final outcome would be sensitive to the order of conciliation. Imagine a revised Jack the Ripper case in which another historian Claire has credence 0.7 in the proposition P that Kosminski is Jack the Ripper. Suppose, as before, that Albert's credence in P is 0.9 while Bridget's is 0.1. If Albert first makes a conciliation with Claire and subsequently with Bridget, his credence in P, according to the Linear EWV, would be 0.45. Let Claire's credence function be $Pr_3(\cdot)$, the process can be formulated as the following:

$$Pr_{1+3}(P) = \frac{0.9 + 0.7}{2} = 0.8,$$

$$Pr_{1+3+2}(P) = \frac{0.8 + 0.1}{2} = 0.45.$$

After Albert conciliates with Claire, they obtain the joint credence function $P_{1+3}(\cdot)$ which assigns P with credence 0.8. When they subsequently make a conciliation with Bridget, the resulting credence function is $Pr_{1+3+2}(\cdot)$ which assigns 0.45 to P. In a slightly different story, Albert first conciliates with Bridget and subsequently with Claire. His credence in P, given the Linear EWV, would be 0.6. That is,

$$Pr_{1+2}(P) = \frac{0.9 + 0.1}{2} = 0.5,$$

$$Pr_{1+2+3}(P) = \frac{0.5 + 0.7}{2} = 0.6.$$

Given this example, we may see that the order Albert conciliates with his peers determines the outcome of conciliation.

Path dependence is undesirable for two reasons. First, the order one makes conciliation with their peers is, in most cases, irrelevant to the proposition in dispute and should not affect the outcome of conciliation in any way. Consider the Jack the Ripper case again. Whether Kosminski committed the Whitechapel murders is a historical event that has already obtained. The temporal order Albert makes conciliation with his peers has nothing to do with the real identity of Jack the Ripper. Hence, it would be absurd to let the order Albert conciliates be a factor determining the outcome of conciliation. Moreover, the result shows that the Linear EWV is diachronically incoherent. The core claim of the EWV is that every individual's credence should be treated equally. If one adopts the Linear EWV and makes conciliation with their peers one at a time, the testimonies that are received at some early stage would be weighted less than the testimonies that come later, as earlier testimonies have been mixed up with new testimonies for more times. In other words, the importance of a single testimony would be gradually diluted as new pieces of testimonial evidence emerge. As this result violates the core idea of the EWV that one should weight all their peers' credences equally, the Linear EWV runs the risk of being self-refuting.

One might attempt to defend the Linear EWV by arguing that people can assign weights in a more sophisticated way. Instead of reassigning equal weight in every single conciliation, an individual should keep track of all the conciliations they have made and assign the correct weight to new peers. In the given example, after Albert conciliates with Bridget, he should be aware of the fact that his credence has already been mixed up with Bridget's prior credence. When Albert subsequently meets Claire and makes another conciliation, he should know that the correct weight to assign to Claire's credence is one-third rather than a half. In short, if Albert is smart enough, he should know the correct weight to assign to his peers' credence. The problem of path dependence only occurs to stubborn individuals.

Although the solution is rather convincing, it is difficult to implement in most cases of peer disagreement. In a simple case which involves only a small number of individuals, it is relatively easy for one to remember the details of all the conciliations they have made. However, in a slightly complicated case, it would be overly demanding to ask an individual to memorise all the details of every conciliation that has taken place. Suppose that Albert, in the searching of the true identity of Jack the Ripper, consults twenty peers at different times. It would be extremely tough for him to remember every peer's prior credence and assign the correct weight to each of them. Hence, it is pragmatically impossible for one to always conciliate with their peers this way.[9] For this reason, path dependence remains a defect of the Linear EWV.

Another way of saving the Linear EWV is to deny that one would ever need to make multiple conciliations with their peers. In the given scenario, after Albert conciliates with Bridget, it can be said that Claire is no longer his peer since, after the conciliating with a peer, Albert's credence has become more likely to be correct. Hence, Albert is no longer Claire's peer and does not have to assign equal weight to her credence. The problem of path dependence, therefore, would not occur in the first place. However, this solution is based on a volatile notion of peerhood. In any ordinary case of peer disagreement, one does not become superior immediately after they conciliate with a single peer. Imagine a case where a panel of scientists aims to make a joint decision concerning government policies. Two among them have a private conversation and, right after they make a conciliation, announce that their credence over the proposition in question is the correct one since they have become superior to all other scientists. If one agrees that this case is absurd, they would have to admit that a proper notion of peerhood should prevent this kind of things from happening. For this reason, we may conclude that path dependence remains a shortcoming of the EWV.

[9]Gardiner (2014) provides a thorough review of the possible solutions to this problem.

	$Pr_1(\cdot\wedge\cdot)$	$Pr_2(\cdot\wedge\cdot)$	$Pr_{1+2}(\cdot\wedge\cdot)$	$Pr_{1+2}(\cdot)Pr_{1+2}(\cdot)$
K, R	0.08	0.33	0.205	0.19
$K, \neg R$	0.12	0.27	0.195	0.21
$\neg K, R$	0.32	0.22	0.27	0.285
$\neg K, \neg R$	0.48	0.18	0.33	0.315

Table 3.3: Albert ($Pr_1(\cdot)$) and Bridget's ($Pr_2(\cdot)$) credences in K and R.

3.4.3 The problem of independence preservation

Another crucial problem of the Linear Conciliatory View is that it fails to preserve their judgement of independence (Bradley et al., 2014; Elkin and Wheeler, 2018b). Let us consider a concrete example provided by Elkin and Wheeler: Albert, based on his evidence, does not find it likely that it will rain in Kyoto tomorrow (K) and has 0.2 credence in the proposition. On the other hand, with some evidence, he thinks that there may be an unpublished volume of *The Lord of the Rings* (R) and has 0.4 credence in its existence. Based on the same background knowledge and evidence, Bridget has 0.6 credence in the former proposition and 0.55 in the latter. Further suppose that Albert judges K and R as mutually independent. His credence in the conjunction of K and R, thus, is equal to the product of his credence in K and his credence in R, namely 0.16. Bridget also judges K and R as mutually independent and has 0.33 credence in the conjunction of K and R. When they conciliate with each other according to the Linear EWV, they come to accept a new credence function $Pr_{1+2}(\cdot)$ which does not assign equal value to the conjunction of K and R and the product of their respective joint credence in K and R. That is, K and R are not independent according to their joint credence function $Pr_{1+2}(\cdot)$. The distribution of their credences is presented in Table 3.3.[10]

Intuitively, this result is strange. As both Albert and Bridget judge K and R as mutually independent, the joint judgement that K and R are relevant, revealed by the joint credence function $Pr_{1+2}(\cdot)$, comes from nowhere. A more serious problem is that failing to preserve ones' judgement of independence may make them

[10]Thanks to an anonymous referee for indicating that this problem should be discussed.

irrational. Suppose that Albert and Bridget do make conciliation according to the Linear EWV and adopt $Pr_{1+2}(\cdot)$ as their new credence function. A witty gambler can sell them two bets: The first bet costs them £20.5 on the condition that they receive £100 from the gambler if K and R are both true. The second bet costs them £33 on the condition that the gambler pays them £100 if K and R are both false. Since, according to the function $Pr_{1+2}(\cdot \wedge \cdot)$, Albert and Bridget's joint credence in both K and R being true is 0.205, the expected return of the first bet for them is 0. Similarly, since their joint credence in K and R both being false is 0.33, the expected return of the second bet is also 0. As both bets are fair for Albert and Bridget, they will accept both bets.

The gambler can go further and sell Albert and Bridget two other bets: One costs them £21 and pays back £100 if K is true and R is false. Another costs them £28.5 and pays back £100 if K is false and R is true. Since both Albert and Bridget judge K and R as independent, their credence in K and R can also be represented by the function $Pr_{1+2}(\cdot)Pr_{1+2}(\cdot)$. Given this credence function, the two new bets are fair. Hence, Albert and Bridget would also accept this proposal. With the four bets, it is guaranteed that Albert and Bridget are going to lose £3 to the gambler. This case shows that adopting the Linear EWV may lead to a sure loss of money, which indicates that the individuals are irrational. The linear EWV, therefore, should be abandoned.

3.5 Opinion pooling and peer disagreement

One way of saving the Conciliatory View is to change the way we conciliate while retaining the core idea that one should make conciliation with their peers. Since it has been proved that there are nonlinear average functions that are free from the three formal problems, we may adopt the nonlinear average functions to make conciliation and thereby derive alternative Conciliatory Views that are free from the three formal deficiencies.[11]

The study of probabilistic opinion pooling aims to answer one central question: Given a profile of credence functions across a set of individuals, what is the proper way of merging them into a single joint credence function which satisfies specific requirements? To correctly respond to this question, philosophers have examined a variety of possibilities and proposed different pooling functions. Since, as we have seen, the Conciliatory View is essentially formal, we may apply the results in the study of opinion pooling to save this view. In the following sections, I will show that nonlinear average functions are free from the three formal deficiencies of the Linear EWV. As the formal results derived can be generalised to other forms of conciliation, the Conciliatory View can be rescued.

It has been proven that the geometric averaging does commute with conditionalisation (Genest, 1984). Hence, we may adopt geometric averaging to make conciliation and thereby derive a rule of conciliation that commutes with conditionalisation:

Definition 3.5.1. The Geometric Conciliatory View

Given a case in which n individuals $1,\dots,n$ disagree over the proposition P, the Geometric Conciliatory View suggests the individuals involved to have credence

$$\sum_{\omega\in P} c\cdot Pr_1(\omega)^{\alpha_1}\cdots Pr_n(\omega)^{\alpha_n}$$

in the proposition P, where the factors $\alpha_1,\dots,\alpha_n$ are the weights assigned to each individual which sum up to 1, and the constant c is a normalisation factor which guarantees that the sum of joint credences across all worlds equals 1:

$$c=\frac{1}{\sum_{\omega'\in\Omega} Pr_1(\omega')^{\alpha_1}\cdots Pr_n(\omega')^{\alpha_n}}.$$

A technical point to be highlighted is that the credence functions here take a single

[11] This idea is first proposed by Martini et al. (2013). However, they did not fully explore the possible outcomes of adopting nonlinear functions to make conciliation.

possible world, rather than a proposition, as its input.[12] The credence each individual has in a proposition can be derived by summing up the credence over each possible worlds included in the proposition.[13]

To show that the Geometric Conciliatory View does commute with conditionalisation, we need to reformulate the rule of conditionalisation.[14] Given any piece of evidence E, the information it carries allows the individuals to derive a likelihood function L which assigns either the value 1 or 0 to each possible world. If a world ω_k is in E, then $L(\cdot)$ assigns 1 to ω_k. An individual i can then update their credence function P as

$$Pr_i^L(\omega) = \frac{Pr_i(\omega)L(\omega)}{\sum_{\omega' \in \Omega} Pr_i(\omega')L(\omega')}$$

which is equivalent to

$$Pr_i^E(\omega) = \frac{Pr_i(\omega)Pr_i(E|\omega)}{\sum_{\omega' \in \Omega} Pr_i(\omega')Pr_i(E|\omega')}.$$

Since this formula is the rule of conditionalisation in a different form, updating with a likelihood function is equivalent to updating with conditionalisation.

With the new form of conditionalisation, we can now show that the Geometric Conciliatory View commutes with conditionalisation. Take $Pr^L_{Pr_1,\ldots,Pr_n}(\cdot)$ as the joint credence function obtained in the case where the individuals first conciliate with all others and subsequently update with a likelihood function L. $Pr_{Pr_1^L,\ldots,Pr_n^L}(\cdot)$, on the other hand, stands for the joint credence function obtained in the case where the individuals first update with the likelihood function $L(\cdot)$ and then conciliate with the others. To show that the two functions yield the same outcome, it suffices to show that the two functions are proportional since any two probability functions that are proportional to one another must be identical. Consider the

[12] This is in line with the definition that each credence function assigns a value to each possible world.

[13] For a detailed explanation, see Dietrich and List (2016, p.8).

[14] The likelihood function, compared to Bayesian conditionalisation, is actually a more general update rule, as it can take any possible value, rather than the discrete values 1 and 0. From this perspective, Bayesian conditionalisation is a limiting case of updating by a likelihood function. See Dietrich and List (2016).

case in which the individuals make a conciliation first. Suppose there are n individuals who disagree over their credence of a single possible world ω. If they adopt the Geometric Conciliatory View and make a conciliation, the result would be $Pr_1(\omega)^{\alpha_1} \cdots Pr_n(\omega)^{\alpha_n}$. When they jointly receive a piece of evidence E and derive a likelihood function $L(\cdot)$ from it, the outcome of updating would then be proportional to $Pr_1(\omega)^{\alpha_1} \cdots Pr_n(\omega)^{\alpha_n} L(\omega)$. On the other hand, when the individuals first update their credence functions with $L(\cdot)$, we have a set of updated credence functions $Pr_i^L(\omega)$. Each function in the set is equivalent to $Pr_i(\omega)L(\omega)$. The individuals later conciliate with the others and get the result $(Pr_1(\omega)L(\omega))^{\alpha_1} \cdots (Pr_n(\omega)L(\omega))^{\alpha_n}$, which is equivalent to $Pr_1(\omega)^{\alpha_1} \cdots Pr_n(\omega)^{\alpha_n} \cdot L(\omega)^{\alpha_1+\cdots+\alpha_n}$. Since $\alpha_1, ..., \alpha_n$ sum up to one, this result is proportional to $Pr_1(\omega)^{\alpha_1} \cdots Pr_n(\omega)^{\alpha_n} L(\omega)$, namely the result in the first case. We may hence conclude that the Geometric Conciliatory View does commute with conditionalisation.[15]

The problem of path dependence can be solved by adopting the Multiplicative Conciliatory View which suggests individuals to conciliate with multiplicative averaging:[16]

Definition 3.5.2. The Multiplicative Conciliatory View

Given a case in which n individuals $1, ..., n$ disagree over the proposition P, the Multiplicative EWV suggests the individuals involved to have credence

$$\sum_{w \in P} c \cdot Pr_1(\omega) \cdots Pr_n(\omega)$$

in the proposition P. The constant c is a normalisation factor which guarantees that the sum of joint credences of all propositions equals to 1.

$$c = \frac{1}{\sum_{\omega' \in \Omega} Pr_1(\omega') \cdots Pr_n(\omega')}.$$

Two points should be noted: First, the Multiplicative Conciliatory View, like the

[15]This simplified proof is presented by Dietrich and List (2016). The original proof that geometric averaging commutes with conditionalisation is provided by Genest (1984).

[16]Dietrich (2010) and Easwaran et al. (2016) both provide detailed analysis of the features of the multiplicative average function.

Geometric Conciliatory View, also takes a single world as the input instead of a proposition. Second, this view assigns equal weight to all the credence functions involved. Hence, we do not have to explicitly assign weight to each credence function.

It need not be proved that the Multiplicative Conciliatory View is path independent. Since multiplication is associative, it trivially holds that any conciliation made this way is also associative.[17] The problem of path dependence, hence, would not occur for anyone conciliating with the Multiplicative Conciliatory View.[18][19]

Both the Geometric EWV and the Multiplicative EWV preserve one's judgement that two events are independent. The proof is also trivial. According to both the Geometric and the Multiplicative Conciliatory View, individuals should multiply their credences to make conciliation. Hence, the joint credence function of a group of peers assign a value which is equivalent to the product of the credences of each individual. The problem of independence preservation can thus be solved. This point can be illustrated with a toy example: Consider a case involving two peers whose credences are represented respectively by the function $Pr_1(\cdot)$ and $Pr_2(\cdot)$. Let their joint credence function be $Pr_{1+2}(\cdot)$. To show that both the Geometric and the Multiplicative Conciliatory View preserve their judgement that two propositions are independent, what we need to prove is that $Pr_{1+2}(\cdot \wedge \cdot)$ is equivalent to $Pr_{1+2}(\cdot)Pr_{1+2}(\cdot)$ when the inputs are independent. According to the

[17]One may wonder whether the process of normalisation make the Multiplicative Conciliatory View path dependent. To show that it does not, it suffices to show that the normalisation factors do not vary with the path. Consider a toy example in which $Pr_1(\omega) = \alpha$, $Pr_2(\omega) = \beta$ and $Pr_3(\omega) = \gamma$. Given the Multiplicative Conciliatory View, one may obtain the result that $Pr_{1+2}(\omega) = c \cdot \alpha\beta$ and $Pr_{1+2+3}(\omega) = c' \cdot c \cdot \alpha\beta\gamma$ where c and c' stand respectively for the normalisation factor at each stage of conciliation. If we change the order of conciliation, we may derive the final result $Pr_{1+3+2}(\omega) = c'' \cdot c''' \cdot \alpha\beta\gamma$. By expanding the normalisation factors, one can see that $c \cdot c'$ is equivalent to $c'' \cdot c'''$. We may see that the process of normalisation does not make the Multiplicative Conciliatory View path dependent.

[18]Easwaran et al. (2016, p.16) provides a different proof to show the same result.

[19]It should be noted that multiplicative averaging does not commute with conditionalisation. Suppose that a group of five individuals conciliate with multiplicative averaging and obtain a result x. Upon receiving a piece of evidence e, they update with conditionalisation and get the final result $xL(e)$. If they change the order and update before conciliation, the result would be $xL(e)^5$, which differs from $xL(e)$.

definition of the Geometric and the Multiplicative Conciliatory View, $Pr_{1+2}(\cdot \wedge \cdot)$ is equivalent to $Pr_1(\cdot \wedge \cdot)Pr_2(\cdot \wedge \cdot)$. When the inputs are independent for both individuals, $Pr_{1+2}(\cdot \wedge \cdot)$ is equivalent to $Pr_1(\cdot)Pr_1(\cdot)Pr_2(\cdot)Pr_2(\cdot)$. Since this formula is equivalent to $Pr_{1+2}(\cdot)Pr_{1+2}(\cdot)$, one's judgement that two propositions are independent can be well preserved.

In sum, the three formal deficiencies of the Linear EWV can be solved respectively by making conciliation with different nonlinear average functions. Since the proofs do not assume that all credence function involved are assigned with equal weight, the same formal result holds for every possible weight distribution. The Conciliatory View, hence, is free from the three formal deficiencies.[20]

3.6 Other features of nonlinear conciliation

Although switching to nonlinear average functions may save the Conciliatory View from the three formal deficiencies, there is a standing worry that both nonlinear average functions introduced are far from ideal. The Geometric Conciliatory View, despite commuting with conditionalisation, is still path dependent. The Multiplicative Conciliatory View, on the contrary, is path independent but not commutative with conditionalisation. Also, it should be noted that both nonlinear Conciliatory Views fail to be *eventwise independent*. That is, if we adopt the nonlinear Conciliatory Views, the collective credences of a group do not depend solely on the conciliating individuals' credences of the proposition but would be influenced by some other factors, such as the content of the agenda (Aczél and Wagner, 1980; McConway, 1981; Stewart and Quintana, 2018; Dietrich and List, 2016).[21] Since

[20]One may think of a case where one needs several features. For instance, one may expect their average function to commute both with the acquisition of new testimonial evidence and the rule of conditionalisation. In such a case, no single average function is proper. This is indeed a problem to the current approach. A possible solution which I have not yet fully explored is to develop a hybrid average rule which mixes the outcomes of different nonlinear average functions and keep the valuable features as much as possible.

[21]A practical consequence of adopting the nonlinear Conciliatory Views is that the result of conciliation would be partly determined by the agenda, namely the set of propositions people disagree on. Since the nonlinear average functions are not eventwise independent, the joint credence a group has in a proposition may differ under different agendas. See McConway (1981) for the

no average function is perfect, some might still consider the Conciliatory View untenable.

