A NON-IDEAL EPISTEMOLOGY OF DISAGREEMENT: PRAGMATISM AND THE NEED FOR DEMOCRATIC INQUIRY

Winston Sun Leung

A thesis submitted in fulfilment of the requirement for the Degree of Doctor in Philosophy

Department of Philosophy

Faculty of Arts and Social Sciences

The University of Sydney

2023

ACKNOWLEDGEMENTS

When I started this dissertation five years ago, I wanted to write something which would explain what we should do when we disagree about who or what is rational to believe. Given my background in formal logic and epistemology, I thought that the answer to the kind of disagreement that I was interested in would be an *epistemic* one – one which explained how the probabilistic confirmation which a body of evidence confers to a proposition or doxastic attitude, or some deductive tests for logical consistency, would lead people from drastically different walks of life to come to find the same rational conclusion. Little did I know that the answer that I would end up at would be an *ethical* and *political* one – one which highlighted the ethical importance of *trusting* other people and institutions as a pre-requisite condition to obtaining the evidence and intelligent methods which make us rational in the first place. To see disagreement as a positive opportunity to learn something from other epistemic agents – as this thesis argues – is therefore a conclusion which only makes sense once we realise how indebted we are to those around and before us in enabling us to be as rational as we are today.

This is why it is only right that I acknowledge my own debts to the people who have enabled me to write this thesis. Without all of the people I mention below, this thesis would not be what it turnt out to be, and neither would I be who I am today.

First, to the weekly reading group which was run by my supervisor Professor David Macarthur, in which I workshopped all of the material in this thesis. Thank you to Daniel Pash, Oliver Gordon, Samuel Khoo, Gavin Kitching, Rory Torrens and Georgia Fagan for all the comments and feedback over the years. Special thanks to Olly for pointing me to the importance of the ideal/non-ideal distinction; Dan, for consistently admonishing me to write what I think; and for Gavin, for the weekly pre-reading group feedback and encouragement.

To Associate Professor Brian Hedden, who was the supervisor for my Honours thesis, and the former auxiliary supervisor of this thesis until he moved to ANU. Not only did I learn practically everything I know about formal epistemology and decision theory from Brian, but any philosophical rigour that you might see in this thesis comes from his many detailed and astute criticisms that he has given me on my work over the years.

To Professor Nicholas Smith, who is the auxiliary supervisor of this thesis. Taking introductory logic with Nick in my undergrad was one of the reasons why I became interested in philosophy to begin with, so I was honoured and extremely excited when Nick agreed to come on board this thesis. Having now worked with Nick on a closer basis, what I have been most thankful for is Nick's patient and even-handed criticism and consideration towards all of the arguments and drafts that I sent to him. Even in the moments where I was disillusioned if not scornful about certain methods or projects in analytic epistemology, Nick would always take my thoughts and opinions in stride, and then point me in the right direction by questioning the reasoning for my claims. I would not be surprised if the conclusion of this thesis – namely, that open and honest intellectual discussion with dissenting parties can lead not only to agreement, but improved epistemic practice collectively – was unconsciously influenced by these discussions with Nick, and I have been very grateful for all the help and feedback.

To my primary supervisor, Professor David Macarthur. From the very first seminar which I took with David – a seminar on Wittgenstein during my Honours – I was hooked by the philosophical picture which David weaved about Wittgenstein, Stanley Cavell, and pragmatism, and knew immediately that David's work was the kind of philosophy that I wanted to pursue and work on. And throughout this postgraduate, I have seen firsthand how David's passion for teaching extends beyond simply teaching ideas and arguments, to teaching *people* to be better critics who pay attention to the fine details – details such as the often "loaded" language which we begin philosophising with; or the ways in which our attempt to clarify our ideas through abstract theorising can lead us to neglect the complexities of

our actual life and practice. Being under David's supervision has therefore not only taught me to be a better philosopher, but to be better at hearing what people are actually saying, and saying what I actually mean. To say the least, I have been personally enriched in this postgraduate experience, and it has been an absolute privilege to work with David.

Next, for some personal words of thank you.

To my close friends who have continually encouraged and supported me throughout this thesis – to Justin, Crista, Psalm, and Hong. Thank you for the constant care and support you have shown me throughout these years – be it your willingness to listen to me rant despite how "crazy" philosophy can often sound, or the ways in which you have each kept me sane during the stress and fatigue of the whole process, I am extremely grateful to call you all friends, and could not have finished this thesis without your support.

To my father. My father has always been something of a hero in my eyes – someone who everybody (including myself) would go to for advice, and someone who always just knew what to do or what to say. Dad, I know you've worried about me whenever I told you how stressed or tired I have been from this thesis, and I know it's been hard not being able to see each other while I have been busy writing. Thank you for always caring about me, and giving me the advice that I need to hear when I am lost, confused, or unsure. I am really grateful for having you in my life, and very proud to be your son.

To my sisters, Emily and Carol. There is no-one in this world that I have more inane disagreements with, and yet no-one in this world who give me greater joy than you two. Thank you for tolerating and supporting me when I was a grouchy recluse; and for constantly reminding me to not take life so seriously, and therefore actually enjoy life. Words cannot express how grateful I am for you two and I can only hope to be as good of a brother to you as well. Or maybe I will just buy you more food.

To my mother. If I do achieve this doctorate degree, and were it permissible for me to do so, I would give half of the degree to my mother. This is because my mother has been *the* unceasing source of comfort and support throughout the entire writing process. There is a Chinese saying that a son with a mother is a treasure, and a son without a mother is grass. (Admittedly, it sounds better as a Chinese rhyme.) But the sentiment remains: truly, I am so lucky to have such a kind and selfless mother, and cannot explain how grateful I am to have her in my life. Mum, thank you for everything you do for me – I love you.

And finally, to the One who has said that He has made everything beautiful in its time. While I have much to learn about who He is and what He has done in my place, I can at least attest to the veracity of this claim through my postgraduate experience – what a beautiful ride this life has been so far.

TABLE OF CONTENTS

1   THE EPISTEMIC PROBLEM OF DISAGREEMENT

1.1   *Introduction: On the Need for an Non-Ideal Epistemology of Disagreement*

Disagreement is an inescapable aspect of human life. In the everyday disputes that we have with

friends, family members, and lovers; or the larger socio-political, economic, and religious debates

which shape our society and the wider world; disagreement is something which we all experience, and

a persistent and intractable problem between us and others.

The intractability which we face in our everyday disagreements has only worsened in the current

Information Age. Rather than aiding us in the resolution of disagreement, the proliferation of

information that is now accessible to us through digital media has created a paradoxical situation

where we as epistemic agents have more access to knowledge and information than any other moment

in history, and yet are often too skeptical or ill-equipped to use this information in any meaningful

way. We have become rightfully skeptical of the information that we have access to because of the

sheer amount of intentional or unintentional misinformation that is shared on the World Wide Web.

This is such that the relative ease with which information can be disseminated on social media

platforms, as well as the minimal accountability which these social media platforms put on its users,

makes it so that even qualified professionals and reputable news sources can unknowingly use and

then propagate false sources of information, let alone the general public.[1] Furthermore, even when we

receive reliable first-hand information to consider for ourselves, such information is often expressed in

the technical jargon and unfamiliar forms of reasoning used by specialised research and bureaucratic

---

[1] For some excellent work on the epistemic problems which social media cause, see Karen Frost-Arnold, "Trustworthiness and Truth: The Epistemic Pitfalls of Internet Accountability," *Episteme* 11, no. 1 (2014): 63–81; Karen Frost-Arnold, "Social Media, Trust, and the Epistemology of Prejudice," *Social Epistemology* 30, no. 5-6 (2016): 513–31.

processes, such that we have little to no chance of understanding the first-hand information, let alone using it within epistemic practice.[2] As such, for all of the incredible advancements in knowledge and understanding which modern science and technology has made and made available to us online, our social epistemic environment still does not have a clear answer as to how we use this information to resolve, or at least reasonably respond, to the prominent disagreements of our time: when faced with disagreement about what is required for a healthy diet, who one should vote for, or what one should do in a global pandemic, we are left at a loss as to who or what is rational to believe. This is true in both the immediate disagreements which we have with the people around us, and in the wider social or global disagreements about issues we are forced to take a stance on.

The aim of this thesis is to provide a full-fledged *pragmatist epistemology of disagreement*, one which provides an alternate epistemic account of disagreement to the one centrally given within analytic social epistemology. Within analytic social epistemology, the well-established literature on the epistemology of disagreement has provided a comprehensive epistemic account of disagreement through the concept of "epistemic peerhood": that is, a relational concept that compares the relative

---

[2] In their famous handbook on how to read well, Mortimer J. Adler and Charles Van Doren note that "until approximately the end of the nineteenth century, the major scientific books were written for a lay audience…[such that] intelligent and well-read persons were expected to read scientific books as well as history and philosophy; there were no hard and fast distinctions, no boundaries that could not be crossed." And yet, Adler and Van Doren lament how this is increasingly rare in modern times. See Mortimer J. Adler and Charles Van Doren, *How to Read a Book* (New York: MJF Books, 1972), 255. Of course, while there are still a great number of bestsellers which explain important scientific subjects and research for the general public, these popularisations are usually written by journalists rather than by the scientists themselves. Elijah Millgram argues that this is an important point because academic writing and journalism are often held to different professional standards of writing, such that certain epistemic features (such as the defeasibility conditions of the information presented) are often obscured within popular journalism. See Elijah Millgram, *The Great Endarkenment: Philosophy for an Age of Hyperspecialization* (Oxford University Press USA, 2015), 28ff. For more information on the potential ways in which the different professional standards of scientists and journalists can lead to miscommunication, Millgram cites the following: Naomi Oreskes and Erik M. Conway, *Merchants of Doubt* (New York: Bloomsbury Press, 2010), 269ff.

likelihood that one is correct about a given proposition to another, such that person A is an epistemic superior, inferior, or peer to person B just in case person A is more likely, less likely, or equally likely to be correct about a given proposition than person B is respectively. In analysing the epistemic problem of disagreement through the concept of epistemic peerhood, the analytic epistemology of disagreement takes the epistemic challenge which disagreement poses to be principally about the epistemic peerhood of a dissenting party. Under this account, the fact that another epistemic agent disagrees with me does not in and of itself pose an epistemic challenge for my beliefs: rather, a disagreement becomes an epistemic challenge to my belief once I learn that the person I am disagreeing with is someone who is capable of showing that I am wrong (i.e. that is, I have reason to believe that I am on the wrong side of a disagreement because they are an epistemic peer or superior). The concept of epistemic peerhood, as it is explicated through the theoretical case of ideal peer disagreement, therefore allows the analytic epistemology of disagreement to provide an overall evidence-based response to disagreement: once I have enough evidence that the person who I am disagreeing with is an epistemic superior, inferior, or peer, I can determine my response to the disagreement by deferring to expertise, moderately changing my beliefs, or remaining steadfast about my beliefs according to my evidence. Such is the general gist of the familiar and well-thought-out response given within the analytic epistemology of disagreement.

What is somewhat less known, however, is the alternate social-epistemic account of disagreement which has been developed contemporaneously within the literature that intersects American pragmatism and democratic theory. Central to this pragmatist account has been the notion of social *inquiry*: that is, the idea of a continuous deliberative process in which a group of epistemic agents collectively come to find an answer about a given problem or question. In drawing from the epistemic features of effective social inquiry, such as the examples of science and democracy, the pragmatists argue that effective social epistemic practice often takes a similar pattern or structure: one in which a group of epistemic agents freely exchange ideas, reasons, and objections within public discussion; form

a collective plan of action in light of this discussion; and revise this course of action by assessing its consequences. What is fascinating about this process of inquiry is its ability to leverage the disagreement within the community of inquirers as a way to *improve* collective epistemic practice: when disagreement is intentionally facilitated and regulated within the public discussion which pervades all stages of social inquiry, disagreement can in fact become an epistemic boon through which a community of inquirers can obtain a diverse pool of ideas, methods, and evidence. The pragmatists therefore reveal how disagreement is a feature rather than an obstacle within inquiry: one which we can use to find valuable information to come up with better plans of action during deliberation, or to improve existing plans of action upon revision. Such is the pragmatist account of disagreement, one which has a proven track record through the examples of experimental science and democratic reform.

In this thesis, I shall argue that the pragmatist theory of inquiry provides a better epistemic account of disagreement than the one that is proposed within the analytic epistemology of disagreement. More specifically, I shall argue that the pragmatist theory of inquiry allows us to outline a robust *non-ideal* epistemology of disagreement, one which is better suited to explain how an epistemic agent should respond to real-life cases of disagreement like the one described in my introductory sketch. To set up the basic premises of this argument, however, two key definitional question must first be addressed: first, what we mean exactly by a "non-ideal" epistemology of disagreement; and second, what we mean by pragmatist epistemology. This will be the focus of this first chapter. In particular, the structure of this chapter will be as follows: in §1.2, I shall draw from the literature on ideal and non-ideal theory within political philosophy to make a distinction between ideal and non-ideal epistemology. In particular, I shall draw from Laura Valentini's distinction between ideal and non-ideal political theory to outline two theoretical differences between ideal and non-ideal epistemology. This means that while ideal epistemic theories are typically theories which are utopian and end-state, non-ideal epistemic

theories are typically theories which are realistic and transitional. Once I have explained the distinction

between ideal and non-ideal epistemology, I shall explain why the epistemic problem which

disagreement poses to our everyday epistemic practice is a distinctly non-ideal epistemic problem.

In §1.3, I shall provide a brief account of the pragmatist theory of inquiry so as to clarify what I mean

by "pragmatist epistemology". After a brief survey of the ways in which the term "pragmatist

epistemology" is misunderstood or made ambiguous within the general philosophical discourse, this

section will give an introductory account for the pragmatist epistemology that thesis is concerned with:

that is, the theory of inquiry as it is outlined by John Dewey, and subsequent Deweyan thinkers such as

Hilary Putnam, Ruth-Anna Putnam, Elizabeth Anderson, Jack Knight, and James Johnson amongst

others. In recontextualising non-ideal epistemology as being a practical matter of settling beliefs within

a community of inquirers, I shall argue that the Deweyan theory of inquiry provides a transition

epistemic account which makes it particularly suitable for answering non-ideal epistemic problems

such as the problem of disagreement. As we shall see in Chapter 5, it is for these reasons which

Dewey's epistemology has been adopted within democratic theory, in order to provide a practical

account of how disagreement can be regulated and resolved within cooperative social epistemic

practice. I shall then conclude the chapter with a brief outline of the thesis and its subsequent

chapters.

1.2 Disagreement and the Ideal/Non-Ideal Distinction

1.2.1 On the Theoretical Differences between Ideal and Non-Ideal Epistemology

What is the distinction between ideal and non-ideal epistemology which this thesis is concerned with?

To answer this question, this section shall consider the distinction made between ideal and non-ideal

theory in political philosophy, and show how it can also be used to highlight two kinds of theoretical

features which distinguish ideal epistemic theories from non-ideal epistemic theories. Let us begin

with a brief overview of the ideal/non-ideal distinction as it is used in political philosophy. The

distinction between ideal and non-ideal theory is commonly traced back to John Rawls, and in particular, his discussion of the concept of justice. This distinction between ideal and non-ideal theory was an essential part of Rawls' methodology to political theorising: Rawls believed that before we can understand what social justice practically requires of us in our actual circumstances, a prior *theory* of social justice must first be given which explicates our concept of social justice.[3] For Rawls, the appropriate theory of social justice for this purpose would be an *ideal theory* – one which outlined principles for the basic institutions of society which comprehensively explained the requirements of social justice as the satisfaction of certain counter-factual and idealised social conditions. From there, one can move from ideal theory into *non-ideal* theory (that is, theorising to do with real-world applications of social justice) by adjusting the principles derived from the ideal theory of justice to match the actual social conditions of one's social and political context.

However, while Rawls saw ideal theory as being a necessary precursor to non-ideal theorising, the subsequent literature in political philosophy has since considered ideal theory and non-ideal theory to be separate theoretical projects with different methods and goals.[4] In her survey of the use of ideal and

---

[3] In the beginning of *A Theory of Justice*, Rawls justifies his focus on distributive justice within a closed society by saying that "it is natural to conjecture that once we have a sound theory for [the basic structure of society], the remaining problems of justice will prove more tractable in the light of it." Very soon after, Rawls justifies his idealised premise that all members within society act justly (over what he calls "partial compliance theory") within his social theory like so: "Obviously the problems of partial compliance theory are the pressing and urgent matters. These are the things that we are faced with in everyday life. The reason for beginning with ideal theory is that it provides, I believe, the only basis for the systematic grasp of these more pressing problems…At least, I shall assume that a deeper understanding can be gained in no other way, and that the nature and aims of a perfectly just society is the fundamental part of the theory of justice." See John Rawls, *A Theory of Justice* (Cambridge, Massachusetts: Harvard University Press, 1999), §2.

[4] This thesis will focus primarily on the distinctions between ideal and non-ideal theory as it is explicated within Laura Valentini, "Ideal Vs. Non-Ideal Theory: A Conceptual Map," *Philosophy Compass* 7, no. 9 (2012): 654–64. However, for more information on the ideal and non-ideal distinction within political theory, see Zofia Stemplowska, "What's Ideal about Ideal Theory?" *Social Theory and Practice* 34, no. 3 (2008): 319–40; Alan Hamlin and Zofia Stemplowska, "Theory, Ideal Theory,

non-ideal theory in political philosophy, Laura Valentini helpfully explains how the distinction between ideal and non-ideal theory has been used to mean different things in different contexts. This has meant that the ideal/non-ideal distinction could refer to two kinds of theoretical differences within political theory: namely, the difference between utopian and realistic political theory, and end-state and transition political theory.[5] Let us consider both in turn.

The first is the difference between utopianism versus realism – ideal and non-ideal theories differ in theoretical assumptions insofar as they differ about the extent to which empirical 'feasibility conditions' constrain normative political theorising. An example of utopian political theory would be G. A. Cohen, who considers justice to be "a timeless (and fact-free) value akin to a Platonic ideal", one which is then taken into consideration along with other values, as well as feasibility constraints, when deciding what we practically ought to do. Cohen is therefore an example of utopian political theory, as his account of justice is "altogether independent of factual constraints" and one which considers an institution to be just or unjust irrespective of any empirical fact which may prohibit one from being

---

and the Theory of Ideals," *Political Studies Review* 10, no. 1 (2012): 48–62; Ingrid Robeyns, "Ideal Theory in Theory and Practice," *Social Theory and Practice* 34, no. 3 (2008): 341–62; A. John Simmons, "Ideal and Nonideal Theory," *Philosophy & Public Affairs* 38, no. 1 (2010): 5–36; Lisa Herzog, "Ideal and Non-Ideal Theory and the Problem of Knowledge," *Journal of Applied Philosophy* 29, no. 4 (2012): 271–88; Robert Jubb, "Tragedies of Non-Ideal Theory," *European Journal of Political Theory* 11, no. 3 (2012): 229–46.

[5] Actually, Valentini's paper considers three theoretical differences between ideal theory and non-ideal theory: namely, between full compliance and partial compliance theories; between utopian and realistic theory; and between end-state and transitional theories. For the purposes of this thesis, however, I have taken the distinction between full compliance and partial compliance theory to be a subset of utopian and realistic theories respectively. This is because I take the difference between full compliance and partial compliance as being just another feasibility condition for theorising – political or epistemic. For more information on full and partial compliance theory, however, see Valentini, "Ideal Vs. Non-Ideal Theory".

just in actuality.[6] By contrast, non-ideal political theory is realistic insofar as it begins with an account for the principles of justice which directly make reference to certain key empirical facts – be it material facts such as moderate scarcity, or psychological facts such as limited altruism.[7]

The second difference between ideal and non-ideal theory is between end-state and transition (or transitional) theories – ideal and non-ideal theories have different theoretical aims insofar as they have different conceptions of what it is that their respective theories aim to achieve.[8] Ideal theories tend to be 'end-state' theories whose theoretical goal is to set a comprehensive account of a political value (e.g. justice, freedom, human right) so as to provide long-term ends for institutional reform. Conversely, non-ideal theories tend to be 'transitional' theories whose theoretical goal is to provide contextual accounts of social reform for the sake of improving rather than perfecting society. A helpful example of end-state and transition theory is the difference in approach which John Rawls and Amartya Sen take in their political theorising: whereas Rawls argued that one must begin with an ideal theory of justice in order to have a systematic way by which to evaluate the importance of social

---

[6] Valentini, "Ideal Vs. Non-Ideal Theory," 657. For more on Cohen's theory of justice, see G. A. Cohen, "Facts and Principles," *Philosophy & Public Affairs* 31, no. 3 (2003): 211–45; and G. A. Cohen, *Rescuing Justice and Equality* (Cambridge, Massachusetts: Harvard University Press, 2008).

[7] Interestingly enough, Valentini points out how Rawls' theory of justice would therefore be considered a realistic (and therefore, non-ideal) theory of justice in its incorporation of facts such as limited altruism. Rawls explicitly claims this in his essay "Justice as Fairness: Political not Metaphysical", where Rawls claims that his conception of justice is not dependent on "claims about universal truth, or claims about the essential nature and identity of persons", but one which is developed in response to the human condition. For Rawls, the need to respond to the human condition means that one cannot theorise about justice without including some facts such as those about moderate scarcity and limited altruism. For more on Rawls' methodology, see John Rawls, "Justice as Fairness: Political Not Metaphysical," *Philosophy & Public Affairs* 14, no. 3 (1985): 223–51; and Rawls, *A Theory of Justice*, §22.

[8] Valentini, "Ideal Vs. Non-Ideal Theory," 660–61.

problems, Sen argued that it is possible (and sometimes, prudent) to begin theorising about how to make society *more* just without needing a vision of society which is *fully* just.[9]

Given Valentini's overview of ideal and non-ideal theory in political philosophy, we can now see how the categories of ideal and non-ideal theory might also be used to explain the theoretical differences between certain projects within epistemology. First, we can see that the distinction between utopian and realistic theory can be used to explain how certain epistemological projects differ about the degree to which one should incorporate feasibility constraints within epistemic theorising.[10] One area of the recent analytic literature which highlights this issue is the debate concerning epistemic guidance: that is, the debate about whether the connection between normativity and guidance requires theories about epistemic norms to provide practical guidance.[11] In a classic paper on this topic, Alvin Goldman argues that the debate between internalism and externalism about epistemic justification has often involved a

---

[9] Valentini cites the following for more information about Sen's theoretical assumptions: see Amartya Sen, "What Do We Want from a Theory of Justice?" *The Journal of Philosophy* 103, no. 5 (2006): 215–38.

[10] In one of the few papers within the current literature on ideal and non-ideal epistemology, Jennifer Rose Carr's paper outlines the distinction between ideal and non-ideal epistemology in a similar way to my distinction between utopian and realistic theory here. Furthermore, Carr outlines this distinction in much greater detail than I do: showing how non-ideal epistemic theories (i.e. 'realistic' theories) depend on epistemic value, epistemic conventions, and contextually contingent constraints. Nevertheless, given that the focus of Carr's paper is to argue in favour of ideal epistemology as the only "normatively robust theory of epistemic rationality", it is clear that Carr is only interested within end-state epistemic theory, and not in transition theory. The difference between Carr's view on the connection between ideal and non-ideal epistemology and the one that this thesis outlines is therefore analogous to the difference between Rawls and Sen on political theory: whereas Carr argues that we cannot approximate ideal epistemic rationality without a prior theory about epistemic rationality, I argue that we can (and sometimes should) theorise about what would be more rational for an epistemic agent to do in a given context without a comprehensive account of epistemic rationality simpliciter. For more information, see Jennifer Rose Carr, "Why Ideal Epistemology?" *Mind* 131, no. 524 (2021): 1131–62.

[11] Nick Hughes provides an excellent series of citations about this topic in Nick Hughes, "Epistemology Without Guidance," *Philosophy Studies*, 2021, fnn. 1-2.

meta-epistemological question regarding the *kind* of epistemic account we are looking for in the first place. In particular, Goldman argues that the debate between internalism and externalism has often rested on a deeper disagreement between regulative and theoretical conceptions of justification: whereas the regulative conception of justification argues that an epistemic account of justification should be "designed specifically to guide a cognizer in regulating or choosing his doxastic attitudes", the theoretical conception of justification states that an epistemic account of justification need only "specify the features of beliefs (or other doxastic attitudes) that confer epistemic status".[12]

At this point, it should be clear how the difference between the regulative and theoretical conceptions of justification is essentially a difference about the degree to which feasibility constraints should be accounted for within an epistemic theory of justification. Whereas the regulative conception of justification claims that an adequate account of justification should explain how an epistemic agent can *feasibly* be justified in their choice of doxastic attitudes, the theoretical conception of justification states that a belief is justified just in case it has certain epistemic features, even if it is not feasible for an epistemic agent to know what those features are in practice. This is why Goldman argues that philosophers who criticise his external "reliabilist" view on the basis of guidance requirements fundamentally misunderstand the kind of theory of justification which reliabilism is, given that reliabilism "is not a *rule* or *prescription* for choosing beliefs or other doxastic attitudes [i.e. a regulative conception of justification]…[but a theoretical analysis of] an *already formed* belief of a cognizer

_____

[12] Alvin I. Goldman, "The Internalist Conception of Justification," *Midwest Studies in Philosophy* 5, no. 1 (1980): 28.

[which] says what features are necessary and sufficient for that belief to count as justified [i.e. a theoretical conception of justification]."[13]

In a recent paper on epistemic guidance, Nick Hughes extends Goldman's point by highlighting how conflicting intuitions about guidance considerations also underlies many of the other debates within contemporary analytic epistemology: such as the debate between Uniqueness versus Permissivism within the epistemology of disagreement, or the debates regarding the assumption of sharp credences and logical omniscience within Bayesian epistemology.[14] In each of these debates, Hughes notes that a key part of the debate has been to show how a given view or claim (e.g. the Uniqueness thesis, or the assumption of sharp credences or logical omniscience) is false or otherwise unacceptable on the basis that it is unrealistically demanding, and therefore unable to provide practical guidance in actual situations. Nevertheless, Hughes argues that these criticisms ultimately fail as theoretical arguments because they are similarly based on an underlying intuition that our epistemic theories should provide practical guidance, one which is neither explicitly not sufficiently defended.[15] As such, a common

_____

[13] That being said, Goldman notes that this does not mean that reliablism cannot play a part within a "regulative epistemology" more generally. Goldman, "The Internalist Conception of Justification," 29. Italics are that of the author. For more information on how reliablism and its role within a regulative epistemology, see §9 of Goldman's paper.

[14] For more information, see Hughes, "Epistemology Without Guidance", §11. While Hughes does not explicitly refer to the particular paper from Goldman which we discuss, Hughes does mention similar papers from Goldman in his citations.

[15] In fact, the majority of Hughes' paper consists of an argument for why guidance considerations do *not* play a role within epistemic theorising. It is this argument which leads Hughes' to his conclusion about "epistemology without guidance", and which leads to Hughes' negative stance towards "non-ideal epistemology" more generally. Hughes' criticism of non-ideal epistemology does not have any bearing for the argument of this thesis, however, because Hughes explicitly states that his argument has focused entirely on the lack of practical guidance given by general epistemic norms (e.g. "be truthful" or "be rational") as opposed to local epistemic norms. See §12 of Hughes' paper. Throughout the rest of this thesis, I shall explain how the pragmatist account of epistemic norms see epistemic norms as methodological maxims of inquiry, maxims which require a community of inquirers to interpret and apply them in a given context. In this way, the pragmatists see epistemic

implication of both Hughes and Goldman's arguments is this: once we realise that epistemic theories can be based on different intuitions and conceptions of epistemology, we can see that many of the conflicts between different views in epistemology might not in fact be different theoretical interpretations about the same epistemic concept, but different *kinds* of epistemic theories which aim to capture different epistemic features entirely. We can therefore use the distinction between utopian and realistic theory to distinguish between these kinds of epistemic theories: whereas a utopian epistemic account explicates an epistemic concept (e.g. justification) as a particular feature or value independent of factual constraints, a realistic epistemic account explicates epistemic concepts in reference to crucial empirical facts which act as feasibility constraints in actual epistemic practice.

In a similar way, we can also use the categories of end-state theory and transition theory to highlight the difference between different theoretical projects within epistemology. Take, for example, the difference between the kind of theoretical account of testimony that is given within the debate between reductionism and non-reductionism about testimony, and that which is given within Miranda Fricker's epistemic account of "testimonial injustice". On the one hand, following the debates between Thomas Reid and David Hume, the ongoing debate between reductionism and anti-reductionism is concerned with the theoretical question about the nature of testimony: that is, whether testimony is an independent source of epistemic justification (i.e. anti-reductionism about testimony), or whether testimonial justification is reducible to other sources of justification such as perception, memory, and inductive inference (reductionism about testimony).[16] On the other hand, the account of testimony

---

norms as general norms which must be contextualised as local norms in epistemic practice. For more information, see §4.5 of this thesis, as well as the entirety of Chapter 5.

[16] The current literature on reductionism and non-reductionism about testimony has largely stemmed from the work of C. A. J. Coady. See Cecil AJ Coady, "Testimony and Observation," *American Philosophical Quarterly* 10, no. 2 (1973): 149–55; Cecil AJ Coady, *Testimony: A Philosophical Study* (Oxford: Clarendon Press, 1992). Jennifer Lackey provides an excellent list of

which is given within Miranda Fricker's case of "testimonial injustice" is one which considers

testimony *in media res*: that is, as testimony is evaluated according to the credibility judgments of a

hearer, and as it is distorted by identity prejudice against particular social groups.[17] Here, it is clear that

the former is an example of end-state theory while the latter is an example of transition theory:

whereas the debate surrounding reductionism and anti-reductionism aims to provide a comprehensive

account of testimony as an epistemic end (i.e. whether testimony confers justification to a belief

outright, or whether its epistemic value is derived from other sources), the discussion around

testimonial injustice addresses epistemic problems surrounding testimony in specific social contexts.

As such, we can see how the distinction between ideal and non-ideal theory within political philosophy

can be applied to distinguish between different epistemological projects with different aims and

---

citations for the recent work on both sides of the debate in Jennifer Lackey, "It Takes Two to Tango: Beyond Reductionism
and Non-Reductionism in the Epistemology of Testimony," in *The Epistemology of Testimony*, ed. Ernest Sosa and Jennifer
Lackey (Oxford University Press, 2006), 160–89, see footnotes 3 and 19.

[17] Of course, it is important to note that Fricker's account is not entirely disconnected from the debate between reductionism
and non-reductionism: in her seminal book on testimonial injustice, Fricker uses the case of testimonial injustice precisely to
advance a virtue epistemology of testimony, one which argues for a moral non-reductionist (or "non-inferentialist", in
Fricker's terms) account of testimony. This shows that, at least on Fricker's own view, there is no reason why transition and
end-state accounts about testimony cannot interact with each other: such that theoretical questions about the nature of
testimony inform and are informed by practical questions about the use and reception of testimony. Nevertheless, it is
interesting to note that Fricker's argument for a moral non-reductionism about testimony is almost entirely absent within the
debate on reductionism versus non-reductionism about testimony. One wonders whether this is because Fricker's argument
does not give a 'purely epistemic' analysis of testimony: Jennifer Lackey, for instance, rejects the "interpersonal view of trust"
as being relevant within the epistemology of testimony, given that the *ethical* considerations which it focuses on precludes it
from having *epistemic* importance. In response to Lackey, I will argue that our everyday epistemic practice actually consists of
judgments which cannot be categorised into purely epistemic or purely ethical reasons. See Appendix A of this thesis. For
more on Fricker's virtue epistemological account of testimony, see Chapter 3 of Miranda Fricker, *Epistemic Injustice: Power
and the Ethics of Knowing* (Oxford: Oxford University Press, 2007). For more on Lackey's criticism of the interpersonal view
of trust, see Jennifer Lackey, "Testimony: Acquiring Knowledge from Others," in *Social Epistemology: Essential Readings*, ed.
Alvin I. Goldman and Dennis Whitcomb (Oxford University Press, 2011), 78–83.

methods. With these distinctions in view, we can now clarify the kind of non-ideal epistemology which this thesis is concerned with. Specifically, the central aim of this thesis is to provide a *transition* epistemic account of disagreement, one which explains how epistemic agents should respond to the problems which disagreement poses in actual epistemic practice.[18] In the next part of this section, we will explore in greater detail how disagreement poses distinctly non-ideal epistemic problems, and why these problems are in need of a transition account of disagreement. Before we move on, however, two more preliminary remarks are needed to clarify the non-ideal epistemic account which I outline within this thesis.

First, while the pragmatist response to disagreement which this thesis develops is a non-ideal account with respect to the distinction between transition and end-state theory, it is also an ideal account insofar as it provides a *utopian* answer to the problem of disagreement. From the outset, it is important to note that the distinction between utopian and realistic theory is a distinction about a difference in degree rather than a difference in kind. This is because there is no defining feature which distinguishes utopian from realistic theory: a theory just is more or less realistic (or utopian) depending on the

---

[18] The best example of transition epistemology within the recent literature is the excellent work done within feminist epistemology on trust and testimony. As we will discuss in the next part of this section, the feminist epistemology on trust and testimony has provided a general non-ideal epistemology of testimony: one which outlines a series of epistemic norms surrounding trust (when receiving testimony) and trustworthiness (when communicating testimony) to explain how we can build reliable epistemic networks of trust. The answer that is given within this literature is extremely thorough and complex: one which traverses the boundaries between feminist philosophy, ethics, social epistemology, political philosophy, and philosophy of science. As a quick aside, I consider the non-ideal epistemology of disagreement which I outline in this thesis to be an epistemic account which complements and supports the non-ideal epistemic account given within feminist epistemology, albeit an account which draws from American pragmatism instead. See fn. 37 for more on the recent literature on the feminist epistemology on trust and testimony. For more on how the term 'non-ideal epistemology' has been defined and substantiated within the feminist literature, see Catharine Saint-Croix, "Non-Ideal Epistemology in a Social World" (PhD thesis, University of Michigan, 2018); Robin McKenna, *Non-Ideal Epistemology* (Oxford: Oxford University Press, forthcoming).

amount of empirical facts it considers within the theoretical account. So, for example, it is possible for

a utopian theory to be made more realistic simply by adding more factual constraints: a physical theory

about motion might be idealised (i.e. utopian) in its assumption of a frictionless plane, and then made

more realistic by including the effect of friction as an additional force.

Once we understand how the distinction between utopian and realistic theory is not a categorical

difference but a matter of degree, we can see how there is a sense in which any non-vacuous normative

theory is, to some extent, *utopian*. David Estlund makes this point when he criticises political theory of

"utopophobia": for Estlund, the continual demand for political theory to be more realistic is

misguided, because "the most realistic normative theory of all, of course, would recommend or require

people and institutions to be *just as they actually are already*."[19] Given that this "complacent realism" is

clearly not the kind of account that is desired by any normative theory, Estlund concludes that every

substantial normative theory is necessarily utopian insofar as it ignores certain empirical facts about

one's actual situation in order to provide an account of what a preferred (i.e. *ideal*) situation would

look like. For Estlund, this is especially important for *transitional* accounts like the one given in this

thesis, because it is precisely because normative theory provides a utopian account of political or

epistemic practice that allows one to understand how they could improve in their given context: that is,

in highlighting what would need to change in a future situation in order for one to be more "just" or

more "rational".

In fact, Estlund's point here about the role in which idealisation plays within normative theory – what

Estlund calls the "aspirational" function of ideal theory – is precisely how John Dewey explains his

_____

[19] David M. Estlund, *Democratic Authority: A Philosophical Framework* (Princeton; Oxford: Princeton University Press, 2008), 263. Emphasis is my own. Estlund continues to provide a much more detailed reflection on the aims and use of normative theory in this chapter, which is worth reading for more information.

own account of democracy and social inquiry, the account on which our pragmatist response to

disagreement is based on. Consider the following passage from Dewey about his theory of democracy:

> It is an ideal in the only intelligible sense of ideal: namely, the tendency and movement of some thing carried to its final limit, viewed as completed, perfected. Since things do not attain such fulfillment but are in actuality distracted and interfered with, democracy in this sense is not a fact and never will be. But neither in this sense is there or has there ever been anything which is a community in its full measure, a community unalloyed by alien elements. The idea or ideal of a community presents, however, actual phases of associated life as they are freed from restrictive and disturbing elements, and are contemplated as having attained their limit of development.[20]

For Dewey, there is no question that the actual social and epistemic conditions of everyday life will

make our ideals about just societies and perfect rationality unattainable. Nevertheless, Dewey insists

that this is besides the point, given that ideals "are not intended to be themselves realized but are

meant to direct our course to realizations of potentialities in existent conditions – potentialities which

would escape notice were it not for the guidance which an ideal, or a definition, provides."[21] In this

way, Dewey explains how ideal (i.e. utopian) theoretical accounts provide a "standpoint…that we can

adopt in the course of problem-solving…[one that] opens up new possibilities and reminds us of how

socially established habits of thinking and acting can appear to be unquestionably natural."[22] The

---

[20] This quote is from Dewey's *The Public and Its Problems*, but taken from Matthew Festenstein, "Ideal and Actual in Dewey's Political Theory," in *Pragmatism and Justice*, ed. Susan Dieleman, David Rondel, and Christopher Voparil (New York, NY: Oxford University Press, 2017), 105. In this paper, Festenstein reveals how Dewey's explanation of the role of ideals is consistent between Dewey's view of social democracy and his epistemic theory of inquiry. For more on the 'aspirational' goal specifically within Dewey's epistemic theory of inquiry, see Hilary Putnam and Ruth Anna Putnam, "Dewey's Logic: Epistemology as Hypothesis," in *Words and Life*, ed. James Conant (Cambridge, Massachusetts; London, England: Harvard University Press, 1994), 198–220, 198 and *passim*.

[21] John Dewey, *Logic: The Theory of Inquiry* (New York: Henry Holt; Company, 1938), 303–4.

[22] Festenstein, "Ideal and Actual in Dewey's Political Theory," 109. As a sidenote, Festenstein's exposition also helpfully explains how Dewey's conception of ideals does not fit nicely within the traditional notions of constitutive or regulative ideals, given that both of these notions are in a sense *ideal* ideals. Instead, Dewey sees idealisations as heuristics which we use

importance of idealisation therefore stems from its ability to change the perspective we have on a given

political or epistemic issue through highlighting the fact that a particular action or decision that is

"unrealistic" or "infeasible" in our present moment need not be so in the future. Dewey therefore

reveals how idealisation can play an important role in transition theory as a model by which we

understand what optimal decision-making looks like, and how to act accordingly. Given that the aim

of this thesis is to give a transition account of disagreement – one which explains how we can be *better*

epistemic agents through disagreement and within disagreement – the epistemic account I give is

therefore utopian insofar as it gives an account of how we could *ideally* resolve disagreement, one

which serves as an example through which we can model our actual attempts to resolve disagreement

off.

Second, it is important to note that the use of idealisation (with respect to utopianism) within the

pragmatist response to disagreement does not mean that the pragmatist response to disagreement is

therefore an end-state epistemology. To illustrate why this is the case, consider the following passage

from Dewey:

> [In transition theory (what Dewey calls the "reconstruction" of philosophy)], the process of growth, of improvement and progress, rather than the static outcome and result, becomes the significant thing. Not health as an end fixed once and for all, but the needed improvement in health – a continual process – is the end and good. The end is no longer a terminus or limit to be reached. It is the active process of transforming the existent situation. Not perfection as a final goal, but the ever-enduring process of perfecting, maturing, refining is the aim in living. Honesty, industry, temperance, justice, like health, wealth and learning, are not goods to be possessed as they would be if they expressed fixed ends to be attained. They are directions of change in the quality of experience. Growth itself is the only moral "end."[23]

---

within our actual (i.e. non-ideal) circumstances, in order to help us understand what we should consider as good epistemic

practice in our given situation.

[23] John Dewey, *Reconstruction in Philosophy* (New York: Henry Holt, 1920), 176–77.

Here, Dewey provides a clear explanation of the difference between transition and end-state theory. For Dewey, what makes transition theory different to end-state theory is its focus on the process of how one improves in social and epistemic practice, rather than on the explication of a particular fixed end (e.g. knowledge, truth, justice, etc.). Dewey's example of health here is an apt one: notice that the concept of health is one which encompasses such a wide variety of issues (e.g. nutritional, mental, immune) that it is not possible to define health as a single fixed end. Furthermore, our pursuit of health within our everyday life is not like the pursuit of a single concrete and realisable goal, but a continual process of maintaining and improving our health.

Of course, at this point, one notes that it is certainly possible to provide specific accounts of what it looks like to be "healthy" within particular domains (e.g. what "mental health" is), just as it is possible to provide specific accounts of what it looks like to be "rational" or "just" (e.g. "logical rationality" or "distributive justice"). This is a possibility which Dewey's focus on transition theory often forgets: on a similar note, Hilary and Ruth-Anna Putnam have criticised Dewey's ameliorative method for obscuring the role and need for conceptual clarification within scientific inquiry.[24] In this thesis, I leave open the possibility that future and further end-state theorising can serve to refine and clarify the account of disagreement I will outline. Nevertheless, to the extent that disagreement is an ongoing epistemic problem in our everyday lives – more like the recurring question of how to maintain one's health than the static question of how to cure a particular disease – Dewey's focus on the process of one's epistemic practice seems to be the best way to address the non-ideal epistemic problem which disagreement poses in our everyday life. What exactly is the non-ideal epistemic problem which disagreement poses? This is the question which we consider in the next section.

―――――――――――――――――――

[24] Putnam and Putnam, "Dewey's Logic," 209.

1.2.2   On the Non-Ideal Epistemic Problem of Disagreement

In this section, I shall highlight how the problem of disagreement is a distinctly non-ideal epistemic problem by comparing the problem of disagreement with the kind of problems which epistemology has typically seen as epistemic challenges for our beliefs: namely, skeptical hypotheses. Throughout the history of epistemology, the justifications for our everyday beliefs have been called into question through the raising of various skeptical arguments, such as Hume's problem of induction and Cartesian skepticism. Crucial to these arguments is the use of a particular class of counter-examples which undermine the legitimacy of certain sources of knowledge and justification such as perception, memory, or reason.[25] For example, in the case of Cartesian skepticism, Descartes famously calls sense perception into 'doubt' in the first *Meditation* by considering the example of dreaming. Given that our experience of dreaming are almost perceptually identical to our experiences of the world when awake, but which are in fact not true experiences of the world, the perceptual experience of dreams

---

[25] As a brief sidenote, one question which I do not discuss is whether skeptical hypotheses are supposed to undermine perception, reason, memory, and the like as sources of knowledge or sources of justification. This is something which also has not been distinguished within the current literature: in the SEP entry on "Epistemology", these sources are described as source of "knowledge and justification", and Robert Audi even goes on to call these sources of "rationality". Furthermore, addressing this question might be important for understanding the implications of skeptical hypotheses in greater detail, given that there is a subtle difference between the claim that skeptical hypotheses show that we can and do not acquire knowledge from sources like perception and induction, and the claim that skeptical hypotheses show that we have no way of justifying our use of these sources in forming our beliefs. However, given that this question is outside the scope of this essay, I shall take skeptical hypotheses undermine these sources as both sources of knowledge and justification. See Matthias Steup and Ram Neta, "Epistemology," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Fall 2020 (Metaphysics Research Lab, Stanford University, 2020); Robert Audi, "Theoretical Rationality: Its Sources, Structures, and Scope," in *The Oxford Handbook of Rationality*, ed. Alfred R. Mele and Piers Rawling (Oxford: Oxford University Press, 2004), 17–44.

undermines our use of perception to justify our beliefs, since we cannot justify why a given perceptual experience is an experience of the external world rather than an experience of a dream.[26]

A similar argument is made in Hume's problem of induction. For Hume, our use of inductive inference is called into question once we consider the contingent possibility of the future being characteristically different to one's past experiences: a famous illustration of this possibility is given by Bertrand Russell with the example of the chicken who believes that they will be fed today as they were yesterday, only to be killed and put on the dining table.[27] For Hume and Russell, the fact that the future can be characteristically different from the past therefore undermines our use of inductive inference to justify our beliefs, since one cannot justify why a given inductive inference wouldn't fail to give us a false belief about the future. As such, in providing cases where our use of perception and inductive inference did not, in fact, give us knowledge or justification, the skeptic therefore undercuts our ability and entitlement to use perception and induction as sources of knowledge and justification.

---

[26] For an excellent philosophical exposition on Descartes' dream argument, see Barry Stroud, *The Significance of Philosophical Scepticism* (Oxford University Press, 1984), 11ff.

[27] Of course, this is merely one half of Hume's argument against the justifiability of inductive inference. More specifically, Hume's problem of induction states that there is no non-circular justification for inductive inference through posing a dilemma against the justification of inductive inference. The dilemma functions as follows: inductive inferences presuppose the Uniformity Principle (that is, the idea that natural phenomena is uniform such that what will happen in the future must resemble what will happen in the past), which must be justified either by deductive argument or inductive reasoning. We cannot justify inductive inference by deductive argument, because no contradiction follows from negating the Uniformity Principle; and we cannot justify inductive inference by inductive reasoning, because this violates the principle of non-circular reasoning. Therefore, there is no non-circular justification for inductive inference. Here, Russell's example of the chicken is a counter-example which shows how the Uniformity Principle cannot be justified by deductive argument, given the contingent possibility that the future does not resemble the past. Bertrand Russell, *The Problems of Philosophy* (London: Williams & Norgate, 1912), 97–98.

However, it is here that we come to the 'two minds' which epistemologists now have towards skeptical arguments. Specifically, while these skeptical arguments are still seen as theoretically important for examining the nature of knowledge and justification, most contemporary epistemologists conclude that skeptical hypotheses are of little to no practical consequences for the justifications of our *actual* beliefs.[28] This is because while skeptical hypotheses – such as being in a dream, or being deluded by a mischievous demon, or being a contestant in something like The Truman Show – may well be theoretically irrefutable, this does not stop us from relying on our perception, memory, and inductive reasoning in everyday life. Let us consider two reasons for why our reliance on these sources is unproblematic when it comes to our everyday epistemic practice. Take the example of perception. First, while one's reliance on perception may lead one to make a mistake occasionally, such mistakes are rare occurrences which are often quite insignificant in their consequence.[29] For example, while one's perception might be faulty at times (e.g. one sees a bent straw inside a glass of water which is

---

[28] The practical irrelevance of skeptical arguments to our everyday epistemic practices as opposed to in philosophical reflection is famously noted by David Hume towards the end of *Enquiry Concerning Human Understanding*. Hume writes: "For as, in common life, we reason every moment concerning fact and existence, and cannot possibly subsist, without continually employing this species of argument, any popular objections, derived from thence, must be insufficient to destroy that evidence. The great subverter of Pyrrhonism or the excessive principles of scepticism, is action, and employment, and the occupations of common life. These principles may flourish and triumph in the schools; where it is, indeed, difficult, if not impossible, to refute them. But as soon as they leave the shade, and by the presence of the real objects, which actuate our passions and sentiments, are put in opposition to the more powerful principles of our nature, they vanish like smoke, and leave the most determined sceptic in the same condition as other mortals." David Hume, *An Enquiry Concerning Human Understanding*, ed. Tom L. Beauchamp (Oxford/New York: Oxford University Press, 1999), §10.21. For other epistemologists who follow this Humean line of thought, see Stroud, *The Significance of Philosophical Scepticism*; and P. F. Strawson, *Scepticism and Naturalism: Some Varieties* (Routledge, 1987).

[29] Of course, this is not to say that there may not be serious consequences to perceptual failings, such as the unfair incarceration of individuals due to faulty eyewitness testimony. Instead, it is to say that such consequences are insufficient to rule out our use of perception: the possibility of faulty eyewitness testimony does not lead us to ban the use of eyewitness testimony, but to attempt to find ways to discern and mitigate unreliable eyewitness testimony through better legal procedure.

actually straight), it is never unreliable to the point where one will stop trusting perception at all when making everyday decisions (e.g. one decides not to drive to Queensland, lest their perception deceives them and leads them to Perth). Second, even in cases where one's reliance on perception leads to incorrect belief, we can often learn from these cases of perceptual failing so as to avoid incorrect belief in future epistemic practice. To take the example of seeing a bent straw in a glass of water which is actually straight, we can learn from this experience by learning how the refraction of light can distort our visual perception objects, such that we can account for refraction in future epistemic practice (i.e. by verifying whether future straws are bent through touch). As such, given the rarity of perceptual failings, as well as our ability to explain and account for perceptual failings in our epistemic practice, the mere possibility of perceptual failings does not proscribe our use of perception as a source of knowledge and justification in everyday life. A similar case can be made for our reliance on other sources, such as memory and of inductive reasoning.

However, it is here where the case of disagreement differs dramatically from the skeptical hypotheses which we have discussed up until now. One way of explaining the epistemic problem of disagreement is to see disagreement as a challenge to both our use of reason and of testimony in the justification for our beliefs. To begin, it is important to note that in addition to the traditional sources of perception, reason, and memory, our everyday epistemic practice also clearly relies on the use of testimony to justify our beliefs. For instance, given that we do not have the ability to conduct research into health science ourselves, nor do we have the time and resources to research the history of each political candidate and party before an election, our epistemic practice is such that we inevitably rely on the testimony of search engines, news companies and social media in order to form our beliefs on health or politics.

However, the reason why disagreement challenges our use of testimony is because different sources of testimony might disagree about a given issue, such that we are forced to choose between these

conflicting sources of testimony: if one scientific study shows that coffee is good for you, and another scientific study shows that coffee is bad for you, it is obvious that you cannot accept the testimony of both studies without being logically inconsistent. Furthermore, the challenge which disagreement poses for our use of testimony can also come to challenge our use of reason to justify our beliefs as well: suppose you are a mathematician who has come up with a complex proof about a mathematical question which has remained unsolved within the literature for years.[30] After publishing your proof, you come to learn that your proof has had mixed reception by your colleagues such that the mathematical community is divided between those who believe your proof to be faulty, and those who believe your proof to be correct. Here, the disagreement between you and your critics forces you to make a decision about what to believe. On the one hand, your critics are also adept mathematicians who you sometimes go to for counsel and for proof-checking, such that it would be reasonable to rely on your friend's answer as testimonial evidence that you have come to the wrong answer. On the other hand, you are feeling confident in the steps of reasoning that you took to arrive at your answer, such that your own reasoning provides strong justification that your answer is correct. Disagreement therefore poses a similar challenge in this case to that between conflicting sources of testimony: in this case, one has to choose between trusting in one's own reasoning, or trusting in the testimony of another's reasoning.[31] As such, disagreement poses an epistemic challenge against our use of reason and testimony as justification for our everyday beliefs.

---

[30] This example is inspired by the real-life example of Shinichi Mochizuki, whose proof for the *abc* conjecture spanned over 500 pages and was too difficult for most number theorists to follow. Given that there have also been criticisms of Mochizuki's proof by some number theorists (which Mochizuki and his advocates have claimed do not affect the proof in any substantial way), Mochizuki's proof currently holds a tenuous place within the literature on number theory. For more information, see Davide Castelvecchi, "Mathematical Proof That Rocked Number Theory Will Be Published," *Nature* 580 (2020): 177.

[31] At this point, one might argue that the two dilemmas which disagreement pose are in fact two different kinds of epistemic problems which ought to be kept separate. This is the prevailing stance, for instance, in the literature on analytic social

Furthermore, unlike the skeptical arguments against perception and inductive reasoning, the case of disagreement clearly has immediate practical consequences on the justification for our everyday beliefs. To explain, consider how the epistemic problem which disagreement poses differs from that of perceptual failings, such that the two reasons which we gave to justify our use of perception despite its limitations do not apply in the case of disagreement. First, disagreements are not only a common occurrence in everyday life, but one where choosing the wrong source of testimony can come with serious real-world consequences. As we have seen in the opening sketch of this chapter, the reason why we have a general skeptical attitude to the information we receive online is because of the very likely possibility that we will receive false if not intentionally deceptive information: one is much more likely to be deceived by a charismatic speaker or an interesting news source in epistemic practice, than to be deceived by instances of false barns or malicious demons. Furthermore, unlike the case of perceptual failings, disagreements among different sources of testimony is such that trusting the wrong source may come with grave epistemic risk – historical examples such as Nazi Germany, the Spanish

---

epistemology: whereas the question of how to decide between conflicting sources of testimony is taken as a problem for the epistemology of testimony (where it is known as the "expert" problem), the question of how to decide between one's own reasoning and that of another is taken as the main problem in the epistemology of disagreement. By the end of this thesis, I aim to show that the two are the same kind of problem. Specifically, I take both problems to be, in essence, a problem about the epistemology of trust, such that the question of when one is rational in their trust of a source of testimony is the same kind of question as the question of when one is rational in their trust of their own reasoning as opposed to trust in others. While this point is not widely discussed within the literature, one person who indirectly makes an argument along these lines is Richard Foley: throughout his publications, Foley has given a consistent line of argument on the role of trust in oneself and others across the disparate issues within the epistemology of trust, testimony, and disagreement. For more on the "expert" problem, see Alvin I. Goldman, "Experts: Which Ones Should You Trust?" in *Social Epistemology: Essential Readings*, ed. Alvin I. Goldman and Dennis Whitcomb (Oxford University Press, 2011), 109–33. For Foley's work on the role of trust in our epistemic practice, see Richard Foley, "Egoism in Epistemology," in *Socializing Epistemology*, ed. Frederick F. Schmitt (Lanham, MD: Rowman; Littlefield, 1994); Richard Foley, *Intellectual Trust in Oneself and Others* (Cambridge: Cambridge University Press, 2001); Richard Foley, "Self-Trust," in *The Routledge Handbook of Trust and Philosophy*, ed. Judith Simon, First (New York; London: Routledge/Taylor; Francis Group, 2020), 231–42.

Inquisition, or cases like bloodletting and lobotomy in the history of medical science show how trusting in a source of testimony (of a political, religious, or scientific authority) may lead one to incorrect beliefs which have serious ramifications to oneself and to society. And such cases still occur when one considers how certain sources of testimony in contemporary society may also lead one to entrenched views that are both harmful and inescapable – cases such as religious cults, extremist political parties, or outspoken media personalities.[32]

Second, as we have seen in the opening sketch of this chapter, the problem of disagreement is an ongoing problem which we have not found an adequate solution for in contemporary society. To take the case of inductive inference, the reason why Hume's problem of induction does not pose a theoretical objection to contemporary research into statistics and machine learning is simply because the kinds of inductive inference which these fields employ have advanced both conceptually and in method.[33] This means that while it is always possible for one to make an invalid inductive inference in their everyday epistemic practice, the question of how we can evaluate or ensure the reliability of an inductive inference is now a question which we do have substantial answers to.[34] However, this is not

―――――――――――――――――――

[32] For an illuminating account of how certain sources of testimony can leave one entrenched within an epistemic "echo chamber", see C. Thi Nguyen, "Echo Chambers and Epistemic Bubbles," *Episteme*, 2018, 1–21.

[33] For example, Gilbert Harman and Sanjeev Kulkarni have proposed a solution to Hume's problem of induction through the use of statistical machine learning theory. Harman and Kulkarni's book is not only a plausible way of addressing Hume's problem of induction, but a clear example of how contemporary statistics and computer science employ notions and methods which have to some extent superseded the traditional categories of inductive inference. See Gilbert Harman and Sanjeev Kulkarni, *Reliable Reasoning: Induction and Statistical Learning Theory*, ed. François Recanati, The Jean Nicod Lectures (Cambridge, Massachusetts; London, England: MIT Press, 2007).

[34] A similar case can be made for simple perceptual failings: consider how the advancements in technology which enhance our senses (e.g. cameras, headphones, etc.) have made the simple perceptual cases which interested Pyrrhonian skeptics largely obsolete. But what about stronger skeptical hypotheses, such as Descartes' evil demon, or a brain-in-a-vat scenario? These skeptical hypotheses are largely orthogonal to the point that I am making here, because my current focus is on skeptical scenarios which question the reliability of a source of knowledge and justification, as opposed to these skeptical hypotheses

the case for the problem of disagreement: despite all the advancements in the transmission of

information in our digital age, we have not yet been able to address the problem of how to discern

between reliable and unreliable information, nor of how we should resolve if not reasonably respond

to the individual and social disagreements we must face in everyday life.[35] As such, given the regular

occurrence of disagreement in our everyday epistemic practice, the severity of the consequences of

incorrectly relying on certain sources of testimony (and of incorrectly relying on one's own cognitive

abilities), and the lack of an adequate explanation or solution for how we should respond to

disagreement and unreliable testimony, the epistemic challenge which disagreement poses is clearly of

a different and more serious kind than that of traditional skeptical hypotheses.

Once we see how the epistemic challenge which disagreement poses is different in kind to those posed

by skeptical hypotheses, the following two questions arise: first, how can we account for the differences

between the problem of disagreement compared to that of skeptical hypotheses; and second, what kind

of epistemology is needed to address the practical challenge which disagreement poses in our everyday

epistemic practice? Here, I shall argue that the distinction between ideal and non-ideal epistemology is

central to answering both questions. First, what is it that makes the problem of disagreement so

different from that of skeptical hypotheses? Given our distinction between ideal and non-ideal

---

which preclude the possibility of knowledge in general. In these stronger skeptical hypotheses, the quote from Hume in fn. 28
seems to apply all the more: our best response to these kinds of skepticism is not further theoretical speculation, but pointing
to the demands of everyday life which call for action that cannot be enacted without a rejection of stronger skeptical
hypotheses. It is at least possible, however, in theory for the kind of skeptical hypotheses that we are looking at to be
addressed and resolved through reason.

[35] That being said, there is excellent progress being made on how to respond to unreliable testimony within the recent
literature on the epistemology of trust. More on this at the end of this section – see also the citations in fn. 37.

epistemology, one way of explaining the difference is to see skeptical arguments as arguments against *ideal* theories of knowledge and justification.

To see why this is the case, consider a scenario in which we took the claim that perception and induction are "sources" of knowledge and justification to be outlining a realistic theory of knowledge and justification, such that we claim that perception and induction are only mediums for knowledge and justification when certain empirical constraints are met. Suppose, for instance, we defined perception and induction with reference to certain empirical facts, such that we defined perception as a sense modality which only gives knowledge when certain physical conditions are met (e.g. conditions about environmental light and sound or about our biological organs), and we defined induction as a form of inference that is only valid when the certain empirical conditions about the relation from the sample size to the general population is met. If we claim that perception and induction are sources of knowledge and evidence under these definitions, then neither Descartes' example of dreams nor Russell's case of the chicken serve as counter-examples to these claims: that is, the case of dreaming does not refute the claim that perception is a source of knowledge because dreaming does not satisfy the physical conditions needed for perception to give knowledge or justification, and Russell's case of the chicken does not refute the claim that inductive inference gives knowledge because Russell's case is not an example where the conditions for a valid inductive inference are met (e.g. because the sample size should not be "all chickens", but "all chickens which reach a certain age and physique").