This worry does not undermine the current approach but instead motivates us to embrace a pluralistic conception of conciliation. Since there does not exist a perfect average function which is applicable in every case, we should, in each specific case, adopt the average function that is most likely to avoid potential problems. For instance, if I am involved in a disagreement where I am sure that no further conciliation would take place but some new evidence may appear, I should adopt the Geometric Conciliatory View. By doing so, I can guarantee that the time I receive the evidence does not determine the outcome of conciliation. Similarly, if I know that someone owns the power of changing the agenda and I do not want the result of conciliation to be manipulated by the agenda setter, I should adopt the Linear Conciliatory View. The next step the conciliationists should take, therefore, is to create a taxonomy of disagreements. By attentively categorising various cases of disagreement, we may apply the right rule of conciliation when a disagreement occurs. The primary aim of this section, hence, is to demonstrate some features of the Geometric and Multiplicative EWV and specify the conditions under which they should be applied.[22]

To see other features of different average functions, we should first compare the outcomes of adopting different average functions in a simple scenario. Suppose

proof that linear averaging is the only function which satisfies the requirement of being eventwise independent. Adopting nonlinear Conciliatory Views, thus, makes the conciliating individuals vulnerable to manipulation by the agenda setters. An agenda setter may decide their joint credences in propositions by setting the agenda in a specific way. The reason is that given different agendas, the underlying set of worlds may change. Imagine a panel of climate scientists negotiating about a set of propositions on an agenda with the intention to decide their joint credences over the propositions. When someone expands the agenda with one more proposition, say whether there will be a hurricane next year, each world ω in the underlying set of possible worlds Ω would have to be replaced by two worlds: one which is a combination of ω and there being hurricane next yest, and another which combines ω and there being no hurricane next year. If one adopts the nonlinear Conciliatory Views, a change of agenda may lead to different outcomes of conciliation. Hence, the agenda setter may manipulate the result by setting the agenda in a specific way. This is another unacceptable result since, as indicated before, an ideal rule of conciliation should not leave space for manipulation.

[22]For the sake of simplicity, here I consider different EWVs, rather than different Conciliatory Views. The formal properties of different EWVs can be generalised to other forms of the Conciliatory View.

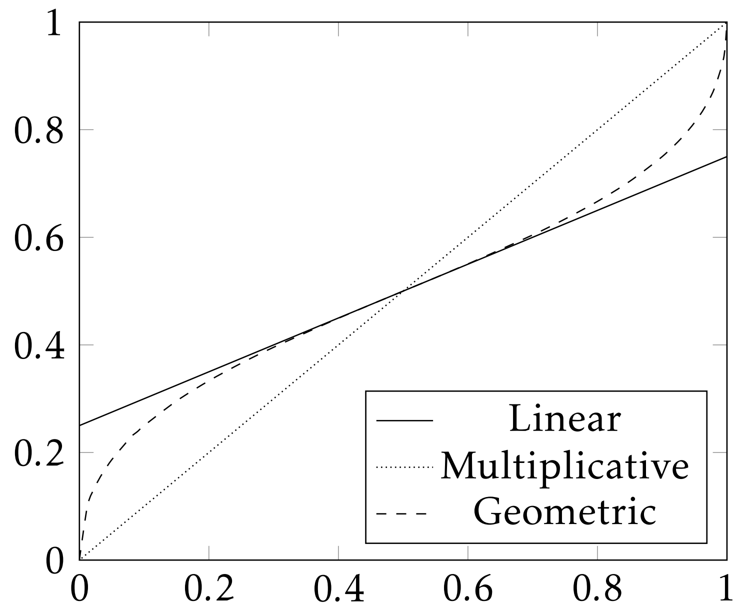

Figure 3-1: Individual with credence 0.5 compromises with a peer whose credence is x

one has 0.5 credence in a proposition and conciliates with a peer whose credence is x.[23] The outcomes of one's conciliation according to different average functions can be seen in Figure 3-1. The value on the x axis stands for one's peer's credence in the proposition under dispute, while the value on the y axis stands for the outcome of conciliation. The solid and dotted lines respectively represent the result of adopting the Linear and Multiplicative EWV, while the S-curve represents the Geometric EWV.

The Linear EWV, compared to the other two average functions, is the most resolute one since the disagreeing individuals who adopt this view never defer, in any sense, to each other. What they do is just split the difference between their credences.

The Multiplicative EWV generates a different result in this case. When one has 0.5 credence in a proposition and conciliates with their peer according to the Multiplicative EWV, the outcome would always equal to their peer's credence. That is,

[23]It should be noted that one having a credence of a proposition is an abbreviation of one having a set of credences in the worlds where the proposition is true.

one always completely yields to their peer.[24]

The outcome of adopting the Geometric EWV is the most intriguing. It behaves like the Linear EWV when the peer's credence is moderate, but gradually deviates from the Linear EWV as the peer's credence gets close to the extreme. Suppose that an individual A has a 0.5 credence in a proposition p and conciliates with geometric averaging. Let A's peer's credence be x, A's credence after to the conciliation can be expressed with the following equation:

$$Pr_A(p) = \frac{x^{\frac{1}{2}} \cdot 0.5^{\frac{1}{2}}}{x^{\frac{1}{2}} \cdot 0.5^{\frac{1}{2}} + (1-x)^{\frac{1}{2}} \cdot (1-0.5)^{\frac{1}{2}}}$$

If one conciliates with geometric averaging, one yields to one's peer to some extent in all the cases where my peer's credence is not 0.5. For example, if we plug $x = 0.3$ in the equation, the outcome is approximately 0.3956 which is slightly closer to 0.3 than 0.5. This feature gets stronger as the peer's credence gets more extremal. If we plug in $x = 0.95$ in the equation, the outcome would be approximately 0.8133. This result, compared to the former one, is closer to the value we plug in for x then 0.5. That is, $|0.8133-0.5|-|0.8133-0.95|$ is greater than $|0.3956-0.5|-|0.3956-0.3|$.

Having 0.5 credence in the disputed claim is a special case. If we relax this assumption and consider other credences, we may find a rather intriguing feature: Based on the Geometric EWV, if one's peer's credence is closer to the extreme, one's credence after conciliating with the peer would be closer to the peer's credence than one's own prior credence. For example, assume that one has a 0.6 credence. If one's peer has a more extreme credence, say 0.2, then the outcome of conciliation would be closer to 0.2 then to 0.6.[25] We may interpret this as a form of deference: When one's peer has a stronger opinion in a dispute, one would defer to the peer.[26]

[24] Easwaran et al. (2016) also mentioned this result.

[25] If we adopt Geometric averaging, the outcome of conciliation would be

$$\frac{0.2^{\frac{1}{2}} \cdot 0.6^{\frac{1}{2}}}{0.2^{\frac{1}{2}} \cdot 0.6^{\frac{1}{2}} + (1-0.2)^{\frac{1}{2}} \cdot (1-0.6)^{\frac{1}{2}}} \approx 0.3797$$

. This result is closer to 0.2 then to 0.6.

[26] A solid proof, however, is unavailable here as cases involving more than two individuals should be considered. Thanks to Catrin Campbell-Moore and Julien Dutant for their comments on this

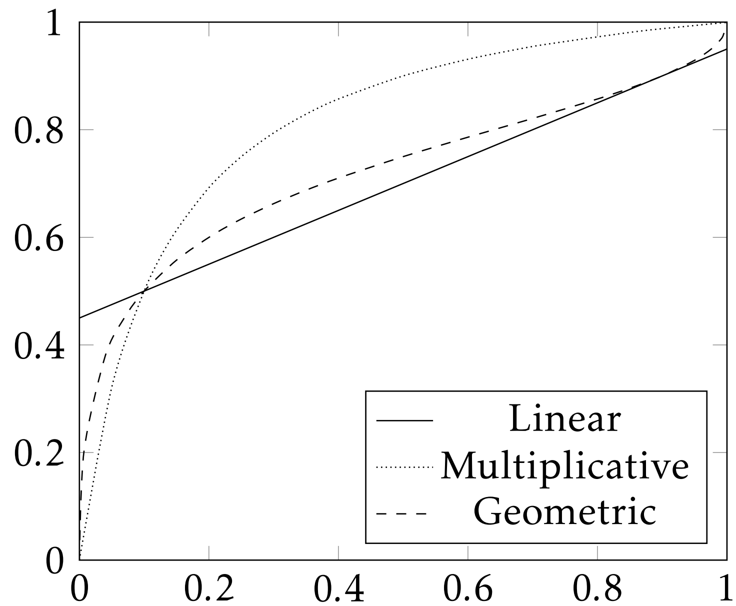

Figure 3-2: Individual with credence 0.9 compromises with a peer whose credence is x

Figure 3-2 presents a case in which one's credence in the proposition under dispute is 0.9, while their peer's credence is again x. In this case, the Linear and the Geometric EWV behave in the same way, while the outcome of adopting the Multiplicative EWV is significantly different. When both the individual and their peer have 0.9 credence in the proposition, the outcome of conciliation, given the Multiplicative EWV, would be greater than 0.9. This is a property Easwaran et al. (2016) call *synergy*. When both the individuals' credences are high, the outcome would be even higher. Because of this property, the result generated by the Multiplicative EWV, compared to the other two EWVs, is always closer to the extreme.

What, then, is the correct way of making conciliation? Should one adopt the Geometric EWV and make a radical change of credence only when the peer is strongly opinionated? Or, should one adopt the Multiplicative EWV and sometimes completely surrender to the peer? As previously indicated, one should pick the EWV which is free from foreseeable problems. Moreover, a general guideline is to pick the rule according to how resolute one wants to be. As I point out, the

point.

Linear EWV, compared to the other two, is the most unwavering one. Individuals who adopt this view never completely surrender to their peers. The Multiplicative EWV, as we have seen, is the least resilient among the three, as it makes the individuals yield to their peers more frequently than any other views. Bearing this feature in mind, one may, in each specific scenario, choose the one that best suits the case.

One may think that in a standard case of peer disagreement, there is a perfect symmetry between the disagreeing individuals. Hence, the individuals involved should never yield to the others' opinion, which implies that the Linear EWV is the only acceptable option. In fact, even in these cases, one may choose to be less resilient about their credence. Here I want to highlight two factors that are crucial in deciding which EWV to adopt. One is the strength of the evidence one possesses which decides how resilient their credence is. The stronger their evidence is, the more unwilling one is to revise their credence. Another factor is the extent the evidence is shared. The more one knows about their peers' evidence before the conciliation, the more likely that one retains their original credence.[27] This point can be illustrated by considering the case in which the individuals do not share their evidence. If one does not know whether their peer has evidence concerning a proposition, when one realises that their peer has some credence different from their own, one should be able to infer that their peer does have some evidence. One may hence be inclined to defer to their peer. If one knows all the evidence their peer possesses, there is no reason to defer. With the two factors explained, we may see how different EWVs capture these intuitions.[28]

Case 1. *An individual has no evidence concerning a proposition p and has* 0.5 *credence*

[27]As noted before, although deviate from the standard cases of disagreement in the literature, cases in which the individuals do not share all their evidence are still worth discussing. See Matheson (2014).

[28]Note that we focus on the evidence each individual possesses here. It is also possible that, upon realising that a peer has a different credence, an individual chooses to deem the peer unreliable. However, this would be a case in which the individual adopts the Steadfast View. If one aims to stick to the Conciliatory View, the more appropriate reaction is to infer that the peer possesses different evidence.

in it.[29] *They do not know if their peer has any evidence.*

From the individual's perspective, their peer is equally likely to have the correct credence as the individual, When the individual realises that their peer has a different credence in the same proposition and hence disagrees with them, it is reasonable for them to think that there are some underlying reasons. Otherwise, their peer would not have a different credence. The most probable reason, given the peerhood between them, is that their peer has better evidence and come up with more definite credence. In such a case, one should completely defer to their peer. After all, the individual has no evidence concerning the dispute. The Multiplicative EWV generates the correct result in this specific circumstance.

Case 2. *An individual has some evidence concerning a proposition p and has* 0.5 *credence in it. They know that their peer has some evidence, but do not know the strength of their peer's evidence.*

When the individual realises that their peer's credence is not radically different from theirs, they could conciliate by moderately deferring to their peer. If one finds out that their peer's credence is very strong, they may realise that their peer's evidence must be rather conclusive. After all, they are equally good in evaluating the strength of the evidence they each possess. Hence, they should yield to their peer. Adopting the Geometric EWV generates the correct result.

One possible challenge to the solution is that their peer might come to have high credence with some weak evidence. If it is so, then it would be wrong for the individual to defer to their peer. However, by assuming the peerhood between them, this kind of cases should not occur. That is, a genuine epistemic peer would not come to have high credence based on insufficient evidence. True epistemic peers should be equally careful in evaluating the evidence available to them.

Apart from the two factors, there are some other important aspects that should be considered when choosing the proper rule for conciliation.

[29] It should be noted that I do not intend to imply that whenever one has no evidence concerning a proposition, one comes to have 0.5 credence over that proposition.

3.6.1 Joint decision making

The purpose of conciliation should be taken as a crucial factor in choosing the rule. On some occasions, the primary purpose of conciliation is to come up with a joint decision on whether to take a certain action. These cases are called *action-disagreement*. Different from belief-disagreement, a true resolution of an action-disagreement is an all or nothing thing, namely that the individuals involved either take action or not. There is no middle ground between the two options. Hence, an ideal rule of conciliation for an action-disagreement should be one which helps the individuals arriving at a consensus about whether to take action. Recall that the Multiplicative EWV has the feature of *synergy*, namely that one's credence enhances another if they point to the same direction. Because of this feature, when all the individuals' credences are above 0.5, the outcome of conciliation with the Multiplicative EWV, compared to the other rules of conciliation, would be much closer to 1. Similarly, when all the individuals' credences are below 0.5, the outcome of adopting the Multiplicative EWV would be very close to 0.

Imagine that a group of people set the following rule: when their joint credence in whether they should perform an act is above 0.7, they will take action. If the joint credence is below 0.7, they would choose to do nothing. The feature of *synergy* would make it more likely for such a group to take an action. Compared to the unanimity preserving Linear EWV, the Multiplicative EWV better tackles some cases of action-disagreement.[30]

3.6.2 Polarisation

A further application of nonlinear average functions is to take them as ways of forming group beliefs. However, doing so may lead to some undesirable consequences. A phenomenon that worries many social epistemologists is *belief polarisation*. Consider a case where two individuals disagree about a controversial fact.

[30]Still, whether this feature is desirable could be doubted. It is possible that a group wants to be conservative when making decisions concerning whether to act. In this kind of case, the individuals involved should not conciliate with multiplicative averaging.

When they are both exposed to some pieces of evidence concerning the disputed fact, it is natural for one to expect that the disagreement between them to be mitigated. However, empirical studies have shown that such expectation differs from what happens in reality.[31] When the individuals are presented with evidence of a mixed character, they tend to strengthen their prior credence on the controversy. That is, one who believes that the disputed fact obtains would become even more certain about the fact, while the other one behaves in precisely the opposite way. Hence, sharing evidence may lead to an increase in the difference between their credences.

Polarisation gets even more severe when we escalate to the level of group disagreement. Suppose that two groups disagree over a proposition p. Members of group A believe that p is more likely to be true than not, while members of group B believe the opposite. When the members of the two groups are exposed to some evidence concerning p, it can be expected that the two groups become more polarised than two individuals. First, what happens in the individual level would occur again: the members of A come to have stronger credence in p, while members of B revise their credences in the opposite way. Second, since the members are now grouped with others who share similar ideas, they would communicate with others and consolidate their credence over the disputed proposition. The two mechanisms make belief polarisation even more intense at the group level.

With the phenomenon of belief polarisation stated, we may now ask the question: Which rule should one adopt when the members of group A intends to come up with a joint credence over p? The Multiplicative EWV is a bad option as its result is comparatively extreme.[32] Suppose that group A consists of four members. After being exposed to the evidence, half of the members have 0.7 credence in p while another half have 0.8. Adopting the Multiplicative EWV, the outcome of their conciliation would be approximately 0.99. Since this result is much higher

[31] See Kelly (2008) and Sunstein (2017) for a full-fledged discussion about polarisation.

[32] Here I assume that the members of a group do not come up with their credences independently. If they do, the Multiplicative EWV can be a good option for them to derive their joint credence as their credences can be amplified.

than each of their prior credence, the difference between the the joint credence of group A and group B becomes greater. Polarisation is further intensified with no substantial reason.[33]

The Geometric EWV performs slightly better in this case. If members of group A conciliate with the geometric EWV, the outcome would be approximately 0.78, which is not very distant from their original credence. However, the Geometric EWV has the feature that when one of the members is strongly opinionated, others tend to defer to their credence. When one of the members have very high credence, the outcome would be dragged toward their credence. Hence, adopting the Geometric EWV may still heighten polarisation in certain situations.

The Linear EWV, compared to the nonlinear ones, is the most conservative rule of conciliation. In cases where people have a strong intention to avoid polarisation, the Linear EWV is the appropriate one to adopt. In sum, both the Geometric and the Multiplicative EWV run the risk of intensifying polarisation. Anyone involved in a disagreement which may lead to polarisation should be aware of the outcome of adopting a rule of conciliation.

Although there are still many other cases that could be discussed, the conclusion I would like to draw has been clearly illustrated by reviewing these possible cases of disagreement. There are cases where one should conditionally defer to their peers. Yet there are also cases where one should not defer in any sense. We may therefore conclude that there is no ultimately correct method of making conciliation. The decision concerning the average function one is supposed to apply must be based upon the specific situation one is involved.

[33] A question arises here: sometimes we do want *synergy*. But in many other occasions, we prefer to avoid polarisation. There is supposed to be a more concrete criteria for us to decide whether to adopt multiplicative averaging or not. I think this is a question about specific types of disagreement. For example, people should avoid polarisation when it comes to political issues. Since the primary aim of this paper is to clarify the formal properties of average functions, I will leave this issue here.

3.7 Conclusion

The formal deficiencies of the Conciliatory View, as we have seen, stem from the misconception that there is only one way, namely linear averaging, that could be adopted to make conciliation. By selecting alternative average functions to make conciliation, the problems dissolve naturally.

Conciliating in a nonlinear way leads to some intriguing results. The most prominent one is that an individual may assign equal weight to all their epistemic peers yet be resilient and makes a marginal revision of their own credence. If we interpret the weight one assigns to their peers as the extent one trusts the peers, we may derive the result that an individual can fully respect their peers' competence in a subject matter, but still retain their credence concerning the proposition in question. From this result, we may see that previous discussion concerning the Equal Weight View and the Conciliatory View are misguided. Before we argue whether we should conciliate, we should elaborate on the notion of conciliation we are using.

The study of different kinds of conciliation motivates us to embrace a pluralistic conception of conciliation. What we should do, hence, is to construct a taxonomy of disagreements carefully. By correctly sorting different cases, we may apply the right function to make conciliation for each case. The discussion over whether to conciliate makes sense only when we are talking about the best way of making conciliation.

Chapter 4

Escaping an Echo Chamber

4.1 Introduction

The core issue in the study of disagreement is whether one should conciliate with one's peer in a case of disagreement. On one side of the debate stand the conciliationists, claiming that one should conciliate with a peer except in extreme cases. On the other side stand the steadfasters, arguing for the position that one should remain steadfast in most cases of disagreement. As both views are supported by various deliberate arguments, epistemologists have not yet reached a consensus about whether one should conciliate in a case of disagreement.

It would be good to have a clear way of defining the views in this debate. But the Steadfast View does not always recommend that one remain steadfast as it sometimes requires one to conciliate with one's peer. Similarly, the Conciliatory View does not always recommend that you conciliate: it allows you to remain steadfast in some extreme cases. Without a precise distinction between the two views, epistemologists may talk past each other while seeking the proper solution to disagreement.

Christensen (2011) points out that the two views are divided by the *Principle of Independence*, which can be reformulated as the following:

> 'In evaluating the epistemic credentials of another's expressed belief

> about *P*, in order to determine how (or whether) to modify my own belief about *P*, I should do so in a way that does not rely on the reasoning behind my initial belief about *P*.' (Christensen, 2011, p.1)

This principle restricts the way one reacts to a disagreement. To see the restriction it sets, consider an example of peer disagreement. Suppose that an individual wants to decide whether to believe *P*. After she reasons with the evidence she has, she comes to believe that *P*. One of her peers, on the contrary, does not believe *P* despite having the same evidence. When the individual realises that her peer does not believe *P*, she can see that her peer reasons in a different way and comes up with a different belief. From her perspective, she may infer that her peer fails to perform the correct reasoning and thus forms an incorrect belief about *P*. Following this line of thought, a natural response for her is to deem her peer unreliable, as her peer fails to reason correctly. Thus, she would hold her original belief concerning *P* in the face of her peer's disagreement. The Principle of Independence prohibits such a line of reasoning. When an individual is involved in a disagreement, what is implied by the very existence of the disagreement is that the individual may have made a mistake in her reasoning about the disputed claim. Since the existence of disagreement implies such a possibility, the individual should not insist that her initial reasoning is correct and deem her peer unreliable for having a different belief. If she does so, she takes what is shown to be possibly mistaken, namely her own reasoning behind *P*, as a reason to dismiss her peer's belief. In other words, she refutes the challenge to her reasoning behind *P* with her reasoning behind *P*. To avoid such circularity, she needs to find a reason independent from her initial reasoning about *P* to show that her peer is unreliable. If she cannot find an independent reason, she should conciliate with her peer by suspending her judgement concerning *P*.

How does the Principle of Independence separate the two views? It is generally agreed that the views conforming to this principle should be categorised into the group of the Conciliatory Views. If an individual accepts this principle, she cannot deem her peer unreliable simply because her peer reasons differently. Thus, in the absence of an independent reason, one who conforms to this principle should conciliate with one's peer. On the other hand, an individual who rejects the Principle of Independence can, when involved in disagreement, deem her peer unreliable without an independent reason. Hence, there is nothing which prohibits the individual to remain steadfast in the face of her peer's disagreement. Rejecting the Principle of Independence thereby allows one to remain steadfast in a disagreement.

As this principle marks the distinction between the two prominent views about disagreement, the debate can hopefully be settled if epistemologists can agree on whether the principle is correct. Like many other tasks in philosophy, this task has not been done yet. The purpose of this chapter is to approach this decade-old problem from a new perspective. Instead of reflecting on the underlying rationale behind the Principle of Independence, I will consider the pragmatic results of adopting and rejecting this principle. More specifically, I will focus on the features of the epistemic communities that will be formed given this principle. A striking discovery is that both conforming to and violating this principle may lead to a defective epistemic network. If an individual follows the Principle of Independence and conciliates with her peer whenever they disagree, it is highly likely for her to end up in an epistemic bubble, an epistemic network where all the members unknowingly share some false belief. On the other hand, if an individual rejects this principle and deems her peer unreliable when they disagree, she would be trapped in an epistemic echo chamber, an epistemic network where members reject any information from sources outside of the network. As both results are undesirable, epistemologists face a potential dilemma concerning the Principle of Independence.

The key notion for dissolving the dilemma is one's estimate of the *reliability* of one's peer. It should be noted that the notion of reliability here differs from the normal understanding. What I refer by the term 'reliability' is the probability one comes up with the correct credence concerning a proposition. That is, if the probability for me to have the correct credence concerning in a proposition p is

0.8, I would assume that the probability for my peer to have the correct credence is also 0.8. Suppose that an individual conforms to the Principle of Independence. When she disagrees with a peer whom she originally recognises as very reliable, she would conciliate with the peer and adjust her belief. Despite the disagreement between them, she still sees her peer as a very reliable person. Therefore, she chooses not to downgrade her estimate of her peer's reliability even when they disagree. That is, she still assumes that it is very probable for her peer to have the correct credence over the dispute. If, on the contrary, an individual rejects the principle, it would then be legitimate for her to remain steadfast and keep her initial belief. What is implied by such an act is that the individual, because of the disagreement, no longer sees her peer as a reliable person and thus downgrades her estimate of her peer's reliability. With this line of reasoning, we may characterise the conflict between the conciliationists and the steadfasters as a problem concerning whether an individual should downgrade her estimate of her peer's reliability in a disagreement. What epistemologists seem to miss is that the two responses are actually compatible. When involved in a disagreement, an individual can both conciliate with her peer and downgrade her estimate of her peer's reliability, as long as she takes the two actions at different times.

The overall project, hence, is to argue for a different type of the Conciliatory View. When one is involved in disagreement, one should on the one hand conciliate with one's peer and revise one's credence over the disputed proposition and, on the other hand, also revise one's estimate of the probability one's peer comes up with the correct credence. If we can develop a systematic method for updating the two values in disagreement, we may have a solution to disagreement which is free from possible negative outcomes.

In the following sections, I will first outline the prominent debate in the study of disagreement, introduce the mainstream solutions and narrow the problem down to the discussion concerning the Principle of Independence. With a clear presentation of the motivation behind the principle, I will illustrate how the principle leads to a potential dilemma for social epistemologists. A possible solution to the dilemma, as I will argue, is to adopt a diachronic strategy for adjusting one's estimate of the peer's reliability. When one is involved in a disagreement, one should make a conciliation and also update one's estimate of the interlocutor's reliability at different times. Given this idea, I develop a new strategy for updating one's estimate of a peer's reliability. We can, with this strategy, avoid being trapped in a defective epistemic network and approach the core problem of disagreement from a different perspective.

4.2 Disagreement

Disagreement, as a phenomenon, has existed for thousands of years. It might be an exaggeration saying that all human beings by nature desire to quarrel, but it is definitely fair to claim that this phenomenon has accompanied human beings since the very beginning of history. Examples are ubiquitous. Aristotle disagreed with Plato on issues in metaphysics. Lavoisier disagreed with Priestley about the cause of combustion. Keynes disagreed with Hayek over the cure for the Great Depression. Disagreement can be found in nearly all disciplines between all kinds of people. One of the reasons for this phenomenon being so common is that it plays a crucial role in the growth of human knowledge. When an individual disagrees with another, she would, on the one hand, reexamine her own view and, on the other hand, scrutinise her interlocutor's claim. Without this process of mutual assessment, the discovery of new knowledge would at best be greatly decelerated and fail at worst.

Despite its importance, disagreement has not been discussed philosophically until recent years. Due to its presence in the philosophy of religion and the rise of social epistemology, disagreement has gradually become a prominent topic in contemporary epistemology.[1] By analysing the very notion of disagreement, philosophers aim to find an ideal way of responding to information about a range of rele-

[1] van Inwagen (1996) is the first one who considers this issue in the context of the philosophy of religion.

vant opinions. To achieve this goal, we must first clarify a fundamental question: what is disagreement?

The simplest form of disagreement involves two individuals and a proposition. One of the individuals believes that the proposition is true, while another believes that it is false. When both individuals realise that the other holds a different doxastic attitude towards the proposition in question, the disagreement between them becomes manifest.[2] This mutual recognition of disagreement motivates both to react in some ways. Without this moment of recognition, people may not think that a solution to disagreement is required.[3] We can also frame this problem in terms of credences: when two individuals differ in their credences over a proposition and realise this fact, they disagree over that proposition. Given this basic form, we may generate different instances of disagreement by altering the basic setting or supplementing specific details. Disagreement can happen between large groups of people, instead of just two individuals. The most noticeable instance of a large scale disagreement is political disagreement, where supporters of different political parties hold extremely different attitudes towards a proposition. Disagreement may also happen between specific types of individuals. Suppose there is a group of individuals having great expertise in a subject. When they have different opinions concerning a proposition in that area of study, a case of expert disagreement arises. If, on the contrary, a group of people lacking expertise in a field disagree over a proposition in that field, we have a case of layperson disagreement. When a group of individuals disagree on whether to jointly believe a proposition, we have a standard case of belief disagreement. If a group of individuals disagree over whether to perform an act together, we have a case of action disagreement. All these different kinds of disagreements are significant and hence deserve further

[2]One may argue that intrapersonal disagreement, namely the kind of case in which an individual has a contradictory pair of beliefs, is also a form of disagreement. As it involves only one individual, it should be taken as the simplest form disagreement. It is without doubt true that intrapersonal disagreement is also a kind of disagreement. However, the solution to intrapersonal disagreement is substantially different from the solution to other kinds of disagreements. For this reason, taking it as the simplest case may bring in unnecessary confusion to the discussion.

[3]There are cases in which only one of the two individuals notices they disagree over a proposition. Still, the one aware of the disagreement needs to react somehow.

discussion.

With the basic understanding of disagreement, we may formulate the core question as the following: what should an individual do when she is involved in a disagreement? More specifically, how should an individual react when she realises that her interlocutor does not share her belief but instead denies it? At first glance, this question is quite simple. If the individual involved in disagreement believes that she is far superior to her interlocutor in the ability to find the truth, she does not need to do anything but retain her original belief. If she believes that she is inferior, she should listen to her interlocutor who performs better in finding the truth. To see this, consider the following example:

Example 4.2.1. Dermatologist

Stephanie, a qualified dermatologist, examines her patient who has a skin condition. After the examination, Stephanie makes the judgement that the patient has atopic eczema. The patient, however, believes that what he has is not eczema but psoriasis. Stephanie is fully aware of the fact that her patient has no expertise in dermatology. The patient also knows that, compared to a dermatologist, his knowledge in dermatology is insufficient.

Since Stephanie knows that the patient has never received any training in dermatology, she can infer that she is superior, compared to the patient, in making diagnoses about skin conditions. Hence, when Stephanie realises that the patient does not believe that he has atopic eczema, it is rational for her to ignore the patient's judgement and keep her original belief that the patient does have atopic eczema. The patient, ideally, should follow the same line of reasoning. Before seeing the dermatologist, the patient believes that he has psoriasis. When the dermatologist tells him that what he has is atopic eczema, the patient, knowing that he has not been trained in dermatology, should give up his belief that what bothers him is psoriasis. The disagreement between the dermatologist and the patient, hence, can be immediately solved.

One's superiority over one's interlocutor may stem from several different sources.

The dermatologist case shows that when all the disagreeing individuals know that one among them performs better in finding the truth concerning the issue, the inferior one should defer and the superior one should remain steadfast. In this kind of case, the superiority comes from the training one has received. Since the dermatologist went to medical school, she is more *reliable* in making judgements about skin conditions. One can also be superior than another by having better evidence concerning the dispute.[4] This point can be illustrated by another example.

Example 4.2.2. Disagreeing dermatologists

Stephanie and Conor are both qualified dermatologists. One day, they assess a patient together. After talking with the patient, Conor comes to believe that the patient has psoriasis. Stephanie also talks with the patient but finds it difficult to make a diagnosis merely with the testimonial evidence collected from the conversation. To play safe, she asks the patient to show her the affected areas. With this additional piece of evidence, Stephanie makes the judgement that what the patient has is not psoriasis but atopic eczema. The two dermatologists later meet up to discuss their findings.

At the very beginning of their meeting, Conor may find it strange that Stephanie holds a belief which differs from his. When Conor later realises that Stephanie disagrees with him because she possesses an extra piece of evidence concerning the patient's skin condition, he should give up his judgement that the patient has psoriasis and accept Stephanie's belief. Based on the fact that they are equally good as dermatologists, the one who has more evidence is more likely to make a correct judgement. On the contrary, Stephanie, upon realising that her colleague Conor makes a judgement without really seeing the affected areas, should retain her judgement that the patient has atopic eczema. Since she possesses a piece of evidence which Conor does not have, she does have a good reason to dismiss Conor's

[4] Some people might think that the notion of evidence can be interpreted broadly so as to include the information one receives from other kind of sources. For example, the training one receives at school could be, in a broad sense, understood as a type of evidence. To be precise, the evidence referred here is the direct evidence concerning the proposition at issue.

judgement. In this case, Stephanie's superiority over Conor stems from the additional piece of evidence she possesses, rather than the training she received.

Given these examples of disagreement, it might be confusing why it has been taken as an important topic in epistemology. As most instances of disagreement can be easily solved, it is natural for one to think that the study of this phenomenon bears little philosophical significance. Indeed, most disagreements that happen in our daily life can be easily solved. When one disagrees with another, one can, in ordinary cases, spot a difference in either their reliability or the evidence they each possesses. When all the disagreeing individuals recognise the difference between them, the disagreement between them can be easily dissolved.

Spotting the difference, however, is not always possible. What makes disagreement a real issue in epistemology is the existence of cases where people cannot find any difference in either the reliability or the evidence yet still have different opinions. Philosophers call such cases *peer disagreement*. To get an idea of what such disagreement looks like, we may consider a modified version of the dermatologist case:

Example 4.2.3. Dermatologists with the same evidence

Stephanie and Conor are both dermatologists. They went to the same medical school, received the same training and had equally good track records in making correct diagnosis. All these factors combined, one can infer that they are equally likely to make a correct diagnosis when it comes to skin diseases. Both dermatologists know that they are equal in all these aspects, therefore recognise each other as an *epistemic peer*. One day, they assess a patient with a skin condition together. There is no difference in the evidence they each possess. However, Stephanie believes that the patient has atopic eczema, while Conor believes that the patient has psoriasis.

When the two dermatologists realise that they hold different beliefs, what are they supposed to do? Since they recognise each other as epistemic peers, they would both think that their own judgement and their peer's judgement are equally likely to be correct.[5] Hence, the peer's judgement cannot be dismissed. They also share all their evidence, which implies that there is no difference between the evidence they each possess. Consequently, there is nothing they can refer to dismiss their peer's judgement.

In this chapter, I choose to focus on this sort of hard case of disagreement for two reasons. First, since they are the hardest ones to solve, these cases have aroused the interest of many epistemologists. Second, focusing on the hardest cases is the most efficient way of solving the problem of disagreement. If one finds the ideal response to peer disagreement, the same response should also be applicable to the easier cases.[6] Even if the solution to peer disagreement is not directly applicable to the easier cases, it may shed some light to an overall solution to disagreement.

We may now go one step further and formulate the most crucial question in the study of disagreement: when a group of equally reliable individuals disagree over a proposition based on the same body of evidence, what is the proper response for each of them? Before we delve into different views about this problem, there are two preliminary remarks about how the responses can be evaluated. First, a response can be evaluated with at least two standards. An ideal solution to disagreement could either be truth-tracking, rational or both. If a solution is truth-tracking, an individual who adopts this solution would have a better ratio of true beliefs to false ones than those who do not adopt this solution. Under a probabilistic framework of belief, one who adopts a truth-tracking solution, compared to the others who do not, would have credences that are closer to the truth. In contrast, if a solution to disagreement is rational, it conforms to the requirement of rationality. Those who accept this solution would be comparatively more rational than those who reject the solution. Although most epistemologists focus on the second ideal when searching for the correct solution to peer disagreement, the first ideal should not be ignored. When we move on to evaluate different solutions

[5]This is a crucial assumption. Even if they actually differ in reliability, as long as they see each other as an epistemic peer, the disagreement between them is a peer disagreement.

[6]Matheson (2014) compares everyday disagreement to idealised cases and points out that the aim of studying the latter is to solve other type of cases, including the former.

to disagreement, we must bear this distinction in mind as some of them focus on the first, while some others focus on the second.

Another important remark is that the traditional framework of doxastic states, under which one either believes, disbelieves or suspends judgement concerning a proposition, is too coarse-grained to capture the subtle differences between different views. Hence, we have to adopt a probabilistic framework and represent one's doxastic states in terms of credences.[7] Here I will adopt the framework constructed in Chapter three.

There are several mainstream solutions to peer disagreement. The most coarse-grained distinction lies between whether an individual should defer to her peer to some extent and make a *conciliation* by revising her credences concerning the disputed proposition (Feldman, 2006; Elga, 2007b; Christensen, 2007). The view that an individual should conciliate with her peers whenever disagreement occurs is the Conciliatory View. On the contrary, the view requiring an individual to retain her initial credence is the *Steadfast View*. Both views are prominent in the literature and supported by many convincing arguments. Here I will begin with an introduction to the Steadfast View.

4.3 The Steadfast View

Some individuals, when confronted by their epistemic peer, choose to retain their original credence in the proposition under dispute. This reaction seems quite intuitive at first glance. People form their credences with some reasons or evidence. As one's credences do not come from nowhere, it is natural for one to defend one's credences when confronted by other people. The Steadfast View provides substantial support to such response.

A deliberate argument for the Steadfast View can be constructed by reformulating the core question of disagreement from a first person perspective: when an

[7]Kelly (2010) provides an argument supporting the idea of dealing with disagreement under a probabilistic framework.

individual, whom I consider as my epistemic peer, disagrees with me over a proposition based on the same body of evidence which I possess, what should I do? In this kind of situation, I know that at least one of us made an incorrect inference concerning the proposition in question.[8] I am also very confident that one of the interlocutors, namely me, derives my credence over that proposition under a normal situation. That is, I am highly confident that I am sober, wide-awake and not suffering from any noticeable cognitive defect. Hence, I am at least as likely to have the correct credence as my peer. However, I cannot make sure that my peer also reasons under a normal condition like I do. My peer may be drunk, drowsy or having hallucinations while forming her credence concerning the proposition we disagree with. Since I have no access to my peer's subjective experience, I am in no position to tell if my peer reasons normally. If I choose to revise my credence, it would be possible for me to end up having an incorrect credence because of my peer's cognitive defect. Hence, sticking with my original credence seems to be the best option from my first-person perspective.

One may object to this argument by claiming that when an individual adopts this line of reasoning, she ignores the possibility that she is the one suffering from cognitive defects while her peer reasons correctly. Whilst it is true that an individual cannot know whether her peer reasons without any cognitive defect, she also cannot know whether she reasons properly either since, if she is the one having a cognitive defect, she could fail to notice it. The existence of this latter possibility implies that an individual should not dismiss her peer's credence when they disagree.

Supporters of the Steadfast View reject this argument by claiming that remaining steadfast, despite being flawed, is the only rational option for one involved in disagreement. An individual involved in disagreement needs to consider two possibilities. First, it may be that her interlocutor reasons in an abnormal way, thereby derives an incorrect credence which results in the disagreement. Second, the individual herself may reason in an abnormal way and obtains an incorrect credence

[8] Here the Uniqueness Thesis is implicitly assumed. I will discuss this thesis later in this section.

in the disputed claim. The steadfasters argue that the two possibilities are not equiprobable from the individual's first-person perspective. Her subjective experience shows that the latter possibility, compared to the former one, is less likely to be true. The individual thinks of herself as reasoning in a normal and hence reliable way. It is without doubt possible that this piece of personal evidence misleads her to consider herself as reasoning normally, but she is in no position to tell whether the evidence is misleading. Since her evidence indicates that she reasons normally, retaining the initial credence is the only rational response to a disagreement. In other words, one would be irrational if she does have evidence that she reasons normally but still thinks that she has made a mistake.

Another argument based on one's subjective reasoning, provided by Plantinga (2000), states that an individual cannot deny what appears as true for her. If an individual sees a proposition P as true, asking her to change her mind seems like a inappropriate requirement. Consider the dermatologist example where Stephanie in very confident that the patient has atopic eczema. Based on the same body of evidence, Conor tells Stephanie that the patient has psoriasis. If Stephanie checks the patient's affected area again, it would not suddenly appear to her that the patient has psoriasis.[9] She should reason with the evidence in the same way and form the same credence concerning the patient's skin condition because of the peerhodd between them. To further illustrate this argument, let us consider a simplified scenario. Suppose that there is a very easy way for one to check whether a person has atopic eczema: if there are exactly three dark red spots in the affected area, then the patient has atopic eczema. On the contrary, if there is only one red spot in the affected area, the patient has psoriasis. In the given example, Stephanie sees three red spots and thus comes to be very confident in the patient having atopic eczema. Upon realising that Conor holds a much lower credence, Stephanie checks the patient's affected area again and still sees three red spots. She thereby remains highly

[9]One might think that this action removes the peerhood between them as Stephanie has more evidence. What I would like to highlight is that Stephanie makes her inference based on the same body of evidence. All she does is simply review the evidence again. Hence, it should not remove the peerhood between the two dermatologists.

confident in the patient having atopic eczema. As the evidence Stephanie sees remains the same, asking her to change her credence implies that she has to adopt a credence which her evidence does not support. Since the requirement is rather counterintuitive, the steadfasters conclude that an individual should not give up her initial credence when involved in a disagreement.