Of course, at this point, one might argue that to define perception and induction in this way would be to misunderstand the epistemic import of skeptical hypotheses and arguments. After all, the purpose of skeptical arguments is not to make an empirical claim about how perception and induction is or is not to be used within everyday epistemic practice, but to make a theoretical claim about the sources

from which we receive knowledge and justification.[36] In other words, the epistemic import of skeptical arguments comes from how skeptical hypotheses show that we cannot justify the claim that perception and induction are sources of knowledge and justification, because we cannot deny that it is possible for perception and induction to fail to provide knowledge and justification. But this is precisely my point: it is only because one desires a certain kind of theory about knowledge and justification that makes skeptical hypotheses important for that kind of epistemic theorising. In particular, it is because we want to claim that perception and induction are sources of knowledge *simpliciter* – that is, that they are sufficient conditions for knowledge and justification irrespective of any additional empirical constraints – that lead skeptical hypotheses to refute this claim. Otherwise, the perceptual experience of dreams would simply be seen as an empirical case in which perception did not give knowledge, rather than a theoretical case against perception *being* a source of knowledge. As such, given the ideal/non-ideal distinction, we can see that the skeptical hypotheses that are traditionally discussed within epistemology are only effective in undermining ideal theories of knowledge and justification: that is, theories which see knowledge and justification (as well as the sources themselves) as epistemic values that are independent of empirical constraints (following utopianism), and as epistemic ends which do not require any reference to a given context (i.e. following end-state theory).

The ideal/non-ideal distinction also helps to explain why skeptical hypotheses have little to no practical consequences for our everyday epistemic practice, and why the case of disagreement clearly does. If

---

[36] This is why the upshot of Hume's problem of induction is not that inductive inference cannot be used, but that there is no non-circular justification for inductive inference. In a similar way, skeptical arguments against perception do not show that perception should not be used in everyday epistemic practice, but that there is no non-circular justification for "perceptual knowledge". For more information about the argument against a non-circular justification of perceptual knowledge, see William P. Alston, "Perceptual Knowledge," in *The Blackwell Guide to Epistemology*, ed. John Greco and Ernest Sosa (Oxford, UK: Blackwell, 1999), 221–42; and Chapters 10 and 11 of Jonathan Dancy, *An Introduction to Contemporary Epistemology* (Oxford: Blackwell, 1985).

what we are looking for in an epistemic theory is an empirical and context-sensitive account of

knowledge and justification, one which explains how we can and do receive justification from certain

sources in certain contexts, then we can give epistemic accounts of perception and induction which

adequately explain how to use perception and induction in practice, why perception and induction

may fail in certain contexts, and how to identify these failures in epistemic practice and adjust

accordingly. But this is the kind of understanding that we lack when it comes to disagreement and

testimony: the reason why the problem outlined in the opening sketch of this chapter seems so

important and urgent for our current age is because we lack any clear way of understanding why we

can fail to reach disagreement with a dissenting party at a given moment, or how we can find reliable

information from sources of testimony about a topic which we are unfamiliar with. To address the

problem of disagreement would therefore require a non-ideal epistemology of disagreement: one

which is sufficiently realistic in its inclusion of empirical constraints so as to explain how epistemic

agents have a *feasible* chance of responding rationally to a given disagreement, and transitional insofar

as it explains how an epistemic agent is to work towards more prudent reliance of testimony when

justifying their beliefs, and better methods of resolving disagreements over matters of fact in their

given context.

Of course, at this point, more still needs to be said about what kind of theory would satisfy this goal of

a non-ideal epistemology of disagreement. Is the kind of theory we are looking for a theoretical

explication of disagreement: one which illuminates disagreement as a epistemic relation between

epistemic agents and some matter of fact? What is the underlying cause of disagreement which we are

trying to account for: is disagreement caused by the different bodies of evidence that each party has, by

different inferential rules which each party employ, by the differences in value (including epistemic

value) that each party holds, or the influence of non-epistemic factors (e.g. "pragmatic" factors such as

time and energy, or non-epistemic factors such as desires and emotions) on one's beliefs? What kind of

epistemic norms would we need to account for the rational response to disagreement: should we focus on epistemic norms which explain how individual epistemic agents ought to revise their belief when faced with disagreement, or should we focus on social norms which guide how an epistemic community ought to engage in and resolve disagreement? These are the kinds of questions which will be answered over the course of this thesis. To conclude this section, however, two remarks are needed to clarify the scope of the non-ideal epistemology of disagreement given in this thesis.

First, while I have discussed the close connection between disagreement and testimony within our epistemic practice, this thesis is primarily interested in giving an epistemic account of disagreement, one which explains how we as epistemic agents can come to a rational resolution (if not, a rational conclusion) to disagreement within our epistemic practice. The main reason for focusing on the topic of disagreement rather than testimony is because, within the literature, much more has been done towards outlining a non-ideal epistemology of testimony as opposed to the non-ideal epistemology of disagreement. In particular, due to the incredibly insightful literature within the feminist epistemology of trust, there is now a substantial answer to the question of how we should evaluate and discern unreliable information and testimony within a given context: an answer that comprises a series of epistemic norms surrounding trust and trustworthiness that describe how epistemic agents should employ trust when receiving testimony, and communicate trustworthiness when giving testimony.[37]

---

[37] For the classic papers which started the literature, see the following: John Hardwig, "Epistemic Dependence," *Journal of Philosophy* 82, no. 7 (1985): 335–49; John Hardwig, "The Role of Trust in Knowledge," *The Journal of Philosophy* 88, no. 12 (1991): 693–708; Judith Baker, "Trust and Rationality," *Pacific Philosophical Quarterly* 68, no. 1 (1987): 1–13; Annette C. Baier, "Trust and Antitrust," *Ethics* 96, no. 2 (1986): 231–60; Onora O'Neill, *A Question of Trust* (Cambridge: Cambridge University Press, 2002); Onora O'Neill, *Autonomy and Trust in Bioethics* (Cambridge: Cambridge University Press, 2002). For more on the recent literature, see the following: Pamela Hieronymi, "The Reasons of Trust," *Australasian Journal of Philosophy* 86, no. 2 (2008): 213–36; Karen Jones, "Trustworthiness," *Ethics* 123, no. 1 (2012): 61–85; Karen Jones, "The Politics of Intellectual Self-Trust," *Social Epistemology* 26, no. 2 (2012): 237–51; Frost-Arnold, "Trustworthiness and Truth"; Frost-Arnold, "Social Media, Trust, and the Epistemology of Prejudice"; Karen Jones, "'But i Was Counting on You!'," in *The*

By contrast, however, the question of how we can come to rationally respond if not resolve the prominent disagreements we face in our everyday practice is one which has not been addressed to the same degree, and therefore in need of more examination. To answer this question will require an independent non-ideal epistemic account of disagreement, one whose existence is justified in its own right: that is, even if we have a clear answer as to how to discern and evaluate the reliability of testimonial sources, there remains a question about what we should do when reliable sources disagree not in information but in opinion. As such, while the role of trust and testimony will be prevalent throughout the epistemic account given within this thesis, the primary concern of this thesis will be to provide an epistemic account of disagreement.

Second, the method by which I arrive at a general non-ideal epistemology of disagreement in this thesis is through an ameliorative strategy. In particular, my argument for a pragmatist epistemology of disagreement is based on a critique of the non-ideal epistemology of disagreement proposed within analytic social epistemology (what I shall call the "evidentialist response to disagreement"), followed by an argument for how the inadequacies of the evidentialist response to disagreement are addressed and resolved by the pragmatist theory of inquiry (culminating in what I call the "pragmatist response to disagreement"). As a result, the non-ideal epistemology of disagreement which I develop throughout this thesis is a synthesis of the ideas about disagreement from analytic social epistemology and pragmatist social epistemology. To consider philosophical accounts about the epistemology of disagreement which reside outside of these two bodies of literature is outside scope of this thesis. As such, I do not examine the topic of disagreement as it is explored within the question of scientific

---

Philosophy of Trust, ed. Paul Faulkner and Thomas W. Simpson (Oxford: Oxford University Press, 2017), 161–76; Naomi Scheman, "Trust and Trustworthiness," in *The Routledge Handbook of Trust and Philosophy*, ed. Judith Simon, First (New York; London: Routledge/Taylor; Francis Group, 2020), 28–40.

consensus within critical social epistemology, nor do I consider the topic of "rational disagreement" as it is discussed within continental political thought.[38]

This brings us to the next and final section of this introductory chapter. In this last section, I will clarify what I mean by 'pragmatist epistemology', so as to lay the final theoretical foundations for the argument of this thesis. Specifically, I will explain how the pragmatist tradition provides a distinct epistemic theory in the theory of inquiry, one whose starting assumptions allow the theory of inquiry to be particularly conducive to theorising about non-ideal epistemic problems.

1.3   What is Pragmatist Epistemology?

1.3.1   On the General Confusion about "Pragmatist Epistemology"

Within the current philosophical discourse, the term 'pragmatist epistemology' is not widely used and one which is likely to breed misunderstanding without further clarification. As an introduction to this

---

[38] The Routledge Handbook of Social Epistemology distinguishes between two seperate accounts of social epistemology within the recent philosophical literature: that is, between the "critical social epistemology" which is based off the work of Steve Fuller within Science and Technology Studies, and "analytic social epistemology" which is based off the work of Alvin Goldman in analytic epistemology. Meanwhile, the interest in the rationality of disagreement comes from Jacques Rancière, and his seminal work on disagreement. For more on the difference between critical social epistemology and analytic social epistemology, see David Henderson, "On the Background of Social Epistemology," in *The Routledge Handbook of Social Epistemology*, ed. Miranda Fricker et al. (New York; London: Routledge/Taylor; Francis Group, 2020), 3–9; Finn Collin, "The Twin Roots and Branches of Social Epistemology," in *The Routledge Handbook of Social Epistemology*, ed. Miranda Fricker et al. (New York; London: Routledge/Taylor; Francis Group, 2020), 21–30. For an overview of analytic social epistemology from Alvin Goldman himself, see Alvin I. Goldman, "The What, Why, and How of Social Epistemology," in *The Routledge Handbook of Social Epistemology*, ed. Miranda Fricker et al. (New York; London: Routledge/Taylor; Francis Group, 2020), 10–20. For more on the topic of rational disagreement within continental political thought, see Chapter 3 of Jacques Rancière, *Disagreement: Politics and Philosophy*, trans. Julie Rose (Minneapolis; London: University of Minnesota Press, 1999). It is entitled "The Rationality of Disagreement". See also Axel Honneth and Jacques Rancière, *Recognition or Disagreement: A Critical Encounter on the Politics of Freedom, Equality, and Identity*, ed. Katia Genel and Jean-Philippe Deranty (New York: Columbia University Press, 2016). For more on Fuller on scientific consensus, see Steve Fuller, *Social Epistemology*, 2nd ed. (Bloomington: Indiana University Press, 2002), 208–10.

section then, let us begin by examining two factors which contribute to the possible misunderstanding about pragmatist epistemology which I outline in this thesis: namely, the lack of research and scholarship on American pragmatism (and in particular, pragmatist thought about epistemic issues) within contemporary epistemology, and the general lack of consensus about philosophical method within contemporary pragmatism. These two factors have greatly contributed to the confusion and disuse of the term 'pragmatist epistemology' within the current philosophical literature.

Let us examine these two factors in greater detail. Within contemporary epistemology, the term 'pragmatist epistemology' is likely to be misunderstood because of two reasons: first, because of a general unfamiliarity with the pragmatist literature about epistemic issues; and second, how this has led the current use of 'pragmatism' within contemporary epistemology to have little to no connection with American pragmatism. To begin, it is important to note that my claim about the general unfamiliarity with pragmatist thought within contemporary epistemology does not refer to an unfamiliarity with pragmatist thinkers, but with an unfamiliarity with pragmatism as a wider philosophical movement. In terms of the work of individual pragmatists, there are many well-known pragmatists whose individual epistemic theories and arguments are still discussed today. These include (but are not limited to) Charles Sanders Peirce's work on abductive inference and dispositionalism about belief; William James on doxastic voluntarism; W. V. O. Quine's epistemological holism and naturalism; and Susan Haack on foundherentism about epistemic justification.[39]

Nevertheless, when considering pragmatism as a general philosophical movement, it is clear that no single epistemic theory or method comes to mind: compare, for instance, how the verificationist

---

[39] A survey of many of these ideas can be found in Robert Talisse and Scott Aikin's definitive reader on pragmatism. See Robert B. Talisse and Scott F. Aikin, eds., *The Pragmatism Reader: From Peirce Through the Present* (Princeton; Oxford: Princeton University Press, 2011).

theory of meaning is seen as a view held by the logical positivists, or the empiricist theory of sense-data

is seen as a view held by early modern empiricists.[40] At best, one might consider the pragmatic maxim

– the method originating from Peirce to evaluate concepts in terms of their effects on our practice –

and its various applications to the concept of truth as a distinctly pragmatist contribution to

epistemology: but even here, both the pragmatist maxim and the pragmatic theories of truth are not

examined in any great detail within the current literature, and prone to misinterpretation without a

greater understanding of pragmatist thought more generally.[41]

The unfamiliarity which contemporary epistemology has with pragmatist thought is especially

pertinent for the purposes of this thesis, because the current literature on analytic epistemology has

since begun to use the term 'pragmatism' to refer to epistemic theories which have little to nothing to

do with the American pragmatist tradition.[42] For example, consider how the term 'pragmatist' is used

---

[40] A rare exception to this trend is the recent literature on a pragmatist epistemology of democracy: the literature which is the focus of my pragmatist epistemology of disagreement in Chapter 5 of this thesis (see also Appendix A). However, it is important to note that this literature on the pragmatist epistemology of democracy is still not widely recognised within mainstream social epistemology because of the recency of the literature and because of the classification of this literature sometimes as political philosophy or democratic theory, rather than as applied social epistemology. For a list of citations on this literature, see fn. 278.

[41] For example, Cheryl Misak has argued extensively about the misunderstanding of Peirce's account of truth, a misunderstanding which Misak argues is based on a misreading of the pragmatic maxim. In particular, Misak argues that the pragmatic maxim is often taken to be a method of *defining* concepts in terms of their consequences on one's interaction with the world, rather than as a method of giving a "pragmatic elucidation" of concepts in terms of their function in doxastic practice. Accordingly, Peirce's claim that truth is "the end of inquiry" is taken as an anti-realist theory about truth, when it is really an account of how the concept of truth functions in our doxastic practice: that is, in the way the concept of truth plays into how we work towards and make expectations about true propositions in ways which we do not for false propositions. For more information, see Cheryl Misak, *Truth and the End of Inquiry: A Peircean Account of Truth* (Oxford: Clarendon Press, 1991); and Cheryl Misak, *Truth, Politics, Morality: Pragmatism and Deliberation* (New York: Routledge, 2000), especially Chapter 2.

[42] One exception to this trend is the recent literature on pragmatist theories of belief: a series of accounts which argue that beliefs should be explained in terms of their affects on one's behaviour, rather than in terms of one's intellectual

to refer to proponents of the "pragmatic encroachment" thesis: that is, the claim that differences in one's pragmatic circumstances can lead to a difference in the epistemic status of one's belief (e.g. as "justified" belief or as "knowledge"). To illustrate the pragmatic encroachment thesis by example, suppose an undergraduate student wrote an excellent essay about the historical circumstances which led to the First World War, such that the marker of this essay commented that this student "knows" the various historical facts about First World War. Yet, if this student requested to be a lecturer on this subject, the difference in pragmatic circumstances (i.e. the level of expertise needed to be a lecturer due to the higher pay and position, as well as influence on the students) may lead the same marker to reject the student's request because they do not "know" the history of the First World War to teach the subject. In cases like this, proponents of the pragmatic encroachment thesis (also known as "epistemological pragmatists") claim that the epistemic status of the student's belief does change because of the difference in pragmatic circumstances, while critics of the pragmatic encroachment thesis (also known as "epistemological purists") claim that the epistemic status of the student's belief does not.

Here, it is clear that the debate around the pragmatic encroachment thesis has little to nothing to do with American pragmatism: while proponents of the pragmatic encroachment thesis do sometimes refer to pragmatist ideas such as fallibilism, their arguments for these ideas are not based on the arguments and thinkers within the pragmatist tradition, but on the implications of epistemic

---

endorsements. While these theories differ from the pragmatist tradition slightly in terms of their attempts to provide metaphysical definitions of the concept of belief, these theories have direct precedence and explicit reference to pragmatist thought and thinkers. See Aaron Z. Zimmerman, *Belief: A Pragmatic Picture* (Oxford: Oxford University Press, 2018); and Eric Schwitzgebel, "The Pragmatic Metaphysics of Belief," in *The Fragmented Mind*, ed. Cristina Borgoni, Dirk Kindermann, and Andrea Onofri (Oxford: Oxford University Press, 2021), 350–75. In his paper, Schwitzgebel also cites the following as an example of a pragmatist account of belief: see Krzysztof Poslajko, "The Lycan-Stich Argument and the Plasticity of "Belief"," *Erkenntnis* 87 (2022): 1257–73.

contextualism about knowledge.[43] The literature on pragmatic encroachment is therefore one of many

examples within contemporary epistemology where the term 'pragmatism' no longer refers to the

American pragmatist tradition, but to any view that considers the role of practical factors on epistemic

concepts: practical factors such as one's goals, interests, circumstances, or action.[44] While research into

particular pragmatist thinkers on various topics in epistemology remains strong, the current literature

on epistemology does not for the most part identify the term 'pragmatist epistemology' with any

distinct epistemic theory or view that is based within the American pragmatist tradition.

At the same time, the term 'pragmatist epistemology' is an ambiguous term within contemporary

pragmatism itself, given the substantial disagreements which pragmatists have over epistemology and

---

[43] Specifically, the recent discussion on pragmatic encroachment is influenced greatly by Jason Stanley's work on how practical interests might justify epistemic contextualism about knowledge. See Jason Stanley, *Knowledge and Practical Interests* (Oxford: Oxford University Press, 2005). For a discussion of fallibilism with respect to the topic of pragmatic encroachment, see Chapter 1 of Jeremy Fantl and Matthew McGrath, *Knowledge in an Uncertain World* (Oxford, England: Oxford University Press, 2009).

[44] Two more examples. First, Timothy Williamson uses the term 'pragmatism' to refer to the attempt to "operationalize epistemology by working only with concepts whose application is always accessible to the agent." An example of this kind of attempt given by Williamson is subjective Bayesianism. We will discuss Williamson, and this passage, more in §4.3 of this thesis. See also Timothy Williamson, *Knowledge and Its Limits* (Oxford: Oxford University Press, 2000), 180. The second example of this is the longstanding debate regarding the kinds of reasons one can have for a belief: whereas evidentialists argue that the only kind of reasons one can have for a belief are epistemic reasons (i.e. reasons about one's 'evidence'), pragmatists argue that there can be practical reasons for holding a belief. While it is true that this debate does have some connection to American pragmatism – given that the history of this debate is typically traced to the interactions between mathematician and philosopher William K. Clifford and American pragmatist and psychologist William James – it is important to note that neither Clifford nor James' original arguments are really discussed within the current literature, and all of the recent arguments for 'pragmatism' about belief have no connection to American pragmatism at all. For more information, see Andrew Reisner, "Weighing Pragmatic and Evidential Reasons for Belief," *Philosophical Studies* 138, no. 1 (2008): 17–27; Andrew Reisner, "The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem," *Philosophical Studies* 145, no. 2 (2009): 257–72; Susanna Rinard, "Believing for Practical Reasons," *Noûs* 53, no. 4 (2018): 763–84.

epistemic issues. Let us consider two ways in which the notion of pragmatist epistemology as a distinct branch of pragmatist thought is complicated by the disagreements pragmatists have over epistemic issues. First, the initial disagreements between the classical pragmatists over the philosophical aim and method of pragmatism highlight how, from the outset, pragmatism begins with a disagreement about how to evaluate epistemic issues. To take but one example, consider the disagreement between Peirce and James over how the pragmatic maxim is used: whereas Peirce used the pragmatic maxim as a way of clarifying conceptual confusion (i.e. as a method for "how to make our ideas clear"), James argued that the pragmatic maxim should be used to consider the effects of believing a proposition on one's life.[45] This in turn reflected on how Peirce and James differed in their evaluation on the justifications for belief: to take the famous example of the Roman Catholic doctrine of transubstantiation, Peirce argued that the pragmatic maxim reveals how the claim that Christ's blood is empirically indistinguishable from wine shows how the doctrine of transubstantiation is meaningless and therefore unjustified as a belief. Conversely, James argued that the pragmatic maxim shows how the psychological effects of the doctrine of transubstantiation on a practicing Catholic was crucial to their faith and life, and therefore justified as a belief.

As such, the differences between Peirce and James over the pragmatic maxim is but one example of the starting differences which pragmatism has about aim and method, differences which clearly do have an impact in the kind of epistemic evaluation that pragmatism gives. These problems regarding the aim and method of pragmatism are only compounded once we introduce Dewey's rejection of the pragmatic maxim in favour for an experimental method based on Darwinian science, as well as the

---

[45] See Talisse and Aikin, *The Pragmatism Reader*, 1–2.

various ways in which pragmatism has been developed by subsequent thinkers such as W. V. O. Quine and Richard Rorty.

This brings us to our second point: the very notion of a pragmatist epistemology has been greatly undermined within contemporary pragmatism because of the strong arguments made by Richard Rorty against the theoretical need and aim of epistemology as a philosophical project. To be fair, the argument against epistemology as a philosophical discipline is not unique to Rorty: amongst the classical pragmatists, Dewey criticised the starting assumptions of epistemology with his argument against the spectator theory of knowledge; and within twentieth century philosophy more generally, the discipline of epistemology was criticised by logical positivists such as Rudolf Carnap, who used the verificationist theory of meaning to claim that substantial epistemological theories are cognitively meaningless.[46]

Nevertheless, what made Rorty's arguments particularly incisive was the fact that Rorty's criticisms did not simply target the assumptions or method of a given epistemic theory, but the metaphilosophical reasons for why one would want an epistemic theory in the first place. In particular, in *Philosophy and the Mirror of Nature*, Rorty argued that the very notion of epistemology as a theoretical enterprise is both unnecessary and presumptuous: unnecessary, because the knowledge claims of scientific research and everyday life progress without any need for an epistemic account which explains what constitutes knowledge in these domains; and presumptuous, because to put forward a "theory" of knowledge is to

---

[46] Dewey argues against the spectator view of knowledge in John Dewey, *The Question for Certainty: A Study of the Relation of Knowledge and Action*, Gifford Lectures (London: G.Allen & Unwin, 1929). Meanwhile, Carnap claims that "all statements belonging to Metaphysics, regulative Ethics, and (metaphysical) Epistemology have this defect, are in fact unverifiable and, therefore, unscientific." See Rudolf Carnap, *The Unity of Science* (London: Kegan Paul, Trench, Hubner, 1934), 22. This quote of Carnap is taken from Hilary Putnam, *The Collapse of the Fact/Value Dichotomy and Other Essays* (Cambridge, Massachusetts; London, England: Harvard University Press, 2002), 18.

assume that one can adjudicate the knowledge claims of another better than they can.[47] While Rorty's

philosophical arguments within the book against the epistemological assumptions of Descartes, Locke

and Kant are no longer discussed in any great detail, Rorty's metaphilosophical arguments against the

the theoretical aim of epistemology – once again, of why we would want an epistemic theory in the

first place – remain influential within contemporary discussions within pragmatism and naturalism.[48]

This has led some pragmatists who accept Rorty's arguments to (unfortunately) hold epistemology in

general as a dubious philosophical project, and even the pragmatists who disagree with Rorty about the

possibility and importance of epistemology are now faced with the question how a pragmatist

epistemology should proceed.[49]

_____

[47] Consider the following quotes. First: "Philosophy as a discipline thus sees itself as the attempt to underwrite or debunk claims to knowledge made by science, morality, art, or religion. It purports to do this on the basis of its special understanding of the nature of knowledge and of mind…philosophy's central concern is to be a general theory of representation, a theory which will divide culture up into the areas which represent reality well, those which represent it less well, and those which do not represent it at all (despite their pretense of doing so)." And: "[Kant] made it possible for epistemology to be thought of as a foundational science, an armchair discipline capable of discovering the "formal"…characteristics of any area of human life [and which] enabled philosophy professors to see themselves as presiding over a tribunal of pure reason, able to determine whether other disciplines were staying within the legal limits set by the "structure" of their subject matters." Richard Rorty, *Philosophy and the Mirror of Nature* (Princeton University Press, 1979), 3, 139.

[48] Here is Cheryl Misak on the influence of Richard Rorty on pragmatism: "Richard Rorty has campaigned over the last two decades to explode an old philosophical picture of truth and objectivity and replace it with his version of pragmatism. In some quarters he has been so successful that the first task for any other kind of pragmatist is to wrest the label from him." And despite Misak's various criticisms of Rorty, Misak concedes that "many of Rorty's negative points are well within what I take to be the real spirit of pragmatism". Misak, *Truth, Politics, Morality*, 11. Meanwhile, Rorty's arguments have convinced naturalists such as Huw Price, about the incompatibility of naturalism with a representationalist account of language or belief on the basis of precisely these reasons about the transcendental claims of representationalism. See Huw Price, *Naturalism Without Mirrors* (Oxford: Oxford University Press, 2011).

[49] While many have criticised Rorty for his claims about truth and objectivity, few have responded to Rorty's specific claims about the theoretical desiderata of epistemology. However, for a convincing response to Rorty about the possibility of a

To summarise then, the general lack of scholarship and research into pragmatist thought within contemporary epistemology, and the general lack of consensus about philosophical method into epistemic issues within contemporary pragmatism, has led to a situation where there exists no clear pragmatist epistemology within the current philosophical discourse. This is an unfortunate turn of events, because as I shall now argue, the pragmatist tradition does contain the philosophical resources for a distinct view on epistemology: one which provides an alternative way of epistemic theorising which is particularly conducive towards evaluating non-ideal epistemic issues. This is the pragmatist theory of inquiry which I outline in the next section.

However, because of the disagreements about the method and aim of pragmatism which we have just seen, a few qualifications are needed for the kind of pragmatist epistemology which I outline in this thesis.[50] First, the pragmatist theory of inquiry that I am focusing on throughout this thesis is the theory of inquiry as it is outlined by John Dewey in his *Logic: The Theory of Inquiry* and *How We Think*, and not the theory of inquiry as it is outlined by C. S. Peirce.[51] The choice to focus on Dewey's theory of inquiry rather than that of Peirce is simply because of the compelling reasons and ideas which Dewey and subsequent Deweyan thinkers give about inquiry and disagreement, rather than any critique or complaints about Peirce.[52] Second, following Hilary Putnam and Ruth-Anna Putnam's

---

pragmatist epistemology, see David Macarthur, "A Kant-Inspired Vision of Pragmatism as Democratic Experimentalism," in *Pragmatism, Kant and Transcendental Philosophy*, ed. Gabrielle Gava and Robert Stern (Routledge, 2015), 67–84.

[50] In stating the assumptions I make about pragmatist epistemology outright, I hope to follow Robert Talisse's lead when he claims that "pragmatists have been in the business of trying to reach agreement about what pragmatism is." Robert B. Talisse, "Saving Pragmatist Democratic Theory (from Itself)," *Ethics & Politics* 12, no. 1 (2010): 12.

[51] See John Dewey, *How We Think* (Boston, New York, Chicago: D. C. Heath & Co., 1910); Dewey, *Logic*.

[52] In fact, an interesting counterpart to this entire thesis is the Peircean view of inquiry and democracy as it is given by Cheryl Misak and Robert Talisse. The Peircean view of inquiry and democracy which Misak and Talisse outlines is interesting because it comes to almost exactly the same conclusion as I do about disagreement in Chapter 5 (and especially in moral and

exposition of Dewey's theory of inquiry as a theory of "epistemology as hypothesis", I take Dewey's

theory of inquiry (in its exposition of common-sense and scientific inquiry, individual and social

inquiry) to be an *epistemic* theory of inquiry, rather than a purely ethical or political theory.[53] Third,

my interpretation of Dewey's theory of inquiry as a non-ideal epistemic theory of inquiry is entirely

my own, as is my use of the term "epistemic practice" to refer to inquiry. I take this to be a fairly

unproblematic reading of Dewey's theory of inquiry. Fourth, when referring to other Deweyan

thinkers on democracy such as Elizabeth Anderson, Jack Knight, and James Johnson, I will take them

as reading the pragmatist theory of inquiry as I do.

And lastly, to move outside the pragmatist literature into contemporary epistemology, any discussion

or engagement on the recent literature on norms of inquiry (or 'zetetic' norms) within contemporary

_____

political disagreements – see Appendix A), but from a completely different basis. One way of explaining the difference
between the Peircean defence of democratic inquiry, and the Deweyan account of democratic inquiry is with an analogy to
the difference in the starting assumptions of political theory from John Rawls and Amartya Sen: recall how Rawls believed
that an ideal theory of justice is required before one can do non-ideal political theory, while Sen argued that one can evaluate
whether a society is *more* just without needing a theory of what is *fully* just. In a similar way, whereas Misak and Talisse see an
epistemic justification for democracy as requiring a prior explanation on the basis of Peirce's constitutive norms of belief, the
Deweyan theory of inquiry that I put forward in this thesis claims that one can argue that "epistemic improvement and
democratic reform go hand-in-hand" (to use Elizabeth Anderson's phrase) without a prior epistemic theory to justify this.
Nevertheless, there is no need to see the pragmatist response to disagreement that I raise in this thesis to be in any conflict
with that given by Misak and Talisse, any more than we need to see Rawl's theory of distributive justice as being in conflict
with Sen's "capabilities approach" to welfare economics. I briefly touch on the Peircean defence in Appendix A of this thesis.
See this for more information. Elizabeth Anderson, "The Epistemology of Democracy," *Episteme* 3, no. 1-2 (2006): 21.

[53] See Putnam and Putnam, "Dewey's Logic". Contrary to Hilary and Ruth-Anna Putnam, Matthew Festenstein has argued
that Dewey's conception of social inquiry is primarily ethical and political rather than epistemological. However, it is
important to note that Festenstein's conception of what epistemology seems to be influenced by Rorty throughout this paper.
As such, given that my interpretation of Dewey's theory of inquiry does not fit Rorty's definition of epistemology, it remains
to be seen whether Festenstein would disagree with the epistemic theory of inquiry I outline in this thesis. For more
information, see Matthew Festenstein, "Inquiry as Critique: On the Legacy of Deweyan Pragmatism for Political Theory,"
Political Studies 49 (2001): 730–48.

epistemology will be outside of the scope of this essay. This is largely because the recent literature on

norms of inquiry has typically focused on giving end-state accounts of extremely abstract norms of

inquiry – for example, the Zetetic Instrumental Principle (or ZIP) states that "if one wants to find out

$Q^?$, then one ought to take the necessary means to figuring out $Q^?$" – the pragmatist theory of inquiry is

interested in how we use a common pattern of problem-solving (taken from both scientific inquiry

and democratic procedure) in order to provide a transition account about how to improve our

epistemic practice.[54] Let us now consider Dewey's theory of inquiry.

1.3.2   Dewey's Theory of Inquiry as Non-Ideal Epistemology

The basic idea of Dewey's theory of inquiry is an extremely simple one: inquiry is the process in which

we as epistemic agents face an initial situation that causes some question or doubt (known as the

"indeterminate situation"), and through various stages of deliberation and action (a process usually

described as "the experimental method"), come to an answer to that initial question or doubt (that is,

ending in a "determinate situation").[55] Taken at face value, this notion of inquiry does not denote any

particularly interesting epistemic phenomenon: being just another way of describing the ordinary

process of problem solving which every human being uses in everyday life. Nor does it appear *prima*

---

[54] In this way, I agree with David Thorstad's recent argument that "there are no epistemic norms of inquiry". Specifically,
Thorstad (taking inspiration from the work of Allan Gibbard) argues that because inquiry is a form of activity – just like
building a house or writing a book – it is not governed by epistemic norms, but must be based on a wider set of all-things-
considered reasons which lie outside the purely epistemic. I agree with Thorstad, but qualify his claim by noting that what I
am calling non-ideal epistemic practice takes into consideration precisely these kinds of all-things-considered reasons. See
David Thorstad, "There Are No Epistemic Norms of Inquiry," *Synthese* 200 (2022). The Zetetic Instrumental Principle is
taken from Jane Friedman, "The Epistemic and the Zetetic," *Philosophical Review* 129, no. 4 (2020): 503.

[55] Or, in Dewey's somewhat convoluted phrasing, inquiry is "the controlled and directed transformation of an indeterminate
situation into one that is so determinate in its constituent distinctions as to convert the elements of the original situation into
a unified whole." Dewey, *Logic*, 138.

facie to lead to a full-fledged epistemology: that is, of the likes of Kant's transcendental argument, or the kind of theoretical analysis given by epistemic logic or Bayesian probability theory.[56] And yet, it is this notion of inquiry which Dewey, and the subsequent Deweyan pragmatists, claim is the overlooked basis of the various cognitive achievements and epistemic success which human beings have made and had throughout history. Throughout the wide span of interests in which Dewey's epistemology covers – ranging from philosophy of science to democratic theory, psychology to philosophy of education – it is this notion of inquiry which explains the way in which an epistemic agent (individually, or within a community of inquirers) can interact with their immediate epistemic environment, find the justification they need to settle their belief (what Dewey calls "warranted belief"), and learn to become rational (although Dewey prefers the term *intelligent*) within their developing epistemic practice.

It is this theory of inquiry that I submit provides the best conceptual framework with which we can understand and evaluate non-epistemic issues. This is because Dewey's theory of inquiry provides a philosophical way of theorising that is recognisably epistemic, but which contrasts with the wider epistemological tradition in its re-interpretation of the traditional epistemic concepts of justification, belief, and knowledge in terms of concepts regarding inquiry and epistemic practice. To use the ideal/non-ideal distinction, we can elaborate by saying that the pragmatist theory of inquiry is able to provide an transition account of epistemic practice which complements and contrasts with the typical end-state accounts given within contemporary analytic epistemology: such that whereas contemporary analytic epistemology analyses epistemic rationality in terms of the epistemic status (e.g. *reasons*) for

---

[56] This is why Dewey uses the terms "judgment" and "reflective thought" to refer to the same process of inquiry: for Dewey, the process of inquiry is not a specialised process which only applies to certain domains of life, but the basic process which describes "how we think". For Dewey on judgment, see Dewey, *Logic*, 283ff. For Dewey on reflective thought, see Dewey, *How We Think*.

one's belief, pragmatist epistemology evaluates epistemic rationality through the *reasoning* one uses in practice. Let us illustrate this with a few examples.

First, suppose a high school student is doing their mathematics homework and has just finished a complicated calculus question, one which involves various steps of reasoning. The high school student is unsure of their answer, and wonders whether the answer they gave was a rational answer. What kind of epistemic theory would explain whether the high school student was rational in their answer, or not? The answer to this question depends on the kind of explanation we want. Perhaps the kind of explanation the student wants is an epistemic evaluation of the answer itself: that is, whether the answer *is* rational or not. This is the kind of explanation which is typically given within contemporary epistemology. Specifically, the typical account given within contemporary epistemology involves providing a post-hoc assessment of the epistemic rationality of the student in terms of the justification for their belief. This means that the student's answer to the question would be rational just in case their justifications for this answer were valid, and irrational just in case their justifications were not valid. The standard by which we assess the validity of the student's justification would then depend on the theory of justification we accept: if we are a reliabilist about epistemic justification, we would assess the student's justification in terms of the process by which they arrived at their belief and its reliability (e.g. through valid deductive inference or not); or if we are an evidentialist, we would assess the student's justification on the basis of the evidence they have about their answer. Such is the standard 'end-state' account given within epistemology.

But suppose the student wanted a different kind of explanation: one which explained *how* the student is supposed to figure out whether the answer is rational. In this case, the standard account given within epistemology would not be the kind of explanation the student wants: to say that the answer *would* be rational if it was formed by a sufficiently reliable process, or if it was based on the evidence they have, would not be helpful if the student is unsure of the reliability of their belief-forming process or perhaps

even unsure of the evidence that they have.[57] What the student wants is a method by which they can assess their answer, and thereby come to a conclusion themselves about whether their answer is rational or not. This is the kind of epistemic account which Dewey's theory of inquiry gives. To give a small elaboration on Dewey's theory of inquiry, consider the following two principles of inquiry given by Putnam: whereas the principle of *fallibilism* claims we we should "not regard the product of any inquiry as immune from criticism", the principle of *experimentalism* claims that we should "try out different ways of resolving problematical situations, or if that is not feasible, observe those who have tried other ways, and reflect carefully on the consequences."[58] Here, the pragmatist theory of inquiry provides a transition account of epistemic norms, which explains what effective inquiry looks like in order for the student to settle the uncertainty about their answer themselves. For example, the student can adopt the principle of experimentalism by trying different methods of answering the question (so as to check if they get the same result), or searching for videos of other people doing similar questions so as to check their understanding of mathematical formulae. In this way, the student can *inquire* into the question about whether their answer is rational or not, and come to find an answer in the particular context that they are in. As such, the pragmatist theory of inquiry provides a transition

---

[57] Of course, within the literature on evidentialism, the question about whether it is possible for one to be unsure of one's evidence depends on the theory of evidence that one accepts: under Feldman's theory of evidence as one's occurrent mental states, for instance, it is not possible for one to be unsure of one's evidence because one's evidence just is the things that one is currently thinking about. But, under Timothy Williamson's theory of evidence as known propositions, it is certainly possible for one to have evidence without knowing that one has evidence. More on this when we discuss theories of evidence in §§4.1-4.3 of this thesis.

[58] Putnam, *The Collapse of the Fact/Value Dichotomy and Other Essays*, 110. While Putnam claims these to be principles of inquiry given by American pragmatism, I take these to be principles given by Putnam to avoid the various disputes about fallibilism and experimentalism within the literature. For more on the dispute between Peirce versus Dewey's interpretation of fallibilism, for instance, see Joseph Margolis, "Peirce's Fallibilism," *Transactions of the Charles S. Peirce Society* 34, no. 3 (1998): 535–69.

epistemic account of rationality, one which explains how one can improve their epistemic practice (i.e. be "more" rational) in their particular context.

Let us consider another example. Suppose a recent university graduate is given a job offer outside his hometown, one which looks to be a promising start to his career but one which he is not sure whether he will enjoy. What kind of epistemic theory can we give to explain whether the university graduate was rational in their decision, or not? In this case, one might argue that the research into decision theory provides both an explanation of what it means for a decision to *be* rational (i.e. to provide an 'end-state' account), and a method for *how* one can figure out what is rational within one's own context (i.e. to also provide a 'transition' account). As a brief explanation of decision theory, the basic premise of decision theory involves determining the rationality of a decision by considering three factors: first, the beliefs one has about the possible states of the world; second, the possible actions they have available to them; and third, the preferences they have about the possible outcomes that may occur as the result of one's action and the actual state of the world.[59] Given one's beliefs, preferences, and possible actions, one way of determining which action is the most rational is the principle of "maximising expected utility": that is, to evaluate each action according to the likelihood of the outcome (based on one's degrees of belief) and one's preference of that outcome, and picking the action which is expected to lead to the most preferred outcome. In this way, the principle of expected utility maximisation not only serves as an account of rational decision-making, but also a method by which one can figure out what the rational decision is within their epistemic practice.

---

[59] For an excellent overview of decision theory, see Michael D. Resnik, *Choices: An Introduction to Decision Theory* (Minneapolis: University of Minnesota Press, 1987).

However, at this point, it is important to highlight how certain conditions must be met in order for one to use the decision theoretic framework to evaluate a decision within one's context. Specifically, the decision theoretic framework begins by assuming that one has determinate beliefs about the relevant states of the world, determinate preferences about the outcomes, and a defined set of possible actions which one can make in a given situation. In simple cases regarding whether one should bring an umbrella given one's beliefs about the weather, such conditions are easily satisfied and the decision theoretic framework gives great results. But when it comes to the more important life decisions, such as the graduate's decision about their career, there are at least two reasons why one might be in a situation where the initial conditions for decision theory are not met.

First, it is very common for one to be forced to make a decision without sufficient information, such that one may not know the likelihood of how the world is, or for one to be unsure even of what one's possible actions are or of one's preferences regarding unknown future outcomes. In the case of the university graduate, for instance, it is entirely possible for them to be so unfamiliar with the workforce that they neither know what actions they can take when seeking employment, or their one's preferences are when it comes to working conditions.[60] Second, the decisions one has to make in

[60] The literature on decision theory has typically considered the topic of uncertain decision-making as affecting one's *beliefs* about the world, such that uncertain decision-making involves making decisions with imprecise credences about the state of the world, or to otherwise make decisions solely on the basis of one's actions and preferences. The topic of uncertainty about one's preferences is indirectly discussed when considering problems of incommensurability or incompleteness: that is, when considering cases where one faces a choice betwen two outcomes where one has no preference for one over the other, and yet one does not believe that both preferences are equally good. As far as I know, there is no discussion on how one should make decisions when one is uncertain about their possible actions. For more on decision-making with imprecise probailities, see Matthias C. M. Troffaes, "Decision Making Under Uncertainty Using Imprecise Probabilities," *International Journal of Approximate Reasoning* 45 (2007): 17–29; Rohan Sud, "A Forward Looking Decision Rule for Imprecise Credences," *Philosophical Studies* 167, no. 1 (2014): 119–39; Susanna Rinard, "A Decision Theory for Imprecise Probabilities," *Philosophers' Imprint* 15, no. 7 (2015); Seamus Bradley, "A Counterexample to Three Imprecise Decision Theories," *Theoria* 85 (2019): 18–30. For more on decision-making without beliefs (also known as decisions under ignorance), see Chapter 2 of

everyday life are often made within a dynamic environment where one's beliefs, preferences, and set of possible actions are constantly changing. This means that long-term decision making (e.g. the ongoing question of what career one should pursue) requires one to also explain how we should evaluate successive decisions over a period of time – a question which has been a particularly thorny question to answer.[61] As such, in situations where we have to make decisions where our beliefs and preferences are either indeterminate or changing, the decision-theoretic framework does not provide a clear transition epistemic account of decision-making, and therefore cannot be used to explain how one should figure out how to make a rational decision in their given context.

To put it in another way, we can say that the decision theoretic framework provides a *transition* account of decision-making which is based on *end-state* accounts of belief, preferences, and action. In situations where one has determinate beliefs, preferences, and set of possible actions, decision theory provides a powerful tool for explaining how one can make a rational decision between various options. But in the cases where one does not have the starting conditions regarding belief, preferences, and actions, decision theory fails to provide any epistemic guidance. By contrast, Dewey's theory of inquiry

---

Resnik, *Choices*. For more on the problem of incommensurability and incompleteness, see Lara Buchak, "Normative Theories of Rational Choice: Rivals to Expected Utility," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Summer 2022 (Metaphysics Research Lab, Stanford University, 2022), §3; Sven Ove Hansson and Till Grüne-Yanoff, "Preferences," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Spring 2022 (Metaphysics Research Lab, Stanford University, 2022), §2.2.

[61] This is largely due to the problem of "diachronic tragedy": the problem when one makes a series of decisions which all seem to be rational decisions when considered individually, but which lead to contradictory results when taken altogether. See Brian Hedden, "Options and Diachronic Tragedy," *Philosophy and Phenomenological Research* 90, no. 2 (2015): 423–51. In fact, it is on the basis of these kinds of problems within diachronic decision making which has led some philosophers to argue for "time-slice epistemology": the view that there are no diachronic norms of rationality, only synchronic norms. See Brian Hedden, "Time-Slice Rationality," *Mind* 124 (2015): 449–91; Brian Hedden, *Reasons Without Persons: Rationality, Identity, and Time* (Oxford University Press UK, 2015); Sarah Moss, "Credal Dilemmas," *Noûs* 48, no. 3 (2014): 665–83.

provides a fuller transition account for decision-making, by situating the moment which we make a decision within a wider context of epistemic practice. In particular, Dewey's account of inquiry reveals how our actual decision-making is often an extended and involved process where an epistemic agent has to partake in smaller acts of deliberation both before and after the moment of making a decision.

Take the case of the recent university graduate. For Dewey, a key part of the structure of inquiry is a prior stage of deliberation in which one brainstorms possible plans of actions with which one can use to inform one's decision.[62] For instance, one way in which the recent university graduate can make a more informed decision is by taking a weekend visit to the place of the job offer in order to see if they enjoy the location, or to figure out ways of transport to and from their hometown. Or, one can ask amongst family and friends to see if they can be connected to recent graduates or other professionals in their field, in order for one to gain a better understanding of the kind of job offers they should expect. In brainstorming and trying possible plans of actions, the university graduate is therefore able to *test* the circumstances of their situation (i.e. following experimentalism) in a way which allows the graduate to make a *better* decision than simply relying on their immediate beliefs and preferences.[63]

---

[62] For more information on Dewey's explanation of deliberation as "imaginative rehearsal", see Steven Fesmire, *Dewey*, ed. Brian Leiter, Routledge Philosophers (London; New York: Routledge, 2015), 88ff.

[63] As a quick aside, Hilary and Ruth-Anna Putnam note that Dewey used this exact point to criticise the hypothetico-deductive (or H-D) model of the scientific method: the view that the scientific method consists of formulating hypotheses which would have certain observable outcomes if true, and then conducting experiments which confirm the hypotheses if one obtains the predicted outcomes. In trying to formalise the scientific method into a definitive procedure, Dewey explains that the H-D model ends up oversimplifying the complexity of scientific inquiry, because taken by itself, the H-D model suggests that "scientific method is merely a combination of guesswork and the fallacy of affirming the consequent." Dewey continues to note that the reason why the hypothetico-deductive pattern is successful within scientific research is because of how this pattern is used within a larger complex process of *inquiry*, one which used "experiment both prior to the formulation of an hypothesis and afterward, [and where] deductions in question always employ a host of so-called auxiliary hypotheses." As such, once again, Dewey's theory of inquiry reveals how the ways in which we as epistemic agents conduct individual or collective deliberation within our epistemic practice does not consist of making a single decision, but often consists of an

Of course, at this point, it is important to clarify what the purpose of the comparison between Dewey's theory of inquiry and decision theory is. It would certainly be unfair to compare how Dewey's theory of inquiry does in capturing long-term decision-making to decision theory, given that the decision-theoretic framework is clearly designed to evaluate the rationality of a single decision made in a single moment in time. Instead, the purpose here is to make a metaphilosophical point about how rational decision-making is typically considered within epistemology: because decision theory is the dominant way of theorising about rational decision-making within contemporary epistemology, contemporary epistemology often evaluates decisions from a position where one already has determined beliefs, preferences, and possible actions. But this is also the reason why decision theory (and other end-state epistemic accounts) cannot explain what an epistemic agent is supposed to do in actual epistemic practice. Because end-state epistemic accounts focus on explicating an epistemic concept in terms of its necessary and sufficient conditions, these accounts cannot explain what an epistemic agent should do when those conditions are not met in actual practice: if one's beliefs do not satisfy the conditions for knowledge or rationality according to a given theory, then one is simply lacking in knowledge or irrational according to that theory. By contrast, Dewey's theory of inquiry aims to capture the part of our epistemic practice where we engage in a process of inquiry so as to arrive at our beliefs and preferences in the first place. This allows Dewey's theory of inquiry to provide a non-ideal account of epistemic practice, one which is severely lacking within the current epistemic literature.

This is the pragmatist theory of inquiry which I will argue provides the best non-ideal epistemic account of disagreement. Nevertheless, in order to defend this claim, we must address the clear competitor to Dewey's theory of inquiry within the analytic epistemology of disagreement: namely,

---

extended process which one tries a series of preliminary and tentative plans of action, in order to arrive at an answer about the decision in question. Putnam and Putnam, "Dewey's Logic," 203.

evidentialism. One of the main appeals of evidentialism as an epistemic theory is its explication of justification in terms of the evidence that one has presently available. This allows evidentialism to provide a unified account of epistemic rationality, such that one can always evaluate what is rational to believe in their particular context by referring to the total body of evidence they have available. Accordingly, evidentialism has had a significant influence within the current research into various non-ideal epistemic issues within the literature: such as questions regarding epistemic self-doubt within individual epistemology, or recent studies into testimony and disagreement within social epistemology.[64]

Furthermore, within the literature on evidentialism, philosophers such as Richard Feldman have explicitly argued that the epistemic rationality of an epistemic agent should be evaluated solely on the ways in which they revise their belief on the basis of their evidence, rather than the actions they take to arrive at their body of evidence (more on this distinction between "methodological rationality" and "current-state rationality" in §4.2 of this thesis). This suggests that the kind of theory which Dewey's theory of inquiry provides is one which is not "epistemic" at all: because it is based on practical considerations around what is prudent rather than epistemic considerations around what is rational. Evidentialism therefore places itself as a strong competitor to Dewey's theory of inquiry: in providing a general non-ideal epistemic response to disagreement which rejects the need for considering practical concerns regarding the acquiring of evidence.

―――――――――――――――――

[64] For more on evidentialist accounts of self-doubt and testimony, see Sherrilyn Roush, "Epistemic Self-Doubt," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Winter 2017 (Metaphysics Research Lab, Stanford University, 2023), §2; Nick Leonard, "Epistemological Problems of Testimony," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Spring 2023 (Metaphysics Research Lab, Stanford University, 2023) §3.

This brings us to the remaining chapters of this thesis. Chapter 2 will outline how evidentialism has been used to develop a general non-ideal epistemic response to disagreement. To begin, I shall explain how evidentialism has developed from a theory about epistemic justification into one of the dominant forms of epistemic theorising within analytic epistemology. In particular, I shall argue that the basic assumptions of evidentialism has been used as a theoretical basis for Bayesian epistemology (i.e. as an interpretation of probability), such that the conjunction of evidentialism and Bayesian epistemology (what I call "Bayesian evidentialism") has led to a revolution within contemporary analytic epistemology: one of similar importance to the introduction of predicate logic into conceptual analysis, or the introduction of possible-world semantics into contemporary metaphysics. It is this broader view of evidentialism which I claim has consistently played a strong influence on the analytic epistemology of disagreement: both in the initial characterisation of the epistemic problem which disagreement poses, and within the various theoretical responses which have been put forward towards this epistemic problem. In fact, I shall argue that the dominance of this evidentialist view has led the literature on the analytic epistemology of disagreement to largely come to a close. This is because of the convincing answer to the epistemic problem of disagreement given by the Total Evidence View: a view which ingeniously uses the case of ideal peer disagreement as an epistemic *criteria* by which one evaluates dissenting epistemic parties as higher-order evidence. The Total Evidence View therefore provides a general evidentialist response to disagreement: one which claims that the rational response to disagreement is to incorporate dissenting parties into one's total body of evidence as higher-order evidence, and then to simply revise one's beliefs accordingly.

Once we have a clear grasp on the evidentialist response to disagreement, Chapters 3 and 4 will put forward a critique of the evidentialist response to disagreement. Chapter 3 will set-up the critique by highlighting how the evidentialist response to disagreement requires a substantial theory about what evidence is in order to defend the norm to follow your evidence. Specifically, I shall explain how the

evidentialist response to disagreement is only valid if evidence has two key epistemic features: first, that evidence determines the epistemic rationality of an epistemic agent by providing doxastic justification towards one's belief; and second, that evidence is commonable in such a way that dissenting parties can come to a *common* body of evidence, and use evidence to adjudicate between competing claims. This brings us to Chapter 4, where we will examine three of the main philosophical theories of evidence within the literature – namely, the empiricist theory of evidence as sense-data, Earl Conee and Richard Feldman's original theory of evidence as occurrent mental states, and Timothy Williamson's theory of evidence as known propositions – and show that no theory of evidence satisfies the theoretical desiderata needed to substantiate the evidentialist response to disagreement. This is because, as Williamson shows, the strong evidentialist claim that evidence determines epistemic rationality fails to account for cases where one is not in the position to know what one's total body of evidence is, or to otherwise know what is rational to believe in a certain context. It is here where we return to the pragmatist theory of inquiry. In particular, I shall explain how the pragmatist theory of inquiry elaborates on Williamson's notion of epistemic context by highlighting the role which other epistemic agents play in constituting our epistemic context. This reveals how our interactions with other epistemic agents within inquiry is not simply as passive sources of higher-order evidence but as active interlocutors which challenge our interpretation of evidence, and provide us with new and different epistemic methods.

This brings us to Chapter 5 of my thesis, where I outline the pragmatist response to disagreement and explain how the pragmatist response to disagreement provides a better non-ideal epistemology of disagreement. Continuing from the last chapter, Chapter 5 will begin by considering Thomas Kelly's epistemic analysis of the psychological research on "belief polarisation". Kelly's epistemic analysis of belief polarisation provides a helpful starting point for connecting the evidentialist response to disagreement to the pragmatist account of social inquiry, since Kelly's account of belief polarisation

reveals how differences within the subjective attitudes of two dissenting parties can lead them to worsen their disagreement even if both dissenting parties started with the same initial body of evidence. For Kelly, the upshot of belief polarisation is to highlight how one's subjective attitude can play a significant *causal* role in one's inquiry, such that it is possible for two epistemic agents to eventually acquire radically different bodies of evidence that oddly allow both parties to be justified in their conflicting opinions.

It is these kinds of subjective problems that epistemic agents face in individual inquiry which lead the pragmatists to insist on the importance of cooperative social inquiry. In particular, Deweyan pragmatists have given substantial epistemic accounts about two kinds of successful social inquiry – namely, science and democracy – in order to show how disagreements over matters of fact can be systematically discussed and resolved (if not, settled) over the course of inquiry. I will focus on two accounts whith have been given within the literature which show how successful social inquiry contains certain structural features which allow a community of inquirers to resolve disagreement. First, Hilary Putnam's account of scientific inquiry highlights the importance of *ethical* norms for the *epistemic* efficacy of scientific research, since ethical norms are required to producing a cooperative discursive environment through which a community of inquirers can mutually exchange ideas, reasons, and objections. Second, Elizabeth Anderson's account of democracy highlights the epistemic functions which *social institutions* play within democratic procedure, given that democratic processes are the best means by which we receive the information needed to form effective decisions, or to revise ineffective decisions. In each account, both Putnam and Anderson reveal how the pragmatist theory of inquiry sees disagreement as an essential epistemic feature within social inquiry, one where the opinions of disenting party become the new information which we use to revise and improve collective epistemic practice.

The pragmatist response to disagreement therefore states that the best epistemic response to disagreement involves a commitment to continue social inquiry, and cooperating with dissenting parties so as to come to greater understanding and better decision-making. A crucial implication of the pragmatist response to disagreement is therefore that the main epistemic problem within disagreement is also the only epistemic solution to disagreement. It is *us*: taken as separate and individual cognitive islands, we are forced to moral conflict and measly cognitive achievements. But taken as a cooperative community, we can aspire to deeper mutual understanding, and greater epistemic potential.

2   DISAGREEMENT AND EVIDENCE

2.1   Introduction: How Evidence Resolves Disagreement

The focus of the following three chapters is the role which evidence plays in resolving disagreement as it is explained within contemporary analytic philosophy. One of the main strategies which we use to resolve disagreement – be it deciding for ourselves between dissenting opinions, or seeking resolution with a dissenting party – is by using a common body of evidence to adjudicate between differing opinions. To take a few examples, evidence is used to resolve disagreements within legal contexts as written, forensic, and testimonial evidence are brought forth by the defendant and the prosecution so that the judge and jury may decide the verdict of the case. Evidence also plays a major role within scientific inquiry – when deciding between competing hypotheses about a particular scientific question, it is the replicable evidence of physical results or statistical data which the scientific community uses to favour one hypothesis over another. And finally, when settling everyday disputes, it is natural for us as epistemic agents to answer questions about what diet to adopt or what political candidate to vote for by collecting more information (and thus, evidence) through the use of search engines or various media platforms.

Given that we often use evidence to resolve disagreement in both professional contexts and our everyday epistemic practice, it is no surprise that the focus on evidence is reflected within the philosophical literature on the epistemology of disagreement. Before we continue to consider this literature, however, it is important to note that what contemporary epistemologists mean by 'evidence' is different to how it is used within ordinary discourse. Specifically, within contemporary analytic epistemology, evidence is used as a technical term which refers to "the kind of thing which can make a difference to what one is *justified* in believing or (what is often, but not always, taken to be the same

thing) what it is *reasonable* for one to believe".[65] This means that whereas our ordinary use of the term

'evidence' refers to different things in different contexts – for example, evidence for historians may

take the form of physical written documents, while evidence in chemistry may take the form of

numerical digits which reflect spectrometer readings – epistemic theories of evidence are concerned

with whatever stands behind these pieces of evidence which allow them to provide justification for our

beliefs, and therefore constitute 'evidence' in the first place. As we shall see, this is why epistemic

theories of evidence have tended to propose psychological or mental items as theoretical candidates for

the referent of evidence as opposed to physical objects – candidates such as sense data, mental states,

experience, and known propositions.

It is this epistemic concept of evidence which will be the focus of the following three chapters. To give

a brief overview of this chapter, the structure of my argument is as follows: in §2.2, I shall explain how

the epistemic concept of evidence has been developed within contemporary epistemology through

evidentialism. While evidentialism began as a theoretical claim about epistemic justification, it has

since been developed into a full-fledged theory about epistemic rationality and one of the key

theoretical approaches to analysing epistemic issues within contemporary analytic epistemology. In

§2.3, I explain how the theoretical strengths of evidentialism are seen in full display within the analytic

epistemology of disagreement, where the concept of evidence is used both to establish the basic (ideal)

theoretical assumptions regarding epistemic problem of disagreement, as well as to develop a general

non-ideal approach to responding to disagreement. With regard to the former, the influence of

––––––––––––––––––––––––––––––

[65] Thomas Kelly, "Evidence," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Winter 2016 (Metaphysics
Research Lab, Stanford University, 2016), §1. The italics are from Kelly. Also, while Kelly admits that this is only one of the
many ways in which evidence is used within philosophy (see the other sections of his SEP entry for more information), it is
clear that this is the main way in which evidence is used within contemporary analytic epistemology.

evidentialism can be seen in the 'same evidence' condition of peer disagreement, in establishing the connection between evidence and epistemic rationality through the Uniqueness Thesis, and in explaining the motivating assumptions behind conciliatory versus steadfast approaches to disagreement. With regards to the latter, the Total Evidence View uses evidentialism to provide a general non-ideal epistemology of disagreement through the notion of 'higher-order evidence': the notion that dissenting parties are epistemically relevant to your beliefs, because they provide evidence about your evidence. Within the analytic epistemology of disagreement, then, a promising non-ideal epistemology of disagreement comes from the *evidentialist* response to disagreement: that the correct epistemic response to disagreement is simply to follow one's evidence. This will bring us to Chapter 3 and 4, where we will critically examine the various attempts to provide a philosophical explication of this epistemic concept of evidence. Let us now move on to examine evidentialism in greater detail.

2.2   *Evidentialism Within Contemporary Analytic Epistemology*

2.2.1   Evidence as Source of Epistemic Justification: On Evidentialism

While the general notion of evidence has been discussed throughout the history of Western epistemology, the recent focus on evidence as a distinct epistemic concept has stemmed from the influence of evidentialism within analytic epistemology.[66] The original account of evidentialism comes from Earl Conee and Richard Feldman, who defined evidentialism as follows:

> **Evidentialism (Justification Thesis):** Doxastic attitude *D* toward proposition *p* is epistemically justified for *S* at *t* if and only if having *D* toward *p* fits the evidence *S* has at *t*.[67]

---

[66] We will examine the historical origins of our modern concept of evidence in greater detail in §4.1 of this thesis.

[67] Earl Conee and Richard Feldman, "Evidentialism," *Philosophical Studies* 48, no. 1 (1985): 15.