Both arguments aim to show that, from one's first-person perspective, remaining steadfast is the only ideal response to peer disagreement. Different from the two arguments, recent discussion on *Permissivism*, the doctrine that a given body of evidence may justify multiple credences, provides a reason to remain steadfast which does not depend on one's first-person experience. Permissivism can be characterised as the rejection of the *Uniqueness Thesis*, which is defined as the following:

Definition 4.3.1. The Uniqueness Thesis

For any body of evidence E and proposition P, E justifies at most one competitor doxastic attitude toward P.

According to the *Uniqueness Thesis*, a body of evidence E supports at most one credence concerning a proposition. If the thesis is right, when two individuals differ in their credences over a proposition, at most one of them is correct. Permissivists reject this thesis and claim that a body of evidence can support more than one credence over a proposition, thereby offering a reason for people to remain steadfast when involved in a disagreement. In a case where two individuals possess the same body of evidence, it is possible, according to Permissivism, that both their credences are justified by the evidence they have. From this fact, it is possible that both individuals are rational in retaining their initial credence. If both are indeed rational, there is no need for either of them to revise their credence.[10]

Still, the fact that disagreeing individuals could all be rational does not im-

[10]It should be noted that Permissivism does not directly support the Steadfast View. The role it plays is that, if it is true, the steadfasters may have another reason to argue that one could be rational in retaining one's credence in a disputed claim. For a detailed discussion concerning Permissivism, see Schoenfield (2014), Ballantyne (2018) and White (2005).

ply that all the disagreeing individuals are rational.[11] Supporters of the Steadfast View, thus, need to provide additional reasons to justify their solution. One of the attempts is the *Right Reasons View* (Kelly, 2005; Titelbaum, 2015, 2019).

4.4 The Right Reasons View

Some epistemologists argue for the Steadfast View by claiming that an individual who makes inference with the *right reason* need not defer to her peer. This claim, although intuitively correct, sounds a bit trivial. Titelbaum (2015) provides a substantive argument for this view based on his *Fixed Point Thesis*:

Definition 4.4.1. The Fixed Point Thesis of Rationality

Mistakes *about* the requirements of rationality are mistakes *of* rationality.

This thesis states that if one gets the requirement of rationality wrong, one is irrational. To understand the scenario this thesis describes, consider a case where an individual makes an inference concerning whether a proposition P is true with a piece of misleading evidence indicating that P is false. As a rational agent, the individual reasons in an impeccable way and comes to believe that P is false. In this kind of situation, we would say that the individual is mistaken for having a false belief. However, since the individual makes her inference correctly, she is still rational. The Fixed Point Thesis shows that there is an exception. If the proposition P is *about* the requirement of rationality, one who makes a mistake about its truth-value is irrational. In other words, if one is misled by a piece of evidence concerning the requirement of rationality, one is irrational.[12]

With the *Fixed Point Thesis*, Titelbaum argues that one should remain steadfast when involved in a disagreement. We may reconstruct his argument with a re-

[11] It should be noted that the Steadfast View can be interpreted in two ways. In the first sense, it is reasonable for an individual to stick to their initial credence upon realising that their peer disagrees, given that their pre-disagreement credence is rational. The second sense of the view says that an individual should stick to their initial credence *even if* the initial credence is irrational. However, an individual may fail to notice that his or her own irrationality and still stick to the initial credence.

[12] Fur further details, see Titelbaum (2015, 2019).

vised version of the dermatologist case. Suppose that Stephanie, based on a body of evidence *E* she shares with Conor, has a very high credence in the patient having atopic eczema. Further suppose that *E* does entail that the patient has atopic eczema and, as a direct result, *E* rationally requires one to have a high credence in the patient having atopic eczema.[13] However, her peer Conor has an extremely low credence in the claim with the same evidence. If, upon realising that Conor has a different credence, Stephanie does not choose to remain steadfast but makes a conciliation, she would change her mind and come to have a moderate credence concerning the patient's skin condition. She would no longer be highly confident in the patient having atopic eczema but would take other possibilities as probable. When Stephanie gives up her initial credence over the dispute, she does not only think that her initial credence concerning the patient's skin condition is false, but also thinks that the evidence *E* does not support that the patient has atopic eczema.[14] She would, because of Conor's testimony, believe that rationality requires her to have a moderate credence in the patient having atopic eczema based on her initial evidence *E*. That is, Conor's testimony let Stephanie believe that she has made a mistake in taking *E* as entailing the patient having eczema. In this scenario, Stephanie's conciliation leads to a mistake about the requirement of rationality concerning *E*. According to the Fixed Point Thesis, Stephanie is irrational for failing to correctly recognise what rationality requires. To avoid being irrational, one should act upon the right reason and remain steadfast instead of changing one's opinion when one is involved in a disagreement.

One may claim that Titelbaum's argument merely shows that an individual should stick to her initial credence when she makes inference with the right reason, but does not tell an individual what is the right reason. As an individual does not always know whether her reason is right, she does not know what rationality requires and thus does not know whether she should conciliate with her peers.

[13] Here I am not assuming that the Uniqueness Thesis is true. One may reject the Uniqueness Thesis yet still admit that, in some cases, there is only one rational response to a body of evidence.

[14] Some people claim that Conor's disagreement alone cannot support not *p* for Stephanie. One may think of a scenario where there are many more dermatologists who share Conor's credence and in turn force Stephanie to make such a revision.

Remaining steadfast, hence, may make one irrational. In response to this criticism, Titelbaum claims that the purpose of the Right Reasons View is to evaluate whether one is rational, not to provide actual guidance for one to judge if her reaction is rational. The Right Reasons View aims to show that acting with the right reason is rational. What is indeed a rational response to a specific body of evidence is beyond its scope. Hence, a single case where an individual fails to recognise the right reason does not undermine the Right Reason view.

Although the Right Reasons View generates the correct verdict that an individual is irrational if she fails to stick to the right reason, it does not tell us how an individual should react to disagreement when she is unsure what the right reason is. One would need some other factors to tell what counts as the right reason. We may turn to other views for some further considerations.

4.5 The Justificationist View

A comparatively moderate view, proposed by Lackey (2008), states that whether one should conciliate depends on whether one has justification for one's credence. It is thus called the *Justificationist View*. Different from the steadfasters who require people to always remain steadfast in a disagreement, Lackey admits that there do exist cases in which one should make a conciliation. As there are also cases in which one should not conciliate, Lackey concludes that there is no single strategy applicable to all kinds of disagreements. What we should do, instead of searching for an ultimate solution applicable to all kinds of cases, is to find the key factor governing our intuition about whether to conciliate. The factor, Lackey claims, is one's justification in one's credence. When one has justification in one's credence before the disagreement, one does not have to conciliate. On the other hand, if one is not justified in having one's credence, one should conciliate when involved in disagreement.

To see how Lackey arrives at her conclusion, we need to first compare cases of disagreement that elicit contradictory intuitions. Consider the most widely discussed example in the study of disagreement:

Example 4.5.1. Restaurant Check

> Suppose that five of us go out to dinner. It's time to pay the check, so the question we are interested in is how much we each owe. We can all see the bill total clearly, we all agree to give a 20 percent tip, and we further agree to split the whole cost evenly, not worrying over who asked for imported water, or skipped desert, or drank more of the wine. I do the math in my head and become highly confident that our shares are $43 each. Meanwhile, my friend does the math in her head and becomes highly confident that our shares are $45 each (Christensen, 2007, p.193).

An implicit assumption in this case is that the disagreeing individuals are epistemic peers on this problem. That is, they are equally competent and have equally good track records in making this kind of calculation. This assumption should be quite natural, as splitting the bill does not involve any advanced knowledge in mathematics. As long as they can do basic addition and division, they should both be regarded as having expertise in this subject. Since, as we have assumed, they are equally likely to be correct, one would be inclined to think that the disagreeing individuals should assign equal weight to both their conclusions. If they do so, they would end up being highly confident that the price for each to pay is $44. After all, they are equally good at doing simple math and make their calculation with the same evidence. There is no factor which indicates an asymmetry between the circumstance they are in. It is thus claimed that making a conciliation is the only rational solution.

The intuitive response to *Restaurant Check* is that one should make a conciliation when involved in a disagreement. However, there exists another case which elicits a different intuition:

Example 4.5.2. Mental Math

Harry and I, who have been colleagues for the past six years, were drinking coffee at Starbucks and trying to determine how many people from our department will be attending the upcoming APA. I, reasoning aloud, say, 'Well, Mark and Mary are going on Wednesday, and Sam and Stacey are going on Thursday, and, since $2 + 2 = 4$, there will be four other members of our department at that conference'. In response, Harry asserts, 'But $2 + 2$ does not equal 4'. Prior to this disagreement, neither Harry nor I had any reason to think that the other is evidentially or cognitively deficient in any way, and we both sincerely avowed our respective conflicting beliefs (Lackey, 2008).

Suppose Harry insists that two plus two equals six, do I have to make a conciliation and conclude that there are five people in the department attending the APA? The answer, intuitively, is no. Since the correlated mathematical fact is extremely basic, it is bizarre for me to respond by conciliating with Harry. Instead of inferring that both of us are making some minor mistakes as we did in *Restaurant Check*, it is more reasonable for me to think that at least one of us is cognitively deficient. As I am extremely certain that two plus two is not six, I have a very high credence in Harry being the one who gets things wrong. The intuitively rational response for me, thus, is to retain my initial credence without making any compromise.

Since the two cases elicit different intuitions concerning the proper response to disagreement, we seem to have a dilemma. If we reject the Steadfast View and insist that one should conciliate whenever one disagrees with one's peers, we would have to, in *Mental Math*, conciliate with Harry and come to be just moderately confident in two plus two equals four. If we take remaining steadfast as the ideal solution, we would have to, in *Restaurant Check*, remain highly confident that each one in the restaurant case should pay $43. As we have seen, both consequences are quite counterintuitive.

Instead of choosing between the two horns of this dilemma, Lackey points out that the two cases are substantially different. In *Restaurant Check*, both individuals have no strong justification for their credences before the disagreement. That is, since I performed the calculation in my head, there is no way I can make sure that I have not made any mistake in my calculation.[15] Hence, I am not strongly justified to have a high credence in the result being $43. Similarly, my friend is not strongly justified in having a high credence in the result being $45. Since both our credences about the final result are not well justified, we ought to make a conciliation in the disagreement. On the contrary, I do have a strong justification that two plus two equals four prior to my disagreement with Harry in *Mental Math*. It is a simple mathematical fact that I have known for years. Due to the fact that I am strongly justified in my belief about the result of two plus two, I do not have to make a conciliation with Harry. By pointing out this difference, Lackey argues that the proper solution to disagreement hinges on whether one is justified in having one's credence. In a case where no one has strong justification for one's credence, the best option is to make a conciliation. In a case where one is strongly justified in one's credence, one should choose to retain one's initial credence.[16]

Although Lackey's explanation is quite convincing, it does not settle the debate as it may collapse into the Steadfast View. In a standard case of peer disagreement where all the individuals involved have the same evidence, each individual thinks that her credence is justified but the opponent's credence is not. When one is involved in this kind of situation, should one remain steadfast? According to the Justificationist View, the answer is yes. When one's credence is justified, one need not make a conciliation. Even in a case where one's peer claims that she is justified in having her credence, one could, given that one's credence is justified, infer that one's peer fails to reason properly and mistakenly takes her credence as justified. The Justificationist View, hence, generates the result that one should remain steadfast in an idealised case of peer disagreement. If one adopts the Justificationist View, one's response to the standard case of disagreement would be the same as a steadfaster. The arguments against the Steadfast View would thus

[15]It can be seen that the supporters of the Conciliatory View are more sympathetic to a form of skepticism. An individual does not have to accept the strong claim that they can never make a mistake with their reasoning. However, when there is strong evidence, namely some peers' testimony, that an individual may have made a mistake in their reasoning, they should not insist that they have not.

[16]Some people take this view as the standard Steadfast View. Here I take the Steadfast View as any view which may violate the Principle of Independence.

also be arguments against the Justificationist View. What the Justificationist View does, hence, is to provide an explanation for one to remain steadfast. Since the action of remaining steadfast is what the conciliationists are dissatisfied with, the Justificationist View does not bring an end to the debate. We still need to explore other possibilities to solve the problem of peer disagreement.

4.6 The Total Evidence View

Another theory which provides reasons for one to remain steadfast is the *Total Evidence View* which says that one should form one's credence with the *total evidence* one has (Kelly, 2010). Consider a revised version of *Equally Reliable Dermatologists* in which Stephanie has 0.9 credence in the patient having atopic eczema while Conor has 0.1 credence in the same proposition. When Stephanie and Conor realise that they disagree over the patient's skin condition, they each possess three pieces of evidence:

E The original evidence E they share.

E_S Stephanie has 0.9 credence in the patient having atopic eczema.

E_C Conor has 0.1 credence in the patient having atopic eczema.

Here the first piece of evidence E is substantially different from the rest since it is a piece of *first-order* evidence concerning the patient's skin condition. It may include the appearance of the patient's affected area, the description of the symptoms provided by the patient and other relevant facts the dermatologists share. The two pieces of testimonial evidence, relative to E, are *higher-order* evidence that reveal each dermatologist's credence based on E. They should be taken as higher-order evidence because they indicate some properties of the dermatologists. With such information, we may have some indirect evidence concerning the patient's skin condition. Given the setting, should one make a conciliation and hold a credence of 0.5 in the patient having atopic eczema? There is one line of reasoning which requires one to do so. Since the two pieces of higher-order evidence point to very different facts, they cancel each other out. One should thus split the difference between the two dermatologists' credences and end up with the average, namely 0.5 credence, concerning the patient's skin condition. We can easily spot the problem of this result: if one does end up splitting the difference between the credences of both, one makes a mistake for letting the higher-order evidence swamp the first-order evidence. That is, the significance of the first-order evidence would be totally dismissed if one ended up having 0.5 credence in the claim.

A more natural way of reasoning is to take all the relevant evidence into account. Let us consider the same case again from Stephanie's perspective. Before the disagreement takes place, she possesses the evidence E and, based on E, comes to have 0.9 credence in the patient having atopic eczema. Upon realising that Conor disagrees with her, she obtains an additional piece of evidence E_C that her peer has 0.1 credence in the patient having atopic eczema. This piece of higher-order evidence, combined with Stephanie's own higher-order evidence E_S, indicates that she should have 0.5 credence concerning the patient's skin condition.[17] Still, she has to take the first-order evidence E into account. For Stephanie, E shows that the patient does have atopic eczema. Thus, her credence should be higher than 0.5, which shows that she considers it more likely than not that the patient has atopic eczema. In sum, because of Conor's testimony that he has 0.1 credence in the patient having atopic eczema, Stephanie's new credence should be lower than her initial credence 0.9, but higher than the result of a full conciliation, namely 0.5.[18]

It should be noted that Kelly does not see dismissing every piece of first-order evidence and split the difference as a wrong response in disagreement. When one is presented with a vast amount of higher-order evidence pointing to different di-

[17] Here the way one updates with higher-order evidence plays is left unexplained since it plays no role in our discussion.

[18] From this case, we can see that the Total Evidence View can only be presented under a probabilistic framework of beliefs. If we adopt the tripartite conception of beliefs which states that one can only believe, disbelieve or suspend judgement concerning a proposition, we cannot correctly capture the result that Stephanie, after acquiring E_C, should have a credence between a full conciliation and her initial credence.

rections, the weight of the first-order evidence would be very small. In such cases, it is acceptable for one to make a conciliation with one's peer and take the significance of the first-order evidence as negligible. Understood this way, the Total Evidence View requires neither to always conciliate with the peers nor to always remain steadfast. Whet one should do, according to the Total Evidence View, is to carefully consider every piece of evidence and act upon the total evidence.

4.7 The Conciliatory View

So far, we have seen several different views that provide reasons against simply making conciliation with a peer. Given these arguments, the Conciliatory View (henceforth the CV), on which one should always conciliate with one's peers, seems very implausible. In fact, quite the opposite. People attack this view exactly because it is one of the most promising solutions to peer disagreement. Broadly construed, any view which recommends one to conciliate with one's peers can be categorised as a version of the CV.[19] We can easily think of the crucial reason which motivates this view. If an individual refuses to conciliate with a peer who disagrees with her, she can be criticised for being overly confident with their own credence. In order to be epistemically modest, one should conciliate with their peers whenever they disagree.

Since there are many different ways one could make a conciliation, equating the CV with any view requiring one to conciliate leads to an overly broad definition. We may take the dermatologist Stephanie for example. When she realises that she disagrees with Conor and wants to conciliate with him, she can either abandon her credence in the disputed matter and come to have an extremely low credence in the patient having atopic eczema, or make a marginal revision and still have a very high credence in the patient having eczema. Although the two responses lead to radically different outcomes, both count as conciliation. If the CV is defined

[19]One may wonder whether the CV really differs from the other views. This point will be addressed later this section.

simply as a view requiring one to make conciliation, it provides little guidance for the ones involved in a disagreement. For the CV to be substantial, one must provide further detail about how to conciliate.

The most significant form of the CV is the *Equal Weight View* which requires one to assign equal weight to all the disagreeing peers' credences. Recall the case where two dermatologists disagree over the patient's skin condition. One of them has a high credence in the patient having atopic eczema while another has a very low credence in the same proposition. It is assumed that they are not only equally good as dermatologists, but also possess the same evidence concerning the patient's skin condition. Both of them are aware of their parity in reliability and evidence, and thus see each other as an epistemic peer. This last assumption is crucial. Even if they are slightly different regarding some factors, as long as they recognise each other as an epistemic peer, this case can be taken as a peer disagreement. When the two dermatologists realise that they disagree over the patient's skin condition, they should, according to the Equal Weight View, make a conciliation on this issue by assigning equal weight to each other's credence. The reason behind this is quite intuitive. Given the symmetry in all these aspects, there is no way for one to tell where the difference between the individuals lies. They are equal in every aspect relevant to their credence. Thus, they must be equally likely to be correct, and their credences should be valued in the same way.

One may wonder whether the CV really differs from the other views. We have seen that the Justificationist View also requires the disagreeing individuals to conciliate on some occasions. When an individual finds that her credence and her opponent's credence are equally unjustified, she should make a conciliation with her peer. On the Total Evidence View, an individual's total evidence may require her to make a conciliation by assigning equal weight to each interlocutor's credence. On these occasions, the views that are normally categorised as non-conciliatory generate the same verdict as the CV. One might hence take the Equal Weight View as a special case of these non-conciliatory views which happens to assign equal weight to all the disagreeing parties in all cases of peer disagreement. If we agree that the difference between these views lies merely in the extent one should conciliate, we can put these views on a spectrum. On the one end, we have the Steadfast View, which requires one to remain steadfast in all possible cases of disagreement. Right next to the Steadfast View is the Right Reasons View, which requires one to act upon the right reason. As one does not always know which reason is the right one, asking one to follow the right reason often leads to the same result as asking one to follow one's initial reason. The response it requires of the disagreeing peers, hence, would be to remain steadfast. The Justificationist View, compared to the Right Reasons View, is somewhat closer to the other end since it requires one to conciliate when one has no strong justification of one's credence. The view which stands closest to the other end is the Total Evidence view. Since one's total evidence includes the higher-order evidence from one's peer, this view usually leads to a conciliation, albeit an unequal one.[20] On the opposite end is the Equal Weight View, which requires the disagreeing peers to make a conciliation by assigning equal weight to the credences of each individual involved.[21]

The spectrum, however, does not correctly portray the relationships between these different views. The difference between the non-conciliatory views and the CV is categorical one, marked by the *Principle of Independence*:

Definition 4.7.1. The Principle of Independence

In evaluating the epistemic credentials of another's expressed credence about P, in order to determine how (or whether) to modify my own credence about P, I should do so in a way that does not rely on the reasoning behind my initial credence about P.[22]

Rejecting the Principle of Independence amounts to accepting the following line of reasoning: given a body of evidence E, I form a credence in a proposition P.

[20] Although I present the non-conciliatory views in this order, I am not claiming that, necessarily, the views that are closer to the CV on the spectrum requires one to make a greater conciliation with one's peer. It is possible that the Total Evidence View requires one to make a conciliation that is smaller than what the Justificationist View requires.

[21] Christensen (2013, p.78) mentioned the idea of aligning the views on a spectrum.

[22] The original formulation of this principle provided by Christensen (2011), as presented in the first section, focuses on the reasoning behind one's belief instead of credences. Here what I present is a revised version.

When I realise that my peer, who also possesses *E*, has a credence that differs from mine, I can infer that her reasoning behind *P* differs from mine. The fact that she disagrees with me shows that either her or my reasoning must be incorrect. Since my reasoning is actually the correct one, I can infer that my peer is wrong in her reasoning. Hence, I can deem my peer unreliable and retain my credence in *P*. This pattern of reasoning, as one can see, is circular. What is shown by the very existence of the disagreement is that my reasoning might be incorrect. If I deem my peer unreliable, I implicitly take my reasoning to be the correct one. Since the correctness of my reasoning is challenged by my peer, I cannot insist that I am correct based on my belief that I am correct. I have to provide some other reason which does not rely on my contested reasoning to show its correctness.

The Right Reasons View violates the principle of independence since it allows one to dismiss one's peer's credence when one has the right reason. If an individual adopts the Right Reasons View and, when involved in a disagreement, does have the right reason for her credence in the disputed claim, the Right Reasons View requires her to retain her credence and refuse to conciliate. When she does so, her action implies that her peer is unreliable for failing to come up with the right reason. Since the Right Reasons View allows the individual to see her peer as an unreliable person if she has the right reason concerning the disputed claim, it does not conform to the Principle of Independence. The Justificationist view also violates the Principle of Independence. If an individual's credence in the disputed proposition is justified, she does not, according to the Justificationist view, have to conciliate with her peer but can retain her credence. Again, as long as she is indeed justified in having her credence before the disagreement, she does not need any reason independent from her initial reasoning behind the disputed claim to dismiss her peer's credence. Thus, the Justificationist View violates the Principle of Independence. Likewise, the Total Evidence View does not conform to this principle. Suppose the total evidence I have supports my credence in a disputed proposition *P*. When a peer of mine disagrees with me over *P*, I am allowed, according to the Total Evidence View, to retain my credence in *P* if that is what my total evidence supports. That is, the Total Evidence View allows me to ignore my peer's contest if that is what my total evidence supports.

Given this brief review of the three non-conciliatory views, we can see a categorical difference between the CV and the other views.[23]

Now we have a clear criteria for distinguishing between the CV and the non-conciliatory views. The question of disagreement, hence, can be narrowed down to whether we should adopt the Principle of Independence. If the principle is right, then the CV is the correct solution to the problem of disagreement. If the principle is wrong, the Steadfast View prevails.

4.8 Challenging the Principle of Independence

Although the motivating idea behind the Principle of Independence is pretty convincing, it is far from uncontroversial. Kelly (2010) points out that following this principle may lead to several counterintuitive results. Consider the following example:

Example 4.8.1. Right Dermatologist

Stephanie and Conor acknowledge each other as an epistemic peer in dermatology. One day, they diagnose a patient together. After viewing a body of evidence E concerning the patient's skin condition, Stephanie forms a 0.8 credence in the patient having atopic eczema (abbreviated as P) while Conor has 0.2 credence in P. The evidence available to them actually supports one to have a 0.8 credence in the patient having atopic eczema. Stephanie and Conor then compare their notes and realise that they disagree.

According to the CV, the rational reaction for both Stephanie and Conor is to make a conciliation. That is, the CV indicates that if both Stephanie and Conor are rational, they should conciliate and end up with 0.5 credence in P. Kelly points out that this verdict is counterintuitive. In this case, Stephanie correctly responds

[23]Both Christensen (2011) and Kelly (2010) accept taking the Principle of Independence as the criteria for the CV.

to the evidence and forms the correct credence in P. Conor, on the contrary, mistakenly takes the evidence as supporting a 0.2 credence in P. When they conciliate with each other, Stephanie moves from the correct credence to an incorrect one, while Conor's credence gets closer to the correct credence. Such difference between the two, according to Kelly, shows that conciliating could not be the rational response for both of them. Hence, the CV is wrong for generating the incorrect verdict that conciliating is the only rational response for both Stephanie and Conor.

Kelly moves on to provide another case which shows that the CV lowers the standard of rationality too much:

Example 4.8.2. Wrong Dermatologists

Stephanie and Conor are mutually acknowledged peers concerning P. After viewing a body of evidence E together, Stephanie forms a 0.02 credence in P while Conor has 0.04 credence in P. The evidence E actually supports a 0.99 credence in P. Stephanie and Conor then compare notes and realise that they disagree. They follow the dictates of the Equal Weight View and compromise at 0.03.

In this case, again, the CV generates the verdict that the rational response for both of them is to make a conciliation. However, their 0.03 credence is still very far from the correct credence given E. If we adopt the CV, we would have to admit that Stephanie and Conor are rational despite having extremely wrong credences in P. This result, as Kelly sees, is absurd. If we accept that people become rational simply by conciliating with others regardless of how accurate their credences are, we are adopting a overly low standard of rationality. Rationality should require something more than a conciliation.

For Christensen, the two cases do not really pose a serious challenge to the CV. We may begin with the first case in which an individual is right while another is wrong. Christensen argues that the CV does not grant them with equal rationality in this case. It is true that, according to the CV, they both react rationally to the disagreement. But both reacting rationally to the disagreement is different from being equally rational. Let the time before their conciliation be t_1 and the time after their conciliation be t_2. At t_1, Stephanie correctly reasons with E and forms the right credence in the disputed claim. When she meets Conor later and compare their notes, she again reacts rationally, according to the CV, with the testimony Conor provides and ends up having a 0.5 credence in the disputed proposition at time t_2. With the reconstruction, we may see that Stephanie's reasoning in this case is perfectly rational. What makes her end up with an incorrect credence is not the reasoning, but the misleading evidence she receives from Conor. Conor, on the other hand, begins with an incorrect credence at t_1 and, after conciliating with Stephanie, gets closer to the correct credence at t_2. Although he is irrational at t_1, he responds to the disagreement rationally and comes to be rational at t_2. Still, since he is irrational at t_1, he is not as rational as Stephanie. With such explication, Christensen argues that the CV does generate the verdict that they are both rational in responding by conciliating, but does not generate the problematic verdict that they are equally rational.

The second case, similarly, poses no real threat to the CV. When both dermatologists have incorrect credences concerning P, they both fall short of being rational. However, this fact does not imply that making a conciliation is not the rational response to them. Supporters of the CV may claim that, by conciliating, they react rationally to the disagreement but are still not fully rational. In other words, the fact that the outcome of their conciliation fails to satisfy the requirement of rationality does not imply that conciliating is not the rational response to disagreement. Hence, the CV does not lower the standard for rationality.

Still, Christensen has to respond to the intuitively correct verdict generated by the Total Evidence View that, in *Right Dermatologist*, Stephanie should have a credence between her initial credence and the credence after a conciliation. Recall that there are three pieces of evidence in this case:

E The original evidence E they share which supports 0.8 credence in P.

E_S Stephanie has 0.8 credence in P.

E_C Conor has 0.2 credence in P.

According to the Total Evidence View, E_S and E_C combined should be taken as a piece of evidence requiring one to have 0.5 credence in P. If we further consider the evidence E, we would have two pieces of evidence. The first one is the combination of E_s and E_C which requires one to have 0.5 credence in P. The second one is the evidence E which requires one to have 0.8 credence in P. The credence one should have, given the two pieces of evidence, should lie between 0.5 and 0.8. This result contradicts the verdict of the CV which requires both dermatologists to have 0.5 credence in P. Supporters of the CV, hence, need to explain what goes wrong with this result.

As a response, Christensen (2011) claims that the significance of a piece of evidence is agent-sensitive. For any third party who does not generate any of E, E_S and E_C, the rational credence to have is, as the Total Evidence View indicates, between 0.5 and 0.8. On the other hand, if we view the three pieces of evidence from either Stephanie or Conor's perspective instead, we may derive a different result. When a person reasons with a set of evidence, she normally takes her first-person psychological evidence as inert. Suppose that an individual comes to have a credence in a proposition with her evidence. When she realises that a peer has the same credence as she does, she would consider her peer's credence as a piece of evidence showing that she assessed her original evidence correctly. She would not, however, take her own original psychological state as a piece of evidence confirming that she has assessed the evidence correctly. In the example given, Stephanie would not take E_S as a piece of evidence supporting the credence she formed. That is, her psychological state is not treated as a piece of evidence.[24] From this abstract example, we can see that an ordinary individual would ignore her first-person psychological evidence while checking if her reasoning is correct. Following this line of thought, in *Right Dermatologist*, the total evidence Stephanie has should involve only E and E_C. Her first-person psychological evidence E_S which shows that E supports 0.8 credence would be inert and does not count as part of her total evidence. With E and E_C, Stephanie would end up having 0.5 credence in P, which

[24]Note here that this line of reasoning is inconsistent with the Steadfast View.

matches the right credence for her to have at t_2 according to the CV. Similarly, Conor would also take his first-person psychological evidence as inert and forms a 0.5 credence in P at t_2, as the evidence he possesses are E and E_S. Given such reasoning, we may see that if they do adopt the CV, they would end up with the correct credence to have in P at t_2.

A more threatening case to the CV, compared to the previous ones, involves individuals with extremely high credences. In *Mental Math*, when Harry makes a very simple mistake in math, should I, as his peer, refuse to make a conciliation with him? It seems that I should refuse, since I am very confident that Henry has made a mistake. However, the only evidence I have is my very simple reasoning in some basic proposition in math. Apart from this, I do not have any other evidence concerning the proposition. If I conform to the Principle of Independence, I can take my reasoning on this simple fact as evidence for my confidence, but cannot take my reasoning as evidence to deem Harry unreliable. As stated, what has been shown with Harry's disagreement is that I might have made a mistake in my reasoning concerning the simple math problem. Thus, I should not insist the I have not made a mistake with the reason that my reasoning about the math problem is correct. Conforming to the Principle of Independence leads to a bizarre result that I ought to conciliate in this kind of case.

To counter this argument, Christensen provides a detailed analysis of cases involving individuals with extremely high credences. When an individual has a very high credence in a proposition, it is unlikely for her peer to have a very low credence in the same proposition. When it does turn out that her peer has a credence that is radically different from her extremely high credence, she can infer that her peer fails to reason in the normal way. Given such fact, she should be allowed to remain steadfast. It should be noted that the original reasoning behind her extremely high credence is not involved in this line of reasoning. In this case, the fact that Harry's credence is radically different from mine may provide me with an independent reason to deem Harry unreliable. Hence, Harry is no longer a genuine epistemic peer for me. By arguing that this kind of case is extraordinary, Christensen claims that the Principle of Independence could be saved.

There remains another problem for the Principle of Independence. Consider the example provided by Moon (2018):

Example 4.8.3. Peggy's Location

Someone I know to be a reliable testifier tells me,

(P) Peggy is at the party

(Q) If Peggy is at the party, then Quinn is unreliable about Peggy's whereabouts. Apart from Q, I do not have any other information about Quinn. Since the testifier is very reliable, I come to have a high credence in $P \wedge Q$. Quinn then comes up to me and tells me that P is very unlikely to be true. Since P can be derived from $P \wedge Q$, I am very confident that P while Quinn is very confident that $\neg P$. Apart from P, I also infer Q from my high credence in $P \wedge Q$, and then give a low evaluation of the epistemic credentials of Quinn's low credence about P on the basis of my high credence in Q.[25]

Moon argues that it should be legitimate for me to remain steadfast in this case. Before I get to know Quinn's credence in P, I already know that she is unreliable concerning Peggy's location if P is true. I am also told by the reliable testifier that P is true. When Quinn reports her credence on Peggy's location, I am very confident that she is unreliable about P. It should thus be legitimate for me to dismiss her credence on this issue. If we do agree with Moon on this point, we would have a counterexample of the Principle of Independence. In this case, Quinn disagrees with me over P, namely Peggy's location. The reason I rely on to dismiss Quinn's credence about Peggy's location is the conjunction of the two things I am told, namely $P \wedge Q$. Since the disputed claim P can be derived from $P \wedge Q$, I do not rely on a reason independent of my reasoning behind the dispute to dismiss Quinn's credence. This case, thus, counts as a counterexample of the Principle of Independence.

[25]Here I replace every occurrence of the notion belief with credence to fit this example in our probabilistic framework of doxastic states.

Although Moon's problem looks threatening, Christensen (2019) does not take it as a serious problem of the Principle of Independence. With a more sophisticated formulation of the principle, he claims, we can get rid of this problem easily.

Definition 4.8.1. The Principle of Independence extended

> When an agent has formed an initial credence c in P on the basis of the first-order bearing of evidence E, and then gets some evidence that bears on the reliability of her reasoning from E to her credence c in P, her final credence in P should reflect the Independent Hypothetical Credence (IHC) it would be rational for her to have in P: that is, the rational credence in P independent of E's first-order bearing on P, but conditional on her having formed credence c in P on the basis of E, and on the reliability evidence the agent has about herself.[26] (Christensen, 2019, p.18)

To see the difference between this new formulation of the Principle of Independence and the rudimentary ones, we need to introduce a distinction concerning the different roles a piece of evidence can play. Christensen points out that a piece of evidence could have two roles. It may, on the one hand, serve as a *first-order* evidence which directly supports a proposition and, on the other hand, be a *higher-order* evidence which provides indirect support to a proposition by indicating the reliability of an individual. We may illustrate the distinction with an example. Recall that in *Restaurant Check*, my peer disagrees with me about the result of splitting the check. Her testimony that the correct result should be $45, as a piece of first-order evidence, supports the belief that the result is not $43. Also, her testimony is a higher-order evidence showing that my reasoning behind the belief is wrong, since I come up with an incorrect answer in this case. By showing that my reasoning is wrong, my peer's testimony indirectly supports the belief that the

[26] Note that this is just one of his several different sketches of the Principle of Independence. He provides other more complicated formulations of this principle in response to other problems. Given the primary purpose of this paper, we only need to consider this sketch of the principle.

result is not $43.[27]

With the distinction between the two roles a single piece of evidence plays, we may see the difference between the extended and the original formulation of the Principle of Independence. What has been highlighted by the new formulation is that in evaluating one's reliability, we should exclude the first-order support of the evidence but not the higher-order support. We do have to consider what the evidence reveals about the peer's reliability. Given this distinction, we may see that I do not violate the Principle of Independence in *Peggy's Location*. At the very beginning, I received the testimony from a reliable source and formed a high credence in $P \wedge Q$. The testimony involves both a piece of higher-order evidence about Quinn's reliability and a first-order evidence about Peggy's location. Later I met Quinn and realised that she has a high credence in $\neg P$. According to the higher-order evidence I gathered from the reliable source, the fact that Quinn is highly confident in $\neg P$ shows that she is unreliable regarding Peggy's location. Based on my assessment of Quinn's reliability, I decide to dismiss her credence about $\neg P$. The reasoning here does not involve my initial reasoning behind Peggy's real location. In other words, the reason I dismiss Quinn's testimony about Peggy's location is not based on my reasoning about Peggy's location, but based on a piece of evidence about Quinn's reliability concerning Peggy's location. My high credence in P is not the reason for me to dismiss Quinn's low credence, but the reason for me to deem Quinn unreliable. Since the Principle of Independence only forbids one to make an inference concerning one's peer's reliability with one's original reasoning behind the proposition in dispute, Christensen does not take the case of *Peggy's Location* as a real threat to the Principle of Independence.

So far, we have considered various arguments against the Principle of Independence. As none of these arguments really knocks the principle down, the Principle of Independence seems to be correct. As a direct result, the CV should be adopted as the solution to all kinds of disagreements. We should conform to the Principle

[27] Christensen admits that the distinction is quite rough. Nevertheless, we do seem to have an intuitive distinction between what a piece of evidence directly supports and what it indirectly supports.

of Independence and conciliate with our peers whenever a disagreement occurs. This conclusion, however, is a bit hasty. There are some other aspects of the principle that have not been properly examined. One of them is the actual influence it brings to the society. To make a more thorough examination, we should turn our focus to an important question: what would happen if everyone in a community conforms to this principle? Put it more precisely: if all the members in a community follow this principle, what kind of epistemic community will be formed? By exploring the pragmatic consequence of conforming to this principle, we may approach the problem in a new way and come up with a complete analysis of this principle. To accomplish this goal, we may begin with a brief introduction to two widely discussed phenomena: epistemic bubbles and epistemic echo chambers.

4.9 Epistemic Bubble or Echo Chamber?

The phenomena of *epistemic bubbles* and *epistemic echo chambers* have both been widely discussed in recent years. In this section, I will introduce Nguyen's (2018) analysis of both phenomena and focus on some crucial features of them that are related to the Principle of Independence. With these features, we may evaluate the Principle of Independence from a different perspective.

4.9.1 Epistemic bubble

To characterise the phenomenon of an epistemic bubble, we need to begin with the notion of an *epistemic network*. An epistemic network can be defined as a group of individuals sharing credences with the others. Members in a network collect information from various sources, form credences with the information in hand and pass the information on to the others. With the process of information exchange, members in an epistemic network may come to have some new information and, based on this information, derive credences in some further propositions. A toy model may illustrate how an epistemic network works: Suppose that I have a piece of evidence supporting a high credence in the claim that a new virus is spreading in East Asia. Having such credence, I meet my two friends Stephanie and Conor and tell them about the new virus. Both of them trust what I say and come to be very confident in there being a new virus spreading in East Asia. They could also share the information they have with me and let me form a credence with the information they provide. As we share our information, the three of us together form an epistemic network.

A crucial point needs to be highlighted for us to see the prominent feature of an epistemic network. When one gathers information from other members in an epistemic network, what one acquires are pieces of *testimonial evidence*. We may illustrate the difference between testimonial and other types of evidence with a simple case. Suppose that a car accident happened near my house. I look out from the window and see a crashed car. After seeing the car, I tell my sister that there was a car accident near our house. My sister accepts what I say and comes to be very confident that a car accident just happened. In this case, we are both confident that a car accident happened near our house. The difference between us lies in the evidence we have to form our credences. I see the crashed car through the window and witness the car accident. With this piece of perceptual evidence, I come to be very confident that there was a car accident. On the contrary, my sister does not see the car accident with her eyes but merely accepts my report on the car accident. Her credence over this, hence, is formed with the testimonial evidence she collected from me. Although we are both highly confident that there was a car accident, we formed our credence with different types of evidence.

Given the distinction between testimonial and non-testimonial evidence, we may elaborate our definition of an epistemic network: an epistemic network is a group of individuals sharing testimonial evidence with the others. By doing so, members of an epistemic network can get to know things without making observations themselves. In the example given, my sister does not have to see through the window herself to form a credence about what happened outside but can do so with the testimonial evidence she gained from me. Moreover, testimonial evidence may provide one with information that is not directly accessible. Consider another scenario where a friend of mine applied for the PhD program in my department and would like to know the result. After consulting an admission committee, I know that his application was successful and inform him about this. In this case, my friend did not have access to the information which he would like to know. He does not know any of the committee members and has no way to gather evidence about the result himself. With the help of another member in the epistemic network he is in, he can reach an information source that is not directly accessible for him, namely the committee member, and get the information he needs. We can therefore see that an epistemic network can provide its members with some information beyond their reach.