As we can see, Conee and Feldman's original formulation of evidentialism sees evidentialism as a theoretical analysis of epistemic justification, one which claims that our doxastic attitude (e.g. belief, disbelief, suspension of judgement, or degree of belief) towards a given proposition is justified just in case it accords with the evidence we have at that point in time. In this way, evidentialism begins as a form of internalism about epistemic justification, one which served as an important alternative to the traditional 'accessibilist' conception of internalism about epistemic justification, and the various criticisms which were raised against it at the time.[68]

However, since its inception, evidentialism has come to have a much greater impact within contemporary epistemology for two interconnected reasons. First, rather than being seen as simply one within a set of competing theoretical analyses about epistemic justification, evidentialism has since been discussed as an independent epistemic thesis about how evidence normatively constrains belief. In the introduction to a recent collection of essays about evidentialism, Trent Dougherty helpfully notes how "the trend in [contemporary analytic] epistemology is to liberate many traditional concerns from the theory of knowledge proper, in recognition of their independent value."[69] Here, Dougherty is

---

[68] Conee and Feldman famously defend evidentialism as an internalist account of epistemic justification in Earl Conee and Richard Feldman, "Internalism Defended," in *Evidentialism* (Oxford: Oxford University Press, 2004), 53–82. This paper is particularly important because of its distinction between two kinds of internalism: namely, between accessibilism, the internalist view that one can have internal access to the basis of knowledge and justified belief; and mentalism, the internalist view that the basis of epistemic justification is one's mental states. In making this distinction, Conee and Feldman were able to distinguish evidentialism from access internalism (and the criticisms it faced), by categorising evidentialism as a mentalist view about epistemic justification. More about this mentalist view in §4.2 of this thesis. For more information about the criticisms of accessibilist internalism, Conee and Feldman cite the following: see Philip Kitcher, "The Naturalists Return," *The Philosophical Review* 101, no. 1 (1992).

[69] Trent Dougherty, "Introduction," in *Evidentialism and Its Discontents*, ed. Trent Dougherty (New York: Oxford University Press, 2011), 2. Preceding this quote, Dougherty qualifies this claim by noting that there is one exception, presumably referring to Williamson's focus on knowledge in his "E=K" thesis. More on Williamson in §4.3 of this thesis.

referring to how contemporary epistemologists have realised that many of the theories proposed

within epistemology are worth considering even without reference to the concept of knowledge, such

that theories of epistemic justification such as process reliabilism, responsibilism, and evidentialism

should be examined as independent epistemic norms of belief rather than as potential (and competing)

candidates for the analysis of the justification condition of knowledge.[70] Accordingly, in place of the

original formulation, Conee and Feldman have since adopted a stronger view on evidentialism in the

form of a supervenience thesis:

> **Evidentialism (Supervenience Thesis):** The epistemic justification of anyone's doxastic attitude
> toward any proposition at any time strongly supervenes on the evidence that the person has at the
> time.[71]

Throughout the rest of our discussion on evidentialism, it is this supervenience thesis which is the

commonly used definition on evidentialism, and the one we will assume to be the fundamental claim

of evidentialism.[72] This brings us to the second way in which evidentialism has influenced

contemporary epistemology: namely, as a way of characterising epistemic *rationality* in general. To

---

[70] That being said, the respective norms which these epistemic theories put forward can and often do conflict in their
implications for our epistemic practice. But it is certainly more fruitful to examine the differences between these norms as
they are, rather than in reference to the concepts of justification and knowledge. For more on process reliabilism, see Alvin I.
Goldman, *Epistemology and Cognition* (Cambridge: Harvard University Press, 1986). For more on epistemic responsibilism
(or the deontological conception of justification), see Laurence BonJour, *The Structure of Empirical Knowledge* (Cambridge,
MA: Harvard University Press, 1985); William P. Alston, "The Deontological Conception of Epistemic Justification,"
Philosophical Perspectives 2 (1988): 257–99; Richard Feldman, "Epistemic Obligations," *Philosophical Perspectives* 2 (1988):
235–56; and Matthias Steup, ed., *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*
(Oxford: Oxford University Press, 2001).

[71] Earl Conee and Richard Feldman, *Evidentialism* (Oxford: Oxford University Press, 2004), 101.

[72] Conee and Feldman state that the supervenience thesis is the "bedrock" claim of evidentialism, and the literature on
evidentialism has since followed suit in taking the supervenience thesis as being the basic assumption required for a theory on
epistemic justification to be an *evidentialist* theory. For more information, see Conee and Feldman, *Evidentialism*, 2004, 101ff;
Dougherty, "Introduction," 7.

explain what I mean, let us begin by comparing the differences between the two formulations of evidentialism listed above. The main difference between the two formulations of evidentialism is this: whereas the original formulation of evidentialism defines epistemic justification as a three-place relation between one's doxastic attitude, one's body of evidence, and the 'evidential fit' between doxastic attitude and body of evidence, the supervenience thesis defines epistemic justification as a two-place relation between one's doxastic attitude and one's evidence. Apart from being a simpler definition of evidentialism, the supervenience thesis also provides two theoretical advantages to the original formulation of evidentialism: first, the supervenience thesis avoids the tenuous notion of 'evidential fit' of the original formulation, and the problems which it causes. Specifically, one implication of the original formulation of evidentialism is that it allowed the possibility that two epistemic agents could have the same body of evidence and be equally justified in taking different doxastic attitudes, simply because it was possible for two epistemic agents to "fit" their evidence to different doxastic attitudes. While this was not a problem in and of itself, the notion of 'evidential fit' was found to be too vague to be explicated in any meaningful way. This meant that the supervenience thesis was generally preferred in its focus on the concept of *evidence* and the justification it provided towards one's doxastic attitude, rather than focusing on an epistemic agent, and the manner in which they *fit* their evidence towards their doxastic attitudes.[73]

Second, in defining epistemic justification directly in terms of one's evidence, the supervenience thesis allows us to compare the relative justification that different epistemic agents have for their doxastic attitudes by reference to their respective evidence. This is because the supervenience thesis implies that

---

[73] As it turns out, this question of whether one could be justified in taking different doxastic attitudes upon the same body of evidence reappears within the epistemology of disagreement as the debate between Uniqueness and Permissivism. More in §2.3.2 of this thesis.

two epistemic agents with the exact same body of evidence would therefore have the same justification

for a particular doxastic attitude. The supervenience thesis therefore allows the concept of evidence to

be used as a standard by which we can assess the epistemic justification one has for their belief, in a

similar way to how the concept of utility is used as a standard by which we assess the reasons behind

one's action, or how we use the rules of deductive logic to assess the validity of one's inferences. This

has meant that evidentialism has become the dominant way of defining *epistemic* rationality within

contemporary epistemology: in addition to the logical rationality of making valid inferences, or the

instrumental rationality of choosing the correct means for a given end, epistemic rationality is typically

defined as when "one believes propositions that are strongly supported by one's evidence and refrains

from believing propositions that are improbable given one's evidence."[74] While there are other rational

norms of belief which are discussed within the literature – such as 'normic support' views, mentalism,

reliabilism, and dispositionalism – evidentialism has stood out as one of the dominant theoretical

approaches to explicating what is epistemically rational in contemporary epistemological issues due to

its connection to Bayesian epistemology.[75] To explain, let us consider what this connection is and why

it matters in turn.

2.2.2 Evidence as Grounds for Rational Belief Revision: On The Evidential Interpretation of Probability in Bayesian Epistemology

First, what is the connection between evidentialism and Bayesian epistemology? To begin, we can

define Bayesian epistemology as a formal epistemology which uses probability theory to theorise about

---

[74] This is how Thomas Kelly, for instance, defines epistemic rationality in comparison to instrumental rationality. See Thomas Kelly, "Epistemic Rationality as Instrumental Rationality: A Critique," *Philosophy and Phenomenological Research* 66, no. 3 (May 2003): 612.

[75] For more information on these other norms of belief, see §11.1 of Hughes, "Epistemology Without Guidance".

epistemic rationality. Specifically, Bayesian epistemology uses probability theory to explicate epistemic rationality by characterising rational belief revision as a form of probabilistic inference. Bayesian epistemologists defend the claim that rational belief revision is a form of probabilistic inference through the use of a series of pragmatic self-defeating arguments known as Dutch Book Arguments – arguments which analogise from betting ratios to conclude that an epistemic agent's degrees of belief must follow the classical axioms of probability at any given point in time, and that an epistemic agent's degrees of belief must be revised upon new information through an application of the Ratio formula (a process which is known as the Bayesian principle of conditionalisation).[76] In showing how one's belief revision should practically follow the probability axioms, Bayesian epistemology therefore justifies its use of the probability calculus to formally represent an agent's degrees of belief, and to model and evaluate an agent's epistemic rationality.

However, in doing so, Bayesian epistemology requires a theoretical interpretation of probability which explains what it is that the probability calculus is modelling. One important debate in which this question about the correct interpretation of probability comes up is the debate between subjective and

---

[76] Dutch Book arguments work by measuring one's degrees of belief through betting ratios, and then showing how certain assignments of degrees of belief will lead one to take a set of bets which entail a sure-loss. The original Dutch Book argument is made by Frank Ramsey, which shows how you are susceptible to a Dutch book if your credences violate the classical probability axioms: see Frank Ramsey, "Truth and Probability," in *Philosophy of Probability: Contemporary Readings*, ed. Antony Eagle (Routledge, 1926), 52–94, and B. Skyrms, *Pragmatics and Empiricism* (New Haven: Yale University Press, 1984). Kemeny proves the contrapositive result: that if your credences conform to the probability calculus, then no Dutch book can be made against you: see J. Kemeny, "Fair Bets and Inductive Probabilities," *Journal of Symbolic Logic* 20 (1955): 263–73. Finally, there is a Dutch Book argument developed by Lewis and Teller for the principle of conditionalisation: see David Lewis, "A Subjectivist's Guide to Objective Chance," in *Studies in Inductive Logic and Probability*, ed. Richard C. Jeffrey, vol. 2 (Berkeley: University of California Press, 1980) and Paul Teller, "Conditionalization, Observation, and Change of Preference," in *Foundations of Probability Theory*, ed. W. Harper and C. A. Hooker (Dordrecht: D. Reidel, 1976). However, Jeffrey raises suspicions about Dutch Book arguments, and in particular, its identification of credences with betting ratios: for more information, see Richard C. Jeffrey, *The Logic of Decision*, 2nd ed. (University of Chicago Press, 1990).

objective Bayesianism. For context, the debate between subjective and objective Bayesianism is concerned with the kinds of norms which govern the prior probabilities within the initial probability assignment (also known as "the problem of the priors"): whereas subjective Bayesians claim that the only norms which prior probabilities must follow is coherence according to the probability axioms, objective Bayesians claim that prior probabilities are also constrained by certain substantial norms of rationality.

Here, it is important to note that the terms 'subjective' and 'objective' Bayesianism is used to characterise this debate because, historically, this debate was between Bayesians who held a subjective interpretation of probability, and Bayesians who held an objective interpretation of probability. Specifically, whereas subjectivists (or personalists) about probability such as Bruno de Finetti argued that there are no rational constraints upon one's prior probabilities other than probabilistic coherence because the probability calculus is meant to represent one's actual doxastic attitudes, objectivists about probability such as E. T. Jaynes argued that there are substantial rational constraints upon one's prior probabilities because the probability calculus is meant to represent how the ideal rational agent should conform their doxastic attitudes to objective chances within the world.[77] However, the recent literature has since moved on from this debate regarding the correct interpretation of probability because of the success of a third alternative interpretation of probability known as the evidential interpretation of

---

[77] For more information on de Finetti's subjectivist account of Bayesian epistemology, see Bruno De Finetti, "Foresight: Its Logical Laws, Its Subjective Sources," in *Breakthroughs in Statistics* (Springer, 1992), 134–74; and Bruno De Finetti, *Theory of Probability: A Critical Introductory Treatment* (Chichester: John Wiley & Sons Ltd, 2017). For more information on how Jaynes uses a Bayesian framework to study objective chance, see Roger D Rosenkrantz, *Foundations and Applications of Inductive Probability* (Atascadero, CA: Ridgeview Publishing, 1981); E. T. Jaynes, *Papers on Probability, Statistics, and Statistical Physics*, ed. Roger D. Rosenkrantz (Dordrecht, Holland: D. Reidel Publishing Co., 1983); and E. T. Jaynes, *Probability Theory: The Logic of Science*, ed. G. Larry Bretthorst (Cambridge: Cambridge University Press, 2003).

probability.[78] The evidential interpretation of probability defines probability as the degree to which "evidence supports or counts against various hypotheses about the world, for example that our world had a beginning or that the butler did it."[79] In taking the probability calculus to be representing the degree to which an agent's available evidence supports a given proposition, the evidential interpretation of probability provides an intuitive and elegant interpretation of probability which avoids many of the key problems which its purely subjective and objective counterparts face.[80] Let us consider the traditional criticisms which proponents of subjective and objective interpretations of probability respectively faced (call this the debate between classical subjective Bayesians and classical objective Bayesians), and how the evidential interpretation of probability avoids these respective criticisms.

On the one hand, classical subjective Bayesianism is criticised because of its extremely permissive view about the rational constraints an agent's degrees of belief has. Since the subjective interpretation of probability takes the probability calculus as representing an agent's actual doxastic attitudes towards a

---

[78] In fact, the evidential interpretation of probability has been so effective that the recent literature on subjective and objective Bayesianism considers the debate to be solely about the degree to which one's prior probabilities should be constrained by their available *evidence*. This is, for example, how Jon Williamson characterises the debate – see Jon Williamson, *In Defence of Objective Bayesianism* (Oxford: Oxford University Press, 2010).

[79] D. H. Mellor, *Probability: A Philosophical Introduction* (London; New York: Routledge, 2005), 80. This is also sometimes known as the epistemic interpretation of probability.

[80] Just as how theories about epistemic justification are now seen as independent norms of belief rather than competing analyses of the concept fo justification, the current literature considers the various interpretations of probability as different *kinds* of probabilities. D. H. Mellor helpfully categorises three kinds of probability which exist in the world: namely, objective chances which exists as ontological features in the world, subjective credences which exist as psychological features within epistemic agents, and epistemic probabilities which exist as epistemic relations between evidence and propositions. One way of framing my argument here is to see objective Bayesianism as developing from conforming one's credences to objective chances within the world, to conforming one's credences to epistemic probabilities. For more information, see Chapter 1 of Mellor, *Probability*.

set of propositions, the only requirements on an agent's degrees of belief are the rational constraints given by the Dutch Book arguments – namely, that an agent's degrees of belief should follow the probability axioms and the principle of conditionalisation.[81] Under classical subjective Bayesianism then, an epistemic agent could have seemingly irrational degrees of belief for a proposition (e.g. they could believe that the phenomenal appearance of the world is in fact an deceptive act of malicious devils) and still be considered as rational insofar as their degrees of belief follow the probability axioms and the principle of conditionalisation. The evidential interpretation of probability, however, avoids this objection of extreme permissiveness, as the evidential interpretation of probability allows one to defend substantial normative constraints towards an agent's degrees of belief through the concept of evidence. This is because, under the evidential interpretation of probability, the probability calculus would therefore represent the degree to which an agent's body of evidence supports a given proposition. As such, if we take the probability calculus to represent an ideal rational agent, the evidential interpretation of probability therefore suggests that an agent's degrees of belief should not only follow the axioms of probability and the principle of conditionalisation, but also be in line with the evidence that is presently available to them.[82]

On the other hand, classical objective Bayesianism is criticised for providing an overly narrow definition of probability which does not capture our intuitions about likelihood statements and

---

[81] At least, that is the standard view. However, within the recent literature, other Dutch Book arguments have been made for David Lewis' Principal Principle (more on this in the next paragraph), and for sharp credences. For more information on these recent Dutch Book arguments, see §5 of Susan Vineberg, "Dutch Book Arguments," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Fall 2022 (Metaphysics Research Lab, Stanford University, 2022).

[82] James Joyce puts it this way: "On this view [i.e. Bayesian epistemology under the evidential interpretation of probability], a person's total, nonincremental evidence regarding a hypothesis H is directly reflected in her level of confidence in H." James M Joyce, "Bayesianism," in *The Oxford Handbook of Rationality*, ed. Alfred R. Mele and Piers Rawling (Oxford: Oxford University Press, 2004), 143.

epistemic rationality. As we mentioned earlier, classical objective Bayesians such as E. T. Jaynes

consider the probability calculus to represent objective chances to which an epistemic agent's degrees

of belief should conform. By objective chances, the classic objective interpretations of probability

considered probability to be real aleatory features of the world – such that probabilities referred to

equipossibilities which exist in logical space (known as the classical interpretation of probability), or to

the relative frequencies of actual physical phenomena (also known as the frequentist interpretation of

probability). Once we discover what these objective chances are – that is, by mathematical calculation

or statistical testing – classical objective Bayesians argue that one's degree of belief about objective

chances should equal the chances themselves.[83]

However, a major problem for classical objective Bayesianism arises when we consider likelihood

statements for which we do not know the objective chances, or which might not refer to objective

chances at all. To take one example, consider the statement "it is likely that the universe has a

beginning".[84] Here, classical objective Bayesianism states that such a statement is either a statement

which refers to some objective chance (e.g. about the logical or metaphysical possibility of universes

with a beginning), or not a coherent statement at all (e.g., by claiming that there is no likelihood

regarding the beginning of the universe). In both of these cases, however, classical objective

Bayesianism provide an implausible account of epistemic rationality: whereas the former provides an

---

[83] In this way, classic objective Bayesians seem to advocate for David Lewis' Principal Principle. Interestingly enough, however, David Lewis proposes this principle in a paper where he explicitly claims to be a subjectivist about probability. David Lewis is therefore a rare example of an objective Bayesian who was a subjectivist about probability. See Lewis, "A Subjectivist's Guide to Objective Chance". See also David Lewis, "Humean Supervenience Debugged," *Mind* 103 (1994): 473–90.

[84] For more information about how these kinds of cases (known as probabilities of singular propositions) pose a problem for classic chance views, see Mellor, *Probability*, 33.

impossibly stringent standard of rationality where epistemic agents must conform their degree of belief

to facts about (inaccessible) possible universes, the latter claims that there are no rational constraints

regarding one's belief about whether the universe has a beginning, given that there is no objective

chance to which this statement pertains. And yet, our intuitions suggest that there are rational

constraints surrounding our beliefs about these kinds of statements: constraints which are not based

upon any chance process within the world, but on what is reasonable to believe given the empirical

research into cosmology and the expansion of the universe. It is here where the evidential

interpretation of probability provides a better explanation about the kind of epistemic norms which

constrains our doxastic attitudes. Specifically, because the evidential interpretation of probability takes

the probability calculus to represent the degree to which a body of evidence supports a given

proposition, the normative constraints on an epistemic agent's degrees of belief are determined by the

evidence that is available to an epistemic agent at a given point in time. This means that the probability

calculus (under the evidential interpretation of probability) can always represent the epistemic

rationality of an individual, because the credence which an epistemic agent should have towards a

proposition is determined by their available evidence, rather than by unknown or non-existent chance

processes. In this way, the evidential interpretation of probability provides a better interpretation of

the probability calculus within Bayesian *epistemology*: while probability theory may still be used in

other areas of research to examine chance processes within the world, the evidential interpretation of

probability provides the best interpretation of probability for evaluating when an epistemic agent is

justified and rational in their belief.[85]

----

[85] Of course, while this exposition has considered, following the literature on Bayesian epistemology, E. T. Jaynes as an objective Bayesian, it is important to note that Jaynes is more accurately described as a mathematician rather than an epistemologist. This distinction is important given that Jaynes seems more interested in developing methods of applying probability theory to scientific reasoning, rather than answering the philosophical question of what epistemic norms governs

For these general reasons, the evidential interpretation of probability has since become the dominant

interpretation of probability for Bayesian epistemology. Timothy Williamson summarises the situation

as follows:

> Given a scientific hypothesis *h*, we can intelligibly ask: how probable is *h* on present evidence? We are
> asking how much the evidence tells for or against the hypothesis. We are not asking what objective
> physical chance or frequency of truth *h* has. A proposed law of nature may be quite improbable on
> present evidence even though its objective chance of truth is 1. That is quite consistent with the
> obvious point that the evidence bearing on *h* may include evidence about objective chances or
> frequencies. Equally, in asking how probable *h* is on present evidence, we are not asking about
> anyone's actual degree of belief in *h*. Present evidence may tell strongly against *h*, even though
> everyone is irrationally certain of *h*.[86]

We can now state the connection between evidentialism and Bayesian epistemology more clearly.

While one might debate as a matter of historical fact whether evidentialism contributed to the

development of the evidential interpretation of probability[87], it is obvious that the use of evidential

----

one's belief. As such, to the extent that Jaynes is interested in using probability theory to theorise about matters outside of
epistemic rationality, it is slightly inaccurate to consider Jaynes to be a Bayesian epistemologist, let alone an objective
Bayesian. For more information, see the citations in fn. 77.

[86] Williamson, *Knowledge and Its Limits*, 209. As an addendum to the two senses of subjective and objective Bayesianism,
Williamson is an example of one who hold the evidential interpretation of probability but someone who is an objectivist
about Bayesian epistemology: that is, Williamson argues for substantial rational constraints on one's prior doxastic attitudes.

[87] For instance, one might object that the attempt to connect evidentialism with the evidential interpretation of probability is
anachronistic, given that the concept of evidence was used in the inductive logics of John Maynard Keynes and Rudolf
Carnap well before evidentialism was originally formulated by Conee and Feldman. Specifically, both Keynes and Carnap
were using a Bayesian framework to develop an inductive logical system which fully determined the degree of implication
which a piece of evidence *e* conferred onto a hypothesis *h*. As such, the use of evidence as an interpretation of probability in
Bayesian epistemology seems to predate evidentialism, and thus it is anachronistic to claim that evidentialism influenced the
development of the evidential interpretation of probability. However, in response to this objection, a deeper examination into
both the inductive logics of Keynes and Carnap shows that Keynes and Carnap did not in fact subscribe to the evidential
interpretation of probability. This is because, for both Keynes and Carnap, the inductive logical system was intended to
determine the degree of implication evidence confers to a hypothesis as a matter of logical *syntax.* As such, the inductive
logics of Keynes and Carnap would be better described as a form of classical objective Bayesianism, rather than as a precursor
to the evidential interpretation of probability. For more on the inductive logics of Keynes and Carnap, see John Maynard
Keynes, *A Treatise on Probability* (London: Macmillan; Co., 1921); Rudolf Carnap, *Logical Foundations of Probability*

interpretation of probability in Bayesian epistemology commits one to evidentialism. Insofar as the

probability calculus is meant to represent the ideal rational agent, Bayesian epistemology under the

evidential interpretation of probability supposes that epistemic rationality requires one's initial degrees

of belief towards a proposition to be the degree to which one's available evidence supports that

proposition, and thus presupposes that evidentialism is true. Of course, this is not to say that Bayesian

epistemology must use the evidential interpretation of probability, nor does it mean that evidentialism

commits one to Bayesian epistemology. As we have mentioned before, one could reject the evidential

interpretation of probability and instead accept a subjective or objective interpretation of probability.

Or, one could accept evidentialism and yet reject the Bayesian claim that there is such an doxastic

attitude which corresponds to the notion of "degree of belief". Such an evidentialist might argue that

the appropriate doxastic attitudes which fit with one's evidence are "full" attitudes such as all-out

belief, all-out disbelief, or complete suspension of belief. However, for the purposes of this thesis, I

shall take it that even if there is no necessary connection between Bayesian epistemology and

evidentialism, it is clear that the conjoined use of Bayesian epistemology with the evidential

interpretation of probability (call this Bayesian evidentialism from now on) has been so influential as

to be one of the dominant modes of theorising within analytic epistemology – a noticeable change akin

to the introduction of predicate logic into conceptual analysis within analytic philosophy, or possible

world semantics in contemporary metaphysics. There is a hardly a single contemporary debate within

contemporary analytic epistemology which has not been discussed through a broadly Bayesian

evidential framework – that is, by theorising about what is epistemically rational in a given issue by

---

(Chicago: University of Chicago Press, 1950); Rudolf Carnap, *The Continuum of Inductive Methods* (Chicago: University of

Chicago Press, 1952). For some rare examples of evidential accounts of probability that do predate evidentialism, see Glenn

Shafer, *A Mathematical Theory of Evidence* (Princeton; London: Princeton University Press, 1976); L. Jonathan Cohen, *The

Probable and the Provable* (Oxford: Clarendon Press, 1977).

considering how an epistemic agent should revise their beliefs (where revision usually means changing

one's degrees of belief following Bayesianism) according to a *body of evidence* (following

evidentialism).

As a conclusion to this section then, it is helpful to note why Bayesian evidentialism has been so

successful as a mode of theorising within contemporary analytic epistemology. This is especially

important for this thesis since, as I shall argue, one of the main strengths of Bayesian evidentialism

which leads to its prevalence within contemporary analytic epistemology is its applicability to *non-

ideal* epistemological questions. In using the probability calculus to model epistemic rationality,

Bayesian evidentialism borrows the explanatory power of the probability calculus to provide an

analysis of epistemic rationality which is suited for evaluating non-ideal epistemic issues. To begin my

defence of this claim, consider the following quote from Williamson on the unique theoretical

strengths of Bayesian epistemology:

> One strength of Bayesianism is that the mathematical structure of the probability calculus allows it to
> make illuminating distinctions which other approaches miss and provide a qualitatively fine-grained
> analysis of epistemological problems, given assumptions about all reasonable prior assumptions.[88]

What are these "illuminating distinctions" which Bayesian evidentialism provides that other

approaches miss? Here, I shall argue that one of the main kinds of distinctions which Bayesian

evidentialism engages with that other classical approaches do not is those of *non-ideal* epistemology.

To illustrate this, consider the theoretical differences in method and aim between Bayesian

evidentialism and a classic epistemological theory such as the JTB analysis of knowledge. First, when

taken in comparison, it is evident that whereas the JTB analysis of knowledge is a *utopian*

---

[88] Williamson, *Knowledge and Its Limits*, 212. Note that Williamson's point here is made on his chapter on 'evidential

probability' – it is therefore reasonable to assume that by Bayesianism, Williamson is referring to Bayesian evidentialism.

epistemological project, Bayesian evidentialism is a *realistic* epistemological project.[89] We can see how

the JTB analysis of knowledge is a *utopian* theory in its claim that knowledge is factive – in claiming

that a proposition is knowledge only if it is true, the JTB analysis of knowledge evaluates whether a

proposition is knowledge in virtue of factors outside of the epistemic agent (specifically, in virtue of

whether the proposition represents some state of affairs in the world in an appropriate way such that

the proposition is "true"). As such, in examples such as the famous Gettier cases, the JTB analysis of

knowledge is not interested in non-ideal questions of whether an epistemic agent is or could be *aware*

of their epistemic luck (let alone whether they can do anything about their own epistemic luck), but is

instead interested in whether epistemic luck precludes one's justified true belief from constituting

knowledge.

---

[89] That being said, recall that in §1.2.1, I explained how the distinction between utopian and realistic theory is a matter of degree. There are certain utopian assumptions within the classical Bayesian framework – such as the implicit assumption of logical omniscience, as well as the assigning of credence 1 to all evidence propositions. But these utopian assumptions are being addressed by contemporary developments within Bayesian epistemology, which aim towards further realism by replacing utopian assumptions with conditions which better reflect actual epistemic cases. One example is the problem of assigning credence 1 to all evidence propositions, which has been addressed within the recent literature through the suggestion of replacing the standard principle of conditionalisation for Jeffrey conditionalisation, which allows for uncertain evidence. Or, in response to the seemingly utopian assumption of sharp credences, theories of "imprecise probabilities" have been proposed which modify Bayesian epistemology as to account for epistemic situations where one's degrees of belief are not precise. For more on Jeffrey conditionalisation, see Jeffrey, *The Logic of Decision*. For more on imprecise probabilities, see Seamus Bradley, "Imprecise Probabilities," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Spring 2019 (Metaphysics Research Lab, Stanford University, 2019). For an argument for a realistic account of sharp credence, see Nicholas J. J. Smith, "Problems of Precision in Fuzzy Theories of Vagueness and Bayesian Epistemology," in *Vagueness and Rationality in Language Use and Cognition*, ed. Richard Dietz, vol. 5 (Switzerland: Springer Nature Switzerland AG, 2019), 31–48.

By contrast, Bayesian epistemology sidesteps the question of whether a proposition is knowledge entirely by focusing on what degree of belief an epistemic agent should have about a proposition.[90] The shift in focus from "knowledge" to "degrees of belief" allows Bayesian epistemology to theorise about an agent's actual doxastic attitudes rather than a epistemic state that an epistemic agent may not be aware of, nor feasibly obtain in their circumstances. The shift to an agent's actual doxastic attitudes is complemented by the evidential interpretation of probability, which provides normative constraints for an epistemic agent's degree of belief for a proposition according to the actual evidence which is available to them in support for that proposition. As such, Bayesian evidentialism shows its applicability to non-ideal epistemological questions by being a *realistic* epistemology: in being sensitive to feasibility conditions with respect to the epistemic agent, Bayesian evidentialism allows one to evaluate epistemic rationality in non-ideal epistemic cases.

Furthermore, in comparing Bayesian evidentialism with the JTB analysis of knowledge, it is also clear that Bayesian evidentialism can be more easily characterised as a *transition* epistemology. This is because the theoretical aim of Bayesian evidentialism is not just to provide an analysis of an epistemic end (e.g. "knowledge", "justification", or "belief"), but rather to give a theoretical framework for agent belief revision (that is, how one's degrees of belief may change or "transition" according to new information). Here, it is important to note that Bayesian evidentialism does contain an end-state theoretical component: specifically, it provides an end-state theory of *evidence* as the source of one's

---

[90] In replacing an all-out conception of "belief" for a graded conception of "degrees of belief", Bayesian epistemology is able to sidestep the question of whether the proposition to be believed is *true* by asking about what level of *confidence* an agent should have on the basis of their evidence. Joyce puts it in the following way: "Bayesians maintain that any adequate epistemology must recognize that beliefs come in varying gradations of strength. They seek to replace the categorical notion of belief as an all-or-nothing attitude of accepting a proposition as true with a graded conception of belief as level of confidence. In general, a person's level of confidence in a proposition X will correspond to the extent to which she is disposed to presuppose X's truth in her theoretical and practical reasoning." Joyce, "Bayesianism," 132–33.

epistemic justification. This is something we will discuss more extensively in the next chapter on

philosophical theories of evidence. Nevertheless, the reason why Bayesian evidentialism is a transition

theory is because it builds upon this end-state theory of justification with the probability calculus,

which allows an epistemic agent to recognise, understand, and evaluate the justification they have for

their degrees of belief in a given epistemic context. As such, rather than simply giving a conception of

justification which serves as a long-term epistemic end that an epistemic agent should aim for,

Bayesian epistemology also allows an epistemic agent to make "short-term reforms" to their beliefs, as

an epistemic agent may use the probability calculus to determine the actual justificatory status of their

beliefs in a given context by referring to the evidence available to them.

Of course, one might argue that the probability calculus will never be used by the general populace in

their everyday epistemic practice, and thus fails to help actual epistemic agents to evaluate their beliefs

in a given moment. But this is placing too high of a demand on non-ideal epistemic theory: insofar as

the principles of Bayesian epistemology can be used in certain contexts to evaluate rational belief

revision, Bayesian epistemology is able to provide a non-ideal theory of epistemic rationality in those

contexts. And this is what we see in how Bayesian evidentialism has been used to evaluate a wide range

of non-ideal epistemological issues within analytic epistemology, as well in other 'applied' fields of

research. These include research within statistics, research into artificial intelligence; psychological

research into learning; studies in confirmation theory within philosophy of science; as well as studies

in risk assessment within decision theory.[91] In conclusion, Bayesian evidentialism provides a non-ideal

---

[91] One might argue that not every Bayesian working within these respective areas accept the evidential interpretation of probability. However, I would argue that insofar as Bayesian epistemology is used as to pick out a distinct view on epistemology and probability (e.g. when Bayesian statistics is defined in contradistinction to frequentist statistics), Bayesian epistemology is often described as a probability assignment for the degree of belief one ought to have on the basis of one's evidence, and thus the evidential interpretation is implicitly assumed. For more on how Bayesian statistics differs frequentist statistics, as well as an argument for why these differences are often negligible in practice, see Daniel Greco, "Significance

theory of epistemic rationality, as it gives an account about the norms (and method) of rational belief

revision as opposed to a theoretical analysis of an epistemic concept. This means that Bayesian

evidentialism is better suited to answering non-ideal questions of *how* an epistemic agent should

update their belief based on available evidence as opposed to simply considering ideal questions of

what constitutes belief, justification, or knowledge for an epistemic agent. The applicability of a

Bayesian evidential framework in analysing non-epistemic issues can be seen in its influence on the

analytic epistemology of disagreement – both in the characterising of the problem of disagreement,

and in the various theoretical approaches which have been proposed in response to this problem.

2.3   On the Role of Evidence in Analytic Epistemology of Disagreement

2.3.1   The Role of Evidence in Characterising Epistemic Disagreement: Ideal Peer Disagreements and the Same Evidence Condition

In order to understand the recent literature on the epistemology of disagreement, it is important to

distinguish between two philosophical questions which we might be interested in when considering

the epistemology of disagreement. On the one hand, there is a theoretical question about the epistemic

nature of disagreements: what is it that makes a disagreement an *epistemic* disagreement (or that leads

a disagreement to have epistemic *dimensions*), and how should we understand the epistemic problem

---

Testing in Theory and Practice," *The British Journal for the Philosophy of Science* 62, no. 3 (2011): 607–37. For more on
Bayesian statistics, see William M. Bolstad and James M. Curran, *Introduction to Bayesian Statistics*, Third (Wiley, 2016); and
Peter M. Lee, *Bayesian Statistics: An Introduction*, Fourth (Wiley, 2012). For more on Bayesian inference in artificial
intelligence, see Glenn Shafer and Judea Pearl, *Probabilistic Reasoning in Intelligent Systems* (San Mateo, CA: Morgan
Kaufmann, 1988); Judea Pearl, *Causality: Models, Reasoning, and Inferences*, Second (Cambridge: Cambridge University Press,
2009). For more on Bayesian learning theory in psychology, see R. A. Jacobs and Kruschke J. K., "Bayesian Learning Theory
Applied to Human Cognition," *Wiley Interdisciplinary Reviews: Cognitive Science* 2, no. 1 (2011): 8–21. For more on
Bayesianism in decision theory, see Resnik, *Choices*; and Christian P. Robert, *The Bayesian Choice: From Decision-Theoretic
Foundations to Computational Implementation*, Second (New York: Springer, 2007).

which these disagreements raise for our beliefs? On the other hand, there is a practical question about the epistemic response to disagreements: what is the appropriate epistemic response to disagreement for an epistemic agent, and how does one know when to revise their beliefs when encountering disagreement?

Here, it is important to note that while there is a distinction between these two philosophical questions, it is clear that both questions are intimately related. In fact, within the analytic epistemology of disagreement, the main strategy has been to use the conceptual question as the starting point for answering the practical question: only once we have understood which kinds of disagreement are epistemically noteworthy can we therefore establish norms about belief revision which explain the appropriate response to the epistemic problem of disagreement.[92] In this way, the general structure of the analytic epistemology of disagreement has followed a similar strategy to Rawls' use of ideal and non-ideal theory: that is, by first providing an *ideal* theory which determines exemplary cases of epistemic disagreement and the appropriate response to these ideal cases, and then using the ideal theory as a model for understanding and evaluating *non-ideal* cases of disagreement. In the remainder of this chapter, I will examine how Bayesian evidentialism has influenced the analytic epistemology of disagreement with respect to both philosophical questions: in conceptualising the nature and problem of epistemic disagreements, and in providing a non-ideal strategy to responding to everyday cases of disagreement. While we will focus mainly on how evidentialism has influenced the analytic

---

[92] For example, this is how Frances and Matheson structure the epistemology of disagreement for the Stanford Encyclopedia of Philosophy, and how Frances structures his own introductory book. See Bryan Frances and Jonathan Matheson, "Disagreement," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Winter 2019 (Metaphysics Research Lab, Stanford University, 2019); Bryan Frances, *Disagreement* (Cambridge, UK: Polity Press, 2014).

epistemology of disagreement, I will return to how Bayesianism contributes to the non-ideal solution to the epistemic problem of disagreement at the end of §2.3.3.

To start, let us consider how the concept of evidence has been used in conceptualising the epistemic problem of disagreement. As we have mentioned, it is clear from the outset that not every case of disagreement will pose an epistemic challenge to my beliefs. For example, if I were to have a disagreement with a friend about whether a particular music album was enjoyable or not, such a disagreement might leave my justifications for my beliefs about the album completely unaffected even if I respected their opinions on music immensely. This is because my belief about whether a certain album is good may be dependent on my personal aesthetic taste in music rather than on any information that is relayed to me by others. As such, in order to specify which kinds of disagreement pose an epistemic challenge to our beliefs, the epistemology of disagreement has focused on disagreements about matters of fact – about the truth (or relative likelihood) of a given proposition.[93]

In addition to restricting the focus of the epistemology of disagreement to disagreements about facts, the literature within the analytic epistemology of disagreement has predominantly focused on the specific case of *ideal peer disagreements*. To explain why, consider the following examples. Suppose

---

[93] That being said, it is always *possible* for my personal aesthetic taste to be influenced by the information I receive from others. For instance, my appreciation of a certain genre of music might increase if I learnt the cultural impact that it had in a certain society, or my enjoyment of a certain album might decrease should I learn that the artist is actually deeply misogynistic. As such, while the epistemology of disagreement is concerned with disagreements about matters of fact, this should not be taken to imply that epistemic problems do not occur within areas which are traditionally assumed to be concerned with matters of value – such as ethics and aesthetics. To the extent that matters of fact can influence a moral disagreement – for example, that debates about the truthfulness of a political candidate can be influenced by empirical facts about when they have lied – such moral disagreements may still be epistemically evaluated. For more on this topic, see Appendix A of this thesis, where I will argue that the pragmatist response to disagreement applies even to disagreements within evaluative areas (such as ethics, politics, and aesthetics), insofar as there are disagreements about facts in these areas.

your three-year old niece argues with you about whether the number "five" is the biggest number or not. The main reason why this disagreement is not one of epistemic concern is because it is natural to assume that your niece is not your epistemic peer, but epistemically inferior in this case. By contrast, if a mathematics professor disagrees with you about the solution to a problem in differential calculus that they specialise in, such a disagreement is also not of epistemic concern since it is natural to assume that the professor is not an epistemic peer but epistemically superior to you. Using these examples, we can see how in disagreements where the dissenting party is epistemically inferior or epistemically superior, disagreement does not pose an epistemic problem because there is a clear answer to how one should respond to those disagreements: to ignore and remain steadfast in one's beliefs, or to defer and change one's beliefs respectively. As such, the type of disagreement which poses a serious epistemic problem are those where you encounter another epistemic agent who is equally as likely to be correct about the proposition in question as you are (i.e. an epistemic peer), and yet disagrees with you.[94]

It is for these reasons in which the epistemology of disagreement has focused almost exclusively on the case of ideal peer disagreement, such that the literature revolves around the following two questions: first, how can I determine whether the dissenting party is an epistemic peer (and thus identify that the disagreement is epistemically noteworthy), and second, what is the appropriate epistemic response to a disagreement with an epistemic peer? In response to the first question, epistemic peerhood has been typically characterised in the literature as the satisfaction of two conditions:

Same Evidence Condition: Epistemic peers have the same evidence as is relevant to proposition *p*.

---

[94] Elgin puts it this way: "Disagreement per se does not jeopardise epistemic standing. More problematic are cases in which opponents are, and consider themselves to be, epistemic peers. Then they have the same evidence, reasoning abilities, training, and background assumptions." Catherine Z. Elgin, "Persistent Disagreement," in *Disagreement*, ed. Richard Feldman and Ted A. Warfield (New York: Oxford University Press, 2010), 53.

Dispositional Condition: Epistemic peers are equally disposed to respond to the evidence regarding proposition *p* in an epistemically appropriate way.[95]

We will discuss the dispositional condition in more detail in §2.3.3 on 'higher-order evidence'. However, setting the dispositional condition aside, it is clear that a key criteria for identifying epistemic peerhood is the respective evidence that each dissenting party has. This is in line with our intuitions about everyday epistemic practice: when disagreeing with another epistemic agent, it is common for us to assess the situation by asking what evidence they have for their beliefs, and comparing that to the evidence we have for our own. As such, within the recent literature, the focus on evidence in characterising epistemic peerhood shows how the epistemic problem of disagreement has been understood through an evidentialist framework. Rather than seeing disagreement as a problem about knowledge — for example, rather than considering Gettier cases where an individual by chance is the correct party within a disagreement — the epistemic problem of disagreement is taken as a problem

---

[95] For formulations of ideal peer disagreement in the literature, see Nathan L. King, "Disagreement: What's the Problem? Or a Good Peer Is Hard to Find," *Philosophy and Phenomenological Research* 85, no. 2 (2012): 252; Jonathan Matheson, "Disagreement: Idealized and Everyday," in *The Ethics of Belief: Individual and Social*, ed. Rico Vitz and Jonathan Matheson (Oxford University Press, 2014), 316–17. A few comments on the formulations of ideal peer disagreement. First, within the literature, some have also included an "acknowledgement" condition such that one needs to acknowledge that the other dissenting parties in fact disagree with you. The acknowledgement condition is not important for the purposes of this thesis. Second, whereas Matheson divides the dispositional condition in terms of faculties (e.g. one's intelligence, perceptual ability, etc.) and virtues (e.g. open-minded, patient, intellectually courageous), King categorises all of these factors as part of the dispositional condition. This difference in nomenclature does not make any substantial philosophical difference to the discussion. However, thirdly, what does make a difference is in how Matheson and King conceptualise peer disagreement: while King formulates ideal peer disagreement as a question about whether a dissenting party *actually* is an epistemic peer, Matheson formulates ideal peer disagreement as a question about justification — that is, whether one has *reasons to believe* that a dissenting party is an epistemic peer. More on why this difference in formulation matters when we discuss Kelly and Matheson on higher-order evidence in §2.3.3.

of how one should respond to disagreement based upon the *evidence* that they and the dissenting party have.[96]

2.3.2 On Steadfast and Conciliatory Approaches to Disagreement: Uniqueness versus Permissivism

The importance of one's evidence in determining the rational response to disagreement has been further accentuated by the Uniqueness Thesis, a claim about the connection between evidence and epistemic rationality. We can define the Uniqueness Thesis as follows:

> **The Uniqueness Thesis:** A body of evidence always justifies at most one proposition in a set of competing propositions. Furthermore, a body of evidence justifies at most one doxastic attitude toward any particular proposition.[97]

Here, it is important to see how the Uniqueness Thesis connects with the supervenience thesis of evidentialism. One way of explaining the connection is to see the Uniqueness Thesis as the basis for the supervenience thesis: the reason why our justification for our beliefs strongly supervenes on our evidence is precisely because a body of evidence only justifies *at most* one doxastic attitude towards a particular proposition.[98] To briefly explain the motivation behind the Uniqueness Thesis, consider a

---

[96] This is not to say that some philosophers (albeit a minority of them) have not considered the problem of disagreement as a problem of knowledge – see for example, John Hawthorne and Amia Srinivasan, "Disagreement Without Transparence: Some Bleak Thoughts," in *The Epistemology of Disagreement: New Essays*, ed. David Christensen and Jennifer Lackey (Oxford: Oxford University Press, 2013), 9–30.

[97] Richard Feldman, "Reasonable Religious Disagreements," in *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, ed. Louise Antony (Oxford University Press, 2006), 205. The two claims within the Uniqueness Thesis are referred to as Propositional Uniqueness and Personal Uniqueness respectively, and a Uniqueness theorist may accept one or both of these claims. For more on Propositional Uniqueness and Personal Uniqueness, see Matthew Kopec and Michael G. Titelbaum, "The Uniqueness Thesis," *Philosophy Compass* 11, no. 4 (2016): 189–200.

[98] Kopec and Titelbaum consider the Uniqueness Thesis to be a necessary condition to the supervenience thesis (what they call "Strong Evidentialism"), even though they consider the two these to be "almost logically equivalent". Meanwhile, Ballantyne and Coffman note that the Uniqueness thesis is clearly stronger than the original formulation of evidentialism,

simple case. Suppose you walk outside and look up to see a blue sky. What propositions would you be justified in believing on the basis of such a perceptual experience? The obvious answer seems to be that the perceptual experience is evidence for the proposition "the sky is blue", and therefore one is justified in believing the proposition "the sky is blue". In this case, there is only *one* way for us to take the evidence. As such, on the basis of this perceptual experience, it would not be rational to say that you now have justification for the proposition "the sky is green". Nor would it be rational to say that your perceptual experience therefore justifies that you disbelieve "the sky is blue", or to conclude that one must suspend judgment about whether "the sky is blue" is true. It is cases like these where our intuitions lead us to accept the Uniqueness Thesis - the evidence which consists of your visual experience of a blue sky justifies at most one proposition in a set of competing hypotheses (i.e. hypotheses about the colour of the sky), and justifies at most one doxastic attitude amongst a set of competing doxastic attitudes about any proposition (in this case, justification for belief as opposed to disbelief, suspension of belief).

Why is the Uniqueness Thesis important to the epistemology of disagreement? If the Uniqueness Thesis is true, it is simply not possible for dissenting parties to disagree upon the same body of evidence and be *equally rational*. Suppose you and a friend of yours were both to look up and see a blue sky. For the reasons stated above, this perceptual experience can only justify belief in the proposition "the sky is blue", and disbelief for any competing propositions (e.g. "the sky is green"). As such, there are only two reasons why you and your friend could possibly disagree on the colour of the sky. One reason is that you and your friend do not in fact have the same body of evidence: while you

---

given that the Uniqueness Thesis implies that there is only one kind of "evidential fit" from evidence to doxastic attitude. See Kopec and Titelbaum, "The Uniqueness Thesis," 193; Nathan Ballantyne and E. J. Coffman, "Uniqueness, Evidence, and Rationality," *Philosophers' Imprint* 11, no. 18 (2011): 2–3.

and your friend share the common visual experience of looking up at the sky, you and your friend may have other bodies of evidence which lead you and your friend to make different conclusions (e.g. your friend might be Japanese, and consider the particular colour of the sky to be more accurately classified as green, rather than blue). However, suppose you and your friend do in fact have the same body of evidence as is relevant to the proposition "the sky is blue". Then, the only other explanation for why you and your friend disagree is that you or your friend (or both) are irrational: since the evidence only justifies the proposition "the sky is blue", the dissenting party or parties who disbelieve the proposition "the sky is blue" is therefore shown to be irrational in their beliefs.

In claiming that dissenting parties cannot disagree on the basis of a given body of evidence and be equally rational, the Uniqueness Thesis therefore highlights why peer disagreements pose an epistemic challenge to our beliefs. This is because if the Uniqueness Thesis is true, a peer disagreement in which you or one of the other dissenting parties must be irrational if not mistaken in some way. Given that you also have no further evidence that the other dissenting parties are more likely to be mistaken than you (because if you did, they would not be your epistemic peer), the Uniqueness Thesis implies that encountering peer disagreement gives you a *prima facie* reason to believe that you are mistaken in your beliefs. This result is formally known as Defeat, and can be precisfied as follows: "Learning that a peer disagrees with you about *p* gives you a reason to believe you are mistaken about *p*."[99]

Of course, at this point, it is important to note that the Uniqueness Thesis has not been without controversy within the analytic epistemology of disagreement. Within the literature, some have argued against the Uniqueness Thesis in favour of Permissivism, the claim that multiple doxastic attitudes

---

[99] See §5.1 of Frances and Matheson, "Disagreement". See also David Christensen, "Epistemology of Disagreement: The Good News," *Philosophical Review* 116, no. 2 (2007): 187–217, *passim*.

might sometimes be *permissible* on the basis of a body of evidence. More specifically, Permissivism has often been distinguished as two distinct positions: Moderate Permissivism and Extreme Permissivism. While Moderate Permissivism states that it is possible for a body of evidence to justify differing degrees of belief about a given proposition, Extreme Permissivism takes the permissivist claim further by stating that it is possible for a body of evidence to be such that an epistemic agent could be justified in believing P, and justified in believing not-P, on the same body of evidence.[100]

Nevertheless, it is important here to note that one can hold either Permissivist positions and still accept the original justification thesis of evidentialism: even in the case of Extreme Permissivism, the claim that it is possible for a body of evidence to permit mutually exclusive doxastic attitudes does not show that one's doxastic attitudes are not justified by evidence to *some* degree. As such, even amongst the critics of Uniqueness, very few philosophers deny evidentialism or the use of the concept of evidence to characterise and evaluate the problem of disagreement. In fact, the question about the degree to which evidence permits multiple doxastic attitudes has been at the heart of the debate between the key theoretical responses to ideal peer disagreement proposed within the literature. More specifically, the debate between Uniqueness and Permissivism has often been seen as the deciding factor between the two main approaches to resolving disagreement proposed within the literature: namely, between conciliatory and steadfast approaches to peer disagreement.[101] In their overview of the literature, David Christensen and Jennifer Lackey explain the difference between the conciliatory

---

[100] The original formulation of Moderate and Extreme Permissivism comes from Roger White. See Roger White, "Epistemic Permissiveness," *Philosophical Perspectives* 19 (2005): 445–59; Roger White, "Evidence Cannot Be Permissive," in *Contemporary Debates in Epistemology*, ed. Matthias Steup, John Turri, and Ernest Sosa, Second (UK: Wiley-Blackwell, 2014), 312–23.

[101] Matthew Kopec and Michael Titelbaum also make this connection in their overview on the Uniqueness Thesis. See Kopec and Titelbaum, "The Uniqueness Thesis," 193.

and steadfast approach to peer disagreement, and states that the theoretical responses to peer

disagreement can be categorised according to where they fall on the spectrum between conciliation

and steadfastness:

> Some philosophers advocate positions toward what might be called the "conciliatory" (or
> "conformist") end of the spectrum. On their views, many of the beliefs people hold on a wide range of
> disputed issues—from the controversial to the mundane—need to be either substantially revised or
> altogether abandoned. Other philosophers advocate positions toward what might be called the
> "steadfast" (or "non-conformist") end of the spectrum. On their views, most of those holding opinions
> on disputed issues need not lower their confidence in the face of disagreement, unless there are non-
> disagreement-related reasons for doing so. Of course, this vastly oversimplifies the discussion. Most
> epistemologists hold that conciliatory responses are appropriate in some cases and steadfast responses
> in others. But there still seem to be clear differences in the overall degree of belief-revision various
> philosophers' positions require.[102]

How does the debate between Uniqueness and Permissivism play into choosing between the

conciliatory and the steadfast approach to disagreement? As Christensen and Lackey allude to in this

passage, a key factor in deciding between conciliatory and steadfast approaches to disagreement is the

relative weight one gives to the epistemic challenge which disagreement poses for the justification of

our beliefs: if one believes that encountering disagreement does not in and of itself provide a reason to

revise their beliefs, then steadfastness seems to be an appropriate response to disagreement; whereas if

one believes that encountering disagreement does provide a reason to revise their beliefs, then

conciliation seems to be the right response instead. Accordingly, the reason why the debate between

Uniqueness and Permissivism affects the choice between conciliatory and steadfast views about

disagreement is because the debate between Uniqueness and Permissivism affects the relative

epistemic significance which peer disagreement has on our beliefs. As we have seen above, if

_____

[102] David Christensen and Jennifer Lackey, "Introduction," in *The Epistemology of Disagreement: New Essays*, ed. David
Christensen and Jennifer Lackey (Oxford: Oxford University Press, 2013), 1. Frances makes a similar point in evaluating
alternative approaches in terms of where they stand in-between conciliation (what he calls the Equal Weight View) and
steadfastness (also known as the Steadfast View). See §5.1 of Frances and Matheson, "Disagreement".

Uniqueness is true, then peer disagreement gives one *prima facie* reasons to believe that they are irrational in their beliefs, which means that seeking conciliation seems to be the appropriate response to peer disagreement. Conversely, if Permissivism is true, it is possible for dissenting parties to have different doxastic attitudes towards a body of evidence and still be equally justified. As such, under Permissivism, encountering peer disagreement does not provide you with *prima facie* reasons to doubt one's evidence or justifications, and remaining steadfast in one's beliefs would therefore be the appropriate response to peer disagreement.[103]

In summary, we can see how the concept of evidence has been used in ideal theorising about the epistemology of disagreement. The concept of evidence has been used to explain the nature of ideal peer disagreements, as epistemic peerhood requires that a dissenting party has the same amount of

---

[103] While I have argued here that the debate between Uniqueness and Permissivism affects the decision between conciliatory and steadfast views about disagreement, this does not mean that one cannot adopt the conciliatory or steadfast view for other reasons. So, within the literature, there have been many arguments for why conciliation does not require one to accept the Uniqueness Thesis. Conversely, the steadfast view has also been defended independently from Permissivism – for example, some philosophers have defended the steadfast view on the basis of self-trust, that one should be confident in one's faculties irrespective of the accusations of others. For more information on why conciliation does not require one to accept the Uniqueness Thesis, see David Christensen, "Disagreement as Evidence: The Epistemology of Controversy," *Philosophy Compass* 4, no. 5 (2009): 756–67; Stewart Cohen, "A Defense of the (Almost) Equal Weight View," in *The Epistemology of Disagreement: New Essays*, ed. David Christensen and Jennifer Lackey (Oxford: Oxford University Press, 2013), 98–117; Matthew Lee, "Conciliationism Without Uniqueness," *Grazer Philosophische Studien* 88 (2013): 161–88; Rik Peels and Anthony Booth, "Why Responsible Belief Is Permissible Belief," *Analytic Philosophy* 55 (2014): 75–88; Benjamin Levinstein, "Permissive Rationality and Sensitivity," *Philosophy and Phenomenological Research* 94, no. 2 (2017): 342–70; David Henderson et al., "Nonconciliation in Peer Disagreement: Its Phenomenology and Its Rationality," *Grazer Philosophische Studien* 94 (2017): 194–225. For more information on how self-trust supports the steadfast view, see David Enoch, "Not Just a Truthometer: Taking Oneself Seriously (but Not Too Seriously) in Cases of Peer Disagreement," *Mind* 119 (2010): 953–97; Robert Pasnau, "Disagreement and the Value of Self-Trust," *Philosophical Studies* 172, no. 9 (2015): 2315–39; Karl Schafer, "How Common Is Peer Disagreement? On Self-Trust and Rational Symmetry," *Philosophy and Phenomenological Research* 91, no. 1 (2015): 25–46; Ralph Wedgwood, "The Moral Evil Demons," in *Disagreement*, ed. Richard Feldman and Ted A. Warfield (Oxford: Oxford University Press, 2010).

evidence about the proposition in question as oneself. The concept of evidence has also been used to explain what exactly is the epistemic challenge which is posed by the disagreement – according to the Uniqueness Thesis, disagreement poses an epistemic challenge due to the fact that it is theoretically impossible for two epistemic agents to disagree upon a body of evidence and be equally rational. As such, under the Uniqueness Thesis, peer disagreement provides *prima facie* reasons to believe that one is irrational if not mistaken, and thus in need of further revision to their beliefs. Finally, the concept of evidence has been used when considering theoretical responses to ideal peer disagreement. In debating whether multiple doxastic attitudes may be justified by a given body of evidence, the debate between Uniqueness and Permissivism affects the relative epistemic significance of peer disagreement, and thus contributes to deciding between conciliatory and steadfast approaches to peer disagreement. Now that we have examined how the concept of evidence has played within ideal theorising about the epistemology of disagreement, let us how the concept of evidence plays into an non-ideal epistemology for how to respond to actual disagreement.

2.3.3   Higher-Order Evidence and the Evidentialist Response to Disagreement

As we have discussed before, the basic trajectory within analytic epistemology of disagreement has been to clarify the epistemic problem of disagreement through the ideal case of peer disagreement, before considering how this ideal case helps us to evaluate and respond to the everyday and often non-ideal cases of disagreement. This additional non-ideal theory is needed because, as it is often noted within the literature, perhaps the only epistemic peer that you will ever encounter in everyday life is yourself.[104] To explain why this is the case, let us re-examine the two conditions of epistemic peerhood:

---

[104] And even then, you might only be an epistemic peer to yourself in the present moment, and not to a past or future self. Both Matheson and King begin their papers by making this point on the extreme rarity of encountering an epistemic peer. See King, "Disagreement"; Matheson, "Disagreement".

in the case of the dispositional condition, it is rare that we will encounter another epistemic agent who shares all the dispositions that are relevant to appropriately responding to the evidence. This is especially once we note the sheer range of dispositions which might be relevant to being a reliable evidence assessor: factors that affect one's reliability as an evidence assessor include one's general cognitive capacities, acquired skills, sensory faculties such as visual acuity, and even character traits (such as patience or cautiousness).[105] Given this wide range of dispositions which may contribute to being a reliable evidence assessor, it is unlikely that one will ever find another epistemic agent who satisfies the dispositional condition, such that they share in all the relevant dispositions as you. The same is the case for the same evidence condition: the idea that dissenting parties should have the same evidence as is relevant to the proposition might unfeasibly require another epistemic agent to have the same past experiences as you. This is because one often draws from one's past experience to make inductive inferences – be they direct inferences about future events from similar past situations or inferences about the reliability of certain media sources, political parties, or the testimony of certain people around you from past experiences. As such, it is almost assured that the same evidence condition is unsatisfiable in actual cases of disagreement.

Nevertheless, just as how Rawls' ideal case of a well-ordered society is used as a theoretical model for evaluating whether actual societies are just, the ideal case of peer disagreement can also be used as a theoretical model through which we evaluate actual cases of disagreement. One strategy that has been proposed by philosophers like Thomas Kelly and Jonathan Matheson is the Total Evidence View: that

---

[105] See King, "Disagreement," 258–61.

the rational response to disagreement depends on the total body of evidence that one has available at the time.[106]

Here, it is important to note how the Total Evidence View has been developed in a way that is different from predecessors such as the conciliatory and steadfast approaches to peer disagreement. Specifically, the original formulation of conciliation and steadfastness were theoretical responses which apply specifically to the *ideal* case of peer disagreement. In focusing only on ideal peer disagreement, the conciliatory and steadfast approach to disagreement is therefore defended solely on the basis of theoretical concerns such as the connection between evidence and rationality (i.e. the debate between Uniqueness and Permissivism). Once these theoretical concerns have been sufficiently weighed, the correct response to peer disagreement as a whole (be it conciliation, steadfastness, or otherwise) follows as a natural consequence from these theoretical considerations. When facing actual cases of disagreement, however, Christensen and Lackey note that most philosophers acknowledge that there are other facts which make it such that conciliation is the appropriate response to some disagreements, and steadfastness is the appropriate response to others.[107] On what basis do we decide whether conciliation or steadfastness (or some other response) is the appropriate epistemic response to actual cases of disagreement? It is here where the Total Evidence View provides an answer: the Total Evidence View states that one's response to a particular disagreement should be determined by the evidence that is available in that context. Crucial to the Total Evidence View is the notion that

---

[106] See Thomas Kelly, "Peer Disagreement and Higher Order Evidence," in *Social Epistemology: Essential Readings*, ed. Alvin I. Goldman and Dennis Whitcomb (Oxford University Press, 2011), 183–217; Matheson, "Disagreement". Note that while Matheson doesn't explicitly endorse the Total Evidence View in his paper, Matheson's strategy of using ideal peer disagreement to assess evidential asymmetries in actual cases of disagreement is in line with the Total Evidence View.

[107] Christensen and Lackey, "Introduction," 1.

dissenting parties should be treated as a part of the total evidence one has about the proposition in question through the notion of "higher-order evidence". Here is an explanation from Thomas Kelly:

> Given that reasonable individuals are disposed to respond correctly to their evidence, the fact that a reasonable individual responds to her evidence in one way rather than another is itself evidence: it is evidence about her evidence. That is, the fact that a (generally) reasonable individual believes hypothesis H on the basis of evidence E is some evidence that it is reasonable to believe H on the basis of E. The beliefs of a reasonable individual will thus constitute higher-order evidence, evidence about the character of her first-order evidence. Of course, such higher-order evidence, like most other evidence, is not conclusive evidence: it does not follow from the fact that a generally reasonable individual believes H on the basis of E that it is reasonable to believe H on the basis of E. In a case in which E does not adequately support H but a generally reasonable individual mistakenly believes H on the basis of E, the fact that the individual believes as she does constitutes misleading evidence about the character of the evidence E. But misleading evidence is evidence nonetheless. In general, then, the fact that a reasonable person believes H on the basis of E constitutes evidence about the character of E.[108]

How does the notion of higher-order evidence help us in determining the correct epistemic response to actual cases of disagreement? Jonathan Matheson argues that we can use the theoretical notions of epistemic peerhood and peer disagreement as a way of assessing another epistemic agent as a piece of higher-order evidence. Given that a dissenting party is higher-order evidence, what is important about the notion of epistemic peerhood is not whether a dissenting party is in fact an exact epistemic peer, but that a dissenting party gives you reasons to believe (that is, *evidence*) that they are an epistemic peer. This evidential shift from facts about a dissenting party to one's evidence about a dissenting party allows us to use the conditions which constitute epistemic peerhood as *criteria* to assess epistemic agents as higher-order evidence. For example, in the case of the same evidence condition, Matheson argues that even if dissenting parties do not have the exact same body of evidence, you can still

_____

[108] Thomas Kelly, "The Epistemic Significance of Disagreement," in *Oxford Studies in Epistemology*, ed. Tamar Szabó Gendler and John Hawthorne, vol. 1 (Oxford: Oxford University Press, 2005), 186. For a sophisticated account about the degree of evidential support one receives from such higher-order evidence, see Richard Feldman, "Evidence of Evidence Is Evidence," in *The Ethics of Belief: Individual and Social*, ed. Jonathan Matheson and Rico Vitz (Oxford: Oxford University Press, 2014), 284–300.

evaluate a dissenting party as an epistemic peer so long as you have reasons to believe that a dissenting party has *equally good* evidence as you do. By equally good evidence, Matheson is referring to cases where you believe the dissenting party has as much evidence to support not-P as you have evidence to support P.[109] Take Christensen's example of two disagreeing parties who are each supported in their beliefs by their self-conducted polls on the issue.[110] Furthermore, these polls are such that they have sample sizes which are equally big, and equally representative of the general population. In this case, while the dissenting parties do not share the same evidence, their evidence is such that they support their respective beliefs equally strongly. This means that each dissenting party has reasons to believe that the other dissenting party is as likely to be correct about the proposition as they are, and therefore that the other is an epistemic peer. As such, in actual cases of disagreement, one can assess whether a dissenting party is an epistemic peer by looking at one's available evidence, and considering whether one has reasons to believe that the dissenting party has equally good evidence about the given proposition.

Similarly, in the case of the dispositional condition, what is important is not whether a dissenting party has the exact same dispositions as you, but whether you have reasons to believe that the dispositions of a dissenting party make them as likely to be as correct about the proposition as you are. To illustrate, suppose you are having a disagreement with another epistemic agent who you know to have considerably different dispositions to you: to keep things simple, suppose that you are more intelligent than them but they have better eyesight than you do. Despite the fact that you have different dispositions to the dissenting party, the dissenting party might still be an epistemic peer depending on

---

[109] Matheson, "Disagreement," 317.

[110] See Christensen, "Epistemology of Disagreement," 211–12. This example is taken from Matheson, "Disagreement," 317.

the relative impact that these dispositions have to making both of you respond to the evidence in an appropriate way. So, if the disagreement was over whether the bird that both you and the dissenting party see is a red-backed shrike, perhaps the advantages of your intelligence and their eyesight cancel out such that you are as equally likely to be correct as they are. Or, if the disagreement is over a prediction about a basketball game, perhaps your intelligence and their eyesight are irrelevant to the prediction such that your other dispositions make you and the dissenting party equally likely to be correct about the proposition.[111] As such, one can assess whether a dissenting party is an epistemic peer by looking at one's available evidence – if one's evidence about the dissenting party shows that their dispositions make them equally as likely to respond to the evidence in an appropriate way as you are, then you therefore have reasons to believe that a dissenting party is an epistemic peer.

In shifting the focus to one's evidence about dissenting parties, the Total Evidence View therefore allows us to provide an evidential analysis of disagreement, one which provides many non-ideal epistemic insights as to the correct epistemic response to a given disagreement. First, in treating dissenting parties as higher-order evidence, Matheson notes that we can also use the same evidence and dispositional condition to evaluate epistemic superiority and inferiority. If we have reasons to believe that a dissenting party has more evidence about the given proposition, or that they are more disposed to respond to the evidence in the correct way, we have reasons to believe that they are an epistemic superior. Similarly, if we have reasons to believe that a dissenting party has less evidence about the given proposition, or that they are less disposed to respond to the evidence in the correct way, we therefore have reasons to believe that they are an epistemic inferior.

---

[111] Matheson and King both point out how the dispositional condition is context-sensitive, as different dispositions are relevant to different disagreements. See King, "Disagreement," 259; Matheson, "Disagreement," 323.

Second, once we have evidence that a dissenting party is an epistemic superior or inferior, we therefore also have evidence that seeking conciliation or remaining steadfast is the appropriate epistemic response in a given disagreement. If we have reasons to believe that a dissenting party is far more likely to be correct about the proposition as we are, then we therefore have reasons to believe that conciliation is the correct response in this disagreement (that is, we should substantially revise our beliefs as to be in line with the dissenting party). If however we have reasons to believe that a dissenting party is far more likely to be incorrect about the proposition, we therefore have reasons to believe that steadfastness is the correct response in this disagreement (that is, we should remain confident in our beliefs in the face of this disagreement). Here, it is important to note that the Total Evidence View is more nuanced than simply recommending a conciliatory approach to epistemic superiors and a steadfast approach to epistemic inferiors. Matheson notes that in actual cases of disagreement, even if a dissenting party is an epistemic inferior, one might still be required to seek conciliation and revise one's beliefs. To explain, Matheson gives the example of two thermometers where though one thermometer is slightly more reliable than the other, both are usually reliable enough for approximate temperature.[112] Suppose that at an earlier time, we read that the temperature is 30°C on the more reliable thermometer while, at a later time, we read that the temperature is 26°C on the less reliable thermometer. Even if we believe that one thermometer is more reliable than the other, the reading on the less reliable thermometer still provides us with *some* evidence for believing that the temperature is no longer 30°C. As such, in evaluating the correct response to disagreement on the basis of the total evidence that one has, the Total Evidence View provides a fine-grained analysis which reveals how the relative epistemic position of a dissenting party (e.g. taking a gradient approach of seeing a dissenting party as epistemically superior or inferior, as opposed to a discrete approach of

---

[112] Matheson, "Disagreement," 322–23.

seeing a dissenting party as an epistemic superior or an epistemic inferior) as determining the degree to which one should revise their beliefs in a given disagreement.

Third, outside of our evidence about the epistemic peerhood (or epistemic superiority or inferiority) of a dissenting party, the Total Evidence View states that the correct response to a given disagreement is also decided on the basis of other evidence that is available to us. To illustrate, consider the following example which is taken from Kelly.[113] Suppose you are having a conversation with a friend of yours. Your friend is someone you met in university, and someone you have become close to from your many discussions about various topics. As a result of these discussions, you have come to greatly admire your friend's intellect and believe them to be your epistemic peer, if not epistemic superior. And yet, over the course of the conversation, you discover that your friend is a Holocaust denier. Obviously, in this scenario, despite your past experiences which cause you to believe that your friend is an epistemic peer (and at times, an epistemic superior), this is not enough evidence for you to believe that the Holocaust did not occur. This is because under the *total* evidence that you have, the overwhelming evidence for the occurrence of the Holocaust outweighs the opinions of your friend. As such, in this case, the Total Evidence View shows how the fact that the dissenting party is an epistemic peer does not necessarily require you to question the evidence you have for your beliefs – rather, in some cases, the evidence for your beliefs might sometimes cause you to *re-evaluate whether a dissenting party is an epistemic peer*.[114] As such, the Total Evidence View shows how the epistemic peerhood of the dissenting party is but part of the total evidence one has when considering the correct response to a disagreement – in actual cases

---

[113] Thomas Kelly, "Disagreement and the Burdens of Judgment," in *The Epistemology of Disagreement: New Essays*, ed. David Christensen and Jennifer Lackey (Oxford: Oxford University Press, 2013), 40.

[114] Lackey makes a similar point when she discusses cases where one's antecedent evidence is enough for one to discount the opinion of a dissenting party. See Jennifer Lackey, "What Should We Do When We Disagree?" in *Oxford Studies in Epistemology*, ed. Tamar Szabó Gendler and John Hawthorne (Oxford: Oxford University Press, 2010).

of disagreement, one's expectation that a dissenting party will be an epistemic peer might be disappointed, such that one rescinds their judgment of epistemic peerhood.[115]

To conclude, the Total Evidence View provides a general non-ideal approach to evaluating the correct epistemic response to disagreement. While the Total Evidence View was originally proposed by Kelly as simply another theoretical approach to ideal peer disagreement, it has since become the dominant approach in analysing epistemic disagreements in general. This can be seen in how much of the more recent discussion on the epistemology of disagreement has moved away from discussing full-blooded conciliation or steadfast views on peer disagreement, and instead opted to analyse the epistemic problem of disagreement as a problem of higher-order evidence following the Total Evidence View. Feldman explains how the Total Evidence View has since made the debate between conciliation and steadfastness obsolete by providing a general *evidentialist* solution to the epistemic problem of disagreement:

> To the extent that there is a general evidentialist answer to questions about what you should do in response to learning of peer disagreement, it is this: follow your evidence. There is no reason to abandon evidentialism simply because there are cases in which seemingly reasonable peers disagree. In fact, evidentialism seems to me to provide exactly the right way to think about disagreement. It instructs us to ask how learning about a disagreeing peer affects one's evidential situation. It asks us to reflect on what one should think, now that one has this new information about the disagreement…The literature on disagreement contains extensive discussion of principles that say such things as that when confronted with a disagreeing peer, one should always suspend judgment, or split the difference, or stick to one's guns. But I think that if evidentialism is correct, then it is almost surely true that there are no special epistemic principles about disagreement. If there were any such special epistemic principles about what is justified in cases of disagreement, then either they undermine evidentialism (because the justified attitude is not always the one supported by the evidence) or they are implications of evidentialism (because the evidential impact of the evidence one gets from a disagreeing peer always yields the outcome specified by the principle).[116]

---

[115] For a similar line of thought, see Kelly, "Peer Disagreement and Higher Order Evidence," 198–201.

[116] Feldman, "Evidence of Evidence Is Evidence," 287.

We are now in a position to summarise how Bayesian evidentialism has influenced the analytic epistemology of disagreement. As a brief note, it is clear that Bayesian epistemology has also influenced both ideal and non-ideal theorising about the epistemology of disagreement. First, the debate between Uniqueness and Permissivism has often been discussed in tandem with a debate in Bayesian epistemology about objectivism and subjectivism about evidential support. This is because the question of whether there is an objective evidential support relation which determines the degree to which a body of evidence supports a hypothesis has clear implications to the Uniqueness Thesis and its claim on the connection between evidence and epistemic rationality.[117]

Next, Bayesian epistemology is often used in discussion about conciliatory and steadfast approaches to disagreement, and in particular, a form of conciliation known as the Equal Weight View. The Equal Weight View is the view that the correct response to peer disagreement is to revise one's belief such that one gives equal weight to peer opinions. One way to give equal weight to each peer opinion is the "splitting the difference" method, which states that each peer should revise their own beliefs such that they assign the same credence for their own hypothesis as they do to their peers.[118] It is clear that both the Equal Weight View and the "splitting the difference" method is plausible (and arguably, intelligible) only if one accepts the Bayesian claim that there is such a doxastic attitude as credences. If one were only to take into account "full" doxastic attitudes such as all-out belief, disbelief, and

---

[117] Hedden argues that objectivism about evidential support is at least a necessary condition to the Uniqueness Thesis. See Brian Hedden, "A Defense of Objectivism about Evidential Support," *Canadian Journal of Philosophy* 45, no. 5 (2015): 717.

[118] The "splitting the difference" method is typically attributed to Christensen - see Christensen, "Epistemology of Disagreement," 193. For an in-depth study of Bayesian formulations of "splitting the difference", see David Jehle and Branden Fitelson, "What Is the "Equal Weight View"?" *Episteme* 6, no. 3 (2009): 280–93.

suspension of belief, it is hard to see what splitting the difference would look like in a disagreement –

for example, between a theist and an atheist.[119]

Finally, Bayesian epistemology is used by Matheson to precisify his non-ideal approach to finding the

correct epistemic response to actual disagreements. Matheson's strategy involves using the Equal

Weight View as a starting point for belief revision: namely, that one should adopt the "splitting the

difference" method to revise their beliefs if one has reasons to believe that a dissenting party is an

epistemic peer. From here, the "splitting the difference" method is adjusted according to whether one

has reasons to believe that a dissenting party is an epistemic superior or inferior: given the extent to

which a dissenting party is an epistemic superior or inferior, one should accordingly give more or less

weight to the dissenting party's credence when revising their beliefs.[120] While a more comprehensive

examination of the use of Bayesianism within analytic epistemology of disagreement (for example, the

many probabilistic formulae which have been proposed as norms for belief revision upon

encountering disagreement) has been outside of the scope of this thesis, it is clear that the theoretical

role of evidence in resolving disagreement has often been cashed out in Bayesian terms: that is, in

requiring belief revision according to Bayesian conditionalisation.

In conclusion, within the analytic epistemology of disagreement, the role of evidence in resolving

disagreement has been theoretically explained through an evidentialist response to the epistemic

problem of disagreement: that is, in the claim that the correct response to disagreement is simply to act

in accordance to the evidence that is available to you in a given moment. In this way, the epistemic

challenge which disagreement poses is no different to any other epistemic situation: disagreement is

---

[119] Kelly makes this point in Kelly, "Peer Disagreement and Higher Order Evidence," 117.

[120] Matheson, "Disagreement," 320–27.

but another situation where one acquires new evidence, and must revise their beliefs accordingly. It is this conclusion which I shall put into question in the next chapter. Is it really the case that disagreement is 'no different to any other epistemic situation', or that it is 'but another situation where one acquires new evidence'? In the following critique of the evidentialist response to disagreement, the answer to both these questions is decisively no: the epistemic challenge which disagreement poses is unique, because disagreements can undermine our ability to use or follow evidence by placing us in an epistemic situation where we do not know what constitutes evidence or rationality in our given context (See §§4.4-4.5 for more details). We now turn to the next chapter, where we will revisit the topic of evidentialism and its need for a substantial account of the concept of evidence.

3   SUBSTANTIATING THE EVIDENTIALIST RESPONSE TO DISAGREEMENT: TWO THEORETICAL DESIDERATA

3.1   *Introduction: On the Need for a Substantive Theory of Evidence*

In the previous chapter, we examined how evidentialism has been used within the analytic epistemology of disagreement to provide a general non-ideal theory to the practical question of disagreement. The story, however, is far from over: in order for evidentialism to work as a general response to disagreement (that is, in order for the correct epistemic response to disagreement to be to follow one's evidence), more needs to be said about what evidentialists mean exactly by "evidence". This is because without some substantive theory about what evidence is, it is difficult to know how to apply the evidentialist norm to "follow one's evidence", or how to assess whether this norm is the correct epistemic response to disagreement.

The need for a substantial theory of evidence is something which is brought up throughout the various critical discussions about evidentialism: notably, critics of evidentialism have argued that if evidentialism is simply taken as the norm that one ought to believe according to the evidence, or as the claim that evidence is that which justifies belief, then evidentialism would be trivially true due to the generality of these claims.[121] To elaborate on this criticism, consider a similar point that is sometimes made about the correspondence theory of truth: that is, the claim that truth is correspondence to

---

[121] See, for example, Dougherty, "Introduction". In a similar fashion, McCain argues that without substantive claims, evidentialism would define either a family of different theories on epistemic justification, or otherwise only provide a formal schema for a theory of epistemic justification. For more information, see Kevin McCain, "Evidentialism: A Primer," in *Believing in Accordance with the Evidence: New Essays on Evidentialism*, ed. Kevin McCain, Synthese Library (Springer Cham, 2018), 1–2; and Chapter 1 of Kevin McCain, *Evidentialism and Epistemic Justification*, New York (Routledge, 2014).

facts.[122] The correspondence theory of truth is sometimes explained in reference to Aristotle's claim

that a statement is true if it says of what is that it is, and of what is not that it is not.[123] However, taken

by itself, it is clear that Aristotle's claim does not provide a correspondence theory of truth but states

an empty platitude which is trivially true for all accounts of truth. As such, in order for the

correspondence theory of truth to be a *distinctive* theory on truth, more needs to be said substantively

about the concepts of correspondence and fact so as to make the claim that truth is correspondence to

facts one which has actual philosophical implications. In the same way, without a theoretical

explication of what we mean by evidence, evidentialism risks conflating the notion of evidence and

epistemic justification such that the claim that one's epistemic justification supervenes on evidence is

trivially true on any account of epistemic justification. This is especially important for the evidentialist

response to disagreement, because if we consider evidentialism as simply the claim that evidence is

that which justifies belief (see Kelly's definition of evidence in §2.1), then the evidentialist response to

disagreement reduces down to the norm "follow whatever you're justified in believing" – a norm

which is too general to be providing any epistemic guidance for how we should respond to

disagreement.

To address this problem of generality, proponents of evidentialism have given various theoretical

explications on the notion of evidence so as to develop evidentialism as a substantial account of

epistemic rationality. Let us consider two of the main theoretical developments within evidentialism –

---

[122] The following points are made in the entry on the correspondence theory of truth in the Oxford Dictionary of Philosophy. See Simon Blackburn, "Correspondence Theory of Truth," in *The Oxford Dictionary of Philosophy*, Second (Oxford: Oxford University Press, 2005), 81.

[123] This claim is taken from Aristotle's *Metaphysics* Γ, iv. 1011. See W. D. Ross, *Aristotle's Metaphysics, a Revised Text with Introduction and Commentary* (Oxford University Press, 1924).

namely, in providing substantial theories of evidence and of evidential support.[124] First, developments in evidentialism have substantiated the notion of evidence by giving different theoretical accounts about the *ontology* of evidence – that is, theories about what constitutes evidence, and about what should be admitted within one's body of evidence. Second, evidentialists have also expounded on the kind of epistemic justification which evidence gives through accounts about *evidential support* – that is, theories which explain how a body of evidence supports a hypothesis, and the nature of this evidential support relation (e.g. whether it is subjectively or objectively determined).

Here, however, it is important to note that there is no consensus amongst evidentialists about what the correct theory of evidence and evidential support is. As such, in order to assess the evidentialist response to disagreement, we must first decide on the theory of evidence and evidential support which best fits the intuitions behind the evidentialist response to disagreement. To introduce the kind of assessment we will be using to find the correct theory of evidence for the evidentialist response to disagreement, consider how different theories of evidence and evidential support may lead to different conclusions about what it means to "follow one's evidence" in a given disagreement. In the case of differing conceptions of evidence, various philosophers have noted that the plausibility of the

---

[124] While this introduction considers only two theoretical components to evidentialism, there are many other issues which are discussed within the literature – for example, questions regarding when a person is said to *possess* evidence, when a person should *gather* evidence, and explanationist accounts of evidentialism which consider the connection between evidence and explanation. For more information, see Earl Conee and Richard Feldman, "Evidence," in *Epistemology: New Essays*, ed. Quentin Smith (Oxford: Oxford University Press, 2008); Thomas Kelly, "Evidence: Fundamental Concepts and the Phenomenal Conception," *Philosophy Compass* 3, no. 5 (2008): 939–41; McCain, *Evidentialism and Epistemic Justification*; McCain, "Evidentialism"; and Dougherty, "Introduction".

Uniqueness Thesis depends greatly upon one's conception of evidence.[125] For example, if one takes a

physical conception of evidence like that used in law courts (e.g. of weapons, photographs, invoices,

and genetic evidence), then the Uniqueness Thesis would plausibly be false. This is because there are

often many ways that one could reasonably interpret physical objects such as photographs, and thus, it

is reasonable to conclude that multiple propositions may be justified according to the same body of

evidence. Conversely, if one was to take a strong mentalistic notion of evidence such that a body of

evidence includes an epistemic agent's mental states, the Uniqueness Thesis would be vacuously true.

This is because the moment two people come to different conclusions about the evidence (that is, one

believes P, while the other believes not-P), they would have different mental states and thus different

"evidence".

The impact that different conceptions of evidence have on the plausibility of the Uniqueness Thesis

might also have consequences as to what evidentialism might recommend as the correct epistemic

response to a given disagreement. For example, if the mentalistic conception of evidence is the correct

theory of evidence for evidentialism, then it is reasonable to assume that steadfastness is the standard

and preferred epistemic response to disagreement. Within the literature on the steadfast response to

disagreement, many have defended the steadfast response to disagreement on the basis of privileging

one's own private experiences over the experiences of others.[126] For example, Peter van Inwagen has

---

[125] The following argument is given by Kopec and Titelbaum, "The Uniqueness Thesis," 191. Kopec and Titelbaum cite
Ballantyne and Coffman as giving an argument for how the truth of the Uniqueness Thesis restricts the options one has for
conceptions of evidence. See Ballantyne and Coffman, "Uniqueness, Evidence, and Rationality".

[126] Frances and Matheson notes this in their overview of the Steadfast View – see Frances and Matheson, "Disagreement",
§5.2. Frances and Matheson also note that while van Inwagen uses private evidence to defend the Steadfast View, Feldman
sees the idea of private evidence as supporting conciliation. This is because the fact that a dissenting party has private
intuitions and experiences might conversely give you reasons to believe that a dissenting party might have some critical piece
of evidence that you lack. See Feldman, "Reasonable Religious Disagreements," 207–8.

argued that one can remain steadfast in their beliefs in the face of peer disagreement by appealing to

the private experiences one has over the dissenting party. Given that one can never fully share their

intuitions and experiences to a dissenting party, one has reasons to believe that a dissenting party

might lack some critical evidence that one has. In a similar line of thought, many philosophers have

also defended the steadfast response by appealing to self-trust.[127] While peer disagreements exhibit an

evidential symmetry from a third-person perspective, it is clear that each dissenting party comes into

any disagreement from a first-person perspective. Given that this first-person perspective gives one

more access to their own evidence than that of others, self-trust is justified due to the fact that one has

more reasons to believe that one has reliably weighed one's evidence than another.[128] As such, under a

mentalistic conception of evidence, the privileged access one has to one's own evidence over the

evidence of a dissenting party gives reason to assume that, *ceterus paribus*, steadfastness is the correct

epistemic response to disagreement. To a lesser extent, if the physical conception of evidence is the

correct theory of evidence for evidentialism, then this might give one reasons to believe that

conciliation is the correct epistemic response to disagreement. This is because a physical conception of

evidence would make it more likely that the same evidence condition is satisfied in a given

disagreement: it is quite possible that within a given disagreement, dissenting parties would have the

same physical objects accessible to them. Insofar as the satisfaction of the same evidence condition

gives you reasons to believe that a dissenting party is an epistemic peer, the satisfaction of the same

evidence condition might therefore give you reasons to believe that conciliation is the correct

epistemic response to disagreement.

---

[127] For more information, see the list of references in fn. 103.

[128] Lackey makes a similar point in Lackey, "What Should We Do When We Disagree?".

As a brief sidenote, while the following discussion will focus specifically on conceptions of evidence, it is also clear that different conceptions of evidential support will also affect what epistemic response will be preferred under evidentialism. As we noted in the previous discussion, the debate between subjectivism and objectivism about evidential support has implications for the debate between Uniqueness and Permissivism. This is because if subjectivism about evidential support is true (that is, if the degree to which a body of evidence supports a hypothesis is relative to some subjective factors, such as a set of evidential standards), then this entails that Uniqueness is false. Accordingly, the discussion around subjectivism and objectivism about evidential support will therefore also have implications for what epistemic response will be preferred under evidentialism. For instance, if subjectivism about evidential support is true, then this makes it more likely for steadfastness to be the correct epistemic response to a given disagreement. This is because, if one has reasons to believe that one's beliefs are supported by their evidence according to an acceptable set of evidential standards, then disagreement from an epistemic peer need not be a reason to re-evaluate the justification for one's beliefs. After all, it is possible for a dissenting peer to have the same evidence as you, but come to different conclusions due to them adopting a different (but equally acceptable) set of evidential standards. As such, since it is possible for dissenting parties to have the same evidence but come to different conclusions - as subjectivism about evidential support allows for the possibility that two parties come to differing conclusions due to a difference in choice of evidential standards - then it is justifiable to remain steadfast in one's belief in the face of disagreement. With this being said, a more detailed examination as to how differing conceptions of evidential support affects the evidentialist response to disagreement will be outside the scope of this essay.

This brings me to my main critique of the evidentialist response of disagreement. The structure of my argument over this and the following chapter is as follows: in §3.2, I will outline the two theoretical desiderata that are required within a theory of evidence in order for the evidentialist response to

disagreement to be a valid epistemic response to disagreement: namely, that the theory of evidence must account for how evidence determines epistemic rationality, and the theory of evidence must explain how evidence is commonable between epistemic agents. I shall explain how these two theoretical desiderata are needed in order to explain how we can use evidence to justify our doxastic attitudes, and to arbitrate between the competing claims of dissenting parties. This brings us to the next chapter, where I will argue that no theory of evidence can satisfy the evidentialist response to disagreement. To show this, in §§4.1-4.3, I shall examine three of the dominant philosophical theories of evidence under these two theoretical desiderata: namely, the empiricist theory of evidence as sense-data, Earl Conee and Richard Feldman's theory of evidence as occurrent mental states, and Timothy Williamson's theory of evidence as known propositions. I shall explain how each of these theories fail to satisfy our two theoretical desiderata, and are therefore unsuitable for the evidentialist response to disagreement.

While the critical examination of these three theories does not prove that no theory of evidence could be suitable for the evidentialist response to disagreement, in §4.4, I shall argue that the common *way* in which they fail to meet the theoretical desiderata highlight the flaws within the evidentialist response to disagreement. In particular, it is the strong evidentialist claim (as exemplified in the Uniqueness Thesis) that evidence determines epistemic rationality which forces one to make unrealistic assumptions about the role of evidence within our everyday epistemic practice. In defining epistemic rationality as a direct connection between one's body of evidence and their doxastic attitudes, the Uniqueness Thesis fails to account for the role of epistemic agents (individually, as well as other dissenting parties) as an important factor within the interplay between one's body of evidence and rational belief formation. To conclude then, in §4.5, I shall briefly sketch how the pragmatist theory of inquiry provides a different non-ideal account of evidence as it is evaluated and discovered within the context of inquiry. In particular, I will draw from Hilary Putnam's critique of Carnap's inductive logic

to explain how the evaluation of evidence within scientific inquiry requires epistemic norms which are contextually and socially interpreted within a community of inquirers. This reveals how the evidentialist norm to "follow your evidence" itself requires a wider context of inquiry, one which involves other epistemic agents which inform and explain how we are to follow evidence within a given context. To explain what this wider context of inquiry entails brings us to the final chapter of this thesis.

3.2   Two Theoretical Desiderata for Evidence: Evidence Determines Epistemic Rationality and Evidence is Commonable

The strategy we will be using to find the correct theory of evidence for the evidentialist response to disagreement is a form of reverse engineering: seeing as we have a general non-ideal epistemology of disagreement in the Total Evidence View, we can consider how evidence is supposed to function within the Total Evidence View, and work backwards in order to find the theory of evidence where evidence does meet those functions. Given our discussion on the Total Evidence View, the first desiderata that we intuitively require from a theory of evidence is that our evidence should determine epistemic rationality. Let us examine what this desiderata entails, and then explain why such a desiderata might be required for the evidentialist response of disagreement.

To begin, we can clarify this desiderata by making a distinction between epistemic justification and epistemic rationality. In particular, the contemporary epistemic literature has typically distinguished between the concept of justification and rationality by making a distinction between propositional justification and doxastic justification.[129] Whereas propositional justification refers to the justification one has about a proposition, doxastic justification refers to the justification one has for a particular

---

[129] See Dougherty, "Introduction", 11ff.

doxastic attitude about a proposition. This distinction is important for distinguishing between justification and rationality, because rationality requires that an epistemic agent has not only propositional justification, but doxastic justification for a proposition. To illustrate, consider the following example. Suppose you have evidence which gives you propositional justification for the proposition "I will pass this exam". Despite this justification, you might still be irrational in believing that you will pass the exam if you choose to believe this proposition for reasons *other* than your evidence. For example, suppose you believe (irrationally) that positive thinking always manifests results into reality, and come to believe that you will pass the exam on the basis of this assumption rather than on your evidence. In this scenario, while your belief in the proposition "I will pass this exam" is propositionally justified, it is not doxastically justified because you have not adopted your doxastic attitude about the proposition for the *right reasons*.

As such, our first theoretical desiderata states that our theory of evidence must provide an account of evidence which shows how evidence provides not only propositional, but doxastic justification towards propositions. To connect this to our previous discussion on Uniqueness, we can say that our first theoretical desiderata requires not only the weaker claim of Propositional Uniqueness (i.e. that evidence justifies at most one proposition in a set of competing propositions), but also the stronger claim of Personal Uniqueness (i.e. that evidence justifies at most one doxastic attitude towards a given proposition). At this point, it should be clear that this desiderata is a substantive constraint on our theory of evidence, because it already precludes some conceptions of evidence from being the correct theory for the evidentialist response to disagreement. To take an example from the previous section, consider the physicalist conception of evidence which takes evidence to be the set of physical objects which are admitted as evidence within a law court. The physicalist conception of evidence does not meet our first theoretical desiderata because while physical objects provide propositional justification, they do not provide doxastic justification. Suppose, for instance, you are a member of a jury for a legal

case in which the evidence consists of DNA samples which justify the proposition "Simon is the criminal". In this scenario, it might still be reasonable for you to suspend judgment or even disbelieve this proposition because you might, for instance, doubt the reliability of DNA sampling and therefore reject the samples as a reason to believe the proposition. As such, while the DNA samples does provide propositional justification (in justifying only one hypothesis about the criminal in a set of competing hypotheses), it does not provide doxastic justification. On the other hand, if we accepted a mentalistic conception of evidence, such that one's evidence also included intuitions about the reliability of DNA samples, then one's evidence provides not only propositional justification, but doxastic justification.

Why does the evidentialist response to disagreement require that evidence determine not only epistemic justification, but also epistemic rationality? Recall the previous quote from Feldman about the evidentialist response to disagreement. Feldman argues that if evidentialism is the correct theoretical approach to disagreement, then

> it is almost surely true that there are no special epistemic principles about disagreement. If there were any such special epistemic principles about what is justified in cases of disagreement, then either they undermine evidentialism (because the justified attitude is not always the one supported by the evidence) or they are implications of evidentialism (because the evidential impact of the evidence one gets from a disagreeing peer always yields the outcome specified by the principle).[130]

In other words, in order for evidentialism to provide a general non-ideal epistemology of disagreement, there cannot be any other reasons for one's belief about a proposition other than *evidential* reasons. If there are other reasons for believing a proposition outside of evidential reasons, then the evidentialist response to disagreement fails as there is no guarantee that following one's evidence is always the correct epistemic response to a disagreement. As such, the evidentialist response

---

[130] Feldman, "Evidence of Evidence Is Evidence," 287.

to disagreement requires that evidence determine not only the justification for one's belief, but the rationality of one's doxastic attitudes.

Furthermore, Roger White provides a powerful argument for why evidence should determine epistemic rationality, in claiming that our evidence needs to determine epistemic rationality in order for us to have any reason to follow the norms of rationality in the first place.[131] While White's argument is raised against Permissivism, his general point can also be used to motivate our proposed desideratum for a theory of evidence. White's argument is built upon the following thought experiment: suppose you are participating within a jury for a legal case. After you receive the relevant evidence from the testimonies and cross-examinations of the trial, you are tasked with making an evaluation over whether the suspect is guilty. However, in this particular trial, you are made aware that your trial is in fact a Permissive case: that is, your trial is such that your available evidence makes it rationally permissible for you to believe either that the suspect is guilty or not.

Now, suppose you are given the option between choosing to believe that the suspect is guilty on the basis of the evidence, or choosing to believe that the suspect is guilty on the basis of a coin flip where heads means the suspect is guilty. White argues that if this trial really is a Permissive case, then the choice between believing that the suspect is guilty on the basis of the evidence and believing on the

---

[131] In particular, I will be focusing on White's first thought experiment in the paper. While White's thought experiment involves the idea of belief-inducing pills, I am going to modify the thought experiment as to make the situation more plausible. See White, "Evidence Cannot Be Permissive," 315. Ballantyne and Coffman cites another passage from Feldman as supporting White's general argument. See Ballantyne and Coffman, "Uniqueness, Evidence, and Rationality," 3; Richard Feldman, "Epistemological Puzzles about Disagreement," in *Epistemology Futures*, ed. Stephen Hetherington (Oxford: Oxford University Press, 2006), 226.

basis of a coin flip is an *arbitrary* choice.[132] In an ordinary non-Permissive case, the reason why we would choose to believe that the suspect is guilty on the basis of evidence rather than by a coin flip is because believing on the basis of our evidence should make us more likely to be correct in our beliefs. However, in a Permissive case, we are given no such reason: if our evidence allows both verdicts to be rationally permissible for belief, then our evidence provides us no indication that one verdict is better than the other. As such, according to our evidence, the choice between believing whether the suspect is guilty or not is *arbitrary*. Given that the presupposition of a Permissive case leads to the absurd conclusion that believing according to one's evidence is as good as believing on the basis of non-epistemic reasons (e.g. a coin flip), White concludes that evidence cannot be permissive.

Of course, it is important to note that White's example of the Permissive case is clearly a case of Extreme Permissivism.[133] However, even if one argues that White's argument would not apply to Moderate Permissivism, his general point is sufficient to support our proposed desiderata that a theory of evidence should explain how evidence determines epistemic rationality. White's argument is built upon "an apparent conflict between rational belief and arbitrariness" such that, if one doesn't form their beliefs according to the norms of rationality, then one's choice of what to believe is ultimately an (epistemically) arbitrary choice.[134] For White, however, what makes rational belief distinct and

---

[132] By "choice", I am not making any claim about whether an epistemic agent can *choose* their belief. The point of this thought experiment is not to make any claim about doxastic voluntarism, but to highlight how Permissivism would imply that evidence is epistemically irrelevant to one's belief revision.

[133] White himself prefaces his paper by acknowledging this fact. However, there are reasons to believe that White's target is not just Extreme Permissivism, but Moderate Permissivism as well. This is supported by White himself, who confidently claims that cases such as this thought experiment can be adapted to apply to Moderate Permissivism as well. See White, "Evidence Cannot Be Permissive," 313.

[134] Ballantyne and Coffman, "Uniqueness, Evidence, and Rationality," 3.

preferable over irrational belief is the notion of evidence, because it is evidence that makes rational

beliefs more likely to be correct than irrational beliefs. The upshot of White's argument therefore

reveals how one of the roles which our concept of evidence plays in our epistemic practices is in

explaining the normative force of epistemic rationality: in being a reliable indicator of truth, evidence

explains why we should choose to believe rationally since believing according to one's evidence will

make us more likely to be correct than if we didn't. This strongly supports our desiderata that a theory

of evidence should explain how evidence determines epistemic rationality – if our theory of evidence

shows how believing on the basis of evidential reasons will lead us to the truth more reliably, then this

supports the evidentialist claim that following one's evidence is always the correct epistemic response

to disagreement.

This brings us to the second theoretical desiderata for our theory of evidence: namely, that our

evidence should be *commonable*. More precisely, what we require from a theory of evidence is an

explanation of how the epistemic justification which evidence gives to our beliefs can be transmitted

between epistemic agents such that we can have a *common* body of evidence. Before we go on to

explain why this is a crucial desiderata for the evidentialist response to disagreement, let us first

consider how and why this desiderata is usually neglected within the analytic epistemology of

disagreement. One reason why the commonability of evidence has not been widely discussed within

the epistemology of disagreement is because of the general association of evidentialism with

internalism about epistemic justification. Because evidentialism claims that the epistemic justification

for one's doxastic attitudes supervenes on one's evidence – and most evidentialists are also mentalists

about evidence (such that one's evidence consists of one's mental states, see §4.2 for more details) –

evidentialism has typically characterised evidence as something that is possessed by an *individual*

epistemic agent rather than a collective group of epistemic agents. At the same time, while the

epistemology of disagreement is often classified as a problem within analytic *social* epistemology, it is

clear that the recent literature has treated disagreement almost exclusively as an individual rather than

social epistemological issue.[135] Unlike other topics within social epistemology – such as the problem of

belief aggregation within social choice theory – the literature on the epistemology of disagreement has

theorised from the perspective of an individual agent, such that the epistemic problem of disagreement

is explained in terms of its effect on the justification of an individual agent's beliefs, and the theoretical

responses to disagreement are explicated as norms for individual epistemic agents, rather than for an

epistemic community.[136] As such, in both the discussions around evidentialism and the epistemology

---

[135] Across his many overviews and introductions for analytic social epistemology, Alvin Goldman has consistently distinguished between three kinds of social epistemology: between social epistemic questions which involve individual epistemic agents (known as interpersonal SE), social epistemic questions a collective group of epistemic agents (known as collective SE), and social epistemic questions regarding social epistemic institutions (known as institutional SE). Under these categories, Goldman sees the literature on the epistemology of testimony and disagreement as part of interpersonal SE. However, one way of explaining the difference between the evidentialist response to disagreement and the pragmatist response to disagreement which we will consider in Chapter 5 is that the literature on pragmatist democratic theory considers disagreement to be a problem for collective and institutional SE, rather than interpersonal SE. For more on Goldman's taxonomy of social epistemology, see Alvin I. Goldman, "A Guide to Social Epistemology," in *Social Epistemology: Essential Readings*, ed. Alvin I. Goldman and Dennis Whitcomb (Oxford University Press, 2011), 11–37; Alvin I. Goldman and Cailin O'Connor, "Social Epistemology," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Winter 2021 (Metaphysics Research Lab, Stanford University, 2021); Goldman, "The What, Why, and How of Social Epistemology".

[136] While the topic of belief aggregation does make reference to individual epistemic agents (i.e. in aggregating the beliefs of *individual* epistemic agents), it is clear that the focus of this literature is on a collective group of epistemic agents: whether it is possible to attribute notions of judgment and belief to collective groups of epistemic agents, and if so, how we represent or measure collective judgments about particular issues. For more on judgment aggregation in social epistemology, see Christian List, "Group Knowledge and Group Rationality: A Judgment Aggregation Perspective," in *Social Epistemology: Essential Readings*, ed. Alvin I. Goldman and Dennis Whitcomb (Oxford University Press, 2011), 221–41; Philip Pettit, "Groups with Minds of Their Own," in *Social Epistemology: Essential Readings*, ed. Alvin I. Goldman and Dennis Whitcomb (Oxford University Press, 2011), 242–68. For introductions on how judgment aggregation is discussed more formally within computational social choice, see Ulle Endriss, "Judgment Aggregation," in *Handbook of Computational Social Choice*, ed. Felix Brandt et al. (Cambridge University Press, 2016), 399–426; Davide Grossi and Gabriella Pigozzi, *Judgment Aggregation: A Primer* (San Rafael, CA: Morgan & Claypool, 2014).

of disagreement, the importance of the commonability of evidence has been neglected because on the focus on individual epistemic agents, rather than collective groups of epistemic agents.

And yet, the commonability of evidence is a necessary theoretical desiderata for the evidentialist response to disagreement because in order for our evidence to determine epistemic rationality, our evidence must be commonable across epistemic agents. To explain, recall how one of the key implications of the Uniqueness Thesis that is discussed within the literature is the idea that two epistemic agents cannot disagree upon the same body of evidence and be equally rational. This intuition is usually highlighted with the oft-discussed Jury Case.[137] Consider the following retelling: suppose you and your friend have both been asked to be part of a jury for a court case. Your friend is someone you met in university, and someone you have become close to from your many discussions about various topics. As a result of these discussions, you have come to greatly admire your friend's intellect and believe them to be equally as rational as you are. As you and your friend participate within the trials, you are both presented with the same testimonies and cross-examinations such that you both receive the same body of evidence with respect to the trial. And yet, at the end of the trial, you and your friend come to different conclusions – whereas your friend believes the suspect to be guilty, you believe them to be innocent.

The Jury Case is often used as a thought experiment to reveal one's intuitions about the debate between Uniqueness and Permissivism: whereas the Permissivist would argue that it is possible for you and your friend to be equally rational in your respective verdicts, the Uniqueness theorist would argue that it is not. Here, however, it is important to notice that a key reason why the Jury Case is an effective thought experiment is because of an assumption about the commonability of evidence: that is, it is

_____

[137] For more about the Jury Case, see Kopec and Titelbaum, "The Uniqueness Thesis," 189.

because you and your friend have a *common* body of evidence as is relevant for the trial which allows us to examine whether it is rational for you and your friend to come to different verdicts. The Jury Case therefore reveals how a key reason why the Uniqueness Thesis is important to the epistemology of disagreement to begin with is because our intuitions suggest that evidence does not only determine what is epistemically rational for individual epistemic agents to believe, but also for collective groups of epistemic agents (e.g. juries).

Thomas Kelly makes a similar point about the important social function of evidence in explaining how one of the key roles which evidence plays in our everyday epistemic practice is as a neutral arbiter between the competing claims of an epistemic community.[138] To take an example, consider the use of evidence within scientific inquiry. Within scientific inquiry, it is clear that the role which evidence plays within scientific research is not only as justification for the beliefs of an individual scientist, but also as a way of settling what hypotheses are justified within the scientific *community*. This means that once an experiment has provided evidence in favour for a particular hypothesis or theory, the findings of that experiment provides justification not only for the scientist who conducts the experiment, but justification which other scientists receive and use to conduct further research.[139] In fact, the commonability of scientific evidence applies not only to scientists within academic research, but also for the general public. Even without first-hand scientific research, the commonability of scientific evidence means that we as epistemic agents are not only justified, but rationally *obligated* to believe modern Western medicine over ancient bloodletting practices, atomic theory over the ancient theory of the four elements, and that the earth revolves around the sun rather than that the sun revolves

---

[138] See §4 of Kelly, "Evidence," 2016.

[139] John Hardwig astutely points this out when examining the role of trust in knowledge – see Hardwig, "The Role of Trust in Knowledge".

around the earth.[140] Here, it is important to see that the mention of "rational obligation" here is meant

to indicate that the scientific evidence provides not only propositional justification for an epistemic

community, but doxastic justification. To take an example, consider a conspiracy theorist who applies

to become a homeschooling parent, and then chooses to ignore the syllabus and instead teach their

child that the sun revolves around the earth, and that bloodletting is more effective than Western

medicine. Here, our intuition that such a parent is not epistemically entitled to teaching their child in

this way is because the overwhelming amount of scientific evidence provides not only propositional

justification about what the correct scientific or medical theory is, but also doxastic justification about

what the parent and child should believe about medicine and the physical world. And the reason why

the scientific evidence rationally constrains the beliefs of the parent (as well as the child) is because

scientific evidence is commonable: even if the parent did not have first-hand experience in practising

medicine or conducting scientific experiment, the commonability of scientific evidence is such that to

learn of these scientific facts is sufficient to determine what the parent *ought* to believe about these

topics.

We are now in a position to see why the commonability of evidence is an important desiderata for the

evidentialist response to disagreement. To defend the commonability of evidence more directly,

consider how White's argument for why our evidence must determine epistemic rationality applies not

only to what is epistemically rational for an individual epistemic agent, but also what is objectively

rational amongst a group of epistemic agents. Once we apply White's argument to the issue of

---

[140] The point here is not to claim that scientific evidence is infallible, but to highlight the social epistemic function of evidence in rationally constraining the beliefs of a collective group of epistemic agents. When and to what extent should scientific evidence rationally constrain our beliefs is an open question: one which I shall argue must be answered in discussion with other epistemic agents within collective inquiry. See §4.5 and the rest of Chapter 5 of this thesis.

disagreement, we can see how White's criticism about the arbitrariness of choosing what to personally believe in Permissive cases also poses a *social* epistemic dilemma in the context of disagreement. As we have seen before, according to the standard view in the literature on Uniqueness and Permissivism, Permissivism has often been seen to undercut the epistemic significance of disagreement due to the fact that Permissivism implies that it is possible for dissenting parties to have different beliefs about the evidence and still be equally rational. Accordingly, if Permissivism is true, then encountering a dissenting peer does not provide you any reason to believe that you have made an error in your beliefs, and it is therefore reasonable for you to remain steadfast in your beliefs even when facing disagreement from other epistemic agents.

However, contrary to this line of thought, White's argument provides a much grimmer picture as to the epistemic implications of Permissivism: if White's argument is valid, then Permissivism doesn't makes disagreement epistemically insignificant, but epistemically *insoluble*. This is because, as we saw above, one of the main conceptual roles which evidence plays in our epistemic practice is in settling competing claims within an epistemic community. As such, within our social epistemic practice, the need to settle what is justified for a collective group of epistemic agents means that we are not allowed to simply discount dissenting parties on the basis of our own evidence. In the context of White's example of a jury, we can see that it is not enough that our evidence only provides justification for a particular member of the jury – our evidence must also settle what verdict is *objectively* rational for the entire jury to believe is the correct verdict.[141] However, if Permissivism is true, then just like the

---

[141] This is not to deny that there are definitely many legal cases in which the evidence does lead the members of the jury to disagree on the verdict – hung juries are a distinct possibility within trials. But the question we are interested in is about whether the evidence allows members of the jury to rationally disagree. One way of explaining the difference is to note that in many legal cases, there is insufficient evidence to judge one to be guilty or innocent with strong certainty. Where there is insufficient evidence, the rational response to a body of evidence may be to suspend belief. But this response is simply not available in legal cases, which require a decisive verdict as to sentence the suspect. As such, the empirical fact that there are

individual case of choosing whether to believe on the basis of one's evidence or on a coin flip, it is once

again hard to see how we can use evidence to arbitrate between the competing claims of dissenting

parties within a given disagreement. If Permissive cases allow different members of a jury to look at the

evidence in a given trial and rationally come to different conclusions, then evidence gives us no

indication as to which member of the jury is more likely to be correct in their verdict. Permissive cases

therefore imply that such disagreements are cases which cannot be *rationally* resolved – for outside of

pointing to some common body of evidence which makes a particular verdict more likely to be correct,

what other epistemic methods are available to us by which we can rationally convince another

dissenting party?

As such, if evidence does not determine what is objectively rational to believe across epistemic agents,

then this undermines the possibility of a rational resolution to disagreement. Conversely, in order for a

theory of evidence to account for our ordinary practice of using evidence to seek a rational resolution

to disagreement, such a theory must presuppose that evidence is commonable such that dissenting

parties can agree on what is most rational to believe upon the evidence. The commonability of

evidence is therefore crucial for the evidentialist response to disagreement because in order for the

norm of "following one's evidence" to be the correct epistemic response to disagreement, we require

that our evidence provides not only adequate epistemic justification for us as individuals, but also

provides a basis by which we can rationally resolve disagreement between epistemic agents. The idea

that we use common evidence as a rational means of resolving disagreement is important not only in

---

disagreements within actual juries should be distinguished from the epistemic desiderata that evidence ought to determine
what is epistemically rational within a legal case. Arguably, the entire process of a legal trial (in its putting forth of witness
testimonies and cross-examinations) is set up for the purpose of this desiderata: that is, to facilitate the evidence needed to
make the appropriate epistemic judgment about the verdict.

our context of the epistemology of disagreement, but also in other areas of epistemology. For instance, in Bayesian epistemology, certain "merging of opinion" theorems are proposed to explain how a *common* body of evidence may allow dissenting epistemic agents to converge in their degrees of belief.[142] These merging of opinions theorems provide a formal explanation of how rational conciliation between two dissenting parties may occur, by showing how two epistemic agents with significantly different prior probability distributions may slowly converge in their degrees of belief upon conditionalising on the same evidence in the long run. While these theorems are not without their criticisms, we can see that the desiderata that is clearly motivating these "merging of opinions'' theorems is the idea that evidence should allow two dissenting parties to resolve their disagreement over time through a rational process of belief revision.[143] As such, to the extent to which we use evidence as a means of finding resolution between parties within a disagreement, the evidentialist response to disagreement requires that our conception of evidence is commonable. Let us now consider two objections to this desiderata.

In order to defend the desiderata that evidence must be commonable, two important objections must be addressed. First, one might object to the desiderata by claiming that the commonability of evidence

---

[142] For classic papers on merging of opinions theorems, see David Blackwell and Lester Dubins, "Merging of Opinions with Increasing Information," *The Annals of Mathematical Statistics* 33, no. 3 (1962): 882–86; Haim Gaifman and Marc Snir, "Probabilities over Rich Languages, Testing and Randomness," *Journal of Symbolic Logic* 47, no. 3 (1982): 495–548. For a more recent version and defence of the merging of opinion theorem, see Simon M. Huttegger, "Merging of Opinions and Probability Kinematics," *The Review of Symbolic Logic* 8, no. 4 (2015): 611–48.

[143] The two main criticisms that have been raised against these theorems are: first, that these merging of opinions hold a controversial assumption that both parties will agree on all likelihoods (i.e. all probabilities of the form P(E|H)); and second, that the convergence may take "too long" to occur. See John Earman, *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory* (Cambridge, MA: MIT Press, 1992). For another criticism of these theorems, see Gordon Belot, "Bayesian Orgulity," *Philosophy of Science* 80, no. 4 (2013): 483–503.

is an implausible assumption about evidence, given that there are many cases in our actual epistemic practice where dissenting parties cannot in fact arrive at a common body of evidence. For example, in §2.3.3, we discussed how the same evidence condition for epistemic peerhood is often unsatisfied within actual epistemic practice because our evidence often includes past experiences which cannot be fully relayed to other epistemic agents. This means that while we can certainly attempt to convey the facts about our past experiences to other epistemic agents through testimony, it is often the case that one has more reason to believe the evidence one has directly from one's perception or memory, than the evidence that one receives from the testimony of another.[144]

To illustrate, suppose you and a friend are witnesses to a bank robbery. However, in this scenario, you happen to notice the occurrence of the crime earlier, and therefore able to catch a glimpse of a mutual friend assisting the robbers in the robbery. Now, suppose you and your friend come to disagree on whether your mutual friend is guilty – no matter how you try to relay your experience of clearly seeing your mutual friend participating in the robbery, your friend does not accept your testimony as they simply cannot accept that your mutual friend could possibly do such a thing. Surely, in this scenario, it would be unrealistic to insist that your evidence *must* be commonable with your friend, since it might be entirely reasonable from your friend's perspective to reject your testimony on the basis of their total evidence, and thus refuse to admit your witness testimony as part of a common body of evidence.[145]

---

[144] That being said, there are also cases where opposite is true: for instance, where one is aware of systematic social factors which lead one to undervalue the evidential import of testimony from certain people groups, it may sometimes be reasonable to correct one's biases by valuing another's testimony more highly than the evidence one receives from one's own perception and memory. For more on this notion of "corrective trust", see Scheman, "Trust and Trustworthiness".

[145] To connect this to a previous discussion, recall how one of the main justifications for the steadfast response within the epistemology of disagreement was by appealing to self-trust, and the primacy of one's direct experience over indirect testimonial evidence. See fn. 103.

From your friend's perspective, then, it is entirely possible that the evidentialist response to disagreement is such that their total evidence undercuts the evidential support which your witness testimony might give to the belief that your mutual friend is guilty. As such, in this case, the desiderata for the commonability of evidence is unnecessary for the evidentialist response to disagreement, given that one's own total body of evidence is enough to determine what is justified and rational to believe in a given situation.

The problem with this objection, however, is that it misunderstands what we mean by wanting the commonability of evidence to be a theoretical desiderata for the evidentialist response to disagreement. To explain what I mean, let us reconsider our first theoretical desiderata that evidence determines epistemic rationality: notice that this desiderata does not imply that one does in fact have evidence which makes one perfectly rational in every circumstance. Instead, our theoretical desiderata is claiming that evidence is a necessary condition to epistemic rationality: insofar as one is epistemically rational in a given situation, one is epistemically rational precisely because one believes upon their evidence which provides the appropriate doxastic justification. This is the same with our desiderata about the commonability of evidence: what we are claiming is that a necessary condition for finding a rational resolution to a given disagreement is that that disagreement is resolved on the basis of a common body of evidence, which is able to adjudicate between the competing claims of an epistemic community. This aligns with our intuitions about the case above: notice, for instance, that it is precisely because your evidence is not commonable with your friend which leads you to be unable to find a rational resolution for this disagreement. After all, if your experience of seeing your mutual friend assisting the robbers cannot be adequately relayed to your friend, how can you possibly convince your friend to revise their beliefs? As such, to the extent that it is possible for there to be a rational resolution to a given disagreement, the desiderata about the commonability of evidence states that our theory of evidence should explain how it is that evidence can be commonable, such that the

norm to "follow your evidence" allows dissenting parties to eventually come to agreement upon a

common body of evidence. The fact that our evidence may not be fully commonable with that other

dissenting parties therefore does not undermine our desiderata that a theory of evidence should

account for how evidence might be commonable between dissenting parties.

The second objection is this: one might argue that while the commonability of evidence is indeed a

feature of our concept of evidence, it is not a feature of the kind of evidence that we are concerned with

for the evidentialist response to disagreement. According to this objection, one might argue that the

two desiderata for a theory of evidence that have been proposed in this section do not pick out

epistemic features about the same *kind* of evidence. Rather, the two desiderata pick out epistemic

features of two *different* kinds of evidence: a personal kind of evidence which is held by an individual

epistemic agent, and a public kind of evidence which is held by a collective group of epistemic agents.

Whereas personal evidence determines the epistemic rationality of an individual epistemic agent,

public evidence is a commonable kind of evidence which settles competing hypotheses for an

epistemic community. Given that there are two different kinds of evidence, there is no reason to

assume that epistemic features of personal evidence need also be epistemic features of public evidence.

Suppose you adopt a theory of evidence (say, Conee and Feldman's theory that one's total evidence is

all of one's current mental states) which accounts for personal evidence, and another (say, that

evidence consists of all the scientific facts known to an epistemic community) which accounts for

public evidence. In this case, it is easy to see why requiring that one's current mental states be

commonable would be a category mistake, since one could argue that commonability is an epistemic

feature of public evidence rather than personal evidence. Alternatively, if one required that their

individual epistemic rationality be evaluated on the basis of every scientific fact known to our

epistemic community, that would also be a category mistake as personal evidence might have different

epistemic features to public evidence (for instance, we might require luminosity to be a feature of

personal evidence). As such, one might object that the desiderata of the commonability of evidence is not a desiderata for the evidentialist response to disagreement, which is mainly involved with the question of how an individual epistemic agent should revise their beliefs when facing disagreement.

Such a response would also account for why the commonability of evidence has been largely neglected within the analytic epistemology of disagreement. Given the complexity of disagreement as an epistemic phenomena, we could arguably see the epistemology of disagreement as addressing two different epistemic issues: first, the individual epistemic problem of what one is justified in believing as an individual epistemic agent when encountering disagreement; and second, the social epistemic problem of what a group of dissenting parties is justified in believing when facing disagreement. Given that these are two distinct epistemic issues, it would therefore be plausible to assume that any evidentialist response to disagreement would involve two seperate theories of evidence as to address these two respective issues. Under this distinction, the analytic epistemology of disagreement has clearly focused on the individual epistemic problem of disagreement, rather than the social problem. To require that evidence be commonable would therefore to impose a restriction which is unfit for the particular epistemic question that the analytic epistemology of disagreement is focused on, since the commonability of evidence would be more pertinent to the social epistemic problem of disagreement.

While such an objection seems convincing at first glance, the main reason why this objection fails is because it is much harder to untangle individual and social epistemology in practice. As a *prima facie* example, it is not clear that the commonability of scientific evidence can be distinguished between "personal" and "public" evidence in this way. This is because the whole point of scientific evidence is that its commonability has *personal* implications for each member within an epistemic *community*. Consider the example above about the overwhelming scientific evidence for the fact that the earth revolves around the sun, rather than the other way around. Given this example, we can ask the following: is the scientific evidence something which is held by an individual epistemic agents, or

something which is held by an epistemic community, or both? Here, the answer is clearly both: scientific evidence is clearly something which is held by individual epistemic agents and by an epistemic community. Furthermore, the distinction between individual and social seems to fail to apply in the case of scientific evidence, as there is no distinguishable difference between scientific evidence which is held by an individual agent and scientific evidence which is held by an epistemic community – scientific evidence is the *same* evidence for an individual and for a community.

Of course, a sympathiser to the objection might not be convinced simply from this example of scientific evidence. After all, we have already seen a kind of evidence which we could say is a distinctly personal or private kind of evidence: namely, perceptual experience. As such, an alternative version of the second objection can be stated thusly: while it is true that the commonability of public evidence is such that public evidence becomes indistinguishable from personal evidence for individual epistemic agents, the converse is not true. In the case of perceptual experience, for example, the failure to make perceptual experiences commonable shows how personal evidence (e.g. one's perceptual experiences) is in fact distinguishable from public evidence. Given that the evidentialist response to disagreement is focusing on how an individual epistemic agent should respond according to their evidence, the appropriate theory of evidence for the evidentialist response is a theory of personal evidence, and thus commonability is not a desiderata for the evidentialist response to disagreement.

Two responses must be given to this new objection. First, it is not clear that perceptual experience cannot be commonable evidence in the first place. Note that, in the original formulation of the desiderata, the commonability of evidence states that what must be commonable is the *epistemic justification* which evidence confers to epistemic agents. Consider two epistemic agents: one who learns that gasoline is flammable by reading a textbook, and one learns that gasoline is flammable by lighting gasoline on fire. While the two epistemic agents might have had different *sources* of evidence for this scientific fact, the epistemic justification which they obtain for this fact justify their belief in

this fact equally. As such, even though it may never be possible for perceptual experiences to be shared from one epistemic agent to another, this does not entail that the epistemic justification that one receives from that perceptual experience might not be shared to another epistemic agent through testimony. In Miranda Fricker's work on epistemic injustice, for instance, the notion of testimonial injustice (that is, that a speaker is given less credibility than they deserve by a hearer due to prejudices about the social group that the speaker belongs to) clearly assumes that perceptual experiences can and should be commonable via testimony, such that a rejection of testimony for epistemically illegitimate reasons is epistemically unjust.[146] As such, given that one does regularly share the epistemic justification of one's perceptual experiences to another via testimony, the commonability of evidence is an epistemic feature which applies not only to public evidence, but also personal evidence.