An epistemic network is, in ordinary cases, beneficial for its members since the number of information sources one can access can be greatly increased with the help of others. If we agree that having access to more information sources and gaining more information is in general better than having access to less sources, we have to accept the direct result that an epistemic network has a positive effect to its members. However, there are some types of epistemic networks which bring about negative effects. The most widely known one is the *epistemic bubble* (Sunstein, 2017; Nguyen, 2018). Consider the following case: Geoff, a university lecturer and a life-long Labour supporter, reads the *Guardian* everyday. Apart from the *Guardian*, he does not gather information from any other source. Needless to say, he has never read the *Daily Mail* or the *Sun* as he stands against the ideology behind these newspapers. He works in a department where all his colleagues share Geoff's political views and, similarly, do not gather information from any source other than the *Guardian*. As a loner, Geoff does not interact with anyone except his colleagues. All his colleagues, like Geoff, are quite unsociable and rarely talk with people outside of their department. The range of sources they gather information from, thus, is very limited. Suppose that the *Guardian* never reports anything about space projects for some unknown reasons. Since Geoff only reads the *Guardian*, he knows nothing about ongoing space projects. Also, he cannot find out about space projects by talking to his colleagues because, like Geoff, they only read the *Guardian*. In this kind of epistemic network, every member lacks information about a specific topic because they collect information from a small range of sources. Consequently, the information spreading in the network is very limited. People describe the members of this network as being trapped in an *epistemic bubble* where some facts are completely left out in the network they are in. Nguyen (2018) defines an epistemic bubble as 'a social epistemic structure which has inadequate coverage through a process of exclusion by omission'. By the term 'coverage', Nguyen refers to the variety of information sources the members consult. When a group of individuals fail to gather information from a sufficiently wide range of information sources, they would miss certain types of information and be trapped in an epistemic bubble. An analogy may illustrate the problem of an epistemic bubble: the members in an epistemic bubble are like picky eaters. They only consume information from some specific sources and thereby know little about things not reported by these sources.

If the only problem of an epistemic bubble is the lack of access to some kinds of knowledge, we do not have to be too worried about this phenomenon. Since it is impossible for one to know everything, there is no blame if one does not know something.[28] Likewise, it should be acceptable for the members of an epistemic network to lack some knowledge. To see this, consider a Buddhist sangha. Although none of its members has any knowledge in meat cooking, we would not say that it is a defective epistemic network. The members of this group simply do not need to know how to cook meat. Hence, it would be absurd to blame them for lacking the knowledge they find useless. The sangha is indeed an epistemic bubble, but not a culpable one. Compared to the lack of knowledge of some category, a more serious problem, as Nguyen points out, is that people in an epistemic bubble may have incorrect credences due to insufficient exposure to a diverse set of information sources. Consider another hypothetical scenario where an article on the

[28]On some occasions, one may be blamed for lacking knowledge. For instance, if a person claims to be an expert in a field, she could be blamed if she does not know the fundamentals in that field. This kind of case, however, does not undermine the point I want to address here. In most occasions, one should not be blamed for failing to know certain things.

Guardian mistakenly reports that the probability for there being more than three hurricanes in the Atlantic this year (abbreviated as P) is 70%. Having read this article, Geoff has a 0.7 credence in P at time t_1 and shares this news with one of his colleagues. His colleague, who has also read the *Guardian* on that day, tells Geoff that she has 0.7 credence in P. With no access to any other information sources, there is no way for Geoff to change his mind and obtain the correct credence in P.

In addition to the problem raised by Nguyen, I here raise an even more problematic feature of an epistemic bubble. The members of an epistemic bubble, compared to members in a normal epistemic network, are more likely to mistakenly hold their credences resiliently. Since they collect information from a limited range of information sources, they are more likely to misjudge the real number of independent information sources they consult and would thus mistakenly hold their credences resiliently.[29] Let us look at the given example again. Upon receiving his colleague's testimony at time t_2, Geoff holds his 0.7 credence in P more *resiliently* at time t_2 than at t_1. That is, it is harder to change Geoff's credence in P at t_2 than t_1. Geoff's reasoning here seems perfectly rational. He gets the information about the expected number of hurricanes in the Atlantic from two different sources. As the evidence provided by both sources indicates that a 0.7 credence is right, 0.7 is very likely to be the correct credence to have on this prediction. Following this line of reasoning, he should hold his credence in P more resiliently. What he does not know is that since he is trapped in an epistemic bubble together with his colleague, his colleague is not an independent information source but one relying solely on the *Guardian* like he does. Failing to notice this, Geoff does not know that the information supporting his 0.7 credence is given by one, instead of two different information sources. Hence, he is wrong in holding his credence more resiliently at t_2. From this case, we can see a much deeper problem of epistemic bubbles. Since the members in an epistemic bubble collect

[29] This problem is similar to what Nguyen calls *bootstrapped corroboration*. Members of an epistemic bubble may come to be overly confident in their beliefs since every other individual they meet share the same belief. Here what I point out is that they may not be overly confident in their beliefs, but hold their credences overly resiliently.

information from a limited range of sources, these sources are less likely to be independent compared to members in a normal epistemic network. This point can be illustrated by comparing an epistemic bubble with an ideal epistemic network. In an ideal epistemic network with n members, every member collects information from a different information source. Thus, the network is linked with at least n different and mutually independent information sources. Suppose that n is sufficiently large so that the information sources are sufficiently diverse. If a member of such a network takes every other member as an independent information source and exchanges information with each of them, she may come to have a very accurate credence over the proposition in question. Also, if a member of such a network holds her credence resiliently because of the testimonies she received, she is quite likely to be right in doing so. On the contrary, in an epistemic bubble with n members, the number of information sources attached to the network m is much smaller than n. It is then inevitable for there to be subgroups that collect information from the same source. Hence, if the members in an epistemic bubble exchange information with the others and take every other member as an independent source of information, they would misjudge the number of independent information sources they consult and mistakenly hold their credences resiliently.

From this observation, we may derive what I consider to be the key problem of an epistemic bubble: its members do not know that they collect information from a limited range of information sources. If the members of an epistemic network do know that their friends and interlocutors collect information from a limited range of sources, they would, when exchanging information with the others, know that it is quite likely that the other members collect information from the same source which they consult. Thus, when they receive information from other members, they would not revise their credence as if these information comes from a new independent source. Although these members may still have incorrect credences, they would not mistakenly amplify the significance of the information they receive from the others and would not hold their credences resiliently.

What makes an epistemic bubble? Nguyen points out two primary causes. If an individual only interacts with like-minded people, it is very likely for her to end up in an epistemic bubble. Like-minded people have similar backgrounds and tend to collect information from similar sources. We may think about the previous example where Geoff, as one who only reads the *Guardian*, knows nothing about space projects. He is trapped in an epistemic bubble not only because he is picky when it comes to choosing newspaper, but also because he only interacts with his colleagues who have the same taste in newspaper. Like-minded people gather information from a limited set and, since a limited set of sources cannot cover every kind of information, it is very likely for a network formed by like-minded people to be incomplete. Put differently, like-minded people tend to omit information of the same sort, thereby forming an epistemic bubble together.

Another major cause for epistemic bubbles is information filtering. This process takes place whenever one intentionally picks the information for another to consume. For instance, if a non-democratic government censors every piece of information about democracy, its citizens would know little about other forms of government and tend to wholeheartedly embrace a non-democratic system. A notable present-day example is *algorithmic personal filtering*, a process which filters information for an individual based on her online browsing history. With modern technology, websites can present its users with only the topics they are interested in and reduce the probability for them to receive different type of information. Such filtering blocks one away from some types of information and thus results in epistemic bubbles.

A possible cause which Nguyen did not mention is the effect of *conformity* (Asch, 1955, 1956). When people get to know the others' credence in a proposition, they tend to conform to the others and have a similar credence in the same proposition. Asch (1955) claims that this effect is undesirable as it undermines the significance of consensus. To see this, compare the following two cases:

Example 4.9.1. A group of n scientists aim to find out whether a substance has the property X. Each of them conducts an experiment and derives an outcome about the claim they aim to verify. After everyone completes the experiment, the scientists meet up to share their credences and, at the end, reach a consensus.

In this kind of epistemic network, each individual verifies the claim independently. If they end up agreeing that it is very likely that the substance does have the property X, we may take this result as very plausible as it has been checked with n independently conducted experiments.[30] However, the effect of conformity contaminates this ideal picture. Consider another epistemic network:

Example 4.9.2. A group of scientists aim to find out whether a substance has the property X. Some of them independently conduct experiments and generate outcomes about the disputed claim, while some others choose to do nothing. After the first group of scientists complete their experiment, the lazy scientists adopt what they say. They come together to share their credences and, at the end, reach a consensus.

In the first network, each member makes an experiment to test the claim separately and hence derive their initial credences without the information provided by the others. The final consensus they arrive at, hence, has undergone careful scrutiny. In the second network, some members choose not to conduct any experiment but simply form their credences based on the credences of the others. If we take the process of making an experiment as consulting an independent information source, we may infer that the second network involves less independent information sources. Other things being equal, the diversity of the information sources in the second network is lower than the diversity in the first. The consensus reached in the second network, thus, would be less accurate compared to the first. However, the members in the second group may not be aware of this fact. When people exchange their beliefs in an epistemic network, they do not always check where their peers collect their information. This result resembles an epistemic bubble as the members misjudge the number of information sources they consult. In brief, the effect of conformity may reduce the number of information sources which the members of a network consult and leads to an epistemic bubble.

[30] The underlying idea here is the same as Condorcet's Jury Theorem.

The observation that conformity creates epistemic bubbles is pretty striking. As the effect of conformity is quite common in real life, such discovery implies that epistemic bubbles appear frequently. Combined with the other two major causes Nguyen pointed out, we can conclude that we are actually very likely to be trapped in an epistemic bubble. We all, more or less, misjudge the number of information sources we consult and make incorrect evaluation about the resilience of some of our credences. Consequently, we are all at high-risk from an epistemic bubble.

4.9.2 Epistemic echo chamber

Fortunately, it is possible to get people out from an epistemic bubble. Recall that an epistemic bubble appears when a group of people gather information from a limited range of sources. What people in an epistemic bubble lack is not the ability to distinguish misinformation from true information, but access to a sufficiently diverse set of information sources. When they are exposed to a greater variety of information sources, they would correct their credences and get out of an epistemic bubble.

Compared to an epistemic bubble, an epistemic *echo chamber* is much harder to break. A archetypal echo chamber consists of a set of core members, call them the gurus, and a set of peripheral members who follow the gurus. The followers take the gurus as extremely reliable sources of information and are very confident in all the gurus' teachings. As one can easily see, this basic structure is very common. For most groups, there exist some key opinion leaders whose words, for the other members in the group, weigh much more than any other person in the world. What makes an echo chamber distinct from other epistemic networks is that in an echo chamber, the core members' authority can never be defeated. Typically, the gurus in an echo chamber tell their followers that all the information sources outside the chamber are unreliable and seldom provide correct information. The followers, with their great confidence in the gurus' reliability, strongly believe what the gurus say and see all external sources as, in most situations, unreliable. With this belief, the followers would never see their gurus as unreliable. To see this, we can think about what would happen when the followers get information from the external source. I will first consider a case where the followers get information that agrees with the gurus, and then information that disagrees with the gurus.

Suppose there is an epistemic echo chamber in which the gurus tell their followers that 5G towers spread a new kind of virus. One day, one of the followers receives a piece of information from an external source, which indicates that 5G towers do spread viruses. The follower, with the information she collected from the external source, would think that the gurus' teaching about 5G towers is right. As the gurus' teaching is confirmed, the follower is further reassured that the gurus are extremely reliable. She may not, however, deem the external sources reliable for providing information that is in accordance with the gurus' teachings. For the followers, the external sources seldom report correctly. Even though one of the external sources provides a piece of correct information about the 5G towers, they are still quite unreliable in general and could provide incorrect information.

If, on the contrary, the incoming piece of information shows that 5G towers are not correlated with the spread of the virus and thus shows the gurus wrong, how would the follower react? Since they see the external sources as unreliable, they would take the incoming information as a piece of misinformation. From their perspective, this line of reasoning is very plausible since the gurus already told them that the external sources are unreliable. That is, when the followers receive a piece of information which contradicts the gurus' teaching, the coherent way for them to understand this fact is to deem the information incorrect and take it as evidence showing that the external sources are indeed unreliable. Again, the gurus' teaching about the external sources is confirmed. Either way, the follower's confidence in the gurus' reliability remains uninfluenced. They are either right about 5G towers, or right about the external sources being unreliable.

Moreover, the followers' beliefs echo with each other. Suppose that a group of followers respectively receive pieces of information which contradicts the gurus' teachings from several different sources outside the chamber. Each of them, based on the assumption that the gurus are extremely reliable, derives the conclusion that the external sources are unreliable. When they meet up to share their finding about the external information sources, their beliefs would confirm each other's belief, as they all infer that the external sources are unreliable. The conclusion they will end up with, hence, is that the external sources are extremely unreliable. The members of an echo chamber will gradually come to be extremely confident that all other members in the chamber, especially the gurus, are very reliable. They will also end up being extremely confident in any information circulating in the chamber, as the information comes from some internal sources. On the contrary, since the external sources are deemed unreliable, members of an echo chamber would not gather any information from them. Hence, it would be extremely hard to convince the members of a chamber that they are wrong. Once such an echo chamber has been established, one cannot expect to correct the members' mistaken credences merely by presenting them with the correct information.

4.9.3 A potential dilemma of Independence

Both epistemic bubble and echo chamber, as we have seen, bring negative effects to their members. Thus, we should avoid being trapped in either of them. This purpose, however, is much harder than it seems. In this section, I will argue that if we conform to the Principle of Independence, we could be trapped in an epistemic bubble. On the other hand, if we reject the Principle of Independence, we could end up in an epistemic echo chamber. Since both results are epistemically undesirable, a dilemma arises.[31]

We may begin by considering the results of violating the Principle of Independence. When an individual disagrees with a peer and chooses to reject the Principle of Independence, she would, instead of revising her credence in the disputed claim, revise her estimate of her peer's reliability. Let us take the two dermatologists for example again. If Stephanie does not accept the Principle of Indepen-

[31] Put it more precisely, it is not the case that the Steadfast View leads to an echo chamber but *trying to follow steadfast norms* leads to a chamber.

dence, she would remain steadfast regarding her judgement of the patient's skin condition. Still, she needs to react to the fact that Conor disagrees with her.[32] As she chooses not to revise her credence, she must revise her estimate of Conor's reliability. The underlying idea is that Conor has made a mistake concerning the patient's skin condition and is no longer Stephanie's peer. In other words, the fact that Conor mistakenly takes the patient as having psoriasis could be a piece of evidence showing that Conor is not as reliable as Stephanie. Thus, Stephanie should downgrade her estimate of Conor's reliability. If she takes this as a general strategy when involved in disagreement, she would end up in an echo chamber. According to Nguyen's analysis, members in an echo chamber respond to disagreement in exactly this way. Whenever they receive information which contradicts what they believe, they see the information source as unreliable and downgrade their estimate of the reliability of the source. Consequently, members of a chamber deem every external source unreliable and only trust the other members in the same chamber who never reports anything that contradicts what they believe. An echo chamber is thus formed.

Since violating the Principle of Independence leads to an unacceptable result, conforming to the principle seems to be the only proper response. A natural response is to turn our attention to a different question: what would happen if everyone in an epistemic network adopted the CV and conciliated with their peers whenever a disagreement occurs? Zollman's (2012) model of an epistemic network provides an answer.[33] Consider a group of individuals $1,\dots n$. Each individual is connected to a group of other individuals, namely their neighbours. Let $\mathcal{N}_i$ stand for the set of i's neighbours (i is also included in $\mathcal{N}_i$) and let $c_i(\cdot)$ stand for the function representing individual i's credences. With this setting, we may simulate the outcome of adopting the CV. When all the individuals conciliate with every other,

[32]One may argue against this point by claiming that Stephanie does not have to do anything. It is true that nothing strictly requires Stephanie to react. However, if we take her as a rational agent, she should react to the disagreement somehow.

[33]Zollman's analysis focuses on the features of different network structures, rather than the practical results of adopting the CV. Nevertheless, his simulation does show how an epistemic network evolves when all its members adopt the CV.

each one takes the average of all the neighbours' credences as their new credence. For example, suppose that an individual w has three neighbours x, y and z. When w conciliates with her neighbours concerning their credences in a proposition P, she takes the linear average of her credence and all her neighbours' credences as her new credence in P, which is $\frac{1}{4}(c_w(P) + c_x(P) + c_y(P) + c_z(P))$. With a computer simulation, Zollman shows that when all the individuals in the network conciliate with their neighbours regarding their credences over a proposition P, every individual in the network comes to have the same credence over P in a few rounds of information exchange. That is, they reach a consensus quickly by making conciliations.[34] With Zollman's simulation, we may obtain the expected result that conforming to the Principle of Independence helps people to reach a consensus.[35]

Unfortunately, this result seems undesirable. When an individual chooses to conciliate with every interlocutor whenever she is involved in disagreement, she implicitly takes every interlocutor as an independent information source. As we have seen, if one takes everyone else as an independent information source, one would misjudge the real number of independent information sources one consults. Making conciliation, thus, leads one into an epistemic bubble. To illustrate, consider an epistemic network with n members which is attached to m independent information sources such that $m < n$. Given this assumption, we may infer that there exists at least a pair of members in the network who consult the same information source. When a member of this network conciliates with the other members, what she does is to sum up everyone's credence and divide the outcome by n. This action implicitly implies that there are n independent information sources, while in fact there are only m sources. We can therefore see that one who adopts the CV runs the risk of misjudging the number of independent information sources.

A direct result of such misjudgement is that the members would mistakenly hold the consensus resiliently. Consider an epistemic network with 10 members

[34]Here I assume that the members conciliate with linear averaging. The result that individuals reach a consensus quickly via conciliation should remain true for other average functions .

[35]Given Zollman's simulation, networks of different structures reach a consensus at different speed. Here what we care about is whether a network reaches a consensus, not the speed it reaches consensus.

(abbreviate as $\mathcal{N}$). Suppose that the first and the second member collect information from the same source and the third and the fourth collect information from a different source. The rest of the group, similarly, form three pairs and collect information from three different sources. When a member k of this network conciliates with every other, her credence would be the sum of all the members' credences divided by 10, which can be expressed as the following:

$$\frac{1}{10}\sum_{i\in\mathcal{N}} c_i(P)$$

.

This formula implies that there are 10 information sources, which is wrong. Now consider another individual who collects information from a new source and meet the then members. Let the new network be $\mathcal{N}'$. In the absence of any additional information about her reliability, one who adopts the CV would assign her credence with the weight $\frac{1}{11}$. That is, the outcome of conciliation would be

$$\frac{1}{11}\sum_{i\in\mathcal{N}'} c_i(P)$$

.

This is again a wrong result. Since the newcomer brings information from the sixth source, the real weight for her credence should be one-sixth, which is much greater than the weight assigned here. From this case, we an see that because of the members conciliate at t_1, they hold their consensus more resiliently at t_2 as it is harder to change their credence. When a new individual having a different credence joins the network and conciliates with the others, her credence would matter much less than it should. This outcome is problematic, as we do want the new information to be assigned with the correct weight. We may arrive at the conclusion that conciliating with others may make the consensus overly resilient and eventually lead to a stubborn epistemic network.

With a review of the results of conforming to the Principle of Independence, we may derive the result that for any network which is not an epistemic bubble, it may evolve into a bubble if all its members conciliate with the others. If it is an epistemic bubble, conciliating would make it worse. Here we can find a potential dilemma which follows from the Principle of Independence. When one conforms to the Principle of Independence, it would be likely for one to misjudge the significance of one's interlocutor's credences and end up in a situation that resembles an epistemic bubble. When one violates the Principle of Independence, one would stick to one's original credences and downgrade one's interlocutors' reliability. Eventually, one would get trapped in an epistemic echo chamber. We are hence trapped in a dilemma of the Principle of Independence.