Second, the fact that personal evidence is commonable is crucial for the evidentialist response to disagreement, because the evidentialist response to disagreement is predicated on using the personal evidence of other epistemic agents – more accurately, the epistemic justification other epistemic agents have from their evidence, and can relay through testimony – as *higher-order evidence*. To illustrate by a familiar example, suppose you are working on a particularly difficult algebraic problem. After checking your working multiple times, you discover no errors in your reasoning and believe you have strong evidence that you have the correct solution. However, when you discuss the problem with your classmates, your classmates all give an answer that is drastically different to yours. This leaves you in a complicated situation: under your assessment of the situation, your evidence suggests that you used the correct method to solve the mathematical problem. And yet, when you attempt to explain your

---

[146] See Fricker, *Epistemic Injustice*.

reasoning to others, no other epistemic agent agrees with your reasoning. What then should you believe is the correct solution to this mathematical problem according to your evidence?

As we have seen before, the Total Evidence View states that what you should do in this situation is to consider your classmates as *higher-order* evidence by assessing their epistemic peerhood, and then to adjust your beliefs according to this new piece of evidence. However, it is important here to recall that one of the criteria by which we assess the epistemic peerhood of another is the same evidence condition: that is, we evaluate whether an epistemic agent is an epistemic peer by considering whether we have reasons to believe (i.e. evidence) that they have an equal amount of relevant evidence for their beliefs which as you. The reason why the same evidence condition therefore requires the commonability of personal evidence is because our evidence about the evidence of another epistemic agent is dependent on personal evidence being commonable. To use the case above, suppose you come to learn that your classmates all worked on the algebraic problem together, and all came to the wrong conclusion by the same mistaken step of reasoning. You therefore judge them to be unjustified in their solution, and remain confident about your own answer. Here, notice that the reason why you were able to evaluate that your classmates were not epistemic peers was based on the fact that you could understand their reasoning, and come to acquire evidence that their reasoning was flawed. To be able to assess the evidence that another epistemic agent has for their beliefs is therefore only possible if evidence is commonable: if mathematical reasoning was such that one's calculations were completely inscrutable to another, then there could be no way of assessing when another person were correct or incorrect in their solutions, or justified or unjustified in their beliefs. Conversely, an implicit assumption in social epistemic practice – be it in education, collective research, or social media – is the commonability of personal evidence: in our everyday epistemic practice, we do not take the intuitions, reasoning, or evidence of an expert tradesman or academic to be exclusive to themselves, but

something which can be taught and transmitted so that other people can come to learn and become knowledgeable expert themselves.[147]

Now that we have put forward our two theoretical desiderata for the evidentialist response to disagreement, we can now continue to critically examine the main philosophical theories of evidence proposed within the literature. While the focus of my critical examination in the next chapter is on whether a given theory of evidence is suitable as a theoretical basis for the evidentialist response to disagreement – that is, whether it allows the norm to "follow your evidence" to be the sole norm which characterises one's epistemic rationality – we shall see that the failure of each of these theories of evidence to meet our two theoretical desiderata will also have further epistemic implications about the nature of evidence, and its connection to epistemic rationality. As such, one can read the following as a genealogy about philosophical theories of evidence: one where our understanding of the nature of evidence and its connection to epistemic rationality has developed through criticisms of each preceding theory of evidence. We begin in §4.1 with the first epistemic theory of evidence: namely, the empiricist conception of evidence as sense-data. After a brief overview, our criticial examination explains how the empiricist theory of evidence fails because of its inability to account for non-sensible evidence. This brings us into §4.2 with Conee and Feldman's mentalist conception of evidence as occurrent mental states. Here, our critical examination reveals how Feldman's mentalist conception of evidence fails to account for the use of evidence to evaluate interpersonal rationality. To address this problem, we examine Williamson's conception of evidence as known propositions in §4.3. Here,

_____

[147] How far does this commonability go? Are there intuitions which one has which cannot be fully conveyed to another? If one has a moral judgment which is based on a particular sentiment or feeling, is that something that is commonable? We will discuss these questions in greater detail in Chapter 5 – see also Appendix A for more about the specific case of moral judgments.

Williamson explains how the use of evidence to evaluate interpersonal rationality necessarily implies

that one can be in an epistemic position where one does not know what one's evidence is. The

implications of Williamson's argument therefore brings us to §4.4, where we consider how the notion

of epistemic position (or, epistemic context) reveals why the evidentialist response to disagreement

fails as a non-ideal epistemology of disagreement.

4   A CRITIQUE OF THE EVIDENTIALIST RESPONSE TO DISAGREEMENT

4.1   The Empiricist Conception of Evidence: Evidence as Thing, and Evidence as Sense-Data

Before we examine the empiricist theory of evidence in further detail, it is important to begin with some preliminary context. By the empiricist theory of evidence, I am specifically referring to the notion of sensory evidence that was developed starting from the empiricism of early modern philosophy and Renaissance science.[148] In his genealogical account of the development of probability, Ian Hacking argues that our contemporary understanding and intuitions around the concept of evidence originates from the early modern period.[149] This is because the history of Western epistemology before the early modern period was primarily concerned with knowledge, and in particular, a particular idea of knowledge as knowledge of universal truths which are derived by demonstrative proofs and therefore true of necessity.[150] As such, before the early modern period, the

---

[148] Some might argue that the notion of sensory evidence in Western philosophy dates back to Aristotle, who differed from Plato by claiming that we can receive knowledge of reality through our senses. However, as we will see shortly, it is not clear that Aristotle did consider sensory experience as *evidence* at all. Given the pre-Renaissance focus on knowledge of universal truths which can be acquired by demonstration, Aristotle saw sense perception as giving us necessary knowledge about the physical world rather than defeasible evidence for scientific hypotheses.

[149] More accurately, Hacking argues that while it is unclear what the exact historical connection is between this new conception of evidence and the experimental method of early modern empiricism, it is clear that our historiographical retelling of this period sees the two as inextricably linked. For more information, see Chapters 4 and 5 of Ian Hacking, *The Emergence of Probability: A Philosophical Study of Early Ideas about Probability, Induction and Statistical Inference*, 2nd ed. (New York: Cambridge University Press, 2006).

[150] Hacking, *The Emergence of Probability*, 20. While Hacking is only referring to medieval epistemology, and in particular Thomas Aquinas, it is clear that the medieval project for *scientia* was a continuation of the Aristotelian project for theoretical knowledge (*episteme*). In the *Posterior Analytics*, Aristotle says that "it is necessary for demonstrative knowledge to depend on things that are true and primitive and immediate and better known than the conclusion, to which they must also be prior and

notion of evidence referred primarily to the *testimonial* evidence one received by witnesses in legal

cases, commonly held opinions by the masses, and the authority of ancient learning from wise teachers

or texts.[151] While such testimonial evidence was seen to give "probable" support for one's opinions,

Hacking helpfully points out that the classical idea of "probable" (*probabilitas*) did not refer to our

contemporary mathematical notion of probability, but to the approbation one could give to an

opinion. As such, for philosophers before the early modern period such as Aristotle and Aquinas, the

support which testimonial evidence gives to an opinion could not be classified as an *epistemic* reason

to believe the opinion, but more an *ethical* sense of credibility and approvability of a person's character

and reputation (i.e. their *ethos*).[152] According to Hacking, it is only in the early modern period where

we begin to see a new concept of evidence as *things*. Rather than conceiving of evidence as referring to

the trustworthy character of books or people, this new kind of evidence conceived of evidence as

originating from *physical objects* which served as "signs" or "indicators" of certain causes – for

example, smoke being a sign and therefore "evidence" of fire.[153]

It is this conception of evidence as things where we first see a distinctly *epistemic* conception of

evidence, one which was developed into the empiricist theory of evidence as sense data.[154] What is the

---

of which they must be explanatory." This quote from Aristotle is taken from Anthony Kenny, *A New History of Western Philosophy* (Oxford: Clarendon Press, 2010), 133, see also 129-134.

[151] Hacking, *The Emergence of Probability*, 21ff.

[152] See Hacking, *The Emergence of Probability*, 22ff. For an alternative account which sees Aquinas as considering testimonial evidence as providing epistemic status to opinions (albeit a status weaker than knowledge), see Richard Cross, "Testimony, Error, and Reasonable Belief in Medieval Religious Epistemology," in *Knowledge, Belief, and God: New Insights in Religious Epistemology*, ed. Matthew A. Benton, John Hawthorne, and Dani Rabinowitz (Oxford: Oxford University Press, 2018), 32.

[153] Hacking, *The Emergence of Probability*, 32–33.

[154] That being said, it is important to note that philosophers within the empiricist tradition differed in their accounts about the epistemic significance of sense-data. While some took sense-data as being sources of evidence which confirmed or

connection between the inchoate conception of evidence as things, and the well-developed empiricist

account of evidence as sense-data? One of the ways in which the conception of evidence as things

influenced Renaissance science is in how the new "experimental" method of Renaissance science used

physical objects and phenomena as the basis of *commonable* evidence for scientific hypotheses and

theories. Within early modern empiricism, both philosophers and scientists accounted for how

physical objects provided commonable evidence by focusing on how different epistemic agents would

receive the same *sense-data* from the physical phenomena produced as the result of any given scientific

---

rejected particular scientific hypotheses, others took sense-data as being sources of *knowledge*. To take a modern example, while Quine considers sense-data to be a form of evidence when stating that he is "interested in the flow of evidence from the triggering of the senses to the pronouncements of science", Russell considers sense-data to be knowledge when describing sense-data as "the things that are immediately known in sensation: such things as colours, sounds, smells, hardnesses, roughnesses, and so on." The notions of sense-data as being a source of evidence or a source of knowledge have both been present throughout the history of empiricism. Nevertheless, it is sufficient for our present purposes to say that there was an account of sense-data as evidence within the empiricist tradition. See W. V. O. Quine, "Comment on Parsons," in *Perspectives on Quine*, ed. Robert Barrett and Roger Gibson (Oxford: Basil Blackwell, 1990), 293; Russell, *The Problems of Philosophy*, 12. For more information, see Gary Hatfield, "Sense Data," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Fall 2021 (Metaphysics Research Lab, Stanford University, 2021).

experiment.[155,156] This allowed the new experimental method of Renaissance science to assume that observations about any physical phenomena which occurred over the course of an experiment were replicable, and thus eligible as commonable evidence to directly confirm or reject scientific hypotheses and theories with. The empiricist theory of evidence as sense-data therefore precisified the conception of evidence as things by claiming that the epistemic justification which one receives from physical objects comes from the information (what the empiricists called "ideas") one acquired through sensory experience.

---

[155] To be more precise, the early modern empiricists provided a sophisticated account of how sensory experience provided evidence for our beliefs through a distinction between the primary and secondary qualities of an object. Whereas primary qualities referred to the properties which an object has independent of human observation, secondary qualities referred to the power which objects have to produce ideas within our mind through sensory experience. Examples of primary qualities included occupying space and being in motion, while examples of secondary qualities included heat, colour, and smell. Once this distinction was made, the empiricists argued that one perceived the secondary qualities of an object through sensation, from which one could know the primary qualities of an object through a combination of multiple sensory faculties and inference. In this way, the secondary qualities of an object acted as *evidence* for the primary qualities of an object and their causal patterns. Note that one wasn't said to have *knowledge* of secondary qualities of an object, since secondary qualities were properties of the *sensations* one has about the object in their mind, rather than any property of the object itself. This strategy was adopted not only by philosophers such as Descartes, Locke, and Berkeley, but also by scientists such as Galileo Galilei and Robert Boyle. For more information, see Martha Bolton, "Primary and Secondary Qualities in Early Modern Philosophy," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Summer 2022 (Metaphysics Research Lab, Stanford University, 2022). Anthony Kenny notes that the distinction between primary and secondary qualities does have some antecedence in Aristotle, who distinguished between "common sensibles" (qualities of an object which are perceptible by more than on sense) and "proper sensibles" (qualities of an object which are perceptible by only one sense). See Kenny, *A New History of Western Philosophy*, 601.

[156] Of course, this is not to say that there were not empiricists who put forward skeptical arguments which questioned the commonability of sense-data – a clear example of an empiricist who engaged with these skeptical concerns was Hume. However, for the purposes of this section, it is enough to claim that in the theorising of the experimental scientific method from empiricists such as Francis Bacon, or in the actual scientific practice of the time, the commonability of sense-data was not questioned.

In a sense, the empiricist theory of evidence as sense-data signified a turn in Western epistemology since it is arguably the first *epistemic* theory of evidence in how it conceived of evidence as giving direct *epistemic justification* towards a particular belief.[157] As we have just explained, Western epistemology before the early modern period focused almost exclusively on knowledge as knowledge of universal truths which are found by demonstration. This focus on demonstrative proof meant that the only epistemic justification one could have for a particular belief was through *deduction*. To take a classic example, the only justification for the particular proposition "Socrates is mortal" is from its logical entailment from a deductive argument – say, via deductive inference from the premises "All humans are mortal" and "Socrates is a human".[158] However, it is in the empiricist theory of evidence where we first see a conception of epistemic justification which allows to epistemically infer from *particular to particular*. For the empiricist theory of evidence, the particular observation that "there is smoke emanating from this matchstick" is sufficient evidence (and therefore epistemic justification) for the particular hypothesis "combustion is occurring on this matchstick". Furthermore, the epistemic justification which this particular observation gives to the particular hypothesis is warranted even without any further recourse to deductive argument (e.g. a deductive inference from a universal generalisation). In taking observations from sensory experience as directly serving as *evidence* for hypotheses, empiricist theories of evidence therefore saw evidence as providing a new form of

---

[157] Hacking, *The Emergence of Probability*, 37–38.

[158] Note that the use of the term "particular" throughout this section is referring to Aristotle's technical use of "particular" and "universal" in his logic. For Aristotle, a universal assertion is one where the subject of the sentence refers to the entirety of a particular group of objects, whereas particular assertions are assertions where the subject of the sentence refers to some member(s) within a particular group of objects. So, in this case, the premise "all humans are mortal" is a universal assertion while the conclusion "Socrates is mortal" is a particular assertion. According to Aristotle, given that "Socrates is mortal" is a particular assertion, the only way in which this assertion can be justified is through its logical entailment from the premises, one of which must be a universal assertion about the predicate of mortals.

epistemic justification to belief – one which was epistemically weaker than the logical entailment of syllogistic reasoning, but nevertheless gave justification to our beliefs.

This feature of the empiricist theory of evidence is important for our purposes, as it reveals how the empiricist theory of evidence was seen to provide a comprehensive account of epistemic justification which accordingly determines *epistemic rationality*. To illustrate how this is the case, let us compare the difference between the notion of induction as outlined within Aristotelian scientific method, and the new conception of inductive reasoning used within the experimental scientific method of Renaissance science. One of the clearest expositions of the experimental scientific method is given by Francis Bacon, who advocated for a revision of the Aristotelian scientific method, and in particular, the role of inductive inference within scientific practice.[159] Whereas Aristotle thought of induction as to be used only to derive the first premises of one's demonstration, whereafter one would use syllogistic reasoning to derive all subsequent scientific knowledge from these first premises, Bacon argued that inductive reasoning from empirical data must be used in *every* stage of the scientific process.[160] For Bacon, this meant a complete revision of our notion of inductive reasoning: whereas Aristotle defined inductive inference as referring to the single logical inference from particular to universal generalisation, Bacon saw inductive inference as the continual use of evidence to support scientific claims. More specifically, Bacon's conception of inductive reasoning involved the use of rigorous and

---

[159] Bacon's theory of induction is expounded in his *Novum Organum*, which by its title suggested a revision of Aristotle's view on logic, and in particular, on induction. For more information on Bacon's theory of induction, see Peter Urbach, *Francis Bacon's Philosophy of Science: An Account and a Reappraisal* (La Salle, IL: Open Court, 1987); Michel Malherbe, "Bacon's Method of Science," in *The Cambridge Companion to Bacon*, ed. Markku Peltonen (Cambridge: Cambridge University Press, 1996), 75–98. Much of this section was also informed by Jürgen Klein and Guido Giglioni, "Francis Bacon," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Fall 2020 (Metaphysics Research Lab, Stanford University, 2020).

[160] Malherbe notes that Bacon's conception of inductive inference such that it "applied to all stages of knowledge, and at every phase the whole precess has to be kept in mind. Malherbe, "Bacon's Method of Science," 76.

continual experimentation to find careful and systematic observations about a physical phenomena, from which one could then infer lower axioms which generalise a subset of these facts, which then (with further experimentation) could be used to infer even more general axioms until one arrived at a complex body of scientific knowledge.[161]

As we can see, critical to Bacon's view on inductive inference was the use of experiments (and the *evidence* which these experiments produced) in each stage of the scientific procedure – for Bacon, experiments were not only to be used within the early stages of the scientific method so as to establish the facts about a particular kind of physical phenomena, but also used when evaluating more general axioms by testing hypotheses and theories through confirming or refuting their predictions. As such, Bacon's conception of inductive inference – one which arguably undergirds the experimental scientific method even as it is used today – uses the concept of evidence as the sole factor in determining epistemic rationality: in every stage of the scientific process, it is one's evidence which determines whether a hypothesis, lower axiom, or general scientific theory is justified and therefore epistemically rational to believe. The empiricist theory of evidence therefore provides a theoretically robust account of evidence which accounts for how evidence can provide common epistemic justification between epistemic agents, and how one's evidence circumscribes what is epistemically rational to believe at any given moment. Given these theoretical strengths, it is unsurprising for the empiricist theory of

[161] Benjamin Farrington notes that Bacon's method was one "which by slow and faithful toil gathers information from things and brings it into understanding." Benjamin Farrington, *The Philosophy of Francis Bacon* (Liverpool: Liverpool University Press, 1964), 89.

evidence as sense-data to be the dominant epistemic theory of evidence until well into the twentieth

century.[162]

However, while the empiricist theory of evidence as sense-data has shown itself to contain many

theoretical strengths, the empiricist theory of evidence ultimately fails to satisfy our theoretical

desiderata for a theory of evidence, as sense-data does not in fact determine epistemic rationality. The

failures of the empiricist theory of evidence are clearly shown throughout the history of twentieth

century analytic philosophy, and in particular, in the fall of logical positivism. Let us consider two

problems which show how the empiricist theory of evidence as sense-data ultimately fails to determine

epistemic rationality. The first problem with the empiricist theory of evidence is that it restricts the

scope of scientific hypotheses and theories which can be justified to only those scientific hypotheses

and theories which refer to entities that are observable by sensory experience. Hilary Putnam explains

how this problem led the empiricist theory of evidence to be deemed scientifically untenable by the

twentieth century. For Putnam, one of the consequences of classical empiricism was the conclusion

that there were no scientifically indispensable predicates which referred to entities not observable with

the human senses. This meant that while the existence of atoms were already a topic of scientific

speculation in the seventeenth century, early modern empiricists such as Locke maintained that atoms

could never be something which we could be said to *know*.[163] Such a conclusion was seen to be clearly

untenable by the twentieth century, since the scientific consensus saw not only atoms to be a known

---

[162] As is noted in fn. 154, the empiricist theory of evidence as sense-data was not only held within the early modern period, but was discussed and widely accepted by twentieth century philosophers such as Russell, the logical positivists, and Quine. For more information, see Hatfield, "Sense Data".

[163] Putnam, *The Collapse of the Fact/Value Dichotomy and Other Essays*, 21–22.

part of the scientific ontology, but subatomic particles such as electrons, protons, and neutrons.[164] As such, the numerous discoveries of modern science which referred to entities that could not be observed through sensory experience revealed the inadequacy of the empiricist theory of evidence, as sense-data could not account for all the scientific facts which one was epistemically rational to believe – be it facts about non-sensible physical entities, or facts about psychological or social properties such as "stress", "anxiety", "authoritarian", or "conservative".[165]

Second, even if we disregard the problem of non-sensible or non-physical objects and phenomena, the theory of evidence as sense-data fails to determine epistemic rationality due to the fact that within actual scientific practice, the (rational) decision between competing hypotheses and theories requires more than sense-data. To illustrate, consider the phlogiston theory propounded by chemists throughout the seventeenth and eighteenth centuries to explain processes such as combustion and rusting. The phlogiston theory claimed that a candle's combustion resulted in the decomposition of phlogiston compounds within a candle's material make-up, whereby the candle would revert back to its base elements as the phlogiston was lost in the air. However, in hindsight, this is an odd conclusion, as the phlogiston theory in fact reversed the correct causal sequence of events, seeing as we now know that combustion is actually the chemical reaction in which the base elements of a candle react with oxygen to produce oxides. The phlogiston theory was therefore not only incorrect, but inhibitive of scientific progress in its delaying of the discovery of oxygen. And yet, while we know in retrospect that

---

[164] Putnam, *The Collapse of the Fact/Value Dichotomy and Other Essays*, 22.

[165] Putnam, *The Collapse of the Fact/Value Dichotomy and Other Essays*, 45–46. Putnam notes Donald Davidson and himself as critics which defend ordinary psychological traits, and he cites the following for further information: see Donald Davidson, *Essays on Actions and Events* (Oxford: Clarendon Press, 1960); Hilary Putnam, *The Threefold Cord: Mind, Body, and World* (New York: Columbia University Press, 1999), particularly Part II.

the phlogiston theory was in fact false, it is not clear that its advocates were irrational, or indeed any less rational than its opponents.[166]

The case of the phlogiston theory serves as an example of how one's sense-data about a physical phenomena is insufficient for determining which competing scientific theory is more epistemically rational to believe. This is because in this case, both the phlogiston and oxygen theory were equally plausible accounts for one's sensory experience of the process of combustion and rusting. In his discussion on the philosophical implications of the case of phlogiston, Hasok Chang argues that what really led to the historical transition from the phlogiston theory to the oxygen theory was not the available sense-data, but due to larger theoretical considerations within the field of chemistry. In particular, Chang notes how the rejection of the phlogiston theory was in fact situated within a paradigm shift in chemistry from a "principalist" model to a "compositionist" model: that is, from "the old chemical notion of 'principles', that is to say, basic substances which actively modified other substances and imparted certain characteristic properties to them…[to the new] building-block ontology, in which all pieces of matter had equal ontological status."[167] For Chang, it is the theoretical virtues of modern compositionism over principalism (which originated from alchemical practices)

---

[166] Jonathon Hricko argues that it would have been equally rationally permissible for chemists to have retained the notion of phlogiston, or to eliminate it. For Hricko, the case of phlogiston shows how scientific rationality should be defined by what is rationally permissible according to science, rather than what is rationally required to believe. See Jonathon Hricko, "Scientific Rationality: Phlogiston as a Case Study," in *Rationality: Constraints and Contexts*, ed. Tzu-Wei Hung and Timothy Joseph Lane, London (Elsevier Academic Press, 2017), 37–59.

[167] Hasok Chang, "The Hidden History of Phlogiston: How Philosophical Failure Can Generate Historiographical Refinement," *Hyle: An International Journal for the Philosophy of Chemistry* 16, no. 2 (2010): 70. For more on this transition in chemistry from principalism to compositionism, see Hasok Chang, "Compositionism as a Dominant Way of Knowing in Modern Chemistry," *History of Science* 49, no. 3 (2011): 247–68; Robert Siegfried, *From Elements to Atoms: A History of Chemical Composition*, vol. 92, Transactions of the American Philosophical Society 4 (Philadelphia: American Philosophical Society, 2002).

which ultimately led to the favouring of the oxygen theory over the phlogiston theory. To tie this into a well-discussed topic within the philosophy of science, the case of phlogiston therefore reflects how the choice between scientific theories is assessed not only according to sense-data, but according to *epistemic values* such as simplicity, explanatory power, and unifying power.[168] Given that one rationally decides between scientific hypotheses requires more than sense-data, but also requires facts about the epistemic values of competing theories, the empiricist theory of evidence does not fully determine epistemic rationality.

Furthermore, the historical case of the phlogiston theory also reveals how prior theoretical commitments affects the *evidential support* which sense-data gives. One might object to the argument above by noting that while theoretical considerations might have influenced the early adoption of the oxygen theory, later scientists were able to definitively show the oxygen theory to be better than the phlogiston theory on the basis of *sense-data* – namely, in how the additional weight of certain metals after being burnt provided evidence in favour of the oxygen theory over the phlogiston theory. However, Chang explains how observations of weight were only seen to be epistemically significant because of the compositionist model of chemistry: it is only because of the compositionist assumption

---

[168] The notion that we evaluate scientific theories according to certain theoretical *values* was introduced by Thomas Kuhn in a paper entitled "Objectivity, Value Judgment, and Theory Choice". In this paper, Kuhn points to five criteria which we use to evaluate and decide between competing scientific theories – namely, the criteria of accuracy, consistency, scope, simplicity, and fruitfulness. Since Kuhn's seminal paper, philosophers such as Michael Friedman and Philip Kitcher have contributed to this discussion by claiming explanatory unification as an important additional criteria which we use to evaluate scientific theories – that is, the ability of a scientific theory to provide an overarching explanation of several localised empirical results and phenomena. For more information, see Thomas S. Kuhn, "Objectivity, Value Judgment, and Theory Choice," in *"The Essential Tension: Selected Studies in Scientific Tradition and Change"* (Chicago; London: University of Chicago Press, 1977), 320–29; Philip Kitcher, "Explanatory Unification and the Causal Structure of the World," in *Scientific Explanation*, ed. Philip Kitcher and Wesley Salmon (Minneapolis: University of Minnesota Press, 1989), 410–505; and Michael Friedman, "Explanation and Scientific Understanding," *Journal of Philosophy* 71, no. 1 (1974): 5–19.

that the basic chemical substances are immutable and equal in ontological status (e.g. as hydrogen and oxygen atoms which could be combined or decomposed without any changes to the atoms themselves) that quantitative changes in weight during a chemical reaction were seen to be evidence of the quantity of the basic chemical substances.[169] By contrast, seeing as the phlogiston theory considered phlogiston to be a substance which actively changed the nature of other substances as to give them combustibility or metallic properties, there was no reason to assume that phlogistication would make a substance any heavier, or that dephlogistication would make a substance any lighter. This therefore meant that there was no reason to assume that observations about weight would provide any evidence about the change in chemical substances during a chemical reaction. As such, the conflict between the phlogiston theory and the oxygen theory reveals how the extent to which certain sense-data counts as evidence may itself depend on prior theoretical commitments. To tie this to another well-discussed (albeit controversial) topic within the philosophy of science, we can say that the case of phlogiston theory reveals how observational evidence is *theory-laden* – to the extent that the epistemic justification which a particular

---

[169] Chang, "The Hidden History of Phlogiston," 71; Chang, "Compositionism as a Dominant Way of Knowing in Modern Chemistry," 257–60.

observation may depend on prior theoretical commitments, observational evidence is not evidence *simpliciter*, but is evidence according to one's prior theoretical commitments.[170,171]

In conclusion, the reason why the empiricist theory of evidence fails to determine epistemic rationality is because sense-data often requires theoretical assumptions in order to interpret whether certain sense-data *is evidence in the first place*. As an aside, this point is not a recent insight at all: in his exposition on the scientific method, Bacon already noted that information one received from the senses needed to be "corrected" by the procedure and assumptions of the scientific method itself in

---

[170] For classical arguments on the theory-leadenness of observation, see Chapter 10 in Thomas S. Kuhn, *The Structure of Scientific Revolutions*, Fourth (Chicago; London: University of Chicago Press, 2012); Norwood Russell Hanson, *Patterns of Discovery* (Cambridge: Cambridge University Press, 1958); Paul Karl Feyerabend, "An Attempt at a Realistic Interpretation of Experience," in *Realism, Rationalism, and Scientific Method*, ed. Paul Karl Feyerabend (Cambridge: Cambridge University Press, 1985); Larry Laudan, *Science and Values: An Essay on the Aims of Science and Their Role in Scientific Debate* (Berkeley/Los Angeles: University of California Press, 1984); Ernan McMullin, "Values in Science," *PSA: Proceeedings of the Biennial Meeting of the Philosophy of Science Association 1982*, 1982, 3–28; Ernan McMullin, "The Virtues of a Good Theory," in *The Routledge Companion to Philosophy of Science*, ed. Martin Curd and Stathis Psillos (London: Routledge, 2009). For more recent discussion, see Jerry A. Fodor, "Observation Reconsidered," *Philosophy of Science* 51, no. 1 (1984): 23–43; Paul M. Churchland, "Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor," *Philosophy of Science* 55, no. 2 (1988): 167–87; Jerry A. Fodor, "A Reply to Churchland's "Perceptual Plasticity and Theoretical Neutrality","" *Philosophy of Science* 55, no. 2 (1988): 188–98.

[171] In the recent literature, Susanna Siegel has questioned whether one's cognition (including theoretical commitments) might affect not only the epistemic significance of a perceptual experience, but also affect the *content* of one's perceptual experience. Siegel list examples like a person who is fearful that their partner is angry at them, and thus perceives their partner's behaviour as expressing anger. Or, the example of spermist preformationists, who favoured the hypothesis that sperm cells contained embryos, and claimed to see embryos in sperm cells through the microscope. For Siegel, these examples of what she calls "cognitive penetration" raises serious doubts towards the use of perceptual content (i.e. sense-data) as a theory of evidence: "perception is often thought to be the ultimate source of epistemic and ethical justification, and so cognitive penetrability calls the foundation of some normative theories into question as well." Zoe Jenkin and Susanna Siegel, "Cognitive Penetrability: Modularity, Epistemology, and Ethics," *Review of Philosophy and Psychology* 6, no. 4 (2015): 532. See also Susanna Siegel, "The Epistemic Impact of the Etiology of Experience," *Philosophical Studies* 162, no. 3 (2013): 697–722; Susanna Siegel, "How Is Wishful Seeing Like Wishful Thinking?" *Philosophy and Phenomenological Research* 95, no. 2 (2017): 408–35.

order for one to arrive at scientific facts.[172] As such, in order to accomodate for how theoretical

assumptions and commitments can constitute part of our total evidence, philosophical theories of

evidence have moved away from physical conceptions of evidence to mentalistic conceptions of

evidence. Let us now turn to the two main mentalistic theories which have been proposed within the

literature.

4.2   *The Mentalist Conception of Evidence: Evidence as Occurrent Mental States*

Before we continue to critically evaluate our two remaining philosophical theories of evidence, a brief

remark on how the empiricist theory of evidence differs from the following two "mentalistic" theories

of evidence.[173] By mentalistic theories of evidence, I am referring to how both of the following theories

take evidence as residing in one's *mental states*: for Conee and Feldman, one's evidence consists of

one's occurrent mental states; and for Williamson, one's evidence consists of one's known

propositions.[174] One of the main differences between the empiricist theory of evidence as sense-data,

---

[172] Malherbe comments that for Bacon, "to move from the sensible to the real requires the correction of the senses, the tables of natural history, the abstraction of propositions and the induction of notions. In other words, the full carrying out of the inductive method is needed." Malherbe, "Bacon's Method of Science," 85.

[173] Conee and Feldman give an argument for why internalist theories about epistemic justification should all be classified as "mentalist" theories about epistemic justification. For Conee and Feldman, internalist theories have often been characterised according to two epistemic features: first, internalist theories are *accessible* since one's epistemic justification is determined by things that an epistemic agent has access to; and second, internalist theories are *mentalist* since one's epistemic justification resides in one's occurrent and dispositional mental states. For more information, see Conee and Feldman, "Internalism Defended," 55–58.

[174] One might wonder if the empiricist theory of sense-data is not also a mentalistic theory of evidence. As we discussed in the previous section, the empiricist theory of evidence shifted away from conceiving of evidence as physical objects to focusing on the informational content one receives from the *perceptual experience* of a physical object. Given that perceptual experience is a *mental* rather than a physical phenomenon, it would seem as if the empiricist theory of evidence is better classified as a mentalistic rather than a physical conception of evidence. However, such a line of thought depends on the conception of sense-data which one is considering. In the SEP entry on sense-data, Gary Hatfield notes the differences between contemporary accounts of sense-data, and the original notion of sense-data used by the early modern empiricists. Whereas

and the following two mentalistic theories of evidence is in what these respective theories see to be the

source(s) of one's evidence. For the empiricist theory of evidence, physical objects are the sole source of

one's evidence, as only physical entities are sensible and thus capable of giving sense-data.[175] However,

for mentalistic theories of evidence, the source of our evidence is not limited to simply physical

objects. If we conceive of evidence as consisting of one's mental states, then we can receive evidence

not only from physical objects but also from a wide range of mental phenomena - such as one's

memories of past experiences, receiving testimony from others about their experiences, or from

deductive and inductive inference from our other beliefs and past experiences. Put in another way, we

can say that mentalistic theories of evidence sees one's sources of evidence as spanning all of the

traditional sources of epistemic justification: that is, of perception, introspection, memory, deduction,

---

the traditional view took sense-data to be direct perceptions of a part of the surface of a material object, Russell took sense-data to be neither mental nor physical, but a *tertium quid* or "third thing" which existed in addition to physical objects and one's mental states. Following Russell, certain sense-data theorists (such as G. E. Moore and H. H. Price) argued that sense-data should be treated as *neutral* on ontology, so as to be compatible with various theories on perceptual ontology. And yet, within contemporary overviews on the philosophy of perception, sense-data theory is contrasted with naïve realism, suggesting that sense-data theory is a form of indirect realism and therefore that sense-data is mental. Despite the debate on the ontology of sense-data within the literature, my point regarding the following difference between the empiricist theory of evidence and mentalistic theories of evidence still stands: even if sense-data is a mental state, the empiricist theory of evidence sees sense-data as originating only from physical objects, and thus restricts one's sources for evidence to only physical objects. For more information on the ontology of sense-data, see Hatfield, "Sense Data".

[175] That being said, while only physical entities are sensible, contemporary accounts of sense-data define sense-data as arising from perceptual experience rather than sensation. This shift to perceptual experiences means that it is possible for one to have sense-data about mental phenomena: that is, that one can *perceive* the mental states of another such as pain or anger in the same way that one can perceive the physical states of another such as height or hair colour. However, the possibility that one can perceive mental phenomena raises epistemic problems for certain accounts about the epistemic justification which sense-data gives. This is because in the traditional view on sense-data, sense-data is supposed to be given in perceptual experience *prior* to any act of cognition. However, as we have seen in the examples given by Susanna Siegel, if we consider mental phenomena as something which is perceivable, then this raises questions as to whether our cognition can affect the content of our perceptual experience. For more information, see fn. 171.

induction, and testimony. By contrast, since the empiricist theory of evidence focuses on sensory

evidence, the empiricist theory of evidence restricts one's evidence as originating only from

perception. It is this shift in sources of evidence which we shall see to be the main strength of

mentalistic conceptions of evidence, a strength which highlights the problems with the evidentialist

response to disagreement. To see why this is the case, let us now turn to the first mentalistic

conception of evidence: namely, Conee and Feldman's theory of evidence as consisting of an epistemic

agent's occurrent mental states.

To begin, consider the following definitions given by Conee and Feldman about their conception of

evidence:

> The justificatory status of a person's doxastic attitudes strongly supervenes on the person's occurrent
> and dispositional mental states, events, and conditions. [176]

> A much more restrictive view about available evidence may be formulated as follows: S has p available
> as evidence at t iff S is currently thinking of p.[177]

A key part of understanding Conee and Feldman's theory of evidence is in the distinction between

occurrent and dispositional mental states. Whereas occurrent mental states refer to the mental states

that one is aware of at a given point in time, dispositional mental states refer to mental states which

one has even if one is not aware of it at a given point in time. An example of an occurrent mental state

is perceptual experience: one's perception of an elm tree outside their window is an immediate

apprehension of an object through one's senses, and thus a mental state that one is aware of at the

time. By contrast, an example of a dispositional mental state is knowledge that one has but is currently

---

[176] Conee and Feldman, "Internalism Defended," 56.

[177] Richard Feldman, "Having Evidence," in *Evidentialism* (Oxford: Oxford University Press, 2004), 232. Some syntactical and
edits were made to this quote.

not thinking of – for instance, one might know that Canberra is the capital city of Australia at any given time, but not be consciously thinking about it currently. We are now in a position to unpack Conee and Feldman's theory of evidence. To begin, notice that Conee and Feldman's theory of evidence is characterised specifically in the *second* formulation: namely, that one's evidence consists only of one's *occurrent* mental states.[178] However, Feldman also explains why, as is stated in the first formulation, one's doxastic justification also strongly supervenes on one's *dispositional* mental states. This is because of the simple fact that "a person knows a thing dispositionally provided the person would know it concurrently if he thought of it".[179] To take our example of the fact "Canberra is the capital city of Australia", Feldman's point is that while such a fact is stored as a dispositional mental state, our intuition that this fact is part of our evidence comes from our experience of how such facts typically become occurrent to us whenever they are pertinent to our epistemic situation. Furthermore, when they do not become occurrent to us, it is commonplace for us to say that such facts are not part of our evidence – for instance, if we cannot recall that Canberra is the capital city of Australia, then our lapse in memory means that it is not part of the available evidence we use to justify our beliefs. As such, Conee and Feldman's theory of evidence sees evidence as consisting of our *occurrent* mental

---

[178] At least, this is how Conee and Feldman's theory of evidence is commonly described within the literature. That being said, it is important to note Conee and Feldman waver in their formulation of their theory of evidence in their 2004 book: while their paper on internalism (from which the first formulation comes from) contains passages which suggest that one's evidence includes dispositional mental states, Feldman's solo paper is adamant that one's evidence only consists of one's occurrent mental states. However, as we shall see throughout this section, it is Feldman who gives a more precise account of their theory of evidence that also accounts for the role of dispositional mental states in determining one's evidence. As such, following the literature as well as Feldman, I shall take this second formulation as the definitive explication of their theory of evidence. For more information, see Conee and Feldman, *Evidentialism*, 2004.

[179] Feldman, "Having Evidence," 236.

states: be they occurrent because of our present environment (e.g. perceptual experience), or because they are dispositional mental states that are made occurrent in our present situation.

Why should we accept Conee and Feldman's theory of evidence as occurrent mental states? Conee and Feldman's strategy for supporting their theory of evidence consists of providing intuitive examples of two epistemic agents who are in similar epistemic situations, but whose difference in evidence (and thus epistemic justification) is clearly due to a difference in their mental states. Let us consider two of the examples which Conee and Feldman give, as well as their epistemic evaluation of these examples:

> **Example:** Bob and Ray are sitting in an air-conditioned hotel lobby reading yesterday's newspaper. Each has read that it will be very warm today and, on that basis, each believes that it is very warm today. Then Bob goes outside and feels the heat. They both continue to believe that it is very warm today. But at this point Bob's belief is better justified.

> **Comment:** Bob's justification for the belief was enhanced by his experience of feeling the heat, and thus undergoing a mental change which so to speak "internalized" the actual temperature. Ray had just the forecast to rely on.

> **Example:** A novice bird watcher and an expert are together looking for birds. They both get a good look at a bird in a nearby tree. (In order to avoid irrelevant complexities, we can assume that their visual presentations are exactly alike.) Upon seeing the bird, the expert immediately knows that it is a woodpecker. The expert has fully reasonable beliefs about what woodpeckers look like. The novice has no good reason to believe that it is a woodpecker and is not justified in believing that it is.

> **Comment:** The epistemic difference between novice and expert arises from something that differentiates the two internally. The expert knows the look of a woodpecker. The novice would gain the same justification as the expert if the novice came to share the expert's internal condition concerning the look of woodpeckers.[180]

According to Conee and Feldman, both of these examples are examples of epistemic agents whose difference in evidence (and thus epistemic justification) can only be accounted for by a difference in occurrent mental states. In the first example, while Bob and Ray both have testimonial evidence for

---

[180] Conee and Feldman, "Internalism Defended," 59. Once again, some syntactical edits were made – Conee and Feldman's numbering of their examples is omitted here, and the example of the novice and expert and bird watcher is actually Example 3 in the original paper.

believing that it will be warm today, the reason why Bob is clearly better justified in his belief is that Bob's perceptual experience of being outside allows him to be more *cognisant* of the heat than Ray. In the second example, while both bird watchers are perceiving the same bird, it is a difference in mental states which accounts for the difference in evidence, and thus epistemic justification. Specifically, it is because of the expert's additional prior beliefs about the appearance of woodpeckers that allows the expert to have a different doxastic attitude about the bird (i.e. the belief that the bird is a woodpecker), one which the novice would not be justified in having. As such, from their analyses of these examples, Conee and Feldman give a strong argument for their theory of evidence: seeing as we typically account for the difference in justification that epistemic agents have as arising from a difference in their mental states at the time, it is natural to assume that one's evidence consists of one's occurrent mental states.

Two further comments on the examples above. First, notice that while the first example can be accounted for by the empiricist theory of evidence, the second cannot. This is because under the empiricist theory of evidence, both the novice and expert bird watcher were perceiving the same bird at the same time, and thus were receiving the same amount of evidence as is relevant to identifying the bird. However, even if the novice and expert bird watcher were receiving the same sense-data about the bird, our intuitions that the expert bird watcher has more evidence stem from the fact that the expert bird watcher has had more experience and knowledge about woodpeckers. As such, this example shows how one's evidence consists not only of one's sense-data, but also the prior knowledge and belief that one has. Second, notice that for the example of the novice and expert bird watcher, the difference in their evidence is also due to a difference in dispositional mental states – that is, that the expert had dispositional beliefs about the appearance of woodpeckers which the novice did not. However, note that the expert's dispositional beliefs in the appearance of woodpeckers only made a difference in this scenario because they were made occurrent during the perceiving of the bird. If the expert could not recall their prior beliefs about the appearance of the woodpecker in the moment, then

they would not be justified in believing that it is a woodpecker.[181] Now that we have given a sufficient account of Conee and Feldman's theory of evidence, let us critically evaluate this theory according to our two theoretical desiderata.

How does Conee and Feldman's theory of evidence fare under our two proposed theoretical desiderata? With regard to our desideratum that our evidence determines epistemic rationality, Conee and Feldman's theory of evidence seems to perform better than the empiricist theory of evidence as sense-data. As we saw in the previous section, the empiricist theory of evidence fails to determine epistemic rationality for at least two reasons: first, it fails to determine what is epistemically rational to believe for non-sensible entities; and second, it fails to fully determine epistemic rationality since it does not account for the epistemic justification one receives from epistemic values and theoretical commitments. However, in both of these cases, Conee and Feldman's theory of evidence is sufficient for determining epistemic rationality. When considering non-sensible entities, Conee and Feldman's theory of evidence doesn't suffer the flaws of the empiricist theory of evidence, as conceiving of evidence as our occurrent mental states allows our evidence to come from sources other than perception.

So, for example, our belief in the existence of atoms can come from many sources which allow us to adopt the relevant doxastic attitudes: we can receive testimonial evidence from scientific textbooks and teachers, or we can come to our belief about the atomic composition of certain substances through abductive inference from scientific experiments which involve chemical reactions. Furthermore, it is

_____

[181] One might still wonder whether one's evidence should not also include one's dispositional mental states. After all, if the expert fails to recall his prior beliefs about the appearance of the woodpecker, one might argue that they have also lost their dispositional beliefs as well. It is unclear where this objection leads though: one could claim that one's evidence consist of occurrent mental states, while also noting that one's dispositional mental states play a large part in determining one's occurrent mental states.

reasonable to assume that the epistemic values and theoretical commitments that one holds is reflected principally in one's mental states. For instance, in the example of the expert and novice bird watcher, it is fair to say that the difference between the expert and novice bird watcher is a form of theory-ladeness in observation: insofar as we can say that the expert's prior reasonable beliefs about the appearance of a woodpecker also commit the experts to certain presuppositions, we can see that the expert's prior beliefs leads to a difference in how the expert perceives the bird, and thus to a different conclusion that is more justified than that of the novice bird watcher. A similar argument can be made for the role of epistemic values in accounting for changes in our mental attitudes: when deciding between two competing theories, our acceptance of epistemic values such as simplicity, explanatory power, and coherence is reflected in our mental judgments of a certain theory as being more simple, coherent, or explaining more of our empirical data as another. As such, in comparison to the empiricist theory of evidence, Conee and Feldman's theory of evidence provides a better account for how evidence determines epistemic rationality.

However, despite its theoretical strengths over the empiricist theory of sense-data, Conee and Feldman's theory of evidence ultimately fails to determine epistemic rationality. Furthermore, the reasons for why Conee and Feldman's theory of evidence fail to determine epistemic rationality are important for the subject of this chapter, as they are intimately connected to how and why Conee and Feldman neglect the commonability of evidence.[182] As I shall argue, the main problem in Conee and Feldman's theory of evidence is its focus on *occurrent* mental states, a focus which forces Feldman to assume that epistemic rationality is "current-state rationality". However, in conceiving of epistemic

---

[182] Once again, while the theory of evidence as occurrent mental states is typically taken to be propounded by both Conee and Feldman, the argumentation for the theory of evidence is largely given by Feldman. As such, the following section will focus specifically on Feldman's defence and explication of the theory of evidence as occurrent mental states.

rationality as "current-state rationality", Feldman neglects the *social* dimensions of epistemic

rationality which play a pivotal role in helping us to identify *which* occurrent mental states are in fact

evidence. To examine these problems in further detail, let us begin with the following case given by

Feldman:

> A professor and his wife are going to the movies to see Star Wars, Episode 68. The professor has in his
> hand today's newspaper which contains the listings of movies at the theater and their times. He
> remembers that yesterday's paper said that Star Wars, Episode 68 was showing at 8:00. Knowing that
> movies usually show at the same time each day, he believes that it is showing today at 8:00 as well. He
> does not look in today's paper. When they get to the theater, they discover that the movie started at
> 7:30. When they complain at the box office about the change, they are told that the correct time was
> listed in the newspaper today. The professor's wife says that he should have looked in today's paper
> and he was not justified in thinking it started at 8:00.[183]

Feldman uses this example to address an objection for why one's occurrent mental states do not

determine one's epistemic rationality. Specifically, the objection claims that one's occurrent mental

states do not determine one's epistemic rationality, since epistemic rationality is purely determined not

only by one's occurrent mental states, but also by evaluating whether one's occurrent mental states are

the result of a process of rational belief formation. To explain, note that in this example, the professor's

occurrent mental states (that is, his knowledge of the contents of yesterday's paper) clearly give him

evidence as to believe that the movie will start at 8:00. Furthermore, the professor's occurrent mental

states give him no evidence to the contrary – according to what he is aware of, there is no reason for

the professor to believe that the movie would start at any other time. And yet, one might argue that the

professor is epistemically irrational in his belief that the movie will start at 8:00, since the professor

failed to act in an epistemically rational way in neglecting to check the movie times in today's

newspaper. As such, this example reveals how one's occurrent mental states are insufficient for

determining epistemic rationality: while the sum of the professor's occurrent mental states serves as

---

[183] Richard Feldman, *Epistemology* (Upper Saddle River, NJ: Prentice Hall, 2003), 47.

evidence in favour of the belief that the movie will start at 8:00, the professor was not epistemically rational in this case because he did not gather enough relevant evidence (i.e. by looking at today's newspaper) in favour for his belief.

In response to this objection, however, Feldman argues that such an objection is mistaken due to the fact that it equivocates between two senses of the term rationality. To explain, Feldman distinguishes between two notions of rationality: namely, between current-state rationality, and methodological rationality. Current-state rationality refers to the type of "epistemic appraisal [that is concerned with the question of] whether believing a particular proposition is rational for a person at a time given exactly the situation the person happens to be in at the time."[184] On the other hand, methodological rationality refers to the "epistemic evaluation of a belief [which] has to do with the methods that led to it…[such that] beliefs are methodologically rational if and only if they are formed as the result of good epistemic methods."[185] For Feldman, the reason why the objection fails is due to the fact that while the example does show that the professor might be methodologically irrational, it does not show that the professor is current-state irrational. Furthermore, since our notion of epistemic rationality is in fact current-state rationality rather than methodological rationality, the professor is therefore epistemically rational in this case.

For what reasons should we believe that current-state rationality is in fact our notion for epistemic rationality? Feldman gives two reasons for this claim. First, Feldman argues that our assessments of epistemic rationality are *normative* assessments, such that what an epistemic agent *ought* to do

---

[184] Feldman, "Having Evidence," 233.

[185] Feldman, "Having Evidence," 233.

rationally is dependent on what they *can* do in their given situation.[186] Put in the context of our example above, Feldman argues that even if the professor is lacking in evidence due to a prior act of methodological irrationality, our assessment of the professor's epistemic rationality should be based on what is reasonable for the professor to believe given the epistemic situation he is in:

> Suppose it is true that the professor should have looked at today's newspaper. He messed up and did not do that. Still, the question remains, given that he has been negligent and not done what he should, what is it most reasonable for him to believe? The answer is that it is most reasonable for him to believe that the movie starts at 8:00. More generally, it is most reasonable to believe what is supported by the evidence one does have. Because one does not know what the evidence one does not have will support, it would be unreasonable to be guided by that evidence. So [the notion that one's epistemic rationality is determined by what additional evidence they could have gathered] is mistaken. Even when one should get more evidence, the thing to do at any given time is to be guided by the evidence one does have.[187]

To put Feldman's point in another way, we can say that in order for us to have norms about epistemic rationality, we require that our norms be satisfiable by an epistemic agent at any given time. However, the fact that our norms of rationality should be satisfiable implies that we should be evaluated only according to the *current* epistemic resources that are available to us. So in this example, the professor's failure to look at today's newspaper might have been a mistake which led the professor to miss a crucial piece of evidence about when the movie starts. And yet, this does not mean that the professor violated any norms of epistemic rationality in believing what he did on his limited evidence: after all, without any evidence to the contrary, it is reasonable for the professor to assume that the movie starts at 8:00. As such, Feldman shows how failures of methodological rationality do not imply failures in

---

[186] In this way, Feldman takes the norms of epistemic rationality to be analogous to ethical normativity, and in particular, to the Kantian notion that "ought implies can": "Analogously, in asking what a person morally ought to do, we look at the situation the person is in and evaluate the options open to him. How he's gotten himself into his current situation is not strictly relevant to the evaluation." Feldman, "Having Evidence," 233.

[187] Feldman, *Epistemology*, 48.

epistemic rationality: one can be placed in a situation with limited evidence (as a result of their own actions, or otherwise), and still be epistemically rational in how they respond to that limited evidence.

Second, Feldman argues that methodological rationality is not a form of epistemic rationality at all, but rather a form of practical rationality: "It is unclear to me whether methodological epistemic rationality is an epistemologically central notion at all. Whether believing something is methodologically rational seems to depend largely on practical matters."[188] By "practical matters", Feldman is referring to the fact that our choice of epistemic method (that is, the method regarding how we gather evidence) is largely based on practical concerns regarding the time, attention, and resources we have, as well as the significance of the belief in question.[189] To give a few examples, suppose you are playing along to a trivia game which is occurring on TV, such that you quickly look up a certain fact on Wikipedia. While such a method might be acceptable for this case, it might be unacceptable for you to use this method when writing a research paper for a peer-reviewed journal. Or, suppose you are an in-person mathematics exam, and you answer the questions through some hasty mental calculations. While this might be acceptable given the limited time you had for an exam, this might be rationally unacceptable

---

[188] Feldman, "Having Evidence," 235. On a similar point, Thomas Kelly has argued that practical rationality is not of the same kind as epistemic rationality: see Thomas Kelly, "The Rationality of Belief and Some Other Propositional Attitudes," *Philosophical Studies* 110 (2002): 163–96; Kelly, "Epistemic Rationality as Instrumental Rationality". For a modern response to these arguments, see Asbjørn Steglich-Petersen and Mattias Skipper, "An Instrumentalist Account of How to Weigh Epistemic and Practical Reasons for Belief," *Mind* 129, no. 516 (2020): 1071–94.

[189] That being said, as we discussed in §1.3.1, one way in which the contemporary literature has diverged from Feldman's view is on the question about whether practical matters might affect one's epistemic status. In particular, the discussion around the "pragmatic encroachment" thesis have led some to believe that practical factors may indeed be relevant in determining whether an epistemic agent is justified in their belief, or is possessing knowledge or not. For more information, see Fantl and McGrath, *Knowledge in an Uncertain World*; Brian Kim, "Pragmatic Encroachment in Epistemology," *Philosophy Compass* 12, no. 5 (2017); Brian Kim and Matthew McGrath, eds., *Pragmatic Encroachment in Epistemology*, Routledge Studies in Epistemology (New York; London: Routledge, 2019).

if you did the same thing for a mathematics assignment in which you were given ample time to complete. According to Feldman, once we notice that methodological rationality is largely based on practical concerns, we can see how methodological rationality refers to a type of rationality which is entirely independent from the question of epistemic rationality. This is seen clearly in the example of the professor:

> In the example, perhaps it would have been a good idea to look at the listings in today's newspaper. However, before drawing that conclusion it is worth noting that it is almost always possible to be even more careful and to look for more evidence. The professor had good reason to think that the movie started at 8:00 and to believe that the newspaper would say that it did. With hindsight, it is easy to criticize him. But if he should have checked today's newspaper, then perhaps he also should have checked the movie listings online, or he should have called the theater to confirm what the newspaper said. Maybe he should have called a second time to get someone to confirm what was said on the recording heard during the first call. Further checking is almost always possible. Depending upon the seriousness of the situation, the likelihood that new information will be helpful, and other factors, it is sometimes in your interest to do some further checking. However, it is surely not always sensible to keep on checking. But all of this is independent of the reasonableness of believing what he did given the situation he actually was in.[190]

As such, in the example of the professor, Feldman's defence of the theory of evidence as occurrent mental states is based on the assumption that epistemic rationality is current-state rationality, and thus principally concerned with what one is rational to believe in a particular moment in time.[191] For

---

[190] Feldman, *Epistemology*, 48.

[191] Within the literature, this notion of current-state rationality is also known as time-slice rationality: the view that every norm of epistemic rationality is a synchronic norm. For more information, see fn. 61. See also Trent Dougherty, "The Ethics of Belief Is Ethics (Period): Reassigning Responsibilism," in *The Ethics of Belief: Individual and Social*, ed. Jonathan Matheson and Rico Vitz (Oxford: Oxford University Press, 2014), 146–68; Richard Feldman, "Epistemological Duties," in *The Oxford Handbook of Epistemology*, ed. Paul K. Moser (Oxford: Oxford University Press, 2002), 362–83; Hedden, "Time-Slice Rationality".

Feldman, current-state rationality is the conception of epistemic rationality that underlies not only evidentialism, but theories of epistemic justification more generally.[192]

However, far from being the concept of epistemic rationality which underlies evidentialism, I shall now argue that Feldman's concept of current-state rationality (and the theory of evidence as occurrent mental states) is in fact unsuitable for the evidentialist response to disagreement. The main problem with current-state rationality is that it fails to account for the *social* dimensions of epistemic rationality, as the focus in current-state rationality on the current situation and current mental states of an epistemic agent construes epistemic rationality as referring only to the rationality of an *individual* epistemic agent. As a starting point for my argument, consider how the example of the professor might play out as an example of an epistemic disagreement. Feldman's description of the scenario ends with the professor's wife claiming that he should have checked today's newspaper, and that he was not justified in believing that the movie will start at 8:00. Given Feldman's defence of current-state rationality, we might say that the rational response for the professor to give in this case is to refute his wife's claims, since he was current-state rational in believing that the movie will start at 8:00.

But now, consider the situation from the wife's perspective. Suppose the wife is not satisfied with this response, as the wife rejects the professor's claim that he *knows* that the movie will usually start at the same time each day. The wife does not believe this, and given that the wife does not believe that movies usually start at the same time each day, the wife claims that she would always check the newspaper for the movie times. As such, given the wife's occurrent mental states, the wife is current-state rational in

_____

[192] "The question relevant to evidentialism, and to theories of epistemic justification generally, is 'What should S believe now, given the situation he's actually in?'" Feldman, *Epistemology*, 48.

believing that the professor is unjustified in believing that the movie started at 8:00. The important

question which this continuation of the example raises is this: where does this leave us in our epistemic

evaluation of this disagreement? Was the professor justified in his beliefs or not? Who is correct in this

disagreement, and what would be the rational resolution to this disagreement? The problem with

current-state rationality is that it gives us unsatisfying answers to these questions. Take the question of

whether the professor is in fact justified in his beliefs. Under current-state rationality, it appears that

the answer would be that it depends on who you're talking about: for the professor, he is rational in

believing that he is justified in his belief; and for the wife, she is rational in believing that he is not

justified in his belief. But this is a unsatisfying answer: what we want out of epistemic rationality is not

just about what is rational *relative* to an individual epistemic agent, but what is *objectively* rational

across epistemic agents. As such, the problem with Feldman's notion of current-state rationality is that

it fails to understand that while the norms of epistemic rationality might be context-sensitive to an

epistemic agent, they are not *determined relative* to an epistemic agent. Even if the professor believes

that he is epistemically rational, this does not mean that other epistemic agents will agree.

On a similar note, we can see that Conee and Feldman's theory of evidence as occurrent mental states

faces the same kind of problem in its failure to account for the commonability of evidence. In the

example of the professor, while the wife might be able to receive some evidence about the movie time

by testimonial evidence from the professor, it is also clear that the professor and his wife can diverge in

mental states such that their occurrent mental states fail to allow them to come to a common body of

evidence. But, as we have discussed in §3.2., this is a problem because without this common body of

evidence, the theory of evidence as occurrent mental states fails to satisfy one of the primary roles

which evidence plays in our everyday epistemic practice: that is, in arbitrating between competing

claims as to decide which dissenting party or parties are rational within a given disagreement.

Before we continue in our critical evaluation of current-state rationality, it is helpful to consider how

Feldman might respond to this criticism. Consider the following rebuttal: just like Feldman's

distinction between methodological rationality and current-state rationality, one might argue that the

criticism of current-state rationality given above is flawed because it equivocates between intrapersonal

rationality and interpersonal rationality.[193] As a brief sketch of the distinction, we can say that whereas

interpersonal rationality is an assessment of one's rationality in relation to a collective group of

epistemic agents, intrapersonal rationality considers one's rationality in relation to one's own personal

epistemic situation. Suppose you study for a science test using an outdated science textbook which

causes you to fail the exam. In this scenario, while you might be interpersonally irrational because you

gave incorrect answers to the exam, you might nevertheless be intrapersonally rational insofar as you

were justified in your beliefs based on the evidence you had. Given this distinction, one might defend

current-state rationality by arguing that the so-called "social" dimensions of epistemic rationality don't

apply to current-state rationality at all, since current-state rationality is a form of *intrapersonal*

rationality while the social dimensions of epistemic rationality refer to *interpersonal* rationality.

Furthermore, seeing as epistemic rationality is only concerned with intrapersonal rationality, the

objection fails to refute the claim that current-state rationality is the correct interpretation of epistemic

rationality.

Let us use the distinction between intrapersonal and interpersonal rationality to examine the example

of the professor again. We can evaluate the professor as being interpersonally irrational insofar as he

---

[193] Such a distinction regarding rationality has been made, for instance, within the literature on the Uniqueness Thesis, and the distinction between intrapersonal Uniqueness and interpersonal Uniqueness. See Kopec and Titelbaum, "The Uniqueness Thesis," 191. For more recent work done on intrapersonal or interpersonal rationality, see Han Li, "How Supererogation Can Save Intrapersonal Permissivism," *American Philosophical Quarterly* 56, no. 2 (April 2019): 171–86; Valerie Soon, "An Intrapersonal, Intertemporal Solution to an Interpersonal Dilemma," *Philosophical Studies* 178 (2021): 3353–70.

made an incorrect assessment about the movie time. But this assessment comes from us as an outsider's perspective: it is because we know that the contents of today's newspaper did in fact contain important evidence that we are able to judge the professor as having made a crucial epistemic mistake in not checking the newspaper, and thus epistemically irrational. Or in our continuation of this example from the wife's perspective, the wife's evaluation that the professor is not rational in his beliefs is also an evaluation of interpersonal rationality, since the wife includes her own beliefs in evaluating the professor rather than simply evaluating the professor according to his own occurrent mental states. Both these assessments are different to the question of intrapersonal rationality: arguably, the professor was intrapersonally rational since the professor was justified in his beliefs based on the evidence that he had.

Why then should we take intrapersonal rationality to be the correct notion of epistemic rationality? For the same reason that Feldman gives for current-state rationality: that we use epistemic rationality as a normative assessment about epistemic agents, such that what an epistemic agent *ought* to do rationally should depend on what they *can* do. Suppose we evaluate epistemic rationality on the basis of interpersonal rationality, such that an epistemic agent is rational only if they meet the standards of what is considered objectively rational across epistemic agents. Such a standard would be an idealised rational norm that is unsatisfiable for any actual epistemic agent: in order to meet the standards of what we consider to be objectively rational for our epistemic community, an epistemic agent would (at the very least) need to be logically omniscient (that is, to be able to infer and know all of the logical consequences of every proposition), as well as to know all the facts (e.g. scientific, historical, mathematical, etc.) that we know as an epistemic community. Given that these standards are unsatisfiable for any actual epistemic agent, the norms of epistemic rationality would therefore cease to be meaningful norms for guiding our epistemic practice. Conversely, insofar as we want the norms of epistemic rationality to be norms which guide our epistemic practice, we should therefore consider

intrapersonal rationality (and thus current-state rationality) to be the correct interpretation of epistemic rationality.

Before we respond to this rebuttal, consider how the rebuttal given above is similar to a previous objection that we examined against the commonability of evidence. In particular, one way of rejecting that the commonability of evidence should be a theoretical desiderata for the evidentialist response to disagreement is by making a distinction between individual and social epistemology. This distinction involves asserting that there are two different kinds of epistemic rationality: one which applies to epistemic agents as individuals, and one which applies to a collective group of epistemic agents as a whole. Once we distinguish between individual and social epistemic rationality, we can see how the commonability of evidence is only applicable to social epistemic rationality for a collective group of epistemic agents, while the evidentialist response to disagreement is solely concerned with the epistemic rationality of an individual epistemic agent, and what it means for an individual epistemic agent to be rational in their response to disagreement. These different kinds of epistemic rationality therefore also suggest that we require two distinct epistemic theories of evidence such that the commonability of evidence applies to a public notion of evidence, while the evidentialist response to disagreement requires a personal notion of evidence. As such, the commonability of evidence refers to a type of evidence and epistemic rationality that is of an entirely different kind to that we require for the evidentialist response to disagreement, and thus should not be a theoretical desiderata for our theory of evidence.

As we discussed in §3.2, however, such a distinction between individual and social epistemic rationality fails in practice, since our conception of what is rational for an individual epistemic agent is informed and shaped by what is considered rational within our epistemic community. This is the same problem with the rebuttal that is given above: such a rebuttal overdetermines the distinction between intrapersonal and interpersonal rationality in suggesting that these are two different kinds of epistemic

rationality. Once we see that intrapersonal and interpersonal rationality are mutually constitutive –

that is, that facts about interpersonal rationality determine facts about intrapersonal rationality and

vice versa – we can see that the notion of epistemic rationality cannot be defined by simply

intrapersonal rationality (or by current-state rationality).

To begin, consider how we develop our own sense of intrapersonal rationality by reference to

interpersonal rationality: that is, we learn and assess our own epistemic rationality by considering what

is rational within our epistemic community. Imagine a child who wishes to stick his hand into the fire,

because all of his past experiences have associated bright colours as being a good thing. Just because

the child's mental states make him believe it is reasonable to stick his hand into the fire, this does not

mean that the child is therefore rational in believing this. Furthermore, we would agree with a parent

who rejects the child's judgment and forbids them to touch the fire, since the child is not in an

epistemic situation where they can make the appropriate decisions and beliefs for themselves. In fact,

even if we would (reluctantly) say that the child is current-state rational or intrapersonally rational in

their response to their limited evidence, we would not want to say that the epistemic rationality of the

child should be evaluated simply by this current-state rationality, or by their current epistemic

situation. As such, this example reveals how current-state rationality (or intrapersonal rationality) does

not encapsulate our notion of epistemic rationality: even if our evaluations of an epistemic agent's

epistemic rationality must be context-sensitive to an epistemic agent's particular situation, this does

not mean that an epistemic agent's epistemic rationality is *determined* only by their current epistemic

situation.

We can make a similar argument against the theory of evidence as occurrent mental states: Conee and

Feldman's theory of evidence as occurrent mental states does not determine epistemic rationality, as

we often evaluate the epistemic rationality of an epistemic agent on the basis of criteria which are not

occurrent to them. To illustrate, consider once again the example of the novice and expert bird

watcher. As we discussed before, Conee and Feldman's analysis of this example involved explaining how the difference in epistemic justification (and thus rationality) between the novice and the expert bird-watcher was due to the difference in their occurrent mental states. But notice that this analysis involves making judgments about the epistemic rationality of the novice that the novice could not possibly make in their situation: that is, that we evaluate the epistemic rationality of the novice on the basis of experiences and beliefs that the novice does not have. Furthermore, notice that is an assessment that the novice should also accept: we wouldn't allow the novice to claim, for instance, that they are just as rational as the expert because they respond to their evidence as rationally as the expert responds to theirs.

At this point, one might point out that while the novice does not have the experiences and beliefs of the expert, the novice should have *higher-order* evidence that the expert is more proficient in bird-watching than they do. As such, the novice *does* have occurrent mental states which suggest that the expert is more justified in their beliefs about the perceived bird than them. However, let us imagine if the novice was extremely ignorant and arrogant about their bird-watching prowess, such that they have never even thought to compare their amazing bird-watching prowess to that of the lowly expert. Given that the novice is so egomaniacal about their bird-watching skill, to the point that the expert's skill is entirely absent from their occurrent mental states, their occurrent mental states do not give any evidence to suggest that they are any less competent in their beliefs about the perceived bird as the expert. Surely it would be an odd result to say that it is precisely the novice's extreme narcissism which caused them not to have any occurrent mental states about the expert's experience and skill, and therefore *enabled* the novice to be rational in their beliefs. But this is the implication of Feldman's argument about the difference between methodological and current-state rationality: given the distinction between methodological and current-state rationality, we should not evaluate the novice's epistemic rationality on the basis of the process by which they arrived at their mental states (i.e. their

narcissism), but should evaluate the novice solely on the basis of what they are currently thinking of.[194]

Given that this is an unacceptable conclusion, we should therefore reject Feldman's conception of current-state rationality, and conclude that occurrent mental states do not determine one's epistemic rationality.

How then should we account for Feldman's argument that our assessment of epistemic rationality is a normative assessment, such that we should evaluate what an epistemic agent rationally ought to do based on what they can do? One way of explaining this intuition is by making a distinction between what is reasonable and what is rational.[195] To explain the difference, consider Kvanvig's account of the connection between epistemic rationality and epistemic excusability.[196] Kvanvig argues that if we conceive of epistemic normativity as analogous with moral or legal responsibility, then we should

------

[194] This example of the egomaniacal novice might also make us wonder whether our analysis of the example of the professor is also based on implicit assumptions about the method with which the professor arrived at his occurrent mental states. Imagine if the professor was delusional, such that he received messages from the Egyptian sun god Ra in his dreams which told him that the movie will start at 8:00. Furthermore, this experience was so emotionally intense that it has been the only thing that the professor has been thinking of for the past week. In this case, surely we cannot say that the method with which the professor arrived at his occurrent mental states is irrelevant to how epistemically rational he is.

[195] Rawls also makes a distinction between reasonableness and rationality in his writings. However, for Rawls, the distinction between reasonableness and rationality is a distinction between an ethical category and an epistemic or logical category. Whereas reasonableness refers to "moral characteristics [which are] related to fairness or justice of persons, principles, and conceptions of the good", rationality refers to "the principles and considerations that are used to determine a person's or society's good." The distinction between reasonableness and rationality that is drawn in the following section, however, is a distinction between two epistemic categories. However, Rawls' distinction is interesting to note here because in the subsequent chapters of this thesis, we will examine how the lines between ethics and epistemology are not as strict as one might assume, such that an ethical sense of reasonableness becomes a pivotal part in reconciling with dissenting epistemic agents. The definitions of "reasonableness" and "rational" is taken from Freeman's glossary on Rawls' technical terms. See Samuel Freeman, *Rawls*, ed. Brian Leiter, Routledge Philosophers (London; New York: Routledge, 2007), 480–81.

[196] For more information, see Jonathan L. Kvanvig, "Perspectivism and Reflective Ascent," in *The Epistemology of Disagreement: New Essays*, ed. David Christensen and Jennifer Lackey (Oxford, UK: Oxford University Press, 2013), 226–29.

make a distinction between epistemic rationality and epistemic excusability. Kvanvig begins his distinction by noting that in legal responsibility, we make a distinction between common criminal laws and strict liability laws. In common cases of criminal law, a conviction of a criminal offence contains a *mens rea* requirement such that one is criminally liable only if they were in a culpable mental state when committing the criminal act. For example, in order for one to be convicted for murder rather than manslaughter, the criminal must be shown to have a "malicious" state of mind rather than simply a "negligent" state of mind.

An exception to this *mens rea* requirement, however, are strict liability laws, which hold a person as criminally liable simply on the basis of causal responsibility. Example of strict liability laws include selling alcohol to underage persons, or statutory rape. In these cases of strict liability laws, one is held guilty of a crime irrespective of their state of mind - that is, we do not qualify our judgment of these crimes on the basis of the mental state of the criminal. Kvanvig's own view is that we shouldn't model epistemic normativity on common criminal law rather than strict liability laws - that is, we should outright include a *mens rea* requirement on epistemic normativity. Kvanvig notes, however, that if we model epistemic normativity on strict liability laws - such that the norms of epistemic rationality do not take into account one's mental states - then we can retain our intuition for wanting a *mens rea* requirement by introducing an additional notion of epistemic *excusability*. To go back to the legal context, it is possible for a person who is convicted under a strict liability law to nevertheless be acquitted on the basis of a "good faith" defence. In *United States v. Kantor*, for instance, the defendants were charged with violating the child pornography law for filming an underage pornographic actress

engaging in sexually explicit conduct.[197] However, the defendants argued that they neither knew nor

had any reason to believe that the pornographic actress was underage. In fact, the defendants argued

that the actress (and her parents and agent) went to considerable effort to deceive the defendants about

the actress' age, which could be seen in their presenting of misleading evidence such as a false birth

certificate and false driver's license. As such, while the defendants were convicted of violating the child

pornography law under strict liability, they were acquitted on the basis of a good faith defence - that is,

that the defendants were found to have committed a reasonable mistake of fact where they took

affirmative steps to comply with the law but were misled in their efforts. Notice here that the good

faith defence does not involve showing that the defendants did not violate the child pornography law,

but acts as a further evaluation that the violation was one should be acquitted for under the

circumstances.[198] In the same way, Kvanvig's notion of epistemic *excusability* involves conceiving of

epistemic rationality as a two-stage evaluation, such that it is possible for an epistemic agent to be

irrational under the strict norms of epistemic rationality, but also *excusable* insofar as their failure to

be epistemically rational was a reasonable mistake of fact.

Putting aside Kvanvig's question of how we should conceive of epistemic normativity, we can use

Kvanvig's distinction between epistemic excusability and epistemic rationality to explain the difference

between what is reasonable and what is rational respectively. To illustrate, consider our previous

example of the person who studies for a science exam using an outdated science textbook such that

---

[197] This section relies heavily on information from Levenson's analysis of strict liability crimes, particularly in the case of
United States v. Kantor. See Laurie L. Levenson, "Good Faith Defenses: Reshaping Strict Liability Crimes," *Cornell Law
Review* 78, no. 3 (2017): 401–69.

[198] As such, Levenson argues that good faith defences reintroduce *mens rea* in the process of acquittal rather than conviction:
"The good faith defense would reintroduce mens rea, not as an element that the prosecution must prove for conviction, but as
an element a defendant must disprove for acquittal." Levenson, "Good Faith Defenses," 405.

they fail the exam. In this scenario, we can say that the person was *reasonable* in their answers: we can acknowledge the fact that they studied hard for the exam, and even if they gave the wrong answers, we could recognise that they did understand the requirements of the exam questions. In fact, in such circumstances, we might even accept their excuse as to why they failed the exam – for instance, the science teacher of the student might show that the student's failure was excusable by allowing them to take a replacement exam. But this is different to saying that the person was *rational* – in this case, for instance, we cannot use the person's reasonable efforts to pass the exam to therefore claim that the person *did* in fact give correct answers to the exam. What it means for one to be rational in their academic performance is that they know the content of the subject that they are studying, and are able to demonstrate that by giving correct answers to the exam. As such, while the student was reasonable in their attempts to pass the exam, they were not rational in their performance during the exam.

That being said, it is important to note that the distinction between reasonableness and rational is in fact quite subtle: note that we can describe both reasonableness and rationality as being dependent on whether an epistemic agent is giving the "correct" epistemic response in their situation. The difference between reasonable and rationality lies in how we define what we mean by the correct epistemic response, and what we mean by the agent's epistemic situation. So, in this example of the science student, we can say that we evaluate the student's reasonableness on the basis of their occurrent mental states, but we evaluate the student's rationality on the basis of their *performance* in the exam. We take the student to be reasonable because if we were put into their context – where "their context" is defined as their situation in which they are crucially unaware that our science textbook is outdated – we would respond in the same way as them. However, we take the student to be irrational because if we were put into their context – where "their context" defines a situation where we were required to take the same science exam as the student – we could respond in a way that is much better than them. More on this in §§4.4-4.5 below.

We can therefore use this distinction between reasonableness and rationality to clarify how and why the mentalist conception of evidence as occurrent mental states fails to be an adequate theoretical basis for the evidentialist response to disagreement. More specifically, we can see how the main problem lies in how Feldman equivocates between reasonableness and rationality: while Feldman's explanation of both evidentialism as a supervenience thesis and the Uniqueness Thesis consider evidence to determine epistemic rationality, Feldman's conception of evidence as occurrent mental states and his conception of epistemic rationality as current-state rationality suggests that evidence only determines epistemic reasonableness. In fact, within Feldman's own explanation of current-state rationality, Feldman himself alludes to this distinction at certain moments. For example, in his own explication of the example of the professor, Feldman notes that we can concede that the professor "messed up" (i.e. made an epistemic mistake, and was thus irrational) in not checking the newspaper, but still ask what is "reasonable" for the professor to believe under the circumstances.[199] Here, Feldman is correct in claiming that it is *reasonable* for the professor to believe what he did in the situation that he was in: if we were placed into the same situation in which we were unaware that the newspaper contained crucial evidence about the movie time, we would respond in the same way as the professor. But Feldman fails to see that this is different to saying that the professor was *rational*: if we define the context as being tasked to arrive at the same movie as the professor at the correct time, then it is possible to have done much better than the professor and thus be more *rational* than him.

As such, some concluding thoughts on Conee and Feldman's theory of evidence as occurrent mental states. First, it is important to acknowledge that the intuition that is motivating Conee and Feldman's theory of evidence as occurrent mental states, and Feldman's notion of current-state rationality, is an

---

[199] Feldman, *Epistemology*, 48.

We can therefore use this distinction between reasonableness and rationality to clarify how and why the mentalist conception of evidence as occurrent mental states fails to be an adequate theoretical basis for the evidentialist response to disagreement. More specifically, we can see how the main problem lies in how Feldman equivocates between reasonableness and rationality: while Feldman's explanation of both evidentialism as a supervenience thesis and the Uniqueness Thesis consider evidence to determine epistemic rationality, Feldman's conception of evidence as occurrent mental states and his conception of epistemic rationality as current-state rationality suggests that evidence only determines epistemic reasonableness. In fact, within Feldman's own explanation of current-state rationality, Feldman himself alludes to this distinction at certain moments. For example, in his own explication of the example of the professor, Feldman notes that we can concede that the professor "messed up" (i.e. made an epistemic mistake, and was thus irrational) in not checking the newspaper, but still ask what is "reasonable" for the professor to believe under the circumstances.[199] Here, Feldman is correct in claiming that it is *reasonable* for the professor to believe what he did in the situation that he was in: if we were placed into the same situation in which we were unaware that the newspaper contained crucial evidence about the movie time, we would respond in the same way as the professor. But Feldman fails to see that this is different to saying that the professor was *rational*: if we define the context as being tasked to arrive at the same movie as the professor at the correct time, then it is possible to have done much better than the professor and thus be more *rational* than him.

As such, some concluding thoughts on Conee and Feldman's theory of evidence as occurrent mental states. First, it is important to acknowledge that the intuition that is motivating Conee and Feldman's theory of evidence as occurrent mental states, and Feldman's notion of current-state rationality, is an

---

[199] Feldman, *Epistemology*, 48.

important epistemic intuition that we should not ignore. Arguably, this is the epistemic intuition that underlies internalist theories of epistemic justification more generally: namely, that our justifications for our beliefs need to be justifications that we are aware of in order for us to be epistemic *agents*. It is only when we are aware of the justifications for our beliefs that we can use these justifications to evaluate the epistemic rationality of ourselves and others, and to come to a mutual understanding by bringing forth these justifications as reasons by which we convince and are convinced by others. Nevertheless, in focusing on the justifications that we are (mentally) aware of, Conee and Feldman's theory of evidence as occurrent mental states fail to account for both the temporal and social dimensions of epistemic rationality. While we do make judgments about the *reasonableness* of an epistemic agent – that is, about the competency of an epistemic agent in responding to the evidence that is available in their immediate epistemic situation – this does not mean that we evaluate epistemic rationality only in reference to an epistemic agent's immediate situation. Rather, our judgments of epistemic rationality (of ourselves and of others) often involve considering how we can be more epistemically rational *over time* (that is, that we evaluate epistemic rationality temporally), and how we can be more epistemically rational by *learning from other epistemic agents* (that is, that we evaluate epistemic rationality socially). As such, Conee and Feldman's theory of evidence fails to satisfy both of our theoretical desiderata for the evidentialist response to disagreement. The theory of evidence as occurrent mental states fails to account for how evidence might be commonable as to allow dissenting epistemic agents to find resolution to disagreement on a common body of evidence. Furthermore, the theory of evidence as occurrent mental states fails to explain why evidence determines epistemic rationality, as it fails to account for our social epistemic practice which allow us to learn what is evidence and what is epistemically rational in the first place. We now turn to Williamson's account of evidence as known propositions.

*4.3   Timothy Williamson's "E=K" Thesis: Evidence as Known Propositions, and
Williamson's Critique of Operationalising Epistemology*

In his influential book *Knowledge and its Limits*, Timothy Williamson argues for his "E=K" thesis:

namely, the claim that "all and only knowledge is evidence".[200] Williamson's argument for this thesis

comes from the defence of three claims: that all evidence is propositional, that all propositional

evidence is knowledge, and that all knowledge is evidence. Seeing as we are interested particularly in

Williamson's theory of evidence, let us examine the first two of these three claims: that is, why

Williamson believes that all evidence is propositional, and that all propositional evidence is knowledge.

As we shall see, Williamson's defence of both claims follows a similar line of argument. In particular,

Williamson's strategy for defending both claims involves considering three epistemic practices in

which we use evidence, and arguing that only propositions fits our concept of evidence in these

practices. These three practices include: inference to the best explanation, probabilistic confirmation,

and the ruling out of hypotheses.[201] To begin, consider Williamson's argument for the claim that all

evidence is propositional from the case of inference to the best explanation:

> Consider inference to the best explanation. We often choose between hypotheses by asking which of
> them best explains our evidence – which of them, if true, would explain the evidence better than any
> other one would, if true. Fossil evidence enables us to answer questions about terrestrial life in this
> way. Even if inference to the best explanation is not legitimate in all theoretical contexts, what matters
> for present purposes is that, where evidence does enable us to answer a question, a central way for it to
> do so is by inference to its best explanation. Thus evidence is the kind of thing which hypotheses
> explain. But the kind of thing which hypotheses explain is propositional. Therefore evidence is
> propositional.[202]

---

[200] Williamson, *Knowledge and Its Limits*, 193.

[201] Williamson, *Knowledge and Its Limits*, 194–97.

[202] Williamson, *Knowledge and Its Limits*, 194–95.

According to Williamson, our inferences to the best explanation implicitly assume that evidence is the kind of thing that hypotheses explains. Williamson elaborates on this point in the subsequent passages by stating that "inference to the best explanation concerns why-explanations, which can be put in the form ' – – because . . .', which is ungrammatical unless declarative sentences, complements for 'that', fill both blanks."[203] For example, Williamson considers the injunction "Explain Albania!", and notes that such an injunction is indeterminate without a given context. What then do we require from the context in order to fill out the meaning of this injunction? For Williamson, the explanandum that is required from this injunction must be a *proposition*: "the context allows us to interpret it as an injunction to explain why Albania exists, or has some distinctive feature. What follows 'why' is a declarative sentence, expressing the proposition to be explained – that Albania exists, or that it has the distinctive feature."[204] Williamson therefore shows how in our epistemic practice of making inferences to the best explanation, we take evidence as that which is to be explained by our hypothesis, and thus to be propositional.

Similar arguments are made when considering our use of evidence in probabilistic confirmation, and in the ruling out of hypotheses. In the context of probabilistic confirmation, Williamson notes that our use of evidence to update degrees of belief through Bayesian conditionalisation requires us to consider the probability of our evidence – in other words, to use the simple version of conditionalisation, any application of Bayes' rule to find P(H|E) requires that we know P(E) amongst other probabilities. In order for us to consider the probability of our evidence, however, Williamson argues that our evidence must therefore be propositional because "what has a probability is a proposition; the probability is the

[203] Williamson, *Knowledge and Its Limits*, 195.

[204] Williamson, *Knowledge and Its Limits*, 195.

probability *that* [i.e. probability about a declarative sentence]."[205] As such, in order for evidence to be

the kind of thing which gives evidential support to our beliefs, evidence must be propositional. Lastly,

when ruling out hypotheses, Williamson notes that we use evidence in order to rule out hypotheses

which are inconsistent with our hypotheses. For instance, when we use the observation of a black swan

as evidence that rules out the hypothesis that "All swans are white", we do so on the basis that the

hypothesis "All swans are white" is inconsistent with the observation sentence "There is a swan that is

black". As we can see in this example, Williamson notes that "only propositions can be inconsistent in

the relevant sense. If evidence *e* is inconsistent with an hypothesis *h* in that sense, it must be possible to

deduce ~*h* from *e*; the premises of a deduction are propositions."[206] To conclude, Williamson's

argument for the claim that evidence is propositional involves examining how we use evidence in our

epistemic practice, and demonstrating that evidence is propositional in these practices. In our

inferences to the best explanation, in probabilistic confirmation, and in the ruling out of hypotheses,

our evidence must be propositional in order for it to be legitimate explanandum, evidential support,

and premises for deductive inference respectively.

This brings us to Williamson's second claim: that all propositional evidence is knowledge. Williamson

notes that his argument for the claim that all propositional evidence is knowledge follows the same

---

[205] Williamson, *Knowledge and Its Limits*, 196. Note that immediately following this quote, Williamson clarifies that this interpretation of probability is only "when 'probability' has to do with the evidential status of beliefs, as now; if we speak in this connection of the probability of an event, we mean the probability that it occurred." The kind of probability that Williamson is referring to in this case are evidential probabilities – see §2.2.2 for more information.

[206] Williamson, *Knowledge and Its Limits*, 196.

strategy as his argument that all evidence is propositional.[207] Just as how our epistemic practice shows

that evidence is propositional, Williamson argues that our uses of evidence in epistemic practice would

not be possible if our evidence was not known. So that we do not rehash the same argument, let us

focus mainly on Williamson's analysis of inference to the best explanation, and then see how it applies

to the other two cases. Williamson's argument for the claim that all propositional evidence is

knowledge from inference to the best explanation is as follows: "When we prefer an hypothesis h to an

hypothesis $h*$ because h explains our evidence e better than $h*$ does, we are standardly assuming e to

be known; if we do not know e, why should h's capacity to explain e confirm h for us?"[208] Imagine a

detective whose investigation over various documents, eyewitnesses, and forensics lead him to the

conclusion that the butler did it. For the detective, this inference was an inference to the best

explanation, such that the conclusion "the butler did it" best explains all of the evidence he has

gathered in his investigation. Williamson's point is this: would the inference to this conclusion be as

warranted if the detective did not *know* his evidence? Suppose the detective testified in court that he

believes his evidence but cannot know that his evidence is true. Surely such a statement would be an

admission of the unsteadiness of one's evidence, and therefore undermine in part our justification for

believing the detective's conclusion. Conversely, it is only when we are certain about our knowledge of

the evidence which allows us to be confident in our conclusion that a particular hypothesis is the best

explanation for the evidence. As such, in making inferences to the best explanation, our decision of

---

[207] "As in the previous section, the argument is from the function of evidence. Indeed, the thesis draws support from the role of evidence cited there, in inference to the best explanation, probabilistic reasoning, and the exclusion of hypotheses." Williamson, *Knowledge and Its Limits*, 200.

[208] Williamson, *Knowledge and Its Limits*, 200.

which hypothesis is the best explanation for the evidence implicitly assumes that the evidence is known.

The same is true for our uses of evidence in probabilistic confirmation and in the ruling out of hypotheses. If our evidence was not knowledge, Williamson argues that our use of evidence both to regulate our degrees of belief and to infer the falsity of certain hypotheses would not be warranted.[209] As such, in defending the claim that all propositional evidence is knowledge, Williamson's theory of evidence conceives of evidence as known propositions. More precisely, Williamson claims that "one's total evidence e_α in a case α is the conjunction of all the propositions which one knows in α."[210] Now that we have given an overview of Williamson's theory of evidence, let us critically evaluate Williamson's theory of evidence under our two proposed desiderata.

How does Williamson's theory of evidence fare under our two proposed desiderata? As a preliminary remark, it is helpful to note the similarities and differences between Williamson's theory of evidence and Conee and Feldman's theory of evidence as occurrent mental states. The main similarity between the two theories is in how both theories considers evidence to be occurrent mental states.[211] To

---

[209] Consider the following quote: ''It is likewise hard to see why the probability of *h* on *e* should regulate our degree of belief in *h* unless we know *e*. Again, an incompatibility between *h* and *e* does not rule out *h* unless *e* is known." Or: "It is hard to see how evidence could discriminate between hypotheses in the way we want it to if it did not have to be known." Williamson, *Knowledge and Its Limits*, 200–201.

[210] Williamson, *Knowledge and Its Limits*, 221–22.

[211] In their discussion of Williamson, Conee and Feldman state that Williamson's theory of evidence as knowledge should fit nicely into their conception of evidence as occurrent mental states. However, as we shall see in the subsequent section, Williamson's criticisms of the phenomenal conception of evidence will shown to be applicable to Conee and Feldman's theory of evidence as well. As such, while there are similarities between Williamson's theory of evidence and that proposed by Conee and Feldman, the differences are substantial enough for the theories to be mutually exclusive. See the afterword in Conee and Feldman, "Internalism Defended," 81.

demonstrate this is true of Williamson's theory of evidence, we can see that Williamson's theory of evidence conceives of evidence as mental states because for Williamson, all propositional evidence is knowledge, and knowledge is a state of mind.[212] Furthermore, as we saw in the formulation above, Williamson's takes one's total evidence as something which is relative to a particular epistemic situation: one's total evidence "e_α" must be described in reference to some case α. For Williamson, any description of an epistemic situation must make reference to a moment in time which one considers as the present moment.[213] Describing epistemic situations as a particular moment in time is important for Williamson because of the fact that one can "lose" evidence by forgetting it in a given moment: "if I observe the truth of e and then forget all about it, my evidence no longer includes e."[214] As such, like Conee and Feldman, Williamson's theory of evidence considers one's evidence as consisting of one's occurrent mental states.

It is therefore no surprise that Williamson's theory of evidence shares some of the theoretical strengths of Conee and Feldman's theory of evidence. For instance, just like Conee and Feldman's theory of evidence, Williamson's theory of evidence also avoids the shortcomings of the empiricist theory of sense-data since, under Williamson's theory of evidence, we can account for non-sensible entities and for epistemic values because we *know* facts about non-sensible entities and epistemic values. However, while Williamson's theory of evidence shares many similarities with Conee and Feldman's theory of evidence, Williamson's theory of evidence is crucially different to Conee and Feldman in how it restricts the type of mental states which are admissible as evidence. Specifically, Williamson's theory of

---

[212] Williamson defends that knowledge is a state of mind in Chapter 1 of Williamson, *Knowledge and Its Limits*.

[213] "A case depends on a subject (referred to by 'one'), a time (referred to by the present tense), and a possible world." Williamson, *Knowledge and Its Limits*, 94.

[214] Williamson, *Knowledge and Its Limits*, 201.

evidence restricts one's evidence to *factive* mental states because of the requirement that one's evidence must be *knowledge*. In the rest of this critical evaluation, I shall argue that Williamson's restriction of one's evidence to factive mental states makes it unsuitable for the evidentialist response to disagreement. In particular, I will argue that while Williamson's restriction of evidence to factive mental states appears to solve the shortcomings of Conee and Feldman's theory of evidence that we examined previously, this requirement actually *exacerbates* the problems which are raised against Conee and Feldman's theory of evidence because of its implication that one can be in a position where one does not know what evidence one has, nor what rationality requires of them. While these implications are Williamson explicitly states and accepts for his own purposes, I shall argue that these implications make Williamson's theory incompatible for the evidentialist response to disagreement.

As a starting point for our critical examination of Williamson's theory of evidence, we will focus primarily on Williamson's argument against skepticism – and in particular, in Williamson's claim that skepticism about the external world is motivated by what he calls the "phenomenal" conception of evidence. Williamson's argument against skepticism is important to our evaluation of his theory of evidence for two reasons: first, Williamson's criticisms of the phenomenal conception help to clarify the differences between his theory of evidence and that of Conee and Feldman; and second, Williamson's defence against external-world skepticism leads to certain implications about the nature of evidence which calls our desiderata that evidence should determine epistemic rationality into question. What does Williamson mean by the phenomenal conception of evidence? While Williamson does not give a precise definition of the phenomenal conception, the passage in which Williamson best describes this conception is as follows:

> Thus [the fact that one's evidence in the good case has the same appropriate properties as one's evidence in the bad case] drives evidence towards the purely phenomenal. We should not assume ourselves to grasp the concept of the phenomenal quite independently [of this fact]. Instead, the phenomenal may be postulated as comprising those conditions, whatever they are, which rational subjects can know themselves to be in whenever they are in them. Such conditions may be supposed to comprise conditions on present memory experience as well as on present perceptual experience.

That such conditions exist is supposedly guaranteed by the argument that rationality requires one to respect one's evidence and cannot require one to respect something unless one is in a position to know what it is."[215]

In this passage, Williamson characterises the phenomenal conception of evidence as a condition for theories of evidence: that evidence, whatever it may be, must be such that one is always in the position to know when one has evidence. Following David Lewis, Williamson notes that present perceptual experiences and present memory experiences are examples of what might count as evidence under the phenomenal conception.[216] Let us consider a few examples. First, for Lewis and Williamson, present perceptual experiences count as evidence under the phenomenal conception because when one is having a perceptual experience, one is always in the position to know that they are having a perceptual experience: you cannot have a perceptual experience of a blue sky without *knowing* that you are having a perceptual experience of a blue sky. By contrast, scientific facts such as the proposition "all vertebrates are chordates" do not count as evidence under the phenomenal conception, because one may not be in a position to know that the proposition "all vertebrates are chordates" is indeed a fact. Furthermore, under the phenomenal conception of evidence, *past* perceptual experiences would also be inadmissible as evidence, since one might not be in a position to know what past perceptual experiences one has had.

Given these examples, the condition set by the phenomenal conception of evidence seems to lead us to a picture of evidence that is strikingly similar to Conee and Feldman's theory of evidence as occurrent mental states.[217] We can argue for this similarity more directly on the basis of two reasons. First, in the

---

[215] Williamson, *Knowledge and Its Limits*, 173.

[216] For more information, Williamson suggests the following: see David Lewis, "Elusive Knowledge," *Australasian Journal of Philosophy* 74, no. 4 (1996): 553.

[217] One of the examples might lead us to question whether we should equate the phenomenal conception of evidence with Conee and Feldman's theory of evidence. Specifically, does this therefore mean that Conee and Feldman's theory of evidence

following discussion on skepticism about the external world, Williamson sees the phenomenal conception of evidence as leading one to conclude that one has the same evidence in normal circumstances and in skeptical scenarios. For Kelly, this suggests that the phenomenal conception of evidence sees evidence as "exhausted by one's subjective, non-factive mental states" – a view that is almost identical to Conee and Feldman's theory of evidence.[218] Second, both the phenomenal conception of evidence and the theory of evidence as occurrent mental states are motivated by a similar line of reasoning. As we saw in the quote above, Williamson claims that the main reason why the phenomenal conception of evidence is accepted is because of our intuition that "rationality requires one to respect one's evidence and cannot require one to respect something unless one is in a position to know what it is."[219] This line of reasoning is almost exactly the same as Feldman's defence for current-state rationality: that the norms of rationality are only satisfiable if one's rationality is evaluated by the evidence that one is currently aware of. As such, even if we can note minor differences between the two theories, the phenomenal conception of evidence is similar enough to Conee and

---

would also consider scientific facts as inadmissible into one's evidence? This is a difficult question to answer. At first glance, if an epistemic agent is currently thinking of a fact then, under Conee and Feldman's theory of evidence, that fact would be part of one's evidence. And yet, it is not clear that Conee and Feldman's theory of evidence could demarcate between believing a true and a false fact, such that believing in a true fact counts as evidence while believing in a false fact does not. As such, it is unclear whether Conee and Feldman's theory of evidence admits facts into one's evidence, or whether the theory admits occurrent beliefs into one's evidence, some of which may be beliefs in true facts. Given that Feldman's defence of current-state rationality suggests the latter, we will assume that the phenomenal conception of evidence leads to conceiving of evidence as occurrent mental states.

[218] Kelly, "Evidence," 2016, §2. We can defend the claim that Conee and Feldman's theory of evidence sees evidence as being "non-factive" mental states for similar reasons to the ones given in the previous footnote. To put it in another way, Feldman's defence of current-state rationality seems to suggest that our present available evidence cannot be determined by external factors (such as those which characterise factivity), which are outside of an epistemic agent's awareness.

[219] Williamson, *Knowledge and Its Limits*, 173.

Feldman's theory that Williamson's criticisms of the phenomenal conception of evidence will be equally applicable to Conee and Feldman's theory of evidence as occurrent mental states.

For Williamson, one of the main problems with the phenomenal conception of evidence, as well as the intuition that rationality requires one to respect their evidence, is that it leads to skepticism about the external world. To see why this is the case, Williamson notes that arguments for skepticism about the external world often proceed by a similar line of reasoning, one which involves the comparison between two cases:

> The sceptic compares a good case with a bad one. In the good case, things appear generally as they ordinarily do, and are that way; one believes some proposition p (for example, that one has hands), and p is true; by ordinary standards, one knows p. In the bad case, things still appear generally as they ordinarily do, but are some other way; one still believes p, but p is false; by any standards, one fails to know p, for only true propositions are known. As far as externalism permits, things appear to one in exactly the same way in the good and bad cases. The sceptic argues that because one believes p falsely in the bad case, one does not know p (even though p is true) in the good case.[220]

To motivate their argument, the skeptic about the external world invites us to compare our everyday experience (the good case) with that of a skeptical hypothesis (the bad case): that we are a brain in a vat, a victim of a malicious demon, or within a dream. The skeptic then argues that our phenomenal experience is the same between the two cases: that "things appear to one in exactly the same way in the good and bad cases". From this, the skeptic concludes from the fact that one believes falsely in the bad case that one does not know what is true in the good case.

Why should we accept the skeptic's final conclusion, that one does not know what is true in the good case because one believes falsely in the bad case? According to Williamson, this is where we see how the phenomenal conception of evidence, and the intuition that rationality requires one to respect one's evidence, promotes skepticism about the external world. To begin, Williamson explains how a hidden

---

[220] Williamson, *Knowledge and Its Limits*, 165.

premise to the skeptic's argument is the "same evidence" condition: that one has the *same* evidence in both the good and the bad case.[221] It is only when we accept that one has the same evidence in both cases that the skeptic's final conclusion is warranted, as one's evidence therefore becomes insufficient for showing that one is in the good case rather than the bad case (thus showing that one lacks the justification to "know" what is true in the good case). But for what reason should we accept the premise that one has the same evidence in both cases? Williamson notes, for example, that an externalist about epistemic justification would reject the claim that one has the same evidence in both the good and bad cases.[222] What leads us to assume that we have the same evidence in both the good and bad cases is the phenomenal conception of evidence: that because things appear to one in exactly the same way in both cases, one therefore has the same evidence in both cases. Furthermore, the phenomenal conception of evidence seems to give the right answer here because of our intuition that rational agents respect their evidence, and one can only respect one's evidence if one knows what evidence is. Even if one's experience in the bad case is misleading - that is, they mislead us in believing that we are in the good case - it seems rational for one to have believed they are in the good case on the basis of this experience. After all, as is often mentioned in the literature, misleading evidence is evidence nonetheless. As such, if it is possible for an epistemic agent to be rational in the bad case, then this is only possible because one's phenomenal experiences count as evidence (albeit misleading

---

[221] Williamson, *Knowledge and Its Limits*, 169. Note that this has nothing to do with the same evidence condition for characterising epistemic peerhood.

[222] More specifically, Williamson argues that under a normal skeptical hypothesis, such as that of the Cartesian demon, an externalist would say that the difference in external factors means that one has different evidence in the two cases. However, if the skeptic supposes a skeptical hypothesis which is sufficiently similar to the good case that it meets the externalist standards, Williamson says that the externalist should still reject the skeptic's claim, as the bad case is now something which is, so defined, impossible. See Williamson, *Knowledge and Its Limits*, 169.

evidence) for their belief that they are in the good case.[223] Put in another way, if we assume that one's

rationality supervenes on one's evidence, then the fact that one can be equally rational in both the good

and bad cases implies that one has equal evidence in both cases.

For Williamson, the problem with the phenomenal conception of evidence lies not only in how it leads

us to the untenable conclusion of external-world skepticism, but also in how it undermines our use of

evidence in our everyday epistemic practices. To see why Williamson thinks that the phenomenal

conception undermines our common intuitions about evidence, consider the following passage:

> That one has the same evidence in the good and bad cases is a severe constraint on the nature of
> evidence. It is inconsistent with the view that evidence consists of true propositions like those
> standardly offered as evidence for scientific theories. For example, the good case in which I see that the
> dial reads 0.407 corresponds to a bad case in which the dial does not read 0.407 but I hallucinate and it
> is consistent with everything I know that the dial reads 0.407. Since the proposition that the dial read
> 0.407 is false in the bad case, it is not evidence in the bad case. If my evidence is the same in the two
> cases, then that the dial read 0.407 is not evidence in the good case either. For similar reasons, [the
> same evidence condition] does not permit my evidence to include perceptual states individuated in
> part by relations to the environment. No matter how favourable my epistemic circumstances, I am
> counted as having only as much evidence as I have in the corresponding sceptical scenarios, no matter
> how distant and bizarre. Retinal stimulations and brain states fare no better as evidence, for in some
> sceptical scenarios they are unknowably different too. Thus [the same evidence condition] drives
> evidence towards the purely phenomenal.[224]

For Williamson, one of the odd results of the same evidence condition is that it downsizes our

evidence to match our evidence in the worst skeptical scenarios. This is a problematic result because if

one really does have the same evidence in good and bad cases, then our common intuitions about the

evidence we have in everyday epistemic practice is undermined simply by the possibility of skeptical

scenarios. To explain, Williamson's example of the dial shows how scientific observations can be

undermined by the same evidence condition. For every scientific observation one makes, there is a

skeptical scenario where one does not in fact make an observation but instead is hallucinating. But

---

[223] Williamson, *Knowledge and Its Limits*, 169–70.

[224] Williamson, *Knowledge and Its Limits*, 173.

hallucinations do not count as evidence towards truth: if one only hallucinates about the dial in the bad case, one does not have evidence that the reading on the dial is true. Under the same evidence condition, this therefore means that scientific observations do not count as evidence at all: if one does not have evidence that the reading on the dial is true in the bad case, one does not have evidence that the reading on the dial is true on the good case.

Taken to its logical conclusion, Williamson notes that the same evidence condition "drives evidence towards the purely phenomenal". One way of understanding what Williamson means here is to say that the phenomenal conception of evidence conceives of all evidence as evidence about our phenomenal *experience* of objects, rather than evidence about the object itself. To take the example of the dial, the phenomenal conception of evidence takes our evidence to only be evidence that we are having an *experience* of reading a dial. Whether we are actually reading the dial, or whether the dial exists, or whether the reading on the dial is actually true – all of these questions are questions which we can have no evidence for, since there can always be some skeptical scenario where we are not reading the dial and the dial does not exist and the reading is not true. This is an unacceptable conclusion: if the phenomenal conception of evidence is true, then this undermines our very use of evidence in epistemic practices. For instance, to take the example of evidence within a court of law, one would not be able to use physical object such as weapons or documents as evidence to the guilt of the suspect, because the only evidence which physical objects could give would be evidence about your phenomenal experience (i.e. that you are seeing something that appears to be a weapon or a document), rather than evidence about the world (e.g. about what happened at the crime scene). It is therefore unsurprising that the phenomenal conception of evidence leads to external-world skepticism: the phenomenal conception gives us *no* evidence about the external world, only evidence about our experiences.

Given the untenable consequences of the phenomenal conception of evidence, Williamson argues that we should reject the phenomenal conception of evidence in favour of his own theory of evidence as known propositions.[225] As we saw in Williamson's discussion of the dial example above, what is missing in the phenomenal conception is the idea that our evidence is constituted by *factive* mental states, such that our evidence consists of *true* propositions. Once we conceive of evidence as known (and therefore true) propositions, we can see how one actually has *different* evidence between good and bad cases. In the dial cases, for example, the difference in evidence between the case where one is reading the dial and the case where one is hallucinating about the dial is a difference in *knowledge*: in the good case, one knows the dial reading while in the bad case one does not. Intuitively, this seems to be the correct answer we were looking for to begin with: if one really did have a hallucination about the dial and realised this fact afterward, the natural response would be to change one's previous assessment of their evidence and admit that they did not *know* that there was no dial. Williamson's theory of evidence as known propositions therefore allows us to avoid the consequences of the same evidence condition, which implicitly downsizes our evidence to match the worst skeptical scenario.

Williamson's theory of evidence as known propositions also allows us to reject skepticism about the external world: if we reject the same evidence condition, then there is no reason for us to infer from the fact that we believe falsely in the bad case to the conclusion that we do not know what is true in the good case. The answer to a skeptic who compares our normal case with a skeptical case such as the brain-in-a-vat scenario is simple: one knows that one has hands in the normal case and therefore has evidence for their beliefs, while one does not know that one has hands in the skeptical case and

---

[225] While Williamson admits that the problems with the phenomenal conception of evidence do not give direct evidential support for his E=K thesis, he notes that the argument here "lays the groundwork for the equation". See Williamson, *Knowledge and Its Limits*, 180–81.

therefore does not have evidence for their beliefs. As such, Williamson's theory of evidence as known propositions gives us a better theory of evidence than the phenomenal conception, since it allows us to account for our use of evidence in epistemic practice, and also to reject the same evidence condition and external-world skepticism.

In addition to providing a substantial response to external-world skepticism, we can now see how Williamson's theory of evidence also addresses and overcomes the very problems we raised towards Conee and Feldman's theory of evidence as occurrent mental states. Rather than evaluating rationality simply on the basis of the occurrent mental states of an individual epistemic agent, Williamson's theory of evidence allows us to evaluate rationality on the basis of what is known in an epistemic community. For instance, in the example of the novice and expert bird watcher, one of the concerns that was raised towards Feldman's notion of current-state rationality is that current-state rationality fails to give us an explanation as to why the expert bird watcher has more evidence than the novice bird watcher. If we suppose that the novice and expert bird watcher are responding in the best way that they can according to their respective occurrent mental states, then under current-state rationality, the novice and expert bird watchers seem to be equally rational. However, under Williamson's theory of evidence, the difference in rationality between the novice and expert bird watcher is clearly due to the fact that the expert bird watcher has *more* knowledge than the novice bird watcher. This seems to be the intuitively correct answer: rather than evaluating the rationality of the novice and expert bird watcher on the basis of how their beliefs fit with their current thinking, we evaluate their rationality from the objective standpoint of how much knowledge they have about bird-watching. Williamson's theory of evidence therefore shows how the concept of knowledge can account for the temporal and social dimensions of epistemic rationality. In conceiving of evidence as knowledge, Williamson's theory of evidence explains how an epistemic agent can become more rational over time, as well as how they can learn from the rationality of other epistemic agents by claiming that one's evidence

grows as one acquires more *knowledge* – knowledge that allows them to be more rational over time and knowledge that is found within a community of other epistemic agents.

Seeing as the concept of knowledge is commonable and able to determine epistemic rationality, Williamson's theory of evidence seems to be the ideal candidate for the evidentialist response to disagreement. However, Williamson continues to explain how the implications of his theory of evidence undermines the very premise of the evidentialist response to disagreement entirely. Let us consider Williamson's concluding comments about the upshot of his argument against skepticism, and how Williamson uses the implications of his theory of evidence to explicitly criticise evidentialism. When summarising his argument against external-world skepticism, Williamson notes that one of the implications of his theory of evidence is that it provides a strong defence against external-world skepticism at the cost of self-knowledge: "by sacrificing something in self-knowledge to the sceptic, we stand to gain far more in knowledge of the world".[226] By self-knowledge, Williamson is referring specifically to one's knowledge about their own epistemic situation: that is, knowledge about one's own knowledge or of one's own mental states.[227] To illustrate in the context of external-world skepticism, take the skeptic's comparison between a skeptical scenario such as being a brain in a vat (the bad case), and our normal epistemic situation (the good case). Williamson's theory of evidence successfully defends against external-world skepticism by explaining how the difference in knowledge between the two cases means that one does have knowledge (and thus evidence) that one has hands in our normal epistemic situation. However, Williamson's defence does not guarantee that one has self-knowledge:

---

[226] Williamson, *Knowledge and Its Limits*, 164.

[227] Williamson argues for this more directly in his anti-luminosity argument: an argument which concludes that our mental states do not have perfect causal sensitivity to the world, and so one is not always in a position to know what mental state one is in. For more information, see Chapters 4 and 5 of Williamson, *Knowledge and Its Limits*.

that is, one may not be in a position to *know* that they have this evidence, or to know that one is in the

good case. Williamson's defence against external-world skepticism therefore explains how we have

knowledge about the external world by distinguishing between knowledge and self-knowledge, such

that we do have knowledge about the external world in the good case, even if our knowledge about our

own epistemic situation (i.e. knowledge that we are in the good case) is imperspicuous and not

guaranteed.

For Williamson, the fact that we lack of knowledge about our epistemic situation also has implications

for our knowledge of evidence, as well as our knowledge of the norms of rationality. Consider the

following passages:

> *Whatever* evidence is, one is not always in a position to know what one has of it.[228]

> Just as one cannot always know what one's evidence is, so one cannot always know what rationality
> requires of one. Just like evidence, the requirements of rationality can differ between indiscriminable
> situations. Rationality may be a matter of doing the best one can with what one has, but one cannot
> always know what one has, or whether one has done the best one can with it. If something is a method
> only if one is always in a position to know whether one is complying with it, then there are no
> methods for learning from experience. [229]

For Williamson, given his E=K thesis, the fact that we can lack self-knowledge (i.e. knowledge about

what we know) implies that we may also lack knowledge about our evidence. Specifically, if it is

possible for us to be in a position where we do not know what we know, then it is possible for us to be

in a position where we do not know what our evidence is, or how much evidence we have at that

moment. However, as we noted before, Williamson also claims that there is an intimate connection

between evidence and epistemic rationality, given our intuition that rational thinkers always respect

---

[228] Williamson, *Knowledge and Its Limits*, 178.

[229] Williamson, *Knowledge and Its Limits*, 179.

one's evidence.[230] As such, for Williamson, the fact that we can be in a position where we do not know what our evidence is also implies that it is also possible for us to be in a position where we do not know what rationality requires of us.

We are now in a position to explain why Williamson's theory of evidence cannot be used as the theoretical basis for the evidentialist response to disagreement. To explain, recall that the evidentialist response to disagreement aims to provide a general non-ideal epistemology of disagreement by showing how the correct epistemic response to disagreement is to "follow your evidence". For the evidentialist response to disagreement, "follow your evidence" is the sole norm of rationality which one needs to satisfy to ensure that one is responding rationally to a given disagreement. However, it is here where Williamson's theory of evidence cannot substantiate the evidentialist response to disagreement, because Williamson's theory of evidence suggests that there are cases in which one cannot follow one's evidence given that *one is not in a position to know what their evidence is*. The implications of Williamson's theory of evidence therefore suggests that the evidentialist response to disagreement fails to be the correct epistemic response to all disagreements, because there are certain

---

[230] Throughout his arguments, Williamson's attitude towards this intuition is complicated to say the least. In the quotes above, for instance, Williamson's conclusion that one might not know what the norms of rationality require of us clearly seem to be derived from the contrapositive of the intuition. This suggests that Williamson seems to agree with the intuition implicitly. However, in the earlier discussion, Williamson clearly saw this intuition as misleading in how it suggests the phenomenal conception of evidence (and thus, motivates the argument for external-world skepticism). Furthermore, in subsequent passages, Williamson calls the intuition as giving a "temptation" that leads us to falsely assume that rationality is cognitively transparent. See Williamson, *Knowledge and Its Limits*, 223. For the purposes of this chapter, I will take it that Williamson accepts this intuition to be a valid intuition about the connection between evidence and epistemic rationality. However, Williamson will predominantly take this connection in the negative sense to show that one's lack of evidence reveals one's inability to be rational.

disagreements wherein one does not know what their total evidence is. More on this in the next section.

To put it another way, we can say that while Williamson's theory of evidence does seem to initially satisfy our two proposed theoretical desiderata, it ultimately does not satisfy these desiderata in the way that we require for the evidentialist response to disagreement. This is because the evidentialist response to disagreement requires that evidence be commonable and determine epistemic rationality in a way that is *cognitively transparent* to us, in order for us to use our available evidence to rationally respond to or resolve disagreement.[231] To explain, recall that the reason why we stipulated that a theory of evidence must account for how evidence determines epistemic rationality is so that we can use new evidence to rationally revise our beliefs. Our desiderata therefore assumes not only that evidence determines epistemic rationality, but that evidence determines epistemic rationality in a way that is cognitively transparent to us – that is, in such a way that once we obtain new evidence, we also thereby gain new understanding as to what we should rationally believe. The same is true of our desiderata about the commonability of evidence: the reason why we stipulated that a theory of evidence must account for the commonability of evidence is so that we can figure out how we can use a common body of evidence to arbitrate between the competing claims of dissenting parties.[232] This desiderata

_____

[231] In fact, proponents of Williamson's theory of evidence have argued that this problem of transparency leads to bleak conclusions about the possibility of finding an epistemic account which explains how we should rationally respond to disagreements. See Hawthorne and Srinivasan, "Disagreement Without Transparence".

[232] In fact, Williamson does consider how our lack of self-knowledge might have implications for the commonability of knowledge, and therefore by the E=K thesis, the commonability of evidence. However, unpacking Williamson's argument for the commonability of knowledge and its implications will be outside of the scope of this essay. This is because Williamson's arguments focus on an extremely demanding definition of common knowledge taken from epistemic logic, where common knowledge refers to knowledge which everyone knows that everyone knows that everyone knows that…etc. By contrast, my notion of the commonability of evidence simply requires that evidence be transmissible within an epistemic community,

therefore assumes that evidence is commonable in such a way that epistemic agents can all recognise

the common body of evidence which they share, and therefore use this body of evidence as the basis

for a rational resolution to disagreement – in other words, our second desiderata assumes that

evidence is commonable in a way that is cognitively transparent for each member of a given epistemic

community.

But it is exactly this cognitive transparency about our evidence and about epistemic rationality which

Williamson denies. For Williamson, the rejection of the phenomenal conception of evidence means

that it must be possible for our beliefs about our evidence to be wrong: once we accept that our

evidence is evidence not only of how things seem to us (or, in Feldman's term, how we think of them)

but how things *actually are*, it necessarily follows that we can be wrong about our evidence when how

things seem to us are *not* in fact how things actually are. Williamson therefore concludes that the lack

of cognitive transparency we have about our evidence and rationality implies that there is no non-ideal

epistemic account that can be given which explains how our available evidence can be used to ensure

that we can attain rational belief.

In his concluding statements about his defence against external-world skepticism, for instance,

Williamson notes that one of the implications of his argument is that the Bayesian framework cannot

be operationalised, because one does not always have access to one's evidence: "There is a pragmatist

and subjective Bayesian project to operationalize epistemology by working only with concepts whose

application is always accessible to the agent. The argument of this chapter [i.e. about external world

---

where the scope of an epistemic community is determined accoriding to a given context. For more on Williamson about
common knowledge, see Williamson, *Knowledge and Its Limits*, 114, 122-23, 227-28.

skepticism] implies that the project is doomed to failure."[233] Within a collection of critical essays on

evidentialism, Williamson criticises evidentialism along similar lines by arguing that "one's evidence

can be radically misleading about one's own present epistemic position. If the rationality of an action

depends on one's epistemic position, then one's evidence can be radically misleading about the

rationality of the various actions available to one."[234] Here, Williamson's criticism of evidentialism

therefore poses a problem for the evidentialist response to disagreement, as the evidentialist response

to disagreement depends on using the notion of epistemic position (i.e. in evaluating epistemic

peerhood) in order to obtain higher-order evidence about dissenting epistemic agents. It is this failure

to account for the role of epistemic position within evidentialism which brings us to the final sections

of this chapter.

4.4   *A Critique of the Evidentialist Response to Disagreement: Evidentialism as End-State, Not Transition Theory*

What conclusion should we draw from our critical evaluations of these three philosophical theories of

evidence? As I have noted earlier, the fact that none of these three theories satisfy our theoretical

desiderata does not necessarily prove that there is no theory of evidence that could satisfy our

theoretical desiderata, nor does it show that no theory would be suitable for the evidentialist response

to disagreement. Nevertheless, in this section, I shall argue that the implications of our critical

---

[233] Williamson, *Knowledge and Its Limits*, 180. For more on Williamson against the possibility that epistemology can be operationalised, see Timothy Williamson, "Why Epistemology Cannot Be Operationalized," in *Epistemology: New Essays*, ed. Quentin Smith (Oxford: Oxford University Press, 2008).

[234] Timothy Williamson, "Improbable Knowing," in *Evidentialism and Its Discontents*, ed. Trent Dougherty (New York: Oxford University Press, 2011), 147. Note that in the same collection, Conee and Feldman do respond to Williamson's argument. However, while Conee and Feldman do successfully respond to the particular cases which Williamson brings up in his argument, they do not rebut the general criticism which Williamson makes about evidentialism – more on this in the next section.

evaluation of the three philosophical theories of evidence do point to a problem in the evidentialist response to disagreement. Specifically, the common way in which all three theories of evidence fail to satisfy the two theoretical desiderata needed for the evidentialist response to disagreement strongly supports Williamson's argument against the use of evidentialism as a *non-ideal* epistemology (in Williamson's terms, to "operationalise epistemology").

To explain, let us continue from where we left off with our discussion on Williamson's criticism of evidentialism. As we saw in the previous section, Williamson's theory of evidence implies that an epistemic agent can be in a situation where they do not know what their evidence is, nor what rationality requires of them. It is this insight about our limited awareness both of our evidence and the norms of rationality which Williamson argues is a major theoretical blindspot within evidentialism. More specifically, Williamson claims that the reason why evidentialism fails to acknowledge this is because it assumes an *idealised* account of rationality – one which does not represent how epistemic agents actually use evidence or the norms of rationality in epistemic practice. To highlight how evidentialism implicitly assumes an implausible and idealised conception of rationality, Williamson offers his own contextualist conception of rationality, one which he believes provides a better account of rationality precisely because it addresses the flaw in the standard view of rationality within epistemology.

It is this contextualist view from Williamson which will be the focus of these concluding sections. Understanding Williamson's contextualist account will be important for the purposes of this thesis for two reasons: first, Williamson's contextualist account provides an illuminating account of the distinction between ideal and non-ideal epistemology, one which allows us to see why evidentialism (and therefore, the evidentialist response to disagreement) is in fact an ideal epistemic theory; and second, Williamson's contextualist account of epistemic practice provides a transition epistemology which sets the scene for the pragmatist response to disagreement which we will outline in the next

chapter. Let us therefore begin this with a brief exposition of Williamson's contextualist account of

rationality. Here is the main passage which we will be focusing on:

> The standard conception of rationality depends on a distinction between the aims and methods of cognitive activity. On that conception, truth is an aim. We cannot attain it directly; we cannot follow the rule 'Believe truly!' when we do not know what is true. Therefore we must use methods to reach the truth. Rationality is a method. We can follow rules of rationality because we are always in a position to know what they require. If the [defence against external-world skepticism] is correct, this picture of rationality is mistaken. Just as one cannot always know what one's evidence is, so one cannot always know what rationality requires of one. Just like evidence, the requirements of rationality can differ between indiscriminable situations. Rationality may be a matter of doing the best one can with what one has, but one cannot always know what one has, or whether one has done the best one can with it. If something is a method only if one is always in a position to know whether one is complying with it, then there are no methods for learning from experience. But that standard is too exacting to be useful. We can use something as a method in contexts in which one is usually in a position to know whether one is complying with it, even if in other contexts one is not usually in a position to know whether one is complying with it. In that sense, we can use even believing truly as a method in contexts in which one is usually in a position to know what is true: for example, when forming beliefs in normal conditions about the spatial arrangement of medium-sized objects in one's immediate environment. In more difficult contexts, believing truly becomes an aim and we fall back on the method of believing rationally. Rationality becomes a sub-goal on the way to truth. That does not require one always to be in a position to know what rationality requires of one; it requires merely that one often knows what rationality requires when one does not know what truth requires. Nothing has been said here to undermine that requirement. In still more problematic contexts, paradoxes throw our very standards of rationality into doubt, and we fall back still further on what workable methods we can find. Cognition is irremediably opportunistic.[235]

Williamson begins the defence of his view on epistemic rationality by considering the distinction

between the means and ends of epistemology (what he calls "aims and methods of cognitive activity").

For Williamson, the standard view of epistemology takes the norms of rationality to be the means by

which we attain the epistemic end of true belief. This standard view is commonly accepted due to our

intuition that true belief is not something we can directly attain – that is, we cannot simply follow the

injunction "Believe truly!" because we often do not know what is true. In order to arrive at a position

where we can know what is true then, the standard conception states that we must therefore arrive at

true belief through some epistemic method – namely, through the norms of rationality. As such,

---

[235] Williamson, *Knowledge and Its Limits*, 179–80.