4.10 Breaking an echo chamber

The solution to the dilemma, ideally, is to adopt the Principle of Independence, conciliate with the peer and always carefully check the number of information sources. If one does so, one would not be trapped in an echo chamber since one conciliates with one's peers in a disagreement. Neither would one end up in an epistemic bubble, as one does know the number of information sources and would not mistakenly amplify the weight of the information one received. By doing so, one may come to have the ideal credence concerning a proposition. However, such a strategy is pragmatically infeasible. It is extremely difficult for one to always check whether one's interlocutors gather information from the same set of sources as each other, as it is hard for one to trace the source which provides a piece of information. As an immediate consequence, it is hard for one to always check the number of sources one gathers information from. With this pragmatic concern, we have to give up searching for this ideal solution and aim at a practically operable one.

To find a feasible solution to the dilemma, we need to achieve two goals. On the one hand, we want to update our estimate of the interlocutors' reliability according to the disagreement, thereby avoiding getting trapped in an epistemic bubble. If we fail to do this, we would have to take every epistemic peer as an independent information source and run the risk of misjudging the number of independent information sources. What we want, thus, is a way to revise the interlocutors' reliability and perceive them less than fully reliable. On the other hand, we still want to make conciliation and revise our credence in order to avoid being trapped in an echo chamber. If we refuse to conciliate, we implicitly deem our interlocutors unreliable. Making such a judgement merely with the evidence that we disagree, as we have seen, leads to an epistemic echo chamber. With these concerns, we may infer that when involved in a disagreement, we should revise both our estimate of our interlocutors' reliability and our credence over the disputed matter.

How should we do this? We have seen that if an individual adopts the CV, she would revise her credence and split the difference when she disagrees with an interlocutor. When she conciliates, she should not also revise her estimate of her interlocutor's reliability since, if she does so, she would contradict herself. When she revises her estimate of her interlocutor's reliability because of the disagreement, what is implicitly shown is that she no longer sees the interlocutor as a peer. Hence, she does not have to conciliate with her interlocutor. Put differently, if an individual's interlocutor is not as reliable as she is, she need not conciliate by splitting the difference between them. Similarly, if an individual chooses to violate the CV and downgrade her estimate of the interlocutor's reliability, she need not also conciliate by splitting the difference. After all, her interlocutor is no longer her epistemic peer after she downgrades her estimate of the interlocutor's reliability. She should not take her interlocutor as an epistemic peer and, at the same time, deny her interlocutor's status as a peer. We may see that making conciliation with an interlocutor is incompatible with downgrading the reliability estimate of the interlocutor. Consequently, there does not seem to be an obvious solution to the dilemma.[36]

The key to untie the knot is to establish a diachronic update strategy which

[36]One may think that a possible strategy is do a bit of both, namely to both conciliate with the interlocutor and downgrade the interlocutor's reliability to a small extent. If one does so, one may get rid of the dilemma. This solution appears to me as unacceptable. The crucial point here is that one should not make the two moves at the same time. Once a factor is changed, there is no reason for one to change another. Hence, the dilemma cannot be dissolved this way.

takes the two moves at different times to avoid self-contradiction. Although the two actions are incompatible, there is no factor which forbids one to take them separately. When an individual disagrees with her peer, she can make a conciliation for the current case and downgrade her estimate of the interlocutor's reliability after the first disagreement.[37] If they disagree again at some later point, she would have a more precise estimate of the interlocutor's reliability and would not take the interlocutor as a peer again. If one follows this strategy, one can make both moves without contradicting oneself.

The question that naturally follows is how, in reality, should we take the two moves at different stages? We have already seen the way an individual conciliates, but have not yet explored the way an individual revises her estimate of her interlocutors' reliability. To develop a way for an individual to revise her estimate of an interlocutor's reliability, we must first reflect on the notion of reliability. A possible way of understanding this notion, albeit not the most widely adopted one, is to take it as the probability one forms the correct credence concerning a proposition. Given this interpretation, when two agents are equally reliable, the probability for them to come up with the correct result is the same. For example, when we say that one's reliability is 0.9, what is meant is that one comes up with the correct belief ninety percent of the time. When two individuals are of equal reliability, they are equally likely to believe in a true proposition and reject a false one. In spite of the intuitive plausibility of this way of understanding reliability, it is insufficient when we adopt a probabilistic framework of doxastic states. According to this definition of reliability, when we say that one's reliability is 0.8, what we mean is that person forms the correct credence for eighty percent of the time. Suppose that, given all the background conditions, the ideal credence for one to have in P is x. According to the notion reliability introduced here, an individual with 0.8 reliability would correctly form a x credence eighty percent of the time when she is asked to eval-

[37] One may think that it is more reasonable to do it the other way round, namely first downgrading one's estimate of the interlocutor's reliability and subsequently conciliate with the interlocutor. However, doing so implies that the peerhood between them disappears when they disagree. If one does so, one does not seriously take one's peer as a genuine peer but one who is likely to be false. If one takes the assumption that one's peer is equally reliable, one should conciliate first.

uate P. A problem arises when we consider the cases where the individual gets her credence wrong. If, in the twenty percent of time in which she gets it wrong, her credence is very close to the ideal credence x, we may still see her as a very reliable person. On the contrary, if her credence is very far from the correct credence, we would find her unreliable. The characterisation of reliability introduced is incomplete for failing to capture such a difference. To illustrate, imagine a more concrete case in which Rachael, one who has rheumatoid arthritis, claims that she can correctly predict the probability of rain tomorrow. Let us suppose that the Met Office provides perfect weather predictions. The ideal credence one should have, thus, should be the same as the report provided by the Met Office. If the Met Office announces that the probability of raining tomorrow is 60%, one should have a 0.6 credence in there being a rain tomorrow. When Rachael is asked to predict the weather, she gets eight correct predictions out of ten tries. For the other two tries, her prediction deviates from the correct number within a three percent range. In such a case, Rachael is very reliable in making weather predictions. Her predictions, although sometimes incorrect, are all very close to the real probability of raining. On the contrary, if Rachael is drastically wrong for twenty percent of time and makes predictions that radically deviate from the reports made by the Met Office, she would be deemed quite unreliable. However, according to the definition of reliability, Rachael is equally reliable in the two cases. As this is a counterintuitive result, the traditional definition should be elaborated if we adopt a probabilistic framework of doxastic states.

Knowing that the notion of reliability should not merely be construed as the frequency one has a correct credence, we should reformulate the notion in a way which takes the magnitude one's credence deviates from the ideal credence into account.[38] We should not only care about the frequency, but also the magnitude one's credence deviates from the correct credence to have over a proposition. More importantly, we need to know the direction one's credence deviates from the cor-

[38]The notion I appeal to here is similar to the notion of *accuracy* which has been discussed by the accuracy-first epistemologists. For an explanation of the notion of accuracy, see Leitgeb and Pettigrew (2010).

rect credence. Following this line of thought, one's reliability could be measured in several new ways. One of the possibilities that I would like to explore here is to consider the average difference between one's credence and the correct credence across a set of propositions. For every proposition in a domain, we may subtract the ideal credence from one's credence and calculate the average of the outcomes. The result may be taken as a factor for us to derive one's reliability concerning propositions in this domain. If we define one's reliability in terms of the average difference between one's credence and the ideal credence for a set of propositions, both the frequency and the magnitude one's credence deviates from the ideal credence could be taken as factors determining one's reliability. For example, consider an individual who has a credence x over the proposition P at time t_1 while the ideal credence to have in P is y. If we know that she gets her credence wrong by the magnitude $x - y$ at t_1, when she forms a z credence over a similar proposition P' at some later time t_2, we can take the value $x - y$ as a factor for our estimate of the magnitude z deviates from the correct credence at time t_2.[39][40] By doing so, our notion of reliability not only involves the frequency one gets thing wrong, but also the extent one gets thing wrong.

How should we formulate this new notion of reliability? We could answer this question in a formal setting. Given a set of possible worlds Ω, we may construct a

[39]Here what I mean by a similar proposition is one concerning a similar subject matter. Since the propositions are about similar subject matters, one can take the track record of the disagreements happened between one and one's peer into account. One may challenge this idea and claim that one may not be systematically biased. The extent one's credence deviates from the correct credence may not be a stable value. One may have an extremely incorrect credence concerning a proposition, yet have a perfectly right regarding another proposition. I agree that this could happen. However, it seems to me that the only possible way of capturing such instability of one's credence is to take the calculate the average difference between one's credence and the correct credence.

[40]It should be noted that the negation of a proposition P, although concerning the same subject matter, should not be counted as a similar one. What we intend to record here is the average difference between one's credence and one's peer's credence over a set of similar propositions. If the set contains both a proposition and its negation, the average difference between one's credence and the ideal credence would always be zero. For example, if one has 0.2 credence over a proposition P, one would have a 0.8 credence over the negation of P. If the ideal credence to have over P is x, the difference between one's credence and the ideal credence over P is $0.8 - x$. We may derive that the difference between one's credence and the ideal credence over $\neg P$ is $0.2 - (1 - x)$. The sum of the two values would be 0. It can thus be seen that if we take both a proposition and its negation into consideration, the information we gather would be useless. Hence, the set we consider needs to be a consistent one which contains no contradictory pairs.

set of propositions based on Ω. Let this set of propositions be 2^{Ω}. Suppose there is an individual A whose doxastic state can be represented by a credence function $c_A(\cdot)$. Further assume that there is an omniscient function c which, for every proposition, generates the ideal credence for one to have. Let $d(\cdot,\cdot)$ be a function which measures the difference between two credences such that $d(x,y) = x - y$.[41] To tell how reliable A is, we need to consider the average difference between the omniscient function $c(\cdot)$ and A's credence function $c_A(\cdot)$ over a set of propositions $S \subseteq 2^{\Omega}$,[42] namely

$$\frac{1}{|S|}\sum_{X\in S} d(c_A(X), c(X)).$$

.

The outcome of this formula is the average difference between A's credence and the ideal credence A should have regarding every proposition in S.[43] Suppose the average difference between A's credence and the ideal credence for all the propositions in S is Δ_A. When A forms a credence concerning a new proposition P' which is not in S, we can anticipate $d(c_A(P'), c(P'))$, namely the difference between A's credence in P' and the ideal credence to have in P', to be Δ_A. Hence, we may take A's credence as an indicator which allows us to derive the ideal credence. Think about a simple case in which A's credence is always 0.3 short of the ideal credence. The average difference between A's credence and the ideal credence Δ_A equals 0.3. When we find out that A's credence in P is x but do not know the ideal credence to have, we could derive that $x-(-0.3)$ is the ideal credence to have based on A's track

[41] Here I take the function as measuring the difference between two values for sake of simplicity, namely the result the first value subtracts the second. One can pick a different distance measure based on one's requirements.

[42] As previously mentioned, this set should not contain any contradictory pair.

[43] It should be noted here that the value is generated by A's credence minus the ideal credence. If we switch the order and calculate the average of the outcome $c(\cdot)$ minus $c_A(\cdot)$, the result would be the additive inverse of the average of $c_A(\cdot)$ minus $c(\cdot)$. Another point that need to be mentioned is that I take the difference between two values, instead of the absolute value of their difference as the factor determining the outcome. By measuring the difference between them, we can keep track of the direction one's credence deviates from the ideal credence. That is, one's credence is sometimes higher and sometimes lower than the ideal credence. If we track the absolute value of their difference, we would not be able to record the direction one's credence deviates from the ideal credence.

record. If it turns out that $d(c_A(P'), c(P'))$ is smaller than Δ_A, we should update our estimate and derive a new estimate Δ'_A which is the updated expected difference between A's credence and the ideal credence. This new estimate is equivalent to:

$$\frac{1}{|S \cup \{P'\}|} \sum_{X \in S \cup \{P'\}} d(c_A(X), c(X))$$

.

Next time when we see A trying to come up with a credence in another proposition, we would expect the difference between A's credence and the ideal credence to be Δ'_A. If $d(c_A(P'), c(P'))$ is greater than Δ_A, we should revise our expectation in the same way and derive another value Δ''_A which is greater than Δ_A. When A faces a new proposition and needs to come up with a credence, we could expect the difference between A's credence and the ideal credence to be greater than Δ_A. With this process, we can gradually revise the estimated difference between A's credence and the ideal credence. We may, with the notion of average difference, derive a new notion of reliability to replace the one we discussed.

To see how the overall process works, consider the following example:

Example 4.10.1. Stephanie is a dermatologist whose credences can be represented by the function c_S. Let P_i stand for the proposition that the i-th patient has atopic eczema and c stand for a omniscient function which generates the ideal credence to have in a proposition. Suppose that Stephanie's credences and the ideal credences in P_i are distributed as the following:

	P_1	P_2	P_3
$c_S(\cdot)$	0.3	0.2	0.5
$c(\cdot)$	0.45	0.15	0.55
$c_S(\cdot) - c(\cdot)$	-0.15	0.05	-0.05

Given this distribution, we may derive the following result:

$$\frac{1}{|\{P_1, P_2, P_3\}|}(d(c_S(P_1), c(P_1)) + d(c_S(P_2), c(P_2)) + d(c_S(P_3), c(P_3))) = \frac{1}{3}(-0.15 + 0.05 - 0.05) = -0.05$$

It can thus be derived that Δ_S in this case is -0.05. That is, the value generated by c_S is on average 0.05 lower than the ideal credence. Suppose we, at some later time, find out that the value of $c_S(P_4)$ is 0.3 but do not know the ideal credence to have over P_4. Given the record of the differences between $c_S(\cdot)$ and $c(\cdot)$ regarding P_1, P_2 and P_3, we can derive the ideal credence to have over P_4 by subtracting the average difference Δ_S from $c_S(P_4)$, which is is 0.35. If we find out later that the ideal credence to have over P_4, as generated by c, turns out to be 0.45, the average difference between c_S and c would become 0.1. With this result, we need to update Δ_S and take its value as -0.1. When we know Stephanie's credence over another proposition P_5, we should take $c_S(P_5) - \Delta_S$ as the ideal credence to have concerning P_5.[44]

Based on this new definition of one's reliability, we may now develop a different response to disagreement which may get us rid of the dilemma of independence. Recall that the purpose of knowing a person's reliability is to derive the correct credence for one to have given that person's credence. To achieve this goal, we need to know the average difference between that person's credence and the ideal credence. With the information that, in general, a person's credence is short of the ideal credence by Δ, we may calibrate a person's credence and gradually get closer to the ideal credence. That person's credence, in such a case, is taken as an indicator for the ideal credence to have over the proposition in dispute.

Based on the idea that we can obtain the ideal credence by taking the interlocutor's credence as an indicator, we can develop a method which calculates the expected difference between an individual's credence and the outcome of the conciliation between that individual and another. Suppose there are two individuals A and B. Let their credence functions respectively be $c_A(\cdot)$ and $c_B(\cdot)$. Both functions, like normal probability functions, assign a value in the interval $[0, 1]$ to every

[44] A potential worry here is that the outcome may violate the axioms of probability. That is, if we subtract one's credence with a value, the outcome may be lower than 0 or greater than 1. In such cases, we should simply take the extreme values, namely 0 and 1, as the outcome. If one's credence is always significantly higher than the ideal credence, when one's credence is very low, it should be acceptable to just take the outcome as 0.

proposition in 2^{Ω}. With this basic setting, we can calculate the average difference between their credences and the outcomes of their conciliations in the same way we calculate the average difference between one's credence and the ideal credence. Suppose that A and B disagree over a set of propositions S. A conciliates with B by splitting the difference and adopts the outcome as her new credence. The ideal credence for them to have in this case, namely the result of conciliation, is $\frac{1}{2}(c_A(P) + c_B(P))$ for some $P \in S$. From A's perspective, the next time B disagrees with her over a proposition, she needs to calibrate B's credence with the average difference between the result of their conciliation and B's credence in the past, which can be formulated as

$$\Delta_{AB} = \frac{1}{|S|} \sum_{X \in S} d(c_B(X), \frac{1}{2}(c_A(X) + c_B(X))).$$

Here I use Δ_{AB} to stand for the magnitude A should calibrate B's credence. The value $\frac{1}{2}(c_A(X) + c_B(X))$ is the outcome of their conciliation concerning X. What we measure with this formula is the average difference between the values of $c_B(X)$ and $\frac{1}{2}(c_A(X) + c_B(X))$ across all propositions in S. This outcome is the magnitude A should calibrate B's credence when they disagree. When A disagrees with B over a new proposition P which is not in S but similar to the elements of S, A should conciliate with B by calibrating B's credence in P with the value Δ_{AB} and take the outcome as her new credence in P. Moreover, she needs to update Δ_{AB} by calculating the value of

$$\Delta_{AB} = \frac{1}{|S \cup \{P\}|} \sum_{X \in S} d(c_B(X), \frac{1}{2}(c_A(X) + c_B(X))).$$

This updated value is the average difference across all propositions they have disagreed upon, including the new proposition P. Next time they disagree over another proposition, A should calibrate B's credence with this updated Δ_{AB}.

As we can see, the value derived this way is not really one's reliability. It is a agent-relative value which varies with the individuals involved. That is, for another individual C whose credences differ from A's, the magnitude C has to calibrate when she disagrees with B would be different from Δ_{AB}. Although it is not the notion of reliability we are familiar with, it does provide an useful guide for an individual to calibrate her credences.

We may consider a concrete example to see how this update strategy is carried out. Suppose that the dermatologist Stephanie disagrees with Conor concerning a patient's skin condition. Stephanie has a 0.7 credence that the patient has atopic eczema, while Conor has only 0.3 credence in the claim. Stephanie, as an open-minded person, wants to make a conciliation. Since they have never disagreed in the past, Stephanie sees Conor as an epistemic peer who is equally reliable as herself. Thus, she conciliates with Conor and forms a 0.5 credence in the patient having eczema after realising that Conor's credence differs from hers. However, this is not the whole story. Since Stephanie knows that Conor's credence concerning the patient's having atopic eczema is 0.2 short of the outcome of conciliation, she realises that she has to calibrate Conor's credence next time when they disagree. Suppose that, at some later point, they face a different patient who has a skin condition similar to the first patient. After checking the evidence, Conor forms a 0.4 credence over the second patient having atopic eczema, while Stephanie's credence is 0.7. In this case, how should Stephanie conciliate with Conor? Stephanie knows that Conor's credence is 0.2 lower than the outcome of conciliation on this kind of problem last time. Knowing that Conor's credence in the second patient having eczema is 0.4, Stephanie should calibrate and take 0.6 as the correct credence. Also, she needs to update the magnitude she has to calibrate. If, in the second case, they do conciliate as if they are epistemic peers, the outcome of conciliation would be 0.55. The difference between Conor's credence and the outcome is thus 0.15. Knowing this, Stephanie should update the magnitude she should calibrate by averaging this value and the old one. The result she gets would thus be 0.175. If they see a third patient, Stephanie should calibrate Conor's credence by this value and take the outcome as the correct credence.

By following this update strategy, Stephanie does conciliate with Conor. She does not retain her original credence, but instead takes Conor's credence into consideration when forming her new credence on the dispute. On the other hand, she also revises her estimate of Conor's reliability based on the disagreement. The action of recalculating the average difference between two credence functions after a disagreement generates the extent one's credence deviates from the result of conciliation, and hence tells us how reliable one is. The process of deriving Δ, hence, could be understood as a process of updating one's reliability. Following this strategy, one may both conciliate with one's interlocutor and update the interlocutor's reliability.

Given this strategy, how should one deal with cases involving more than two individuals? Suppose an individual A disagrees with two individuals B and C, how should A conciliate with both them and revise her estimate of the reliability of them? In such a case, A should first calibrate both B and C's credences separately. That is, A should calibrate B's credence with the average difference between B's credences and the results of their conciliation, namely Δ_{AB}. Also, A should calibrate C's credence in the same way with Δ_{AC}. The calibrated credence of both B and C, supposedly, are the credences for A to adopt. Since both B and C's credences are calibrated, the two values should be equally close to the ideal credence for A to have. Hence, A should adopt the average of the two values as her credence over the disputed matter. By expanding the result derived from this case, we may see how to apply the update strategy to a case involving multiple individuals.