Williamson claims that this standard view is what leads to the assumption that the norms of rationality are the *sole* means of epistemology, and therefore to assume that rationality is such that an epistemic agent is always in a position to know what is rational to believe. After all, echoing White's claim for Uniqueness, given that the truth is not directly apprehensible to us, what other method *can* we use to discover the truth outside of the norms of rationality?

In response to this objection, however, Williamson argues that the standard view on epistemic rationality is mistakenly *idealised* (in his words, "too exacting to be useful") because it fails to account for the fact that we *do* use a plurality of epistemic means and ends in our actual epistemic practice depending on our *epistemic context*. As a response to the standard view, Williamson puts forward his own contextualist view where different epistemic means and ends are chosen by triage according to the felicitousness of one's epistemic situation. Williamson begins by noting that it is possible in extremely favourable epistemic circumstances for the act of believing itself to be an epistemic method by which we attain true belief. For example, suppose a child is asked to learn the mathematical statement "2+2=4" by a teacher, who gives a visual proof of the mathematical statement "2+2=4" by placing two balls in front of the child, and then two more. In this case, the injunction "Believe truly!" is indeed satisfiable by simply believing the mathematical statement.

Of course, Williamson notes that most of the epistemic situations which we encounter in everyday life are not so favourable, such that we do not have an obvious method of attaining true belief. In these situations, the standard view is correct in pointing out that we use the norms of rationality as a means towards attaining true belief. To take an example, suppose you look up to see cloudy skies, and consider whether you should bring an umbrella. In this case, since you do not know how the weather will unfold over the course of the day, simply believing does not suffice as a method of deciding *what* to believe or of ensuring that your belief will be *true*. What you do know in this situation, however, is that cloudy skies like the ones that you see typically indicates rain. As such, while you might not know

what is true, you know what is rational to believe: namely, that you should bring an umbrella because cloudy skies mean that it is likely to rain. Your decision of what to believe therefore uses rationality as the means which gives you the best chance of attaining the epistemic end of true belief. More specifically, Williamson claims that in these cases, we use the norms of rationality to reach rational belief as an *epistemic end* which itself is a means to the end of true belief: in his words, that "rationality becomes a sub-goal on the way to truth".

However, Williamson continues to explain that there is a third kind of situation which the standard view of epistemology does not account for: namely, unfavourable epistemic situations where one does not even know what rationality requires of them. To give an example, suppose you are a philosopher who is attending an academic conference on an area of research in which you have absolutely no expertise – say, in particle physics. During question time, your colleague sitting next to you quietly asks you what you think about a particular debate which has been present through the respective presentations. Perhaps in such a situation, your lack of understanding about physics is such that you have no idea what the true physical theory is, nor do you know what would be rational to believe about this debate.

Nevertheless, Williamson argues that even in these situations, the "irremediably opportunistic" nature of our cognition means that still find other means to decide on what is best to believe in a given situation. This is because, irrespective of how unfavourable our epistemic circumstances are, circumstances within our everyday life often force us to make a decision about action, and therefore a decision about what to believe: in our everyday decision-making, suspension of judgment and belief are often not options simply because *inaction* is not an option. As such, even in unfavourable circumstances, Williamson argues that our epistemic practice (in Williamson's word, "cognition") involves finding the best epistemic means available to us. To return to our previous example, while you as a philosopher may have absolutely no idea what is rational to believe about particle physics, it is still

possible for you to do *something* about your decision on what to believe: you can, for instance, consider possible objections to the position of the respective presenters, or you can evaluate which position you think is most plausible by analogising from your own field of expertise. To use our previous distinction between reasonable and rational, it is important to note that your actions do not ensure that you will be *rational* in your belief, seeing as it is quite possible that your attempts to formulate possible objections or analogies from your own research fail to be valid objections or analogies. But you can be at least *reasonable* in your beliefs insofar as you have attempted to give *some* reasoning for your beliefs – reasons which may even lead one to rational belief depending on whether one's reasons are in fact valid reasons (e.g. that the objections that we raise about a particular physical theory are in fact valid objections).

Williamson's contextualist view therefore provides a richer account of our epistemic practice than the standard view of epistemology: one which highlights our choice of epistemic methods on the basis of the felicitousness of our epistemic situations. In extremely favourable epistemic circumstances, one might simply have a direct method of attaining true belief such that no norms of rationality are necessary. However, in circumstances where one does not have a direct method of attaining true belief, one is often in a position where one knows what rationality requires of us, and can therefore use the norms of rationality as a means to attain true belief. And in circumstances where one does not know what rationality requires of us, one can still resourcefully use whatever means is available in one's present situation, in the hopes of attaining rational belief, and thus possibly attaining true belief.

We are now in a position to explain Williamson's criticism of the standard view of rationality within epistemology, and how this connects to Williamson's criticism of evidentialism – and specifically, in the attempt to use evidentialism as a *non-ideal epistemology*. To begin, we can clarify Williamson's criticism of the standard view of rationality by referring to our familiar distinction between ideal and non-ideal epistemology – and more specifically, between *transition* and *end-state* epistemic theories.

Recall the distinction between transition and end-state theory: whereas end-state theories provide a theoretical explication of a concept which serves as a long-term end for our practice (ethically, politically, epistemically, etc.), transition theories provide a contextual account of progressive reforms one can make to improve our practice in a given context. Given this distinction, Williamson's criticism of the standard view about epistemic rationality can be understood as claiming that the standard view is making a category error in assuming that theories of epistemic rationality are *transition* rather than *end-state* theories. In conceiving of rationality as the sole epistemic means by which we arrive at true belief, the standard view of rationality sees epistemic theories of rationality as *transition* epistemic theories: that is, as epistemic theories which explicate context-sensitive norms by which we can use to arrive at true belief. To take the example of Bayesian epistemology, the standard view (which Williamson sees as being held by subjective Bayesians)[236] takes the classic Bayesian principle of conditionalisation as a *transition* norm by which an epistemic agent can always use conditionalisation to revise their degrees of belief to be in line with the evidence they have in a given context, and thus ensure that their beliefs are rational.

For Williamson, however, the standard view of rationality is mistaken because it fails to recognise how any theoretical explication of the norms of rationality is an *end-state* as opposed to a transitional epistemology. At this point, it is important to note that one of the key differences between end-state and transition theories is in the kind of assumptions it makes about *epistemic context* within epistemic theorising. In particular, end-state theories necessarily make idealised assumptions about epistemic context in order to provide conceptual clarification about a given concept. To illustrate with our familiar example of Rawls' theory of justice, Rawls argued that before we consider the practical

---

[236] For more on Williamson's criticisms of subjective Bayesianism, see the citations given in fn. 233.

question of *how* we can make our actual society more just (i.e. how we conduct transition political theory), we must first answer the conceptual question of *what* a just society looks like in the first place (i.e. we require an end-state theory of justice). As such, Rawls' end-state account of the just society makes certain idealised assumptions about the social conditions of the just society, since the purpose of his account is to serve as a long-term political end for how our society *ought* to be were it to be just, rather than as a descriptive account of what our actual societies are like. Rawls' political theory therefore reveals how in order to provide a theoretical account of a particular concept (such as justice or rationality), end-state accounts must therefore make idealised assumptions about epistemic context in order to show what it would look like to be in a position where one can attain justice and rationality as a political or epistemic end respectively.

Returning back to Williamson, we can see that Williamson seems to be making the same point when he claims that "rationality is a sub-goal on the way to truth": while the norms of rationality do serve as a means by which one can increase the likelihood that their belief will be true, Williamson argues that rational belief is also an epistemic *end* insofar as certain conditions within one's context must be met in order for one to be in a position to follow the norms of rationality in the first place. Williamson therefore reveals how theories of epistemic rationality are necessarily end-state theories because any account of the norms of rationality implicitly assumes that one is in a position to know or act in a way which satisfies the norms of rationality. To take our previous example of Bayesian conditionalisation, notice that the Bayesian principle of conditionalistion requires that one's epistemic context meets certain conditions in order for it to be possible to follow principle of conditionalisation at all: conditions such as having certain evidence such that the probability of their evidence is 1, or that an epistemic agent knows and is able to calculate conditional probabilities. To the extent that an epistemic agent is in a position where these conditions are not fulfilled – that they have evidence which they are uncertain of, or that they cannot calculate conditional probabilities – the Bayesian principle of

conditionalisation is a norm of rationality which they cannot follow.[237] Bayesian conditionalisation therefore acts as a principle for how we might arrive at rational belief as an epistemic *end*, an end which might be unattainable depending on our epistemic situation. (For a similar point on decision theory, see §1.3.2.)

By contrast, it is only in Williamson's overarching contextualist account about epistemic means and ends where we find a transition epistemic theory which does account for how we attempt to improve our epistemic practice within a given context. For Williamson, the way in which we improve our epistemic practice does not come from the norms of rationality themselves, but from our cognitive ability to choose the epistemic means which we believe will give us the best chance at attaining true belief in our epistemic situation. Williamson therefore highlights how it is our *cognition*, not the norms of rationality, which allows us to make context-sensitive decisions about what we should do in our epistemic practice. Just as transition political theories reveal how reforms in distributive justice involve different policies and objectives in different social and political contexts (e.g. how the strategies and objectives that a developing nation might use to contribute to global re-distribution of wealth would differ to those of a developed nation), Williamson's contextualist account reveals how our attempt to attain true belief involves different epistemic methods in different epistemic contexts. Within Williamson's contextualist account, theories about epistemic rationality still play an important role in explicating how different norms of rationality might apply or be used in certain epistemic situations.

---

[237] In the previous chapter, we discussed briefly the possibility of whether further research in Bayesian epistemology might modify ideal assumptions (like the assumption that one assigns probability 1 to their evidence) in order to address non-ideal cases. Given Williamson's contextualist account, however, we can see that while further research in Bayesian epistemology might allow us to provide more nuanced norms of rationality which apply to a greater range of epistemic situations, Bayesian epistemology would still be an end-state epistemology insofar as it determines norms of rationality which are only applicable in certain epistemic situations. This is nevertheless a fine result, as we saw in the diverse applications of Bayesian epistemology in various fields of research. For more information, see fn. 91.

However, the fact that these theories about epistemic rationality are end-state rather than transition theories means that the satisfiability of a norm of rationality (and thus, the attainment of the end of rational belief) is not guaranteed, but depends on the felicitousness of one's epistemic situation.

We can now see how Williamson's critique of the standard view of rationality also applies to the evidentialist response to disagreement. This is because we can similarly criticise the evidentialist response to disagreement for making a category error in its assumption that evidentialism is a *transition* account of epistemic rationality, rather than an end-state account. In claiming that the correct epistemic response to any given disagreement is to follow one's evidence, the evidentialist response to disagreement assumes that evidentialism is a *transition* account of rationality: one which sees the evidence that we have at any given moment as the sole means by which we evaluate and acquire rational belief.[238] For Williamson, however, this is an incorrect assumption because evidentialism is clearly an *end-state* account of rationality, given that it assumes that one is already in a position where one knows what one's total evidence is, and what doxastic justification one receives from their evidence (i.e. what rationality requires of them).[239] Once we see how this assumption is not

_____

[238] Recall, for instance, the evidentialist conception of epistemic rationality from Kelly which we discussed in §2.2.1: "By epistemic rationality, I mean, roughly, the kind of rationality which one displays when one believes propositions that are strongly supported by one's evidence and refrains from believing propositions that are improbable given one's evidence." Kelly, "Epistemic Rationality as Instrumental Rationality," 612.

[239] The following argument reveals how the aim of finding a theory of evidence which substantiates the evidentialist response to disagreement is wrongheaded, because having a substantial account of evidence does show that one is always in a position to know what one's evidence is in actual epistemic practice. As a quick sidenote, Williamson's claim that one might not know what rationality requires of us suggests that a substantial account of evidential support would not help the evidentilaist response to disagreement either. This is because the discussion around evidential support within the literature is also a discussion about ideal epistemology: given that the debate between subjectivism and objectivism about evidential support is purely interested in the conceptual question about the degree to which a body of evidence supports a given proposition, and not the practical question of what one should do if one is not in a position to know the degree to which their evidence supports a given proposition. As such, even if we have a substantial account of how a body of evidence supports a hypothesis,

guaranteed in actual epistemic practice, the evidentialist norm to "follow your evidence" is in fact an end-state norm of rationality, one which is only applicable when one's epistemic context satisfies certain conditions.

It is here where we see how our critical evaluation of the three philosophical theories of evidence in §§4.1-4.3 supports Williamson's view. This is because the reason why each of these theories of evidence fails to meet the theoretical desiderata of the evidentialist repsonse to disagreement is because in each theoretical account of evidence, there remains cases where one's evidence does not provide doxastic justification for one's beliefs because one does not know what is rational to believe on the basis of their available evidence. The fact that each of these theories of evidence fail to meet the theoretical desiderata in the exact same way therefore suggests that the problem lies with the theoretical desiderata – and, in particular, that our assumption that evidence determines epistemic rationality (an assumption which is motivated both by the supervenience thesis of evidentialism and the Uniqueness Thesis) is false, because our evidence does not always provide doxastic justification in our actual epistemic practice.

To illustrate this in further detail, let us briefly revisit the conclusions of each of our critical examination into the three theories of evidence. In the case of the empiricist theory of evidence, we saw that the reason why the empiricist theory of evidence as sense-data is widely rejected within contemporary philosophy is because the empiricist theory of evidence is unable to account for cases (like the phlogiston theory) where one has a sufficient amount of sense-data, but does not know what

─────────────────────────

the evidentialist response to disagreement still fails to be the correct epistemic response in situations where one does not know what proposition is supported by one's evidence. Refer back to §3.1 for more on the topic of evidential support. For more specifically from Williamson about the cognitive transparency of evidential support, see Williamson, "Improbable Knowing".

is rational to believe on the basis of the sense-data. The critical discussion around the empiricist theory

of evidence therefore reveals how our evidence does not come directly from our sense perception, but

requires our cognition to evaluate what data we do receive from our sensory faculties.

To address this problem, Conee and Feldman's theory of evidence as occurrent mental states provides

a more liberal conception of evidence which explains how evidence can be obtained not only from

one's perception, but also from sources like induction and memory. However, just like the empiricist

theory of evidence, the main criticism against Conee and Feldman's theory of evidence as occurrent

mental states comes from cases where one's occurrent mental states do not provide doxastic

justification for one's belief, because one's occurrent mental states are radically mistaken about how

the world actually is. In cases like flat-world theorists or proponents of eugenics, we are reluctant to

say that their occurrent mental states are sufficient for determining what is rational for them to believe

because one's evidence should not be determined solely by what one is currently thinking of, but also

whether one's current thoughts are in correspondence to how the world actually is.[240]

It is in response to these kinds of concerns which Williamson proposes his theory of evidence as

known propositions, such that we can evaluate which occurrent mental states are evidence by

considering which of our mental states are true, and therefore knowledge. But, as Williamson himself

explains, it is still possible under Williamson's theory of evidence for one to not receive doxastic

justification from one's evidence, given that it is possible for one to possess evidence without knowing

that one possesses evidence. When relying on testimonial evidence, for instance, it is entirely possible

---

[240] This therefore suggests that the intuition that "misleading evidence is evidence nonetheless" – which is oft discussed in the literature on evidentialism – is in fact *false*. Just as how a knowledge claim is different from knowledge – that is, claiming that *p* is knowledge is different to *p* being in fact knowledge – so the claim that something is evidence (even though it might be in fact misleading) should not be confused with something *actually* being evidence.

for one to receive evidence from a source of testimony without knowing that one has evidence

(e.g. because one does not know when testimony is not actually true, and therefore not actually

evidence).[241]

As such, the general upshot which we learn from our critical evaluation of these three theories of

evidence is that our ability to obtain doxastic justification from evidence depends on factors which

reside *outside* of our evidence: in the empiricist theory of evidence, our ability to obtain doxastic

justification from sensory information is dependent on our ability to intepret sensory information;

under Conee and Feldman's theory of evidence, our ability to obtain doxastic justification from our

occurrent mental states is dependent on whether our thinking is correct; and under Williamson's

theory of evidence, our ability to obtain doxastic justification from our knowledge is dependent on

whether we are a position to know what we know. This means that while the evidentialist claim that

one's doxastic justification supervenes on one's evidence might be true as a theoretical claim about the

doxastic justification one *ought* to have from receiving evidence, we can fail to obtain this doxastic

justification from evidence in *actual* epistemic practice because our epistemic context may not satisfy

the conditions which allow us to know what our evidence is, and or what doxastic justification we

should receive from evidence. As such, our critical evaluation suggests that the evidentialist norm to

---

[241] On a similar note, within the literature on the episteology of testimony, the claim that testimony is the *transmission* of knowledge has been criticised on the basis of cases where the testifer may relay a piece of knowledge to a hearer that the testifier themselves does not know. For example, Jennifer Lackey gives the case of the devout creationist who works as a science teacher, and is forced to teach evolutionary theory to their students. In this case, while the devout creationist is able to reliably convey knowledge about evolutionary theory to the students, the devout creationist does not possess this knowledge themselves because they do not believe it. Here, I argue a similar point from the perspective of the hearer: just as it is possible for a testifier to convey knowledge which they themselves do not know, it is possible for a hearer to receive knowledge which they do not know. For more on the transmission view of knowledge, see Jennifer Lackey, *Learning from Words: Testimony as a Source of Knowledge* (Oxford University Press, 2008).

"follow your evidence" is not a transitional norm, but an *end-state* norm: one which does allow us to rationally revise our beliefs in certain situations, but one which is not always available to us in practice.

This brings us to our main critique of the evidentialist response to disagreement. As we have seen, the reason why the evidentialist response to disagreement can fail to determine the correct epistemic response to a given disagreement is because one may be in a position where one does not know what one's evidence is, or what doxastic justification one receives from one's evidence. As I shall now argue, the evidentialist norm is therefore particularly ill-suited to addressing the non-ideal epistemic problem of disagreement because *disagreements are precisely the kind of unfavourable epistemic situation where one is placed in a position where they do not know what their evidence is or what is rational to believe*.

To explain, let us return to Feldman's explanation of the evidentialist response to disagreement:

> To the extent that there is a general evidentialist answer to questions about what you should do in response to learning of peer disagreement, it is this: follow your evidence…evidentialism seems to me to provide exactly the right way to think about disagreement. It instructs us to ask how learning about a disagreeing peer affects one's evidential situation. It asks us to reflect on what one should think, now that one has this new information about the disagreement.[242]

Here, notice that Feldman's explanation of the evidentialist response to disagreement implicitly assumes that one knows what one's evidence is (that is, that one knows "how learning about a disagreeing peer affects one's evidential situation"), and that one knows what rationality requires of us (that is, that one knows "what one should think…now that one has this new information about the disagreement"). But once we accept that these assumptions are not guaranteed in epistemic practice, it is not clear that Feldman's conclusion that evidentialism gives us the right way to think about disagreement is necessarily true.

---

[242] Feldman, "Evidence of Evidence Is Evidence," 287.

Take, for example, Feldman's conclusion about the ideal case of peer disagreement. In the same paper, Feldman argues that peer disagreement is a situation where one does know what one's evidence is and what is rational to believe: given that encountering a dissenting peer is always higher-order evidence that one that one is mistaken, the rational response to encountering any dissenting peer is therefore to suspend judgment.[243] But it is not clear that this is always the correct epistemic response to peer disagreement. For example, recall that in our exposition of the Total Evidence View, the case of the Holocaust denier revealed how it is possible for one's total evidence to outweigh the higher-order evidence one receives about the epistemic peerhood of a dissenting party, such that one's total evidence may overturn one's judgment that a dissenting party is a peer at all. As such, contrary to Feldman's conclusion about the peer disagreement, there is no single epistemic response to peer disagreement because, depending on one's total evidence, one might decide to revise one's belief on the basis of learning about a dissenting peer, or revise one's judgment about a dissenting party being an epistemic peer on the basis of one's previously held beliefs.

It is this decision between believing in one's original evidence for a particular belief, and believing one's evidence about the epistemic peerhood of a dissenting party which highlights the deeper epistemic problem which disagreement poses. To explain, suppose we encounter a disagreement where we do not know whether to acknowledge a dissenting party as an epistemic peer (and thus to take their disagreement as higher-order evidence against our current beliefs), or whether our present available evidence is sufficient to judge the dissenting party to *not* be an epistemic peer (and thus to reject their

---

[243] "The conclusion of the previous section is that in virtually every case of peer disagreement, where the peers are generally reasonable people who are responsive to evidence, each peer gets evidence against his or her prior view…This suggests that, at least for these cases, [disagreeing peers] cannot justifiably believe different things [and that disagreeing peers] should respond by suspending judgment." Feldman, "Evidence of Evidence Is Evidence," 293–94.

disagreement as providing contrary evidence to one's belief). How then should we follow our evidence in this situation?

This is the deeper non-ideal problem about disagreement which we discussed in §1.2.2: namely, that disagreement poses an epistemic challenge to our beliefs by placing us in a situation where we are divided between trusting the opinions of others, and trusting the opinions of ourself. With regards to this non-ideal epistemic problem, the evidentialist response to disagreement seems to beg the question because the epistemic problem of disagreement comes precisely from the fact that we have evidence both for our own belief about the given proposition, and for our belief in the epistemic reliability of the other. As such, the epistemic problem of disagreement cannot be resolved from following one's evidence because part of the problem is that one has conflicting pieces of evidence.

Furthermore, it is important to notice that in many of these problematic cases of disagreement, one actually does not have as much evidence about the epistemic peerhood of the dissenting party as the evidentialist might believe. This is because there is an important difference difference between the "absence of evidence" and the "evidence of absence" when it comes to epistemic peerhood: that is, there is a difference between a situation where you have *no* evidence that a dissenting party is more or less likely to be correct about the topic as you do, and a situation where you *do* have evidence that a dissenting party is *not* more or less likely (and therefore equally as likely) to be correct about the topic as you do.[244]

---

[244] In fact, in the latter situation, it is not clear that having evidence that a dissenting party is not more or less likely to be correct as you is the same as having evidence that a dissenting party is *equally* likely to be correct as you (and therefore an epistemic peer). To explain, within formal epistemology, proponents of the Dempster-Shafer theory of evidence have criticised the Bayesian assumption that one's evidence should be represented as a probability function, claiming that the assumption that one's evidence always follows the probability axioms is an idealised assumption that overdetermines the evidence one has in actual epistemic practice. By contrast, the proponent of the Dempster-Shafer theory claims that the

Let us illustrate the difference through the example of the Jury Case which we examined in §3.2. Part of the reason why the Jury Case is often used as the starting case within the literature on the epistemology of disagreement is because the Jury Case is supposed to be an example of the kind of disagreement which best approximates the case of peer disagreement in our everyday life. Given that the Jury Case is one where you and your friend are attending the same trial proceedings, and one where you are supposed to have known this friend for many years, the Jury Case approximates the case of peer disagreement because it is a situation where you have evidence that your friend has the same relevant body of evidence as you, and you have evidence that your friend is equally disposed to respond to the evidence in the right way as you.

However, upon further examination, it is not clear that the Jury Case should be considered a case of peer disagreement because it is unclear whether you do in fact have evidence for believing that your friend has the same amount of evidence as you, or that your friend is equally disposed to respond to the evidence in the right way as you. This is because in both cases, it is unclear that you have actual

---

evidence that we obtain in our actual epistemic practice often does not conform to the probability axioms, because there are many cases where the fact that we obtain a piece of evidence that $P(H) = 0.6$ does not allow us to conclude that $P(\sim H) = 0.4$. In this way, the Dempster-Shafer theory claims that evidential support does not follow the probability axioms, and in particular, finite additivity. Here, the Dempster-Shafer theory helps to formalise the claim that the "absence of evidence" is not the "evidence of absence": if one's evidence does not conform to the probability axioms, then the fact that one does not have evidence that a dissenting party is not an epistemic inferior or superior does not entail that one therefore has evidence that a dissenting party is an epistemic peer. To have no evidence is to have no evidence. Furthermore, as we saw in the starting claim of this footnote, even if we do have positive evidence that a dissenting party is not an epistemic inferior or superior, this does not mean that we have evidence that a dissenting party is an epistemic peer: once we reject the axiom of finite additivity, we cannot claim that our degree of belief $bel(\sim A \text{ and } \sim B)$ is the same as our degree of belief $bel(C)$, even if A, B, and C are the only outcomes within the sample space. For more on the Dempster-Shafer theory of evidence, see A. P. Dempster, "Upper and Lower Probabilities Induced by a Multivalued Mapping," *The Annals of Mathematical Statistics* 38, no. 2 (April 1967): 325–39; Glenn Shafer, *A Mathematical Theory of Evidence* (Princeton; London: Princeton University Press, 1976).

evidence for these claims as opposed to no evidence to the contrary. To begin with the latter, the reason why we accept the claim that you have evidence that your friend is equally disposed to respond to the evidence in the right way as you is because of the fact that in this case, your friend is someone which you have known for many years, and whom you know to be a generally reasonable person. But to have evidence that your friend is a generally reasonable person is not the same as having evidence that your friend is equally likely to respond to courtroom evidence in the right way as you: for to have evidence about a person's general proficiency in one area (e.g. mathematical reasoning) is not to have evidence about a person's general proficiency in another area (e.g. interpreting psychological motivation). To the extent that you do not have any actual evidence about your friend's proficiency in legal reasoning, it is not clear that knowing that your friend is a reasonable person constitutes having evidence that they are equally likely to respond to the courtroom evidence in the right way as you. Instead, the fact that your friend has always been a generally reasonable person only serves to indicate that you do not have evidence that your friend is an epistemic inferior: that is, you do not have evidence that your friend is generally unreasonable, and therefore less likely to respond to the evidence in the right way as you.

This is also true when considering the same evidence condition of epistemic peerhood. Here, part of the reason why the Jury Case is supposed to satisfy the same evidence condition is because you and your friend have attended the *same* legal proceedings, are therefore privy to the same testimonies and cross-examinations as each other. However, once we acknowledge that one's evidence can also come from one's past experiences and memory, it is not clear that you do have evidence that your friend has equally good evidence about the verdict as you do. To take but one example, to the extent that verdict of a trial might involve judgments about human character and motivation, it is hard to see how you can assess whether you and your friend have the same evidence or not, seeing as our judgments about human character and motivation are dependent on evidence we've obtained from our personal

experiences and relationships in ways which are too complex to conceptualise, let alone communicate to another person. As such, the fact that you and your friend are both members of the jury does allow you to conclude that your friend is not missing out on some pertinent information about the trial that you have, and that your friend does not have access to additional information about the trial that you do not have. But from this fact, it does not necessarily follow that you have evidence that your friend does not have more or less relevant evidence about this case as you do.

The point of these criticisms is therefore to show that just as there are situations in which one does indeed have higher-order evidence about the evidence and dispositions of another epistemic agent, so too are there situations in which one does *not* have evidence about the evidence and dispositions of another epistemic agent. And in the latter case of situations, one cannot infer from the fact that one has no reason to believe that a dissenting party is any more or less likely to be correct about the topic as oneself, that one therefore *has evidence* that a dissenting party is an epistemic peer. In this way, the evidentialist notion of "higher-order evidence" does not help to solve the non-ideal epistemic problem of disagreement: to the extent that a disagreement poses an epistemic challenge to our belief, such a disagreement is often problematic because we do not have sufficient evidence to arbitrate between the competing claims of ourselves and the dissenting party, and therefore are shown to not have the higher-order evidence needed to evaluate epistemic peerhood. Evidentialism therefore cannot provide an adequate answer to the question of how we should respond to the epistemic problem which disagreement poses: given the evidentially deficient situation which disagreement places us in, the way to respond or resolve disagreement cannot depend on norms about revising our belief upon our available evidence, but must be based in norms about how we acquire better evidence, or arrive at a better epistemic position. In other words, what we need is a *transition* epistemology of disagreement: one which allows us (and the parties which we disagree with) to *improve* our epistemic practice, by allowing us to come to a better epistemic position (both as individual epistemic agents, and collectively

as parties within a given disagreement) by which we do have the resources to know what the rational response to disagreement is, and how to settle collective belief.

4.5   The Pragmatist Account of Evidence: On the Importance of Social-Epistemic Context for Determining Evidence

This brings us back to the pragmatist theory of inquiry. As we discussed in §1.3.2, the pragmatist theory of inquiry provides an epistemic framework which is especially conducive to evaluating non-epistemic issues. This is because the pragmatist focus on the *process* of individual and social inquiry allows the theory of inquiry to provide an epistemic account which re-contextualises the epistemic (end-state) concepts of justification, rational belief, and knowledge from the perspective of an individual epistemic agent or a community of epistemic agents as they acquire these ends within epistemic practice.

In this concluding section, I will explain how the pragmatists use the theory of inquiry to provide a non-ideal epistemic account of evidence – one which shifts the focus from evidence as an epistemic end which one has in a given situation, to evidence as it is gathered, evaluated, and revised over the course of social inquiry. To illustrate this, I will draw from Hilary Putnam's criticism of Rudolf Carnap's inductive logic, because Putnam's argument serves as a helpful bridge between our previous criticism of the evidentialist response to disagreement, and the pragmatist theory of inquiry. More specifically, Putnam's own explication of how evidence functions within scientific inquiry provides a contextualist account that is both similar and more elaborate than the one given by Williamson. Similar to Williamson's claim about our use of cognition in choosing epistemic means within a given context, Putnam's account sees inquiry – as exemplified within scientific inquiry – as the *cognitive* process by which we choose a series of epistemic means to experimentally attempt, enact, and evaluate until we arrive at an answer which settles the initial question or doubt. However, where Putnam's account starts to distinguishes itself from Williamson's account is in the pragmatist emphasis on the

social dimensions of our cognition – and specifically, the ways in which social inquiry uses a

community of epistemic agents to *improve* our cognitive choice of epistemic means, as we obtain *better*

epistemic means from each other through cooperative discussion and deliberation.[245] In this way, the

pragmatist theory of inquiry will be shown to provide a better transition epistemology than

Williamson's contextualist account of epistemic rationality, because the theory of inquiry is able to

account for the important role which epistemic context plays within our epistemic practice, while also

explaining how can *improve* our epistemic context through cooperation with other epistemic agents in

social inquiry.

Let us now consider Putnam's critical comparison of the pragmatist theory of inquiry to Carnap's

inductive logic. From the outset, we can see how Putnam's criticism of Carnap parallels our previous

criticism of evidentialism, given that the main focus of Putnam's criticism is how Carnap's inductive

logic provides an *idealised* conception of scientific rationality. More specifically, Putnam begins his

criticism by noting that while Carnap claims that his inductive logic is a rational reconstruction of the

scientific method, his inductive logic does not make any reference to the role of experiments within the

---

[245] Of course, the literature on the pragmatist theory of inquiry has also considered inquiry as it is conducted by a single epistemic agent in as much detail as it has social inquiry. However, there are two reasons for why we will focus on social inquiry for the remaining parts of this thesis. First, the pragmatist response to disagreement focuses almost exclusively on social inquiry because, as I shall argue in the next concluding chapter, it is this notion of social cooperative inquiry which the pragmatists argue provide the best epistemic means by which a collective group of epistemic agents can come to settle disagreement. Second, at least within the Deweyan account of the pragmatist theory of inquiry, the notion of social inquiry takes precedence over individual inquiry because Dewey highlights how our ability to conduct effective individual inquiry is something which we learn from our social epistemic environment, and the knowledge and understanding which is available in our social epistemic environment is the result of social inquiry. This means that, for Dewey, "effective intelligence is not an original, innate endowment. No matter what are the differences in native intelligence (allowing for the moment that intelligence can be native), the actuality of mind is dependent upon the education which social conditions effect." John Dewey, *The Public and Its Problems: An Essay in Political Inquiry*, ed. Melvin Rogers (Athens, Ohio: Swallow Press, 2016), 226. For more information from Dewey on social inquiry, see Chapter 24 (entitled "Social Inquiry") in Dewey, *Logic*.

scientific method, such that Carnap's inductive logic does not explain why we should conduct experiments within scientific inquiry at all.[246] Of course, in Carnap's defence, one might argue that the reason why Carnap's inductive logic does not consider the role which experimentation plays within the scientific method is because Carnap's inductive logic is not interested with the methodological question of how scientific inquiry is done, but the epistemic question of what makes the hypotheses and theories of science more justified and rational to believe than other theories, such as those of pseudoscience or myth. It is this focus on the epistemic question which leads Carnap to characterise scientific reasoning (and arguably, epistemic rationality) in terms of the confirmation of scientific hypotheses by empirical observation – a characterisation which remains dominant within contemporary philosophy of science and, in particular, Bayesian confirmation theory (see §2.2.2 for more details).[247]

---

[246] In fact, Putnam notes that Carnap's seminal work on inductive logic – that is, *The Logical Foundations of Probability* – does not contain a single reference to experiment at all. See Hilary Putnam, "Pragmatism and Moral Objectivity," in *Words and Life*, ed. James Conant (Cambridge, Massachusetts; London, England: Harvard University Press, 1994), 170. For more discussion on Carnap's inductive logic, see Appendix A of this thesis.

[247] Why should we think that Carnap's inductive logic explains not only the rationality of scientific practice, but also of epistemic rationality more generally? According to Carnap himself, the goal of his inductive logic was to provide an alternative probabilistic explication of the notion of "confirmation", an explication which did not suffer the same consequences as the logical positivist principle of verification: "But I tried a new approach. I believed that the logical concept of probability should supply an exact quantitative explication of a concept which is basic in the methodology of empirical science, viz. the concept of the confirmation of a hypothesis." As such, just as how the verifiability criterion of meaning was used not only to explain scientific statements, but empirical statements more broadly, Carnap's inductive logic applied to how empirical statements in general were confirmed, not simply to scientific hypotheses and theories. This is especially clear when we note that Carnap's previous focus on the verificationist criterion of meaning was not simply to distinguish scientific from non-scientific hypotheses, but to show how ethical, metaphysical, and religious statements were cognitively meaningless. As such, Carnap's inductive logic should be taken as a theory of epistemic rationality insofar as it is meant to provide a general account of inductive inference, rather than a specific account of scientific reasoning. For more on Carnap's non-cognitivism about value statements, see Appendix A of this thesis. Rudolf Carnap, "Carnap's Intellectual Autobiography," in *The Philosophy of Rudolf Carnap*, ed. Paul Arthur Schilpp (LaSalle, IL: Open Court, 1963), 72. For more information on how

However, Putnam astutely points out that one of the by-products of this focus on the question of confirmation is that while "scientific theories are confirmed by 'evidence' in Carnap's systems of inductive logic…it is immaterial (that is to say, there is no way to represent the difference in the formalism) whether that evidence – those 'observation sentences'– is obtained as the result of intelligently directed experimentation, or whether it just happens to be available."[248] In other words, Putnam explains how the focus on confirmation within Carnap's inductive logic implicitly assumes an *idealised* (i.e. end-state) account of evidence: one which does not explain how the scientific method allows us to arrive at a certain body of evidence, but considers scientific reasoning to begin from a position where one already has this evidence. For Putnam, this idealised assumption of evidence as (potentially incidental) observation statements leads Carnap's inductive logic to a flawed conception of the scientific method, because it neglects the considerable amount of reasoning which scientists must partake in in order to formulate plausible hypotheses and valid experiments – reasoning which plays a pivotal part in determining the reliability of the evidence which results from a given experiment, and therefore in determining whether the observation statements one receives is *evidence* at all (i.e. as opposed to misinterpreted data or misleading evidence). Within Carnap's inductive logic, however, all of these methodological concerns become unproblematic if not trivial, such that "the scientific method is reconstructed as a method of computation, computation of a function like Carnap's famous 'c*.'"[249] It is here where Putnam compares Carnap's inductive logic to the pragmatist theory of inquiry. For the

---

Carnap's "logical" concept of probability, and its connection to the contemporary Bayesian evidential interpretation of probability, see Rudolf Carnap, "The Two Concepts of Probability: The Problem of Probability," *Philosophy and Phenomenological Research* 5, no. 4 (1945): 513–32; and Hannes Leitgeb and André Carus, "Rudolf Carnap," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Fall 2022 (Metaphysics Research Lab, Stanford University, 2022), §8.2.

[248] Putnam, "Pragmatism and Moral Objectivity," 170–71.

[249] Putnam, "Pragmatism and Moral Objectivity," 171.

pragmatists, the claim that epistemic rationality is reducible to a computational algorithm is an unacceptable conclusion, because the pragmatist theory of inquiry considers the role of epistemic agents to be essential in determining not only evidence, but epistemic rationality itself. To explain, consider the following exposition of the pragmatist theory of inquiry from Putnam:

> [For] Dewey's conception of inquiry: the model of an algorithm, like a computer program, is rejected. According to the pragmatists, whether the subject be science or ethics, what we have are maxims and not algorithms; and maxims themselves require contextual interpretation. Furthermore, the problem of subjectivity and intersubjectivity was in the minds of the pragmatists from the beginning. They insisted that when one human being in isolation tries to interpret even the best maxim for himself and does not allow others to criticise the way in which he or she interprets those maxim, or the way in which he or she applies them, then the kind of "certainty" that results is always fatally tainted with subjectivity…Notions such as "simplicity", for example, have no clear meaning at all unless inquirers who have proven their competence in the practice of inquiry are able to agree, to some extent at least, on which theories do and which theories do not possess "simplicity." The introduction of new ideas for testing likewise depends on cooperation, for any human being who rejects inputs from other human beings runs out of ideas sooner rather than later, and begins to consider only ideas which in one way or another reflect the prejudices he has formed. Cooperation is necessary both for the formation of ideas and for their rational testing. [250]

Let us unpack Putnam's pragmatist account of the scientific method. Putnam begins by explaining that the reason why the pragmatists reject a formal explication of scientific rationality is because the scientific method is a series of contextual maxims, not a computational algorithm. More specifically, the pragmatists see rational theory-choice within actual scientific practice as requiring the intelligent application of methodological maxims by epistemic agents within a given context – a process which is not reducible to any computational procedure.

To illustrate by example, Putnam explains how the notion of simplicity in scientific practice does not refer to any ostensible attribute about a scientific theory, but refers to a guiding epistemic value which is intelligently used by epistemic agents to *evaluate* competing scientific theories in a given context. This is because in order to even follow the principle that one should choose the simpler scientific

---

[250] Putnam, "Pragmatism and Moral Objectivity," 172.

theory, an epistemic agent must begin by interpreting for themselves what *counts* as simplicity. To take

an example, while the hypothesis that we are being deluded by an evil demon might be more simple

than the claim that we are in the normal world in terms of the number of entities posited (i.e. being

deluded by an evil demon posits you and the demon, while being in the normal world posits millions

of individuals), it is also less simple in terms of the number of ad-hoc explanations which are posited

(i.e. being deluded by an evil demon posits an intentional act of delusion by the demon every time one

interacts with the environment or another individual, while being in the normal world does not

require any such assumption).[251] The norm "choose the simplest explanation" is therefore a maxim

which requires an epistemic agent to play an active role in interpreting the maxim in a given context

(i.e. in deciding what we *mean* by simplicity in a given context) before acting accordingly.

Here, we can see how Putnam's claim that the scientific method is a series of contextual maxims has

some similarities to Williamson's contextualism about epistemic means and ends: just as how

Williamson notes that the injunction "Believe truly!" is satisfied differently (or unsatisfiable)

depending on the epistemic situation, Putnam notes that a methodological principle such as "Choose

the simplest explanation" is a maxim which must be interpreted by an inquirer in their given context,

and therefore potentially unsatisfiable if one does not know which explanation is simplest in their

particular situation.

However, there are also important differences between the account of the scientific method which

Putnam outlines here and Williamson's contextualist account. For example, while Williamson claims

that an epistemic agent must decide on what epistemic method they should follow depending on their

---

[251] This example is inspired from Putnam's arguments how the concept of 'simplicity' can be misused when explaining theory-choice within the philosophy of science. See Hilary Putnam, "Language and Philosophy," in *Mind, Language and Reality* (Cambridge: Cambridge University Press, 1975), 1–32, 26ff.

situation, Williamson's account seems to assume that injunctions such as "Believe truly" are norms which have clearly defined meanings, such that one can *know* an epistemic norm in such a way that one's knowledge of an epistemic norm entails knowing how to follow these norms (or whether a norm is followable) as an epistemic method in any given context.[252] Consequently, Williamson's contextualist account characterises epistemic rationality solely in terms of one's choice of the best epistemic method according to the epistemic norms one knows. By contrast, Putnam argues that a maxim like "choose the simplest explanation" does not have a determinate meaning in Williamson's sense, because while one might know a maxim insofar as one understands the semantic meaning of each word within a maxim, this does not mean that one knows what a maxim means insofar as one might not know how to *apply* this maxim within their context. As such, contrary to Williamson, the pragmatist theory of inquiry does not see one's epistemic practice as simply the act of choosing between a set of pre-determined epistemic methods according to the epistemic norms one knows, but the act of interpreting a methodological maxim in order to understand how to follow the maxim in a given context.

The reason why this difference between Putnam and Williamson is important is because one of the main aspects of our epistemic practice which Williamson's contextualist account does not explain is how we *arrive* at a particular epistemic situation. Just as how evidentialism begins theorising from a point where one already has a distinct body of evidence, Williamson's contextualist account begins theorising from a point where one is in an epistemic situation. This means that Williamson's contextualist account does not consider how an epistemic agent arrives at an epistemic situation, or of

---

[252] This view on the determinate meaning of epistemic norms would at least be consistent with Williamson's epistemicist view on vagueness, and his claim that "ignorance is the real essence of the phenomenon ostensively identified as vagueness." See Timothy Williamson, *Vagueness* (London: Routledge, 1994), 202.

how an epistemic agent might arrive at a *better* epistemic situation: instead, one's epistemic situation comes with a fixed body of knowledge (which includes one's knowledge of epistemic methods), and one's cognition involves choosing the best epistemic method depending on what one knows within their particular situation.

By contrast, the reason why the pragmatists highlight the role of the epistemic agent in interpreting a methodological maxim is because the notion of interpretation highlights the role which epistemic agents play in *creating* epistemic methods in the first place: once we see how our epistemic practice involves an act of interpretation where we figure out how to apply the steps of reasoning or action which a epistemic agent has used within a previous epistemic situation (i.e. their epistemic method) to our current situation, it naturally follows that epistemic methods are not simply acquired in virtue of being in an epistemic situation, but *invented* by epistemic agents in a particular situation. In this way, the pragmatists reveal how the position which we reside in our epistemic practice is situated within a broader historical and social epistemic context – one where our knowledge of how to use certain epistemic methods to acquire rational or true belief in certain contexts comes from us having *learnt* those methods from other epistemic agents, or having come up with these methods in the past ourselves. And this is not simply the case with our knowledge of epistemic methods: Putnam's account of the scientific method suggests that our very understanding of certain concepts – that is, of what counts as "simplicity", "evidence", or "rational" – are themselves the result of our past inquiries, and the inquiries of other epistemic agents. [253]

---

[253] Consider the following quote: "The claim that on the whole we come closer to truth about the world by choosing theories that exhibit simplicity, coherence, past predictive success, and so on…are themselves complex empirical hypotheses that we choose…because we have been guided by the very values in question in our reflections upon records and testimonies concerning past inquiries…the records and testimonies that we have good reason to trust by these very criteria of 'good reason'." Putnam, *The Collapse of the Fact/Value Dichotomy and Other Essays*, 32.

It is this notion of the *social* dimension of our epistemic context which the pragmatists see as explaining not only how we come to have certain epistemic methods within a given situation, but how to find *better* epistemic methods by which we can *improve* our epistemic situation. In particular, the pragmatists argue that it is the social dimensions of inquiry as a *cooperative* activity which plays this crucial role: while an individual epistemic agent might be in a position where they do not know what rationality requires of them, this does not mean that *other* epistemic agents cannot be in a position where they *do* know what is rational to believe. The pragmatist theory of inquiry therefore highlights how, in our everyday epistemic practice, we as epistemic agents do not simply rely on our own individual cognitive capacities, but can and do use the cognition of other epistemic agents as the means by which we obtain *better* methods of arriving at rational, and even true, beliefs. Within the process of inquiry, an epistemic agent improves their epistemic situation by allowing their epistemic practice (be it their application of a particular maxim, or their understanding of what counts as evidence or rational) to be criticised and revised by others so as to arrive at a position where one can better assess their epistemic practice. In this way, disagreement serves not as an epistemic challenge, but as an epistemic *asset*: one which enhances our epistemic rationality by allowing us to learn from how other epistemic agents respond to their epistemic situation in order to more intelligently adapt to our epistemic situation.

It is this picture of social inquiry which I shall argue provides a better account of how we as epistemic agents can come to resolve the non-ideal epistemic problem of disagreement. From the perspective of the pragmatist theory of inquiry, the problem with the evidentialist response to disagreement is that it neglects most of the rational decision making and reasoning one needs to do to arrive at evidence in

the first place.[254] Because the evidentialist defines epistemic rationality in terms of how an epistemic

agent proportions their degree of belief to their evidence, the evidentialist begins their theorising from

a starting position where one already has a distinct body of evidence, but just needs to know how to

rationally revise their beliefs based on this body of evidence. Given that the real epistemic problem is

characterising the norms of rational belief revision, the evidentialist characterisation of epistemic

rationality ignores the question of how one arrives at evidence – instead, evidence is generally

conceived of as an epistemic good one possesses by virtue of being in a specific epistemic situation.

By contrast, the pragmatists sees the revising of one's belief upon their evidence as the last step in an

extended process of collective reasoning – a step which is comparatively easy once we have dealt with

the epistemic problem of how we determine and find evidence in the first place. More specifically, the

pragmatists sees epistemic rationality as predominantly concerned with how one arrives as evidence,

since evidence is the result of laborious experimentation, argumentation, and reasoning from an

epistemic community – that is, as the result of *inquiry*.[255] Consequently, the pragmatist conception of

evidence is not that of a distinct body of evidence where new pieces of evidence are admitted, and old

---

[254] An instance of this which we have discussed extensively is Feldman's explicit rejection of methodological rationality as a
form of epistemic rationality at all. We will see how the pragmatists' rebuttal to the evidentialist rejection of methodological
rationality in §§5.2-5.3 of the next chapter.

[255] Consider the following quote from Dewey: "For the stock and staple of common sense inquiries and judgments are of this
sort. The deliberations of daily life concern in largest measure questions of what to make or to do…unless the decision
reached is arrived at blindly and arbitrarily it is obtained by gathering and surveying evidence appraised as to its weight and
relevance; and by framing and testing plans of action in their capacity as hypotheses: that is, as ideas…[similarly], the conduct
of scientific inquiry, whether physical or mathematical, is a mode of *practice*; the working scientist is a practitioner above all
else, and is constantly engaged in making practical judgments: decisions as to what to do and what means to employ in doing
it." Dewey, *Logic*, 160-61, see also 320. Italics are Dewey's and not my own.

pieces of evidence are defeated over time – but one where the very question of what constitutes evidence is posed, debated, and decided over the course of a given inquiry.

Of course, given that our understanding of what counts as evidence is something which itself may be the subject of inquiry, it is possible for a particular line of inquiry to begin precisely because of a disagreement over what constitutes evidence in a particular context. In situations where a particular line of inquiry is still undecided on what the total evidence is, or what is rational to believe upon this evidence, evidence obviously cannot play a role in resolving this disagreement. However, given that the process of inquiry forces epistemic agents to let their epistemic practice be criticised by other inquirers, the pragmatists claim that this form of revisionary self-criticism leads to the eventual settling of belief within a particular line of inquiry.[256] What does this process of social inquiry entail, and how can we ensure that the criticisms raised by dissenting parties within public discussion is constructive rather than destructive? To answer these questions will require an account of the kind of social inquiry which the pragmatists are envisioning, and the steps which we must take to ensure cooperation within social inquiry. This will be the focus of our next and final chapter.

---

[256] This is also how the pragmatist theory of evidence would also qualify our discussion on the commonability of evidence: while it is true that one of the conceptual roles which evidence plays in our epistemic practice is that of arbitrating between competing beliefs within a community, the pragmatist theory of inquiry shows how our very use of evidence in arbitrating between competing beliefs is predicated on prior reasoning which has settled what counts as evidence within earlier stages of inquiry. For example, within scientific practice, the use of scientific results as common evidence for particular beliefs about our health or our environment is only possible once these results have been firmly established through a prior process of experimentation and systematic review. As such, while the results of inquiry may be used as evidence which arbitrates between future disagreements, the pragmatist theory of inquiry highlights how this function of evidence is itself facilitated because of the *prior agreement* from the community of inquirers on the results of inquiry in the first place.

5   THE PRAGMATIST RESPONSE TO DISAGREEMENT

5.1   *Introduction: How does Social Inquiry Regulate Disagreement?*

As the concluding chapter of this thesis, the aim of this chapter is to outline what I shall call the

pragmatist response to disagreement, one which draws from the various accounts proposed within the

literature on pragmatism and democratic theory in order to show how the pragmatist account of social

inquiry can provide a better response to the epistemic problem of disagreement than evidentialism.

By way of introduction, recall how the previous chapter concluded with a preliminary sketch of the

pragmatist theory of inquiry, and how it provides an alternative account of how epistemic agents

should respond to epistemic disagreement to the evidentialist response to disagreement. Specifically,

the pragmatist theory of inquiry describes inquiry as a social epistemic practice in which the discursive

process of inquiry allows an epistemic community to improve in its epistemic practice through a

mutual exchange of reasons, methods, and ideas. Within this discursive process, the pragmatist theory

of inquiry claims that disagreement is not an epistemic challenge to our beliefs (like it is commonly

framed within the analytic epistemology of disagreement), but an epistemic boon by which we obtain a

diverse range of new epistemic methods (that is, by learning new ways of applying methodological

maxims, or by learning new methods entirely). In this way, the pragmatist theory of inquiry is

therefore able to provide an answer to the non-ideal epistemic problem of disagreement: one which

shows how a community of inquirers can incorporate disagreement within inquiry in a way that

improves rather than obstructs our epistemic practice.

However, it is at this point where more needs to be said about how social inquiry is able to regulate

disagreement in such a way that it becomes an epistemic asset to us. What is it about the underlying

structure of certain social inquiries that allow the very process of inquiry to channel the discourse in a

certain way, such that disagreement within inquiry can be addressed and settled in a productive

manner? The focus of this chapter is to unpack the substantial answer to this question that has been developed within the pragmatist literature. The structure of this chapter is as follows: by way of introduction, I will begin §5.2 by using Thomas Kelly's epistemic analysis of the psychological case of belief polarisation to connect our previous discussion on the analytic epistemology of disagreement with the subsequent discussion on the pragmatist theory of inquiry. In particular, I will show how the upshot of Kelly's epistemic analysis about the case of belief polarisation is that the process of rational belief revision – the process which the evidentialists see as encapsulating what we mean by epistemic rationality – is in fact situated within a larger epistemic practice of inquiry. This is because, according to Kelly's epistemic analysis, our epistemic practice does not consist only of revising our beliefs upon a body of evidence, but also on our decisions about which lines of inquiry we should conduct – decisions which have a strong *causal* impact on the total body of evidence we arrive at in a given epistemic situation. The social-psychological case of belief polarisation therefore reveals how epistemic disagreements are not only caused by the different ways in which epistemic agents can revise their beliefs on the basis of their evidence, but also caused by how epistemic agents can arrive at different bodies of evidence according to their differing intuitions about which lines of inquiry are worth pursuing.

It is these kinds of differing intuitions about potentially fruitful inquiries that epistemic agents have in their individual inquiries which lead the pragmatists to insist on the importance of cooperative social inquiry. Once we acknowledge that our individually biased intuitions about which lines of inquiry are worth pursuing can lead us to fail to obtain crucial evidence about particular issues, the pragmatists argue that the best solution to this problem of subjective bias can only be found in intersubjective discussion. More specifically, in §5.3, I shall explain how the literature on the pragmatist theory of inquiry focuses on developing a transition epistemic account as to how a community of inquirers can best cooperate in order to make the best decisions within inquiry. At the centre of this epistemic

account are the similarities which the pragmatists see between science and democracy: two kinds of inquiry which the pragmatists see as the best examples of successful social epistemic practice. In comparing the examples of science and democracy – that is, in seeing democratic norms within scientific inquiry and in seeing the scientific experimental attitude within democracy – the pragmatists therefore substantiate the pragmatist theory of inquiry by pointing to a particular set of ethical norms and social institutions which allow political and scientific inquiry to leverage disagreement in an epistemically productive way. As such, in focusing on select arguments made by Hilary Putnam and Elizabeth Anderson, I will outline the pragmatists' account of the *democratic* structure of inquiry, and how these structural features allows social inquiry to move from a position of initial disagreement to improved epistemic practice.

This brings us to §5.4, where I will draw from the implications of our critical examination of the arguments of Kelly, Putnam, and Anderson respectively, and use them to outline the pragmatist response to disagreement. In particular, one of the main upshots of Putnam and Anderson's respective accounts of democratic inquiry is that the pragmatist account of effective social inquiry is *itself* the product of social inquiry: that is, that the substantial reasons which Putnam and Anderson give for adopting certain ethical norms and social institutions within social epistemic practice are not *a priori* theoretical reasons, but empirical reasons about the proven track-record of certain methodological features and principles. In this way, just as how the evidentialist response to disagreement can be summed up by the norm to follow your evidence, the pragmatist response to disagreement is characterised by the norm to continue inquiring cooperatively: a norm which allows us as epistemic agents to provide new ways of improving our epistemic practice, and therefore new ways of responding to the disagreements we encounter. I will then conclude this thesis in §5.5, by showing how the pragmatist response to disagreement is an epistemic response to disagreement that uses this empirical basis to argue for an ethical attitude: one which invites us to cooperate with a greater and

more diverse number of dissenting parties in our epistemic practice, because choosing to do so will only improve not only our own epistemic practice, but that of our wider epistemic community.

5.2   The Impact of Subjective Bias on Individual Inquiry

In this section, I will explain how our subsequent exposition of the pragmatist theory of inquiry connects to our previous critique of the evidentialist response to disagreement. To begin, let us continue from Putnam's explanation of the pragmatist theory of inquiry from the previous chapter. In the exposition of Dewey's theory of inquiry which we examined last chapter, Putnam points out that the democratic structure of inquiry is crucial in avoiding what they saw to be one of the main causes of insoluble disagreement: namely, the subjective limitations and biases of individual inquirers. Consider the following excerpt:

> Furthermore, the problem of subjectivity and intersubjectivity was in the minds of the pragmatists from the beginning. They insisted that when one human being in isolation tries to interpret even the best maxims for himself and does not allow others to criticize the way in which he or she interprets those maxims, or the way in which he or she applies them, then the kind of 'certainty' that results is always fatally tainted with subjectivity.[257]

Here, Putnam explains how, for the pragmatists, one of the main reasons why inquiry has to be a *social* epistemic enterprise is due to the subjective limitations and biases of individual inquirers.[258] In particular, the social dimensions of inquiry serve as a means of *correcting* epistemic agents on their use of epistemic methods by subjecting their use of epistemic methods to the criticism of other epistemic

---

[257] Putnam, "Pragmatism and Moral Objectivity," 172.

[258] In a footnote to this passage, Putnam lists the following as works in which pragmatists have discussed this particular issue: see Karl-Otto Apel, *C. S. Peirce: From Pragmatism to Pragmaticism* (Amherst: University of Massachusetts Press, 1981); William James, "The Moral Philosopher and the Moral Life," in *The Will to Believe and Other Essays in Popular Philosophy* (Cambridge: Harvard University Press, 1979); John Dewey, "Nature, Communication and Meaning," in *Experience and Nature* (Mineola, N.Y.: Dover, 1958). For the original footnote, see Putnam, "Pragmatism and Moral Objectivity", fn. 44.

agents. Without this corrective function of inquiry, the pragmatists claim that epistemic agents who attempt to investigate the truth by themselves tend to come to a flawed subjective certainty about their own beliefs and epistemic practice, a certainty which becomes the central impediment to inquiry.

But how exactly does the subjective bias of individual epistemic agents impede individual and collective inquiry, and how does this relate to the issue of finding resolution in epistemic disagreements? In order to connect the discussion on the problem of subjective bias within the pragmatist literature to our topic of epistemic disagreement, what we first need is a better understanding of the ways in which cognitive biases can affect our epistemic practice, both in terms of our individual and social inquiry and in how we as epistemic agents respond (or ought to respond) within epistemic disagreements. As a preliminary remark, it is important to note that the psychological literature does not see the role which cognitive biases play within our epistemic practice as *purely* negative: if we take cognitive biases simply to be our natural psychological predispositions to favour one conclusion over another, then the psychological research affirm certain biases as necessary heuristics within our everyday life, as we require these heuristics in order to make everyday decisions

under practical constraints such as limited time and energy.[259,260] Nevertheless, given that some

cognitive biases clearly do have negative impacts on our epistemic practice, an epistemic account is

needed to explain the role and effect that cognitive biases have on our epistemic practice.

---

[259] Of course, while some might argue that the negative connotation of the word "bias" is inherent to the contemporary meaning of the word, I am referring to an older sense of the word for a particular purpose. Consider the following definition (Definition 3.a.) of the word from the Oxford English Dictionary: "A tendency, inclination, or leaning towards a particular characteristic, behaviour, etc.; a propensity." The reason why I am using this sense of the word bias is because, in the upcoming discussion of Kelly's analysis of the phenomenon of belief polarisation, a key part of Kelly's argument is to show how the infelicitous outcome of belief polarisation might arise despite the rational belief revision of the respective epistemic agents. As such, rather than seeing rational belief revision and cognitive bias as mutually exclusive, Kelly (and the researchers into belief polarisation) seem to assume that the real distinction is between rational and irrational biases, or biases with good outcomes versus biases with bad outcomes. "Bias, n., Adj., And Adv." In OED Online. Oxford University Press, December 2022. https://www.oed.com/view/Entry/18564. (accessed December 26, 2022.) For a defence of biases as important heuristics in everyday decision-making, see the work of Gerd Gigerenzer: in particular, see Gerd Gigerenzer and Daniel G. Goldstein, "Reasoning the Fast and Frugal Way: Models of Bounded Rationality," *Psychological Review* 103, no. 4 (1996): 650–69; Gerd Gigerenzer and Reinhard Selten, eds., *Bounded Rationality: The Adaptive Toolbox* (Cambridge, Massachusetts; London, England: MIT Press, 2002); Gerd Gigerenzer and Wolfgang Gaissmaier, "Heuristic Decision Making," *Annual Review of Psychology* 62, no. 1 (2011): 451–82.

[260] It is also important to note that Kelly has very recently published an extensive philosophical analysis of bias – one which aims to outline a "norm-theoretic" account of bias. While Kelly's own philosophical analysis leads to similar conclusions as the one which this thesis gives, it is also important to note that there are noticeable differences in the aim and subject of our respective epistemic accounts. To explain in a bit more detail, the following pragmatist account of biases would agree with Kelly's conclusion about a "thoroughgoing externalism about bias", and Kelly's idea that biases can be corrected in inquiry through the norm of "following the argument wherever it leads". However, our pragmatist account differs differ from Kelly's account of bias because Kelly's norm-theoretic account considers norms which apply to individual epistemic agents in their individual inquiries (i.e. Kelly provides an account of *individual* epistemology), whereas the pragmatist account of bias considers social-epistemic norms which apply to a collective *group* of epistemic agents (i.e. the pragmatist theory of inquiry is an account about *social* epistemology). In this way, the pragmatist account of bias would qualify Kelly's norm of "following the argument wherever it leads", in showing how the satisfaction of this norm requires *other* epistemic agents who challenge one's own argument, or provide new counter-arguments within social inquiry. See Thomas Kelly, *Bias: A Philosophical Study* (Oxford: Oxford University Press, 2022).

To provide an epistemic account of this sort, I shall now draw from Thomas Kelly's epistemic analysis of the phenomenon of belief polarisation. For Kelly, the empirical phenomenon of belief polarisation is of epistemic interest precisely because it demonstrates how two epistemic agents who have a high degree of certainty about competing hypotheses can be placed into an epistemic situation where obtaining further common evidence counter-intuitively exacerbates the extent of their disagreement. As an example of how two dissenting parties can be *more* convinced of their conflicting beliefs upon receiving the same body of evidence, belief polarisation therefore raises questions about how our initial doxastic attitudes can affect the use of evidence in resolving disagreement and in inquiry more generally.

Let us examine Kelly's epistemic analysis of belief polarisation and its implications in turn. To begin, consider Kelly's key example of belief polarisation: suppose two epistemic agents have conflicting opinions about whether capital punishment is a deterrence for crime, such that one believes that capital punishment is a deterrence while the other is not.[261] Furthermore, suppose both epistemic agents come to receive a common but mixed body of evidence – in this case, both parties learn about a series of statistical studies where some studies suggest that capital punishment is a deterrence while others suggest that it does not. The counter-intuitive result which occurs in cases of belief polarisation is this: rather than experiencing any convergence of opinion from revising their beliefs on the same body of evidence, cases of belief polarisation are such that the same body of evidence causes the two epistemic agents to be more convinced of their respective beliefs, and thus to *diverge* further in their

---

[261] Thomas Kelly, "Disagreement, Dogmatism, and Belief Polarization," *Journal of Philosophy* 105, no. 10 (2008): 611–12. Kelly notes that this example of capital punishment is not only hypothetical, but is based off an actual psychological study into belief polarisation. For more information, see Charles Lord, Lee Ross, and Lepper Mark, "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence," *Journal of Personality and Social Psychology* 37, no. 11 (1979): 2098–2109.

opinion of the given issue. So in Kelly's example, it is learning about the same series of statistical

studies which somehow causes the two epistemic agents to be more certain about their respective

beliefs, such that the same body of evidence causes one to be more convinced that capital punishment

is a deterrence for crime and the other to be more convinced that capital punishment is not a

deterrence. Belief polarisation therefore acts as an example of an epistemic situation where dissenting

parties do not find any resolution from receiving a common body of evidence – on the contrary, a

common body of evidence only serves to *worsen* the extent of their disagreement.

At this point, one might question why we should take belief polarisation to be an epistemic puzzle

rather than an undesirable psychological phenomenon. Given the psychological research on how

actual epistemic agents can be prone to fallacious reasoning in everyday life, one might dismiss belief

polarisation as simply the result of irrational belief revision, and therefore unrelated to the epistemic

question of how one *ought* to respond to evidence in their epistemic practice were they a rational

epistemic agent.[262] However, what makes empirical cases of belief polarisation an epistemic puzzle for

Kelly is the fact that, upon further investigation on behalf of the psychologists researching the

phenomenon and Kelly's own perspective as an epistemologist, it is unclear whether epistemic agents

are really irrational in their belief revision within cases of belief polarisation.[263] As such, some

---

[262] Kelly, "Disagreement, Dogmatism, and Belief Polarization," 625. For some classical psychological studies on common fallacies and irrational biases which actual epistemic agents employ in everyday life, see Daniel Kahneman and Amos Tversky, "Judgment Under Uncertainty: Heuristics and Biases," *Science* 185, no. 4157 (1974): 1124–31; Daniel Kahneman and Amos Tversky, "On the Reality of Cognitive Illusions," *Psychological Review* 103, no. 3 (1996): 582–91. Also see fn. 259 for more on the psychological literature on bias.

[263] Kelly notes that amongst the papers he discusses on belief polarisation, many of the psychologists emphasise how they do not believe that the epistemic agents within a case of belief polarisation are not behaving in an irrational manner. See Kelly, "Disagreement, Dogmatism, and Belief Polarization," 625.

epistemic account needs to be given for why belief polarisation occurs despite the seemingly competent reasoning of the epistemic agents involved.

To show why epistemic agents within belief polarisation are not in fact irrational in their epistemic practice, Kelly begins by considering a misleading but potential explanation for belief polarisation which would make belief polarisation epistemically insignificant: namely, that belief polarisation occurs as a result of dogmatism from epistemic agents.[264] Suppose the two epistemic agents are dogmatists, such that their prior belief about whether capital punishment is or is not a deterrence to crime leads them to reject any counter-evidence towards their belief (i.e. their belief in p or not-p leads them to reject any evidence for not-p or p respectively). Given that the two epistemic agents are dogmatists, one could easily explain the phenomenon of belief polarisation as the result of the two parties receiving the common body of evidence, rejecting any counter-evidence within that body of evidence, and updating only on the subset of the evidence which increases their level of confidence for their pre-existing belief. Furthermore, if the underlying cause of belief polarisation was dogmatism of this sort, then we can safely dismiss belief polarisation as being of no epistemic challenge to our use of evidence within disagreement: belief polarisation would simply be a case of irrational belief revision, and an example of what one should *not* do when encountering evidence that challenges rather than confirms your pre-existing beliefs.

---

[264] Kelly, "Disagreement, Dogmatism, and Belief Polarization," 613–17. Kelly notes that this example of the dogmatist comes from Saul Kripke's discussion on the dogmatism paradox. However, whereas Kelly's example here is about a dogmatist about belief, Kripke's discussion was originally on dogmatism about knowledge. Kelly goes on to note that Kripke's dogmatism paradox about knowledge comes from an unpublished lecture delivered to the Cambridge Moral Sciences Club, and that the first published discussion of the paradox is from Gilbert Harman. See Gilbert Harman, *Thought* (Princeton University Press, 1973), 148–49.

However, it is here which Kelly compares this hypothesis about belief polarisation with the actual

psychological research. While the hypothesis that belief polarisation occurs as a result of dogmatism

might seem to be initially plausible, this hypothesis is ultimately unsupported by the psychological

research because the studies do not show epistemic agents within cases of belief polarisation as

discounting or ignoring counter-evidence on the basis of their prior confidence in a hypothesis. On the

contrary, epistemic agents within belief polarisation actually paid a *greater* attention to counter-

evidence, because their high level of confidence in a particular hypothesis led them to a "heightened

sensitivity to methodological problems in studies when the results of those studies seemed to tell

against their beliefs."[265] Furthermore, according to the researchers on belief polarisation, it is unfair to

claim that the scrutiny which epistemic agents gave to counter-evidence was irrational, because

epistemic agents within belief polarisation were able to correctly identify "genuine limitations and

weaknesses in studies that conflict with their prior beliefs".[266] As such, one of the key differences

between actual epistemic agents within belief polarisation and dogmatists is the fact that epistemic

agents within belief polarisation reject counter-evidence not on the basis of their high level of

confidence in a particular hypothesis, but on the basis of additional *reasons* about the unreliability of

the counter-evidence received.

What types of reasons can an epistemic agent have for rejecting the counter-evidence without

acquiring any additional evidence? According to Kelly, one way in which an epistemic agent can

reasonably reject counter-evidence without any further evidence is by considering alternative

explanations which undercut the counter-evidence itself. To illustrate, Kelly uses the following

---

[265] Kelly, "Disagreement, Dogmatism, and Belief Polarization," 618.

[266] Kelly, "Disagreement, Dogmatism, and Belief Polarization," 618.

continuation of the example of capital punishment: suppose the two epistemic agents both come to

learn that State A and not B has capital punishment, and that State A has a lower murder rate than B.[267]

Now, consider the epistemic agent who has a high level of confidence in the belief that capital

punishment does not deter crime. Kelly argues that one way in which the epistemic agent can

reasonably reject that learning these two facts should lower their confidence in their belief is by

considering other hypotheses which serve as alternative explanations of these facts. In particular, if the

epistemic agent in question can find plausible alternative explanations which show how these two facts

are not causally connected (i.e. that there is no causal connection between the capital punishment of

State A and its lower murder rate relative to State B), then there is no need for them to lower their

degree of belief after learning these two facts. This is because one is only irrational in refusing to lower

their beliefs upon learning these facts if these facts jointly count as counter-evidence for their beliefs in

the first place. However, if one were to come up with plausible alternative hypotheses in which the fact

that State A has a lower murder rate than B has no causal connection to the fact that State A has capital

punishment and State B does not, then the two facts no longer serve as counter-evidence for the belief

that capital punishment does not deter crime. As such, the epistemic agent who has a high level of

confidence in the belief that capital punishment does not deter crime might reasonably discount what

initially appeared to be counter-evidence to their beliefs as a result of considering alternative

hypotheses which give plausible explanations as to how the given facts are not counter-evidence at all.

Kelly therefore argues that the psychological research suggests an alternative explanation to the cause

of belief polarisation than that of dogmatism: because of the greater scrutiny that the respective

epistemic agents place on counter-evidence, each respective agent receives greater evidential support

---

[267] Kelly, "Disagreement, Dogmatism, and Belief Polarization," 619–20.

for their own belief upon receiving a common mixed body of evidence as the subset of evidence which goes against one's prior beliefs is undercut by discovering methodological flaws within or alternative explanations for said evidence. Furthermore, given that each epistemic agent is not violating any norms of rationality during their belief revision – since their belief revision is dependent on legitimate reasons rather than blind dogmatism about their pre-existing beliefs – belief polarisation therefore serves as an example of an epistemic situation where dissenting parties may rationally respond to a common body of evidence, and become more convinced of their respective beliefs as a result.

Still, one might question whether the act of paying greater attention to counter-evidence is nonetheless irrational. Even if an epistemic agent is rational in their criticisms of whether a certain piece of information should be counted counter-evidence, the disproportionate attention an epistemic agent pays to counter-evidence as opposed to confirming evidence suggests some violation of rational belief revision, given that one chooses to respond to some evidence more than others.[268] In response to this objection, Kelly uses the example of scientific inquiry to show how the decision to pay greater attention to counter-evidence over confirming evidence is not a violation of the norms of rationality,

---

[268] Throughout this paper, Kelly evaluates whether the belief revision of an epistemic agent is rational according to whether it violates the principle of the Commutativity of Evidence or not. Considering Kelly's discussion on how the Commutativity of Evidence connects with his proposed explanation of belief polarisation is outside the scope of this thesis. However, here is a brief summary of the Communtativity of Evidence in order to understand why one might think paying greater attention to counter-evidence is just as irrational to the dogmatic rejection of counter-evidence. The principle of Commutativity of Evidence states that one's method of belief revision is rational only if the historical order in which one acquires evidence is irrelevant to how we should revise our belief according to that evidence. It is clear that the dogmatic rejection of counter-evidence violates the principle of Commutativity of Evidence: if a dogmatist initially receives evidence which confers a high degree of certainty in a particular hypothesis, then this affects the response which the dogmatist has to all subsequent counter-evidence. However, the question in this section is whether giving greater scrutiny to counter-evidence violates the principle of Commutativity of Evidence in the same way as the dogmatist, and thus whether greater scrutiny to counter-evidence is an instance of irrational belief revision. As we will see shortly, Kelly is going to argue for the negative: that paying greater scrutiny to counter-evidence is not an instance of irrational belief revision.

but part of our intuitions on what good reasoning looks like.[269] In particular, Kelly's focus is on how scientific inquiry responds to anomalous data which is unexplained by or inconsistent with a commonly held scientific theory. When encountering anomalous data, the standard approach in scientific inquiry is this: rather than spending time and resources trying to find alternative explanations of phenomena for which the commonly held theory already provided plausible explanations, focus your attention and effort on finding plausible alternative hypotheses which explain how the anomalous data might be reconciled with the commonly held theory.

To illustrate this type of scientific reasoning, suppose a scientist finds that an experiment they conducted has given them data which contradicts the laws of thermodynamics. In response to this data, the scientist decides to first check their experimental method, rather than to give up the commonly held laws of thermodynamics and start finding other explanations for thermodynamic processes from scratch. A few points to note. First, notice that the scientist's decision to check their experimental method is not a dogmatic rejection of counter-evidence: in this case, the scientist is not making the (invalid) inference from their prior belief that the laws of thermodynamics is true to the conclusion that the data from their experiment is false. Rather, the scientist is making a *practical* judgment that this situation requires one to find additional evidence (i.e. additional evidence about the reliability of the experiment) due to the overwhelming evidence one already has for the laws of thermodynamics.[270] Second, it is precisely the high level of confidence that the scientist has in the laws

---

[269] Kelly, "Disagreement, Dogmatism, and Belief Polarization," 624–25.

[270] Kelly claims that the difference between the question of what one should infer from one's evidence (which Kelly sees the dogmatist as giving an answer to) and the question of when one should find additional evidence (which Kelly sees as the question concerning the scientist who is responding to anomalous data) is a difference between two kinds of rationality: "We should, I think, distinguish carefully between (i) questions about the rationality of devoting greater scrutiny to apparent counter-evidence in the relevant ways, and (ii) questions about the rationality or epistemic status of the beliefs that result from doing so." Kelly, "Disagreement, Dogmatism, and Belief Polarization," 622. Kelly's distinction here seems to be making

of thermodynamics which allows the scientist to come to this decision: if the scientist was conducting an experiment on a new scientific hypothesis which is not firmly established within the literature, then receiving data which contradicted this hypothesis would not lead one to immediately conclude that the data is anomalous, or that one should immediately check for flaws in one's experimental method.

As such, the scientist's decision to pay greater attention to anomalous data does not stem from an irrational bias towards one's pre-existing belief – rather, the scientist pays greater attention to anomalous data precisely because their total body of evidence suggests that it is more likely that the anomalous data rather than the laws of thermodynamics is wrong. This is similar to a case we examined when discussing the Total Evidence View (see §2.3.3): when assessing a dissenting peer's

_____

the same distinction as Feldman on the difference between methodological and current-state rationality. Consequently, one may wonder what Kelly's own view is about the connection between methodological and current-state rationality. On the one hand, Kelly has written a paper which argues that epistemic rationality and instrumental rationality are distinct forms of rationality, such that epistemic rationality cannot be seen as a form of instrumental rationality. Seeing as Kelly's use of the terms "epistemic" and "instrumental" rationality roughly correspond to current-state and methodological rationality respectively, this suggests that Kelly would agree with Feldman that epistemic rationality is only concerned with (probabilistic or deductive) belief revision on the basis of evidence. On the other hand, in this paper on belief polarisation, Kelly seems to suggest that methodological rationality plays a part in whether an epistemic agent is "reasonable" in how they evaluate what counts as evidence in their epistemic practice. For instance, in the previous example about scientific inquiry, Kelly shows how what is reasonable for the scientist to do within their inquiry affects what the scientist believes about their data (whether it is faulty or not), and vice versa. To re-iterate, my own view which I put forward in the previous chapter is to claim that the connection between these two forms of rationality changes depending on what type of epistemic theorising we are talking about – that is, whether we are talking about ideal or non-ideal epistemology. In the case of ideal epistemology, it is often helpful to make a distinction between methodological and current-state rationality so as to explicate between different cognitive processes which are occurring within belief revision. However, when it comes to non-ideal epistemology, the distinction between methodological and current-state rationality blurs simply because both forms of rationality are necessary to one's attempt to find adequate and better justifications in their particular epistemic position. In other words, if what we are interested in is a practical question such as that of how we can actually come to resolve disagreement in our everyday epistemic practice, then the answer will naturally involve both practical norms of how to gather evidence, and inferential principles of what to infer on the basis of one's evidence. For Kelly's discussion on methodological and current-state rationality, see Kelly, "Epistemic Rationality as Instrumental Rationality".

claim in light of our total evidence, our total evidence might be such that the rational response is not to lower our confidence of our belief, but to lower our confidence that the dissenting party is in fact a peer. Given that the scientist is reasonable in how they respond to anomalous data in scientific inquiry, Kelly concludes that it should also be reasonable for individuals to pay greater attention to counter-evidence when their pre-existing body of evidence suggests a high level of confidence for their belief.[271]

Does the case of belief polarisation therefore undermine the use of evidence in resolving epistemic disagreement, seeing as rational epistemic agents can revise their beliefs on a common body of evidence and come to be further entrenched in their conflicting opinions? For Kelly, the answer is no: as it turns out, the upshot of belief polarisation is not in fact about how an epistemic agent's doxastic attitudes can affect the norms of rational belief revision. Instead, belief polarisation reveals how the initial doxastic attitudes of an epistemic agent can have a *causal* impact on the subsequent body of evidence we acquire during inquiry:

---

[271] Kelly doesn't argue for this directly, but merely states that there is no reason to assume that "what is reasonable in the context of scientific inquiry is unreasonable at the level of individual thinker". However, Gilbert Harman has directly argued for this point in his work on the influence of practical factors on theoretical reasoning, and specifically on the epistemic value of conservatism. Given the cognitive and time limits of epistemic agents, Harman argues that conservatism is an epistemic value that guides theoretical reasoning in much the same way as simplicity and coherence. In our epistemic practice, the epistemic value of conservatism states that we should "start with our present view and try to improve it by getting rid of inconsistency and by increasing its coherence in ways that help us answer questions in which we are interested". For Harman, the fact that the epistemic value of conservatism tells us to begin our reasoning from our present view means that we inevitably "[favour] beliefs that we already have over propositions that we do not already accept". The epistemic value of conservatism therefore shows how our conception of good reasoning is such that we ought to privilege the beliefs that we currently hold (provided that they are held on the basis of good evidence) over propositions which we do not yet have strong evidence for and which we do not currently accept. Kelly, "Disagreement, Dogmatism, and Belief Polarization," 624; Gilbert Harman, "Practical Aspects of Theoretical Reasoning," in *The Oxford Handbook of Rationality*, ed. Alfred R. Mele and Piers Rawling (Oxford: Oxford University Press, 2004), 54. For more information on conservatism as a normative value of theoretical reasoning, see Gilbert Harman, *Change in View: Principles of Reasoning* (Cambridge, Massachusetts; London, England: The MIT Press, 1986).

> For any given body of total evidence –where total evidence is understood as evidence in the broad sense –the order in which the constituent pieces of evidence are acquired makes no difference to what it is reasonable to believe [i.e. to the principles of rational belief revision]…on the other hand, historical facts about when one acquires a given piece of evidence [and what doxastic attitude results from that evidence] might very well make a causal difference to which body of total evidence one ultimately ends up with. One acquires a given piece of evidence at an early stage of inquiry; this might very well influence the subsequent course of inquiry in various ways, by way of making a difference to how one subsequently thinks and acts (which possibilities one considers, which routes get explored as the most promising and fruitful, and so on). And this in turn can make a difference to what evidence one ends up with. In such cases, there is an undeniable element of path-dependence.[272]

Kelly's epistemic analysis of the psychological phenomenon of belief polarisation therefore provides a sophisticated account on the role which our subjective doxastic attitudes plays within our individual inquiry.[273] As it turns out, belief polarisation is not an example of how epistemic agents can receive different evidential justifications from the same body of evidence. Rather, belief polarisation is an example of how one's total body of evidence can be causally influenced by the subjective beliefs which one holds. The case of belief polarisation therefore demonstrates how one's subjective certainty in a

---

[272] Kelly, "Disagreement, Dogmatism, and Belief Polarization," 628.

[273] On a related tangent, one might wonder if the problem of belief polarisation might be solved if we remove epistemic subjects from the process of inquiry. Would it be possible, perhaps, to use some sort of artificial intelligence or machine learning program to filter through empirical data so as to avoid the interference of the epistemic agents and their subjective doxastic attitudes altogether? Research into machine learning algorithms seem to respond in the negative: just as how cognitive bias is an embedded part of our epistemic practice, so too is "algorithmic bias" an embedded part of machine learning algorithms as well. One argument within the recent literature is given by Johnson, who provides an epistemic analysis of cognitive bias by examining the relationship between human cognitive bias with the "algorithmic bias" of certain machine learning programs. Johnson claims that the similarities between algorithmic and cognitive biases show how the causes of biases arise from "seemingly innocuous patterns of information processing" whose attributes are "difficult to identify, mitigate, or evaluate using standard resources in epistemology and ethics". As such, Johnson's conclusion that there is "no purely algorithmic solution" for algorithmic bias seems to parallel Putnam's criticisms of the possibility of a project like Carnap's inductive logic. That being said, given the recency of the topic, the question of whether there is an algorithmic solution to algorithmic bias in machine learning is still a live question. See Gabbrielle M. Johnson, "Algorithmic Bias: On the Implicit Biases of Social Technology," *Synthese* 198, no. 10 (2020): 9941–61. For an overview of the current literature on algorithmic bias, see also Nima Kordzadeh and Maryam Ghasemaghaei, "Algorithmic Bias: Review, Synthesis and Future Research Directions," *European Journal of Information Systems* 31, no. 3 (2022): 388–409.

belief can have consequences on one's subsequent body of evidence by affecting one's judgment about which lines of inquiry seem to be promising and worthy of pursuing. As a result, the differences in the initial doxastic attitudes of two epistemic agents within belief polarisation can influence the subsequent lines of inquiry which they respectively choose to pursue, such that the epistemic agents come to different total bodies of evidence that justify divergence in their respective beliefs.

We can now return to our question of how the problem of subjectivity which is outlined within the pragmatist literature connects to the pragmatist theory of inquiry and to our topic of epistemic disagreements. To begin, notice how Kelly's analysis of belief polarisation serves to elaborate Putnam and the pragmatists' claim that an epistemic agent's subjective attitudes can impede their individual inquiry. Without other epistemic agents, an epistemic agent may never notice the influence that their prior subjective attitudes can play in guiding the course of their inquiry, which in turn might lead them to bodies of evidence which disproportionately favours rather than challenges their pre-existing beliefs. This means that if an epistemic agent comes to initially acquire enough misleading evidence such that one has a high level of confidence in a false belief, their subsequent inquiry might be obstructed insofar as one fails to pursue the lines of inquiry needed to acquire evidence which defeats and corrects one's false beliefs, or the evidence necessary for one to question their own belief so as to seek conciliation with disagreement.

However, for the pragmatists, the problem of subjectivity that obstructs epistemic agents on the level of *individual* inquiry is one which requires a solution within *social* cooperative inquiry. To raise a quick criticism of Kelly's account of belief polarisation, Kelly's concluding explanation of how to avoid belief polarisation remains unconvincing because of his failure to consider any form of social interaction between epistemic agents within their respective inquiry. At the start of Kelly's paper, Kelly notes that the reason why he does not consider cooperative social inquiry – or any social interaction between agents in their epistemic practice for that matter – is because Kelly sees the epistemic problem

of disagreement as a problem that is distinctly for *individual* epistemic agents, rather than a problem for an epistemic community.[274] However, by not considering the social dimensions of our epistemic practice, Kelly's concluding solution for how to avoid belief polarisation is problematic, because it relies on an implausible assumption about the sensitivity of epistemic agents to their epistemic situation.[275] Specifically, Kelly argues that the solution to belief polarisation is simply to be more aware of the ways in which our cognitive biases can skew our epistemic practice, such that our understanding of the influence of cognitive biases on our epistemic practice will lead us "to correct for the operation of the relevant psychological mechanisms, by being less confident of those belief that are likely to have been the past beneficiaries of the mechanisms."[276]

However, the outstanding problem which Kelly's response does not address is how exactly is one supposed to know which beliefs are most likely to have inappropriately benefited from the

---

[274] Kelly begins this paper with a footnote stating that his description of belief polarisation as a "disagreement" is taking disagreement in a relatively weak sense, such that epistemic agents can disagree without ever being aware that the other disagrees, or that the other exists at all. See Kelly, "Disagreement, Dogmatism, and Belief Polarization", fn. 1. For a paper in which Kelly explicitly argues for why disagreement should be taken as an individual rather than a social epistemic problem, see Kelly, "The Epistemic Significance of Disagreement".

[275] I am using "sensitivity" in Williamson's counterfactual sense of the word. Williamson gives a counterfactual notion of sensitivity to the truth such that one is sensitive to the truth if and only if it was the case that if a proposition is false, one would not believe it. As we shall see, the following criticism of Kelly's response to the problems of belief polarisation parallels our reading of Williamson in our criticism of the evidentialist response to disagreement. It appears that Kelly's response to belief polarisation would require a similar sort of sensitivity, such that if a belief was a beneficiary of a skewed body of evidence which disproportionately confirmed it, one would know and adjust one's level of confidence accordingly. As such, conversely, if one is in an epistemic position where they do not have this sort of sensitivity towards their body of evidence, Kelly's response to belief polarisation fails to be applicable in practice. As a quick sidenote, we can see that Kelly's subsequent work on bias seems to acknowledge this problem of sensitivity in how Kelly's account of bias leads to a thoroughgoing externalism about bias. See fn. 260 for more information about Kelly's recent epistemic account of bias. For more information on sensitivity, see chapter 7 of Williamson, *Knowledge and Its Limits*.

[276] Kelly, "Disagreement, Dogmatism, and Belief Polarization," 629.

psychological mechanisms undergirding belief polarisation, or to what degree one is supposed to correct one's level of confidence in a given belief. After all, Kelly's suggestion is not that we should lower our levels of confidence for every belief – only the ones which have been affected by the "relevant psychological mechanisms". As such, Kelly's suggestion faces the same problem as the one raised by Williamson's contextualist criticism of the evidentialist response to disagreement: just as how we may be placed in an epistemic situation where we do not know what our evidence is nor of what rationality requires of us, so too may we be in a position where we do not know which of our belief are in fact past beneficiaries of our psychological biases, and thus affected by a skewed body of evidence. In other words, if we are in an epistemic position where we do not know (or have evidence about) whether a given body of evidence is the product of biases within one's psychology, then we consequently cannot follow Kelly's suggestion to change our belief revision in response to these psychological biases.

5.3   The Democratic Structure of Inquiry

5.3.1   The Democratisation of Inquiry: Putnam on the Norms of Discourse Ethics

It is precisely these types of problems that epistemic agents face in individual inquiry which lead the pragmatists to highlight the epistemic strengths of social cooperative inquiry. Rather than relying on an epistemic agent's ability to identify one's own subjective biases and its effects on the justifications for one's beliefs, the pragmatist points to our reliance on other epistemic agents in our everyday epistemic practice as the best and most natural way of identifying and correcting for our subjective biases.[277] In particular, ever since Dewey, the pragmatists have used the joint examples of democracy

---

[277] As an addendum to the previous section, consider what would happen if two epistemic agents who are in a case of belief polarisation were to interact with each other in their epistemic practice such that they come to share a common body of evidence. This kind of social interaction between epistemic agents would seem to fix many of the epistemic problems which Kelly raises about belief polarisation: even if our psychological make-up as human beings is such that we are naturally more adept at finding problems in evidence that is against rather than for our beliefs, this does not mean that we cannot recognise

and science to substantiate the pragmatist theory of inquiry, such that the pragmatist theory of inquiry

has been described as a form of "democratic experimentalism".[278] To explain how the pragmatists use

the examples of democracy and science to develop the theory of inquiry, I will focus on two arguments

made by Hilary Putnam and Elizabeth Anderson respectively. In examining Putnam's argument for

the democratisation of scientific inquiry, and in Anderson's argument for Dewey's experimentalist

model of democracy, I will show how the pragmatist theory of inquiry uses two structural features

which allow inquiry to address and regulate disagreement within inquiry. In particular, the pragmatist

_____

the problems in the evidence for our beliefs if they were relayed to us by another. Or, even if our doxastic attitudes can cause us to pursue different lines of inquiry and obtain different bodies of evidence as a result, this does not mean that we cannot share our evidence so as to reach a common body of evidence which combines the result of our respective inquiries.

[278] Specifically, in a paper which responds to Eric MacGilvray's criticisms, Cheryl Misak and Robert Talisse note that the literature on the intersection between pragmatism and democratic theory can be roughly categorised into two separate topics: that is, on an epistemic argument that defends democratic norms on the basis of Peirce's constitutive norms of belief, and an epistemic account of social inquiry that is developed on the basis of Dewey's experimentalist model of democracy. Under Misak and Talisse's distinction, this thesis focuses solely in the Deweyan account of social inquiry, given that the Peircean defence of democracy is a philosophical argument in favour of democratic norms, while the Deweyan account is an *explanatory* account of democratic inquiry. The Deweyan account is therefore more important for the purposes of this thesis. For more on the Peircean defence of democracy, see Misak, *Truth, Politics, Morality*; Robert B. Talisse, *Democracy After Liberalism: Pragmatism and Deliberative Politics* (New York: Routledge, 2005); Robert B. Talisse, *A Pragmatist Philosophy of Democracy* (New York: Routledge, 2007); Cheryl Misak and Robert B. Talisse, "Debate: Pragmatist Epsitemology and Democratic Theory: A Reply to Eric MacGilvray," *The Journal of Political Philosophy* 22, no. 3 (2014): 366–76. For more on the Deweyan account of democratic inquiry, see Anderson, "The Epistemology of Democracy"; Hilary Putnam, "A Reconsideration of Deweyan Democracy," in *The Pragmatism Reader: From Peirce Through the Present*, ed. Robert B. Talisse and Scott F. Aikin (Princeton; Oxford: Princeton University Press, 2011), 331–52; Jack Knight and James Johnson, *The Priority of Democracy: The Political Consequences of Pragmatism* (Princeton, NJ: Princeton University Press, 2011); and Putnam and Putnam, "Dewey's Logic". For more information on democratic experimentalism, and the theory of inquiry as a distinctly epistemic theory, see Macarthur, "A Kant-Inspired Vision of Pragmatism as Democratic Experimentalism". And for some pragmatists which criticise the pragmatist literature on epistemic democracy, by arguing that democracy ought to be defended on the basis of moral and political rather than epistemological reasons, see Eric MacGilvray, "Democratic Doubts: Pragmatism and the Epistemic Defense of Democracy," *Journal of Political Philosophy* 22 (2014): 104–23; Festenstein, "Inquiry as Critique". In particular, responding to MacGilvray's arguments will be the central focus of Appendix A of this thesis.

theory of inquiry explains how inquiry uses *ethical norms* which are instated within *social institutions* to produce an epistemic environment that is most conducive towards social cooperation, and therefore to channel disagreement in an epistemically productive way.

Let us examine each argument in turn. First, in response to the problem of subjectivity, Putnam argues that the preferred solution of the pragmatists is to look to social cooperation, and especially social cooperation of the kind found within scientific inquiry. In particular, the pragmatists argue that the epistemic success of scientific inquiry stem from a series of ethical norms which best allow for the sharing and evaluation of information, what Habermas and Apel refer to as "discourse ethics". Let us unpack this notion of discourse ethics, as well as how it connects to the pragmatist notion of the democratic structure of inquiry. Consider the following passage from Putnam:

> But–and this is the crucial point–that cooperation must be of a certain kind in order to be effective. It must, for example, obey the principles of "discourse ethics." Where there is no opportunity to challenge accepted hypotheses by criticizing the evidence upon which their acceptance was based, or by criticizing the application of the norms of scientific inquiry to that evidence, or by offering rival hypotheses, and where questions and suggestions are systematically ignored, the scientific enterprise always suffers…Moreover, it is not just that, on Dewey's conception, good science requires respect for autonomy, symmetric reciprocity, and discourse ethics–that could be true even if scientific theories and hypotheses were, in the end, to be tested by the application of an algorithm, such as the inductive logic for which Carnap hoped–but that, as already observed, the very *interpretation* of the *non-*algorithmic standards by which scientific hypotheses are judged depends on cooperation and discussion structured by the same norms. Both for its full development and for its full application to human problems, science requires the *democratization of inquiry*.[279]

---

[279] Putnam, "Pragmatism and Moral Objectivity," 172–73. The italics is the author's. It is important to note that in this passage, Putnam is using discourse ethics to refer to specifically the norms of communication which Habermas and Apel identify as guiding public discussion and debate, and not to how Habermas and Apel use these norms to provide an argument for democracy as a political order. In the discussion of pragmatism within democratic theory, there is debate amongst pragmatists about how one should defend democracy: while some point to Habermas and Apel's transcendental argument for democracy on the basis of our discursive practices, others give an epistemic defence of democracy on the basis of Peirce's constitutive norms of belief, while others still argue that democracy should be defended on the basis of moral and political rather than epistemological reasons. For more on Habermas and Apel's discourse ethics, see Jurgen Habermas, *The Theory of Communicative Action*, 2 vols. (Boston: Beacon Press, 1985/1989); Jurgen Habermas, *Moral Consciousness and Communicative Action*, trans. Christian Lenhart and Shierry Weber Nicholson (Cambridge: MIT Press, 1990); Jurgen

Let us consider two important points from the passage above. First, Putnam explains how the norms of

discourse ethics are *ethical* norms which establish the social conditions that are most conducive to

cooperation, and thus effective social inquiry. As a brief introduction, Habermas and Apel's discourse

ethics aims to provide a theory of moral justification which sees the validity of moral claims as being

assessed according to its acceptance by the relevant parties within a "reasonable" public discussion. To

explicate the notion of "reasonable" discourse, Habermas outlines what he calls the "ideal speech

situation", which is governed by a series of discursive norms such that each participant is able to

question previous assertions within the discussion, introduce new assertions to the discussion, and

express any attitudes towards assertions without hesitation (that is, without fear from internal or

external coercion).[280]

For Putnam, the norms that Habermas see as characterising reasonable public discourse about moral

claims are the same norms which govern reasonable discourse within scientific inquiry. In the above

passage, Putnam helpfully distinguishes between the norms of discourse ethics and the maxims of the

scientific method by highlighting how the norms of discourse ethics serve not part of the scientific

method itself (insofar as they are not norms about reasoning and inference), but ethical norms which

---

Habermas, *Justification and Application: Remarks on Discourse Ethics*, trans. Ciaran Cronin (Cambridge, Massachusetts; London, England: MIT Press, 1993); and Karl-Otto Apel, *Diskurs Und Verantwortung: Das Problem Des Übergangs Zur Postkoventionellen Moral* (Frankfurt am Main: Suhrkamp, 1985). For more on Habermas and Apel's epistemic justifications of democracy, see Putnam, "A Reconsideration of Deweyan Democracy". For a critique of Habermas and Apel's arguments, see Misak, *Truth, Politics, Morality*, 35–47.

[280] A greater examination of Habermas' notion of ideal speech situation, let alone Habermas and Apel's larger project of discourse ethics, would require much more space, and thus outside the scope of this thesis. For Habermas' original formulation of the ideal speech situation, see Jurgen Habermas, "Discourse Ethics: Notes on a Program of Philosophical Justification," in *Moral Consciousness and Communicative Action*, trans. Christian Lenhart and Shierry Weber Nicholson (Cambridge: MIT Press, 1990), 43–115.

govern over how different interpretations and applications of the scientific maxims are to be discussed over the course of scientific inquiry. And yet, this does not mean that the norms of discourse ethics are any less important to the success of scientific inquiry. As Putnam explains, if the norms of discourse ethics which properly regulate discussion and criticism (that is, discussion about what future lines of scientific inquiry are most promising and worthy of being pursued, or criticism as to the reliability and validity of a particular experimental method) within scientific practice are not followed, the "scientific enterprise always suffers" because it does not have the social cooperation needed for scientific progress.

This brings us to Putnam's second point: that though the norms of discourse ethics are *ethical* norms, they are employed within the theory of inquiry for distinctly *epistemic* reasons.[281] For Putnam, the reason why the norms of discourse ethics are important for scientific inquiry is not simply because of *moral* reasons: that is, the norms of discourse ethics are not implemented and followed within scientific inquiry for the sole purpose of giving each scientist the appropriate respect they deserve as moral agents. Nor are the norms of discourse ethics used only to test the validity of moral claims, as they are in Habermas and Apel's original project.[282] Rather, the norms of discourse ethics are integral

---

[281] In the seminal paper on the Deweyan account of democratic inquiry, Putnam makes this connection between the ethical and the epistemic more explicit, in stating that Dewey's account of democracy provides an epistemic justification for democracy as a form of social life. For more information, see Putnam, "A Reconsideration of Deweyan Democracy".

[282] This caveat is given to point out that there might be small theoretical differences between Habermas and Apel's discourse ethics as a theory of moral justification, and Putnam's explication of scientific inquiry here. However, given that both Habermas and Apel are largely sympathetic to pragmatism (if not, being full-fledged pragmatists themselves), it is not clear that they would make any strong distinctions between using discourse ethics to assess epistemic justifications versus moral justifications. This is especially the case if one accepts Putnam's arguments against a strong dualist notion of the fact/value distinction: if Putnam is right that matters of fact inform matters of value and vice versa, then there does not seem to be any problem in using norms about evaluative discourse to govern scientific inquiry into facts. That being said, within the literature on epistemic democracy, Estlund has argued against the use of Habermas' ideal speech situation to explain and justify the epistemic value of democracy. For Estlund, the use of Habermas' ideal speech situation to defend democracy would

to the *epistemic* success of scientific inquiry because of reasons which we discussed at the end of the last chapter: that is, because the scientific method consists of a series of contextual maxims rather than a universally applicable algorithm.

Here, Putnam helpfully points out that if Carnap's hope for an inductive logic was realised, then scientific inquiry need not be a cooperative activity at all: a scientist could simply compute by themselves which scientific theory or hypothesis is best confirmed by the evidence through the use of a universally computable algorithm.[283] This would mean that while the norms of discourse ethics might

---

require an analogy between democracy and a social contract. Estlund criticises this analogy, arguing that a social contract does provide an helpful model of the ethics of democracy, but does not accurately represent the epistemic features of democratic deliberation. As such, Estlund provides his own "epistemic departure view," which replaces Habermas' ideal speech situation for a model of ideal epistemic deliberation. While considering Estlund's account in greater detail will be outside the scope of this essay, I would submit that Estlund's "epistemic departure view" is not too dissimilar to Dewey's ameliorative account of democracy. See §1.2.1 for more on the connection between the similarity between Estlund and Dewey's view on democracy. For more information specifically on Estlund's epistemic departure view, see Chapter 9 of Estlund, *Democratic Authority*. See Appendix A of this thesis for more on Putnam's argument against the fact/value distinction.

[283] It is interesting to note how the difference between Putnam and Carnap persists in debates surrounding theory choice within contemporary philosophy of science: contrary to Kuhn's claim that there is "no neutral algorithm for theory choice" in science, Okasha has argued for the possibility of an algorithm for theory choice by using the formal apparatus for preference aggregation developed within social choice theory. This is possible because, according to Okasha, the problem of theory choice which Kuhn posed – namely, the problem of how we are to (rationally) choose between scientific theories when different theoretical virtues (such as simplicity, coherence, and explanatory power) recommend different scientific theories – is formally identical to the problem in social choice theory of how we are to aggregate a set of individuals' preferences into one social preference. Okasha therefore concludes that – barring problems related to Arrow's impossibility theorem – it is theoretically possible to use the formal apparatus of preference aggregation within social choice theory to aggregate the performance of scientific theories under the various theoretical virtues so as to provide a rational ordering of which scientific theory is to be preferred. Examining Okasha's argument as well as other criticisms and elaborations within the literature on rational theory choice is outside the scope of this essay – suffice to say, Okasha's argument is not decisive, and it is still a point of contention within the literature as to whether one can provide a formal algorithm for rational theory choice. In contrast to this discussion on rational theory choice, contemporary pragmatism has followed Putnam in criticising the use of formal analyses of the epistemology of democracy. As we will see in the following discussion on Anderson's analysis of the epistemology of democracy, one way of reading Anderson's argument is as a criticism of formal epistemic analyses such as

still be morally appropriate when communicating the results of this algorithm between scientists, they would not be epistemically important as they would not have an effect on the truth of these results, nor on the justifications one would have for a scientific theory or hypothesis. However, if Putnam's criticism of Carnap is right – that is, if the scientific method cannot be formalised into a universally applicable algorithm, because the scientific method consists of maxims which must be interpreted by scientists in a given context – then scientists play an ineliminable role within scientific practice. In particular, scientists play an active role in choosing the epistemic method by which we evaluate scientific hypotheses and theories, such that inquiry requires the judgment of scientists who decide which lines of inquiry are promising in giving us the relevant evidence for or against a scientific theory, or who decide whether one's experimental method is sufficiently reliable and valid as to ensure the acceptability of the empirical results of an experiment as evidence for a theory or hypothesis. To the extent that such judgments are susceptible to the kind of problems surrounding subjectivity which we saw in Kelly's analysis of belief polarisation, the norms of discourse ethics therefore play a distinctly *epistemic* role in leveraging the diversity of opinions and ideas of a community of inquirers to overcome the subjective limitations of any one inquirer. In allowing and empowering each participant to freely express their opinion and ideas, the norms of discourse ethics create a social epistemic environment which allows the diversity of opinions and ideas to be assessed and evaluated in a cooperative way.

---

the Condorcet Jury Theorem and the Diversity Trumps Ability Theorem, which fail to account for how democracies *actually* function as epistemic institutions. More below. For more information on Okasha's original argument, as well as the subsequent discussion, see Samir Okasha, "Theory Choice and Social Choice: Kuhn Versus Arrow," *Mind* 120, no. 477 (2011): 83–115; Jacob Stegenga, "Theory Choice and Social Choice: Okasha Versus Sen," *Mind* 124, no. 493 (2015): 263–77; Michael Morreau, "Theory Choice and Social Choice: Kuhn Vindicated," *Mind* 124, no. 493 (2015): 239–62; Samir Okasha, "On Arrow's Theorem and Scientific Rationality: Reply to Morreau and Stegenga," *Mind* 124, no. 493 (2015): 279–94; Seamus Bradley, "Constraints on Rational Theory Choice," *British Journal for the Philosophy of Science* 68, no. 3 (2017): 639–61.

In summary, Putnam's explanation of the use of the norms of discourse ethics within scientific inquiry helps to further characterise the pragmatist theory of inquiry. For the pragmatists, it is not enough to simply allow each participant within a given inquiry the right to speak, since unrestrained public discussion can often lead to certain participants to oppress the opinions of others through verbal coercion. In order for inquiry to function as an effective form of social epistemic practice, the pragmatists argued that inquiry must be conducted according to ethical norms which ensure that the diversity of opinions and ideas within the community of inquirers are properly received and fairly evaluated. Putnam claims that these ethical norms lead to the *democratisation* of inquiry, since the function of these ethical norms are democratic insofar as allow every participant within public discussion the ability to exchange and evaluate information: "both for its full development and for its full application to human problems, science requires the democratization of inquiry."

5.3.2   The Epistemology of Democracy: Anderson on the Epistemic Function of Democratic Institutions

Within contemporary pragmatist thought, Putnam's exhortation for the democratisation of inquiry has been taken up by subsequent Deweyan thinkers, who build upon Dewey's theory of inquiry in highlighting the importance of certain social institutions in realising and regulating these norms within a given inquiry.[284] In particular, within the contemporary literature on the intersection between pragmatism and democratic theory, many pragmatists have developed substantive accounts of democracy which see democracy not only as a political system but as a social epistemic institution,

---

[284] A few comments are in order here. First, this claim that subsequent Deweyan thinkers are building upon Putnam's account of inquiry is taken from Eric MacGilvray's summary of the literature – see MacGilvray, "Democratic Doubts," 105. Second, it is important to remember that my use of the term "regulate" here does not mean that the pragmatist theory of inquiry involves providing "regulative" norms. Instead, it is merely referring to how public discussion within inquiry is channeled in a productive manner. For more on why Dewey's theory of inquiry involves neither constitutive nor regulative norms, see fn. 22.

such that "democratic decision processes make better use of the distributed knowledge that exists in society than do their rivals."[285]

In this section, I will focus on Elizabeth Anderson's epistemic analysis of Dewey's experimentalist model of democracy for three reasons. First, Anderson's paper provides the clearest exposition of Dewey's epistemic account of democracy, and how this pragmatist account of democracy differs from the other accounts of epistemic democracy within the literature. Second, Anderson's epistemic analysis of the pragmatist account of democracy complements and elaborates on Putnam's argument about the norms of discourse ethics, since Anderson's analysis reveals the role which *social institutions* play in instating and regulating the norms of discourse ethics within a given inquiry. Finally, a large part of Anderson's argument involves explaining how Dewey's experimentalist account of democracy reveals the "epistemic import of dissent". In explaining how democratic institutions use disagreement as an epistemic resource for better collective decision-making, Anderson's analysis helps to clarify the structural features of inquiry, and how they leverage disagreement in an epistemically productive way.

Given these preliminary remarks, let us now begin with a synopsis of Anderson's argument. The main aim of Anderson's paper is to compare Dewey's experimentalist model of democracy with two other epistemic models of democracy – namely, the Condorcet Jury Theorem, and the Diversity Trumps Ability Theorem (DTA Theorem for short) – and to show how Dewey's experimentalist model of democracy provides the best explanation of the epistemic features of democratic institutions. Throughout this argument, Anderson's focus is on the following three constitutive features of democracy: "the epistemic diversity of participants, the interaction of voting with discussion, and

---

[285] Knight and Johnson, *The Priority of Democracy*, p. 151. This quote has been used by various pragmatists as representative of the central role which deliberative democracy plays within pragmatist social epistemology: for places where other pragmatists have affirmed this quote, see MacGilvray, "Democratic Doubts," 114; Misak and Talisse, "Debate," 375.

feedback mechanisms such as periodic elections and protests."[286] Anderson argues that only Dewey's experimentalist model of democracy provides an explanatory *account* of these features of democracy because only Dewey's model of democracy explains how these features allow democracy to function as an *epistemic* institution.

This is especially the case when we compare Dewey's model of democracy with other models such as the Condorcet Jury Theorem and the DTA Theorem. As a brief primer, within the literature on democratic theory, the Condorcet Jury Theorem and the DTA Theorem are often cited as formal mathematical results which provide justification for why democracy is a desirable form of government from an epistemic standpoint.[287] The upshot of these two theorems is as follows: whereas the Condorcet Jury Theorem states that a group of reasonably correct voters (i.e. voters who are correct more than 50% of the time) will rapidly approach probability 1 in yielding the right answer through majority vote as the group increases in size, the DTA Theorem concludes that groups of epistemically diverse non-experts will consistently outperform groups of experts in solving difficult problems. While such mathematical theorems do provide *prima facie* justifications for why democracy is epistemically desirable, Anderson's critique of these theorems centres around the fact that these results do not in and of themselves provide "a particularly illuminating account of the epistemic powers of

---

[286] Anderson, "The Epistemology of Democracy," 8.

[287] For an extensive list of the various ways in which the Condorcet Jury Theorem has been used within the literature on epistemic democracy, see the bibliography in Christian List and Robert E. Goodin, "Epistemic Democracy: Generalizing the Condorcet Jury Theorem," *Journal of Political Philosophy* 9, no. 3 (2001): 277–306. See also Jeremy Waldron, "The Wisdom of the Multitude: Some Reflections on Book 3, Chapter 11 of Aristotle's Politics," *Political Theory* 23, no. 4 (1995): 563–84; Hélène Landemore, *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many* (Princeton: Princeton University Press, 2013).

democracy".[288],[289] In other words, while the Condorcet Jury Theorem and the DTA Theorem give us reasons to believe *that* democracy will lead to good results for an epistemic community, the theorems are insufficient to explain *how* democracy leads us to obtain these epistemic results.[290] Consequently,

---

[288] Anderson, "The Epistemology of Democracy," 11. While Anderson is only referring to the Condorcet Jury Theorem here, the following discussion will also outline the theoretical shortcomings of the DTA Theorem as an epistemic model of democracy by considering criticisms from Anderson as well as the contemporary literature.

[289] The reason why I state that these theorems give "*prima facie*" justifications for democracy is because there are considerable criticisms about both theorems which challenge the epistemic and political implications of both the Condorcet Jury Theorem and the DTA Theorem. For criticisms of the Condorcet Jury Theorem, see Duncan Black, *The Theory of Committees and Elections*, Second (Cambridge: Cambridge University Press, 1963); Krishna K. Ladha, "The Condorcet Jury Theorem, Free Speech, and Correlated Votes," *American Journal of Political Science* 36, no. 3 (1992): 617–34; David Estlund, "The Epistemic Dimension of Democratic Authority," *The Modern Schoolman* 74, no. 4 (1997): 259–76; Daniel Berend and Jacob Paroush, "When Is Condorcet's Jury Theorem Valid?" *Social Choice and Welfare* 15, no. 4 (1998): 418–88; David M. Estlund, "The Irrelevance of the Jury Theorem," in *Democratic Authority: A Philosophical Framework* (Princeton; Oxford: Princeton University Press, 2008), 223–36; and Franz Dietrich and Kai Spiekermann, "Jury Theorems," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Summer 2022 (Metaphysics Research Lab, Stanford University, 2022), §4. For criticisms of the DTA Theorem (at least as it is originally derived in the Hong-Page model), see Abigail Thompson, "Does Diversity Trump Ability?" *Notices of the AMS* 61, no. 9 (2014): 1–24; Sameer Bajaj, "Review of Democratic Reason: Pollitics, Collective Intelligence, and the Rule of the Many, by hélène Landemore," *Ethics* 124, no. 2 (2014): 426–31; Paul J. Quirk, "Making It up on Volume: Are Larger Groups Really Smarter?" *Critical Review* 26, no. 1-2 (2014): 129–50; Patrick Grim et al., "Diversity, Ability, and Expertise in Epistemic Communities," *Philosophy of Science* 86, no. 1 (2018): 98–123; Bennett Holman et al., "Diversity and Democracy: Agent-Based Modeling in Political Philosophy," *Historical Social Research* 43, no. 1 (2018): 259–84. For a defence of the DTA Theorem, see Daniel J. Singer, "Diversity, Not Randomness, Trumps Ability," *Philosophy of Science* 86 (2019): 178–91.

[290] In their excellent work on a pragmatist justification for democracy, Knight and Johnson claim that the "priority of democracy" is best justified by a focus on the formal institutions of democracy, and an explanation of *how* these institutions enhance and improve our collective decision making. Consider the following quote: "Other advocates of democracy may place their primary emphasis elsewhere, but we would argue that any compelling attempt to justify democracy must give pride of place to the formal institutions that facilitate democratic governance. As we argue, the formal institutional rules set the terms and conditions of democratic participation and practice and, in doing so, vitally influence the quality and effectiveness of our collective decision making. When, in our efforts to defend democratic politics, we neglect to account for how we will actually implement our most basic values and commitments, for how we will institutionalize the abstract democratic ideal, then we fail to make the most persuasive case for the priority of democracy. Too often we focus on the question of why we should act democratically at the expense of the equally important explanation of how we should do so." In

Anderson argues that if we take the Condorcet Jury Theorem and the DTA Theorem as epistemic "models" or "accounts" of democracy, what we end up with is an inaccurate and incomplete account of what actually allows democracy to function as an epistemic institution.

To see how these theorems fail to capture the actual epistemic features of democracy, let us begin with Anderson's critical examination of the Condorcet Jury Theorem. For Anderson, the Condorcet Jury Theorem does not provide an adequate account of the epistemic functions of democracy, seeing as *none* of the three aforementioned constitutive features of democracy are properly accounted for by the Condorcet Jury Theorem. First, while the Condorcet Jury Theorem suggests that a larger group of epistemic agents within a majority vote would lead to the correct solution to a social problem, the Condorcet Jury Theorem does not adequately account for why a large group of epistemic agents is necessary for democratic decision-making. Here, Anderson explains how the epistemic justification for democratic principles such as universal suffrage does not come from an appeal to the size of the group of participants that is included, but from an appeal to the *epistemic diversity* of the participants included within democracies.[291] More specifically, Anderson explains how the epistemic diversity of voters is a constitutive feature of democracy due to the complexity of the social problems which democracies address through policy making.[292] Given the complexity and scale on which political

the remainder of this section, I shall explain how this focus on the question of "how" rather than "why" will reflect how the pragmatist account of democracy is a *transition* epistemology of democracy. More below. Preface of Knight and Johnson, *The Priority of Democracy*, x.

[291] For example, see Waldron, "The Wisdom of the Multitude"; Landemore, *Democratic Reason*.

[292] For more information about how the complexity of social problems leads to asymmetrically distributed information that requires epistemic diversity in democracy, see Dewey, *The Public and Its Problems*; Knight and Johnson, *The Priority of Democracy*, 159–61; Landemore, *Democratic Reason*, 82–89. In addition, Anderson, Landemore, and Knight and Johnson point to Friedrich Hayek as being the first and most influential theorist on how epistemic institutions require a vehicle by which they acquire the socially dispersed information needed for effective policy making. However, while Hayek, being

problems must be addressed, political problems and policies often have consequences whose effects are asymmetrically distributed between different demographics. This in turn means that the *information* of these effects on different demographics is itself asymmetrically distributed within society, such that different groups of individuals can have considerably different experiences about the consequences of a political issue or policy.

It is for this reason that democratic theory has pointed to the epistemic diversity of voters as a constitutive feature of democracy, as it is only through an epistemically diverse group of voters that we can gather the information that is asymmetrically distributed across different demographics which is necessary for effective policy making. However, Anderson points out that the Condorcet Jury Theorem works even if voters are epistemically homogenous, such that Condorcet's original proof even assumed homogeneity.[293] In claiming that effective policy-making is guaranteed from a sufficiently large group of voters, the Condorcet Jury Theorem – taken as an independent epistemic account of democracy – fails to account for how democratic decision-making requires not only a large, but an epistemically diverse group of voters in order to function as an effective epistemic institution.

---

principally an economist, focused solely on market prices as a vehicle for obtaining this form of information, Anderson, as well as Knight and Johnson, notes that institutions can also obtain and transmit information through public forums, social media, and votes. For more on Hayek's original analysis of socially dispersed knowledge, see F. A. Hayek, "The Use of Knowledge in Society," *The American Economic Review* 35, no. 4 (1945): 519–30; and F. A. Hayek, "Economics and Knowledge," in *L. S. E. Essays on Cost*, ed. James M. Buchanan and G. F. Thirlby (New York: New York University Press, 1981). For critical interpretations of Hayek's work from a pragmatist perspective, see Anderson, "The Epistemology of Democracy," 8–9; and Knight and Johnson, *The Priority of Democracy*, 52–55.

[293] Anderson, "The Epistemology of Democracy," 11. In fact, Berend and Paroush show that once we move away from the homogenous case, then the theorem is unlikely to hold. This suggests that proponents of the Condorcet Jury Theorem can only appeal to this theorem in cases where there is *substantive* evidence that a group of voters is in fact reasonably competent. See Berend and Paroush, "When Is Condorcet's Jury Theorem Valid?".

Second, the initial premises of the Condorcet Jury Theorem obscures the important epistemic role

which public discussion plays as a constitutive feature of democracy. In particular, Anderson argues

that the problem lies with how the Condorcet Jury Theorem assumes that voters vote independently of

one another, an assumption which is at odds with the emphasis which democracies place in the need

for public discussion prior to and in conjunction with voting.[294] While the assumption that voters vote

independently of each other might seem to be as a reasonable assumption to ensure that one's vote is

not unduly influenced by another, Anderson explains how this assumption is untenable in practice

because public discussion is necessary prior to voting in order for epistemic agents to be informed in

their voting decisions. Without public discussion in the forms of a free press and public forums,

epistemic agents lack any mediums by which they can share information and opinions, and thus fail to

acquire the information needed to make their vote meaningful.[295] As such, the assumption of

---

[294] While Anderson notes that the condition of independent voting within the Condorcet Jury Theorem does not rule out all influence of voters on one another, the condition of independent voting is at least in tension with the emphasis for public forums for sharing and discussing information in democracy. Franz Dietrich and Kai Spiekermann claim that the setup of the Condorcet Jury Theorem leads to a dilemma about how to interpret deliberation: "Two intuitions compete: does deliberation primarily threaten collective epistemic success, by reducing judgmental independence, or primarily increase epistemic success, by raising individual competence?" Anderson sees this kind of dilemma as reflective of the problems with the Condorcet Jury Theorem as a model of democracy, because an adequate account of democracy should show how public discussion works in *tandem* with the voting process, not in *opposition*. See Anderson, "The Epistemology of Democracy," 11; Dietrich and Spiekermann, "Jury Theorems", §4.5.

[295] In support of this claim, Anderson notes that a shareholder vote would not be an accurate representation of the opinions and interests of the respective shareholders, if few to none of the shareholders have any of the necessary information about the records and positions of the nominees for corporate boards of directors. To the extent that the rationality of collective decision-making is dependent on the evidence that each respective party has, Anderson shows that any epistemic account of democracy must therefore take into consideration the role of public discussion in allowing each citizen to make an informed vote which accurately represents their opinions and interests. Of course, Anderon's account does not guarantee that the media platforms and public fora which are available to us in our actual context *do* facilitate public discussion in this way – there is, for instance, good reason to believe that in many places in the world that the free press is not as "free" as it purports to be. Nevertheless, to the extent that we do want the epistemic features which Dewey's experimentalist account of democracy outlines, institutions which facilitate public discussion in the way that Anderson outlines here are a necessary condition to

independent voting within the Condorcet Jury Theorem is in direct tension with the important epistemic role in which public discussion plays within democratic decision-making.

Finally, the Condorcet Jury Theorem does not explain the role of the various feedback mechanisms within democracy, such as periodic elections and protests, which occur *after* the voting process. As a preliminary remark, it is important to note that if the Condorcet Jury Theorem is taken as a model which is specifically targeted towards understanding the epistemic strengths of majority voting in sufficiently large groups, then there is obviously no need for the Condorcet Jury Theorem to account of the stages of democratic decision-making which occur after the voting process. However, Anderson argues that the use of the Condorcet Jury Theorem as the *sole* model of the epistemic strengths of democracy tends to lead epistemic democrats to forget how democratic decision-making assesses the effectiveness of public policy by looking at the *consequences* of a given policy rather than its "*ex ante* popularity" within a vote.[296] In representing the epistemic strengths of democratic decision-making solely in terms of majority voting, epistemic democrats who defend democracy on the basis of the Condorcet Jury Theorem fail to explain why feedback mechanisms are necessary within democratic decision-making at all: if a sufficiently large group of voter is almost guaranteed to yield the right answer in a majority vote, what need is there for democracies to allow processes such as periodic elections or protests which occur after the voting process? By contrast, Anderson argues that feedback mechanisms such as periodic elections and protests are constitutive features of democracy precisely because in actual political practice it is common for "majorities [to] converge on an inefficient solution

---

effective collective decision-making. More on this when we explain how Dewey's utopian vision of democracy connects to our actual epistemic practice in the next section.

[296] Anderson, "The Epistemology of Democracy," 12.

because they fail to anticipate certain consequences of the policies they adopt."[297] Given that actual

policy-making do not lead us to always yield the right answer by majority vote (as the Condorcet Jury

Theorem suggests), feedback mechanisms such as periodic elections and protests are a necessary part

of democratic decision-making so that we "can learn how to devise better solutions and correct [our]

course in light of new information about the consequences of policies."[298] The Condorcet Jury

Theorem therefore fails to account for the various feedback mechanisms within democracy, which play

an important epistemic role due to the fallibility of actual political decision-making.

At this point, it is important to reiterate that Anderson's criticism against the Condorcet Jury Theorem

as an epistemic account of democracy does not mean that the Condorcet