How does this new way of conciliating get rid of the dilemma of Independence? It should be clear that, if an individual follows this strategy, she conciliates with her interlocutors when they disagree as she changes her credence. Since she adopts the others' view concerning the dispute and does not immediately see others as unreliable, she would not be trapped in an epistemic echo chamber. On the other hand, an individual following this strategy would not easily end up in an epistemic bubble, since she does not see everyone as a fully reliable and independent information source. She admits that her interlocutors are less than fully reliable and sometimes provide biased or inaccurate information. Hence, she would not adopt her interlocutors' credences without doubt but would calibrate them. If every member of an epistemic network adopts this strategy, the probability for them to misjudge the number of independent information sources would be reduced. As a consequence, the probability for this epistemic network to evolve into an epistemic bubble could also be reduced.

One may challenge this approach by pointing out that those who adopt this strategy may still misjudge the number of information sources. That is, it is still possible for there to be several individuals in the same epistemic network who collect information from the same source. In such a case, the problem that the individuals mistakenly amplify the importance of some pieces of information may arise. Indeed, it is possible for there to be multiple individuals collecting information from the same source. However, since what we consider are not the credences of these individuals but the calibrated credences of them, it is not the case that all of them are treated as fully reliable individuals collecting information from mutually independent sources. We may illustrate this with a toy example. Consider a group of five individuals A, B, C, D and E where A, B and C have the same credence x over a proposition P since they collect information from the same source. On the other hand, D collects information from another source and has the credence y over P. Let their credence functions respectively be $c_A(\cdot), c_B(\cdot), c_C(\cdot), c_D(\cdot)$ and $c_E(\cdot)$, we may derive that

$$c_A(P) = c_B(P) = c_C(P) = x,\ c_D(P) = y.$$

If they simply conciliate with each other normally, the result would be

$$\frac{1}{4}(3x + y).$$

Such a result is far from ideal. A, B and C collect information from the same source and form the same credence in P. Their credence x should be counted only once. When all the members conciliate by adopting the average as their new credence, the credence x is assigned with an incorrect weight and thus mistakenly amplified. This, as we have seen, is a typical case where conciliation leads to an epistemic bubble.

What would happen if they calibrate each other's credence? We may consider this from the fifth individual E's perspective. Suppose that E comes in and calibrates the others' credences according to the record of the disagreements that have happened between E and the other individuals. Her credence in this dispute would be

$$\frac{1}{4}(3x - \Delta_{EA} - \Delta_{EB} - \Delta_{EC} + y - \Delta_{ED}).$$

E forms her credence in P with several factors, including the others' credences in P, her own credence in P and the average difference between her credences and the others' credences over other propositions. Since such an outcome involves several factors which need to be derived from E's perspective, it is unlikely for the other members of this network to have the same credence as E does.[45] Consider D's credence. Suppose that E's credence over P is z. If D adopts the strategy I proposed, her credence would be

$$\frac{1}{4}(3x - \Delta_{DA} - \Delta_{DB} - \Delta_{DC} + z - \Delta_{DE}).$$

Since there are many factors involved, the probability for D and E to have the same credence is low. For this reason, it is quite unlikely for the members in the group to reach a consensus. If all of them, instead of making a simple conciliation, choose to calibrate the others' credences, all their credences would more or less be different. A direct consequence is that the members would not be overly confident in their credences. Recall that when every member in a network has the same credence over a proposition, it is likely for them to very confident in holding their credence. In such a group, it is quite natural for one to think that one's credence is the correct one. After all, everyone else shares exactly the same credence. Such a case would easily collapse into an epistemic bubble, as the members are overly confident in

[45]Here what I meant by E's perspective is the difference between E's credence and the others' credences in the past, namely Δ_{EA}, Δ_{EB}, Δ_{EC} and Δ_{ED}. It is possible that the sum of these values happen to coincide with the average difference between another individual and the other members. Hence, it is possible that another member calibrates the others' credences and get the same outcome as E. However, the probability for such cases is low as too many factors are involved.

their credences. If, on the contrary, every member's credences always differ from others' credences to some extent, members in the group would be more cautious and less confident in holding their credences. They do find the others having similar credences over some propositions, but their credences always deviate to at least a small extent. Thus, the reason for one to hold fast one's credence is much weaker. The probability of forming an epistemic bubble, therefore, would be lower.[46]

4.11 Conclusion

Epistemologists in the past have aimed to find an ultimate solution to disagreement which deals with all kinds of cases in a impeccable way. This project looks impractical. No matter whether one conciliates, it is always possible for one to end up with the incorrect credence. That is, there does not seem to exist a single solution which completely gets us rid of the probability of having a wrong credence. A possible and more practical goal is to develop a diachronic strategy which leads to the least problematic result in the long run. In other words, for the study of disagreement to be genuinely fruitful, what we should do is not to aim for an one-shot solution, but a solution which gradually leads us to an ideal outcome.

Apart from getting us out of the dilemma of Independence, the strategy I propose here has an additional advantage. It provides a better response to cases of idealised disagreement. Traditionally, one's reliability is taken as a static notion determined by factors that are independent from the disagreement, such as the training one has received. On the contrary, my strategy does not need these external factors. One's reliability can be determined solely with the track record of previous disagreements between the interlocutors. Hence, this strategy provides a better response to idealised cases of peer disagreement where all the disagreeing peers are equal regarding the external factors. For the same reason, the strategy

[46]A standing worry is that one may not be systematically biased. That is, we are in no position to claim that when an individual A is biased regarding a proposition p_1, she is also biased regarding another proposition p_2. A possible solution is to claim that all these propositions involved are in the same category and hence similar in their content. As a result, it is reasonable for us to assume that one is systematically biased.

also tells us how, in the absence of additional information, we should calibrate a person's reliability.

Since this strategy gets people out of the dilemma between an epistemic bubble and an echo chamber, it should be adopted as a general strategy for one to update her credence upon knowing the credences of others. The next step for us to take, therefore, is to explore the features of this strategy and examine it by simulating the practical results it brings. Moreover, we can review previous discussions concerning peer disagreement in light of this new strategy. By doing so, a huge proportion of issues surrounding peer disagreement can hopefully be settled.

Chapter 5

A pluralistic View of Formal Methods

In this last chapter, I critically review the projects introduced in this thesis and draw a conclusion concerning the use of formal tools in philosophical research.

The three projects share a common theme: philosophers should consider a sufficiently wide range of formal tools when facing philosophical problems. What is shown, substantially, is that philosophers sometimes make hasty generalisations when deriving normative claims with a formal approach. Recall that there are two primary uses of formal philosophy. On the one hand, we take formal apparatuses as handy tools for describing a philosophical position. With the help of formal apparatuses, philosophers can deal with problems under a more fine-grained framework and obtain results that are more accurate. On the other hand, when we model a philosophical position with a formal theory, the constraints of the formal theory can be taken as norms governing the philosophical position modelled. By applying formal tools, philosophers can derive normative results concerning the philosophical position modelled. However, there exists an asymmetry between the two uses of formal apparatuses. When a philosopher aims to formally describe a philosophical position, she only needs to find a single theory which correctly captures the features of the position. On the contrary, when a philosopher intends to derive a normative claim about a position, she should consider a variety of formal theories which can model the position in question.

All three chapters serve as examples showing the point I intend to make. In the second chapter, we see that Koscholke and Schippers implicitly argue against the idea of measuring coherence in terms of relevance. They claim that there is no relevance-sensitive measure of coherence which does not suffer from the problem of common cause. Hence, all these measures fail to correctly capture the notion of coherence.

The problem of their reasoning lies in the assumption that they have examined all the relevance-sensitive measure of coherence. With the modified version of these measures, I have shown that the conclusion to be drawn is not that coherence should not be measured by relevance, but that an important factor, namely the number of confirmation relations, should be considered when we measure the coherence of a set. Thus, what Koscholke and Schippers argue for is the normative claim that the notion of relevance should not be taken to measure the coherence of a set, but what I show, in opposition to their claim, is that the notion of coherence can be measured in terms of relevance if we adopt a different averaging function. We may therefore see a case in which philosophers fail to consider alternative averaging functions and end up with a fallacious normative result.

The third chapter, similarly, shows that a normative conclusion needs to be supported by a complete exploration of different formal tools. In the debate over whether to conciliate in a disagreement, some philosophers mistakenly take linear averaging as the only method for us to conciliate. In other words, they equate conciliating to conciliating with linear averaging. Based on this incorrect assumption, they conclude that conciliating leads to some significant formal deficiencies. What I point out is that there are at least three averaging functions which people can adopt to conciliate. Hence, one should not take linear averaging as the only method for conciliating but should consider every possibility. If we adopt the nonlinear averaging functions to make conciliation, we can avoid some of the problems philosophers pointed out.

The fourth chapter illustrates a different point that if one attempts to use a formal theory descriptively and models a philosophical notion, one need not consider all the possibilities. In the project, I introduced a dilemma concerning the Principle of Independence. The key to avoiding the problem, I argue, is to formulate the notion of reliability in a new way and, based on this new formulation of reliability, develop a new strategy for updating our estimate of our interlocutor's reliability. By doing so, one can avoid the potential pitfalls.

Although the two uses are intertwined and sometimes hard to clearly separate, we may still draw a conclusion from the case studies provided in this thesis: philosophers should embrace a pluralistic view of formal methods and explore, to the maximal possible extent, different formal apparatuses before they can reach a normative conclusion.

Bibliography

Aczél, J. and Wagner, C. (1980). A characterization of weighted arithmetic means. *SIAM J. Matrix Analysis Applications*, 1:259–260.

Akiba, K. (2000). Shogenji's probabilistic measure of coherence is incoherent. *Analysis*, 60(4):356–359.

Arntzenius, F. (2003). Some problems for conditionalization and reflection. *Journal of Philosophy*, 100(7):356–370.

Asch, S. E. (1955). Opinions and social pressure. *Scientific American*, 193(5):31–35.

Asch, S. E. (1956). Studies of independence and submission to group pressure: I. a minority of one against a unanimous majority. *Psychological Monographs*, 70(9).

Bacchus, F., Jr, H. E. K., and Thalos, M. (1990). Against conditionalization. *Synthese*, 85(3):475–506.

Ballantyne, N. (2018). Is epistemic permissivism intuitive? *American Philosophical Quarterly*, 55(4):365–378.

BonJour, L. (1985). *The Structure of Empirical Knowledge*. Harvard University Press.

Bovens, L. and Hartmann, S. (2003). *Bayesian Epistemology*. Oxford: Oxford University Press.

Bovens, L. and Olsson, E. J. (2000). Coherentism, reliability and bayesian networks. *Mind*, 109(436):685–719.

Bradley, R., Dietrich, F., and List, C. (2014). Aggregating causal judgments. *Philosophy of Science*, 81(4):491–515.

Briggs, R. and Pettigrew, R. (2020). An accuracy-dominance argument for conditionalization. *Noûs*, 54(1):162–181.

Carnap, R. (1950). *Logical Foundations of Probability*. Chicago]University of Chicago Press.

Chisholm, R. M. (1966). *Theory of Knowledge*. Englewood Cliffs, N.J., Prentice-Hall.

Christensen, D. (1999). Measuring confirmation. *Journal of Philosophy*, 96(9):437–461.

Christensen, D. (2007). Epistemology of disagreement: The good news. *Philosophical Review*, 116(2):187–217.

Christensen, D. (2011). Disagreement, question-begging and epistemic self-criticism. *Philosophers' Imprint*, 11.

Christensen, D. (2013). Epistemic modesty defended. In Christensen, D. and Lackey, J., editors, *The Epistemology of Disagreement: New Essays*, page 77. Oxford University Press.

Christensen, D. (2019). Formulating independence. In Skipper, M. and Steglich-Petersen, A., editors, *Higher-Order Evidence: New Essays*, pages 13–34. Oxford University Press.

Crupi, V., Tentori, K., , and Gonzalez, M. (2007). On bayesian measures of evidential support: Theoretical and empirical issues. *Philosophy of Science*, 74(2):229–252.

Dietrich, F. (2010). Bayesian group belief. *Social Choice and Welfare*, 35(4):595–626.

Dietrich, F. and List, C. (2016). Probabilistic opinion pooling. In Hajek, A. and Hitchcock, C., editors, *Oxford Handbook of Philosophy and Probability*. Oxford: Oxford University Press.

Douven, I. and Meijs, W. (2007). Measuring coherence. *Synthese*, 156(3):405–425.

Easwaran, K., Fenton-Glynn, L., Hitchcock, C., and Velasco, J. D. (2016). Updating on the credences of others: Disagreement, agreement, and synergy. *Philosophers'Imprint*, 16:1–39.

Elga, A. (2007a). Reflection and disagreement. *Noûs*, 41(3):478–502.

Elga, A. (2007b). Reflection and disagreement. *Noûs*, 41(3):478–502.

Elkin, L. and Wheeler, G. (2018a). Resolving peer disagreements through imprecise probabilities. *Noûs*, 52(2):260–278.

Elkin, L. and Wheeler, G. (2018b). Resolving peer disagreements through imprecise probabilities. *Noûs*, 52(2):260–278.

Ewing, A. C. (1934). Idealism: A critical survey. *Philosophy*, 9(36):476–482.

Feldman, R. (2006). Epistemological puzzles about disagreement. In Hetherington, S., editor, *Epistemology Futures*, pages 216–236. Oxford University Press.

Fitelson, B. (2003). A probabilistic theory of coherence. *Analysis*, 63(3):194–199.

Fitelson, B. and Jehle, D. (2009). What is the 'equal weight view'? *Episteme*, 6(3):280–293.

Gardiner, G. (2014). The commutativity of evidence: A problem for conciliatory views of peer disagreement. *Episteme*, 11(1):83–95.

Genest, C. (1984). A characterization theorem for externally bayesian groups. *Ann. Statist.*, 12(3):1100–1105.

Greaves, H. and Wallace, D. (2005). Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind*, 115(459):607–632.

Hild, M. (1998). The coherence argument against conditionalization. *Synthese*, 115(2):229–258.

Kelly, T. (2005). The epistemic significance of disagreement. In Hawthorne, J. and Gendler, T., editors, *Oxford Studies in Epistemology, Volume 1*, pages 167–196. Oxford University Press.

Kelly, T. (2008). Disagreement, dogmatism, and belief polarization. *Journal of Philosophy*, 105(10):611–633.

Kelly, T. (2010). Peer disagreement and higher order evidence. In Goldman, A. I. and Whitcomb, D., editors, *Social Epistemology: Essential Readings*, pages 183–217. Oxford University Press.

Kemeny, J. G. and Oppenheim, P. (1952). Degree of factual support. *Philosophy of Science*, 19(4):307–324.

Keynes, J. M. (1921). *A Treatise on Probability*. Dover Publications.

King, N. L. (2012). Disagreement: What's the problem? or a good peer is hard to find. *Philosophy and Phenomenological Research*, 85(2):249–272.

Klein, P. and Warfield, T. A. (1994). What price coherence? *Analysis*, 54(3):129–132.

Koscholke, J. and Schippers, M. (2019). Coherence and common causes: Against relevance-sensitive measures of coherence. *British Journal for the Philosophy of Science*, 70(3):771–785.

Kuan, K. (2020). Beyond linear conciliation. *Synthese*, pages 1–22.

Lackey, J. (2008). A justificationist view of disagreement's epistemic significance. In Haddock, A. M. A. and Pritchard, D., editors, *Social Epistemology*, pages 145–154. Oxford University Press.

Leitgeb, H. and Pettigrew, R. (2010). An objective justification of bayesianism i: Measuring inaccuracy. *Philosophy of Science*, 77(2):201–235.

Lewis, C. I. (1946). An analysis of knowledge and valuation. *Review of Metaphysics*, 2(7):99–115.

Martini, C., Sprenger, J., and Colyvan, M. (2013). Resolving disagreement through mutual respect. *Erkenntnis*, 78(4):881–898.

Matheson, J. (2014). Disagreement: Idealized and everyday. In Vitz, J. M. R., editor, *The Ethics of Belief: Individual and Social*, pages 315–330. Oxford University Press.

McConway, K. J. (1981). Marginalization and linear opinion pools. *Journal of the American Statistical Association*, 76(374):410–414.

Moon, A. (2018). Independence and new ways to remain steadfast in the face of disagreement. *Episteme*, 15(1):65–79.

Mortimer, H. (1988). *The Logic of Induction*. Halsted Press.

Mulligan, K., Simons, P., and Smith, B. (1984). Truth-makers. *Philosophy and Phenomenological Research*, 44(3):287–321.

Nguyen, C. T. (2018). Echo chambers and epistemic bubbles. *Episteme*, page 1–21.

Nozick, R. (1981). *Philosophical Explanations*. Harvard University Press.

Olsson, E. J. (2002). What is the problem of coherence and truth? *Journal of Philosophy*, 99(5):246–272.

Pettigrew, R. (2019). What is conditionalization, and why should we do it? *Philosophical Studies*, pages 1–37.

Pettit, P. (2006). When to defer to majority testimony - and when not. *Analysis*, 66(3):179–187.

Plantinga, A. (2000). *Warranted Christian Belief*. Oxford University Press USA.

Reichenbach, H. (1956). *The Direction of Time*. Dover Publications.

Rescher, N. (1958). A theory of evidence. *Philosophy of Science*, 25(1):83–94.

Roche, W. (2013). Coherence and probability: A probabilistic account of coherence. In Araszkiewicz, M. and Savelka, J., editors, *Coherence: Insights from philosophy, jurisprudence and artificial intelligence*, pages 59–91. Dordrecht: Springer.

Rorty, R. (1979). Philosophy and the mirror of nature. *Philosophy*, 56(217):427–429.

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.

Schupbach, J. N. (2011). New hope for shogenji's coherence measure. *British Journal for the Philosophy of Science*, 62(1):125–142.

Shogenji, T. (1999). Is coherence truth conducive? *Analysis*, 59(4):338–345.

Shogenji, T. (2001). Reply to akiba on the probabilistic measure of coherence. *Analysis*, 61(2):147–150.

Shogenji, T. (2012). The degree of epistemic justification and the conjunction fallacy. *Synthese*, 184(1):29–48.

Siebel, M. (2005). Against probabilistic measures of coherence. *Erkenntnis*, 63(3):335–360.

Stewart, R. T. and Quintana, I. O. (2018). Probabilistic opinion pooling with imprecise probabilities. *Journal of Philosophical Logic*, 47(1):17–45.

Sunstein, C. R. (2017). Polarization. In *#republic: Divided Democracy in the Age of Social Media*, chapter 3, pages 59–97. Princeton University Press.

Titelbaum, M. G. (2015). Rationality's fixed point (or: In defense of right reason). In Gendler, T. S. and Hawthorne, J., editors, *Oxford Studies in Epistemology, Volume 5*, pages 253–294. Oxford University Press.

Titelbaum, M. G. (2019). Return to reason. In Skipper, M. and Steglich-Petersen, A., editors, *Higher-Order Evidence: New Essays*, pages 226–245. Oxford University Press.

van Inwagen, P. (1996). It is wrong, everywhere, always, for anyone, to believe anything upon insufficient evidence. In Jordan, J. and Howard-Snyder, D., editors, *Faith, Freedom and Rationality*, pages 137–154. Savage, Maryland: Rowman and Littlefield.

White, R. (2005). Epistemic permissiveness. *Philosophical Perspectives*, 19(1):445–459.

Wilson, A. (2010). Disagreement, equal weight and commutativity. *Philosophical Studies*, 149(3):321–326.

Zollman, K. J. S. (2012). Social network structure and the achievement of consensus. *Politics, Philosophy & Economics*, 11(1):26–44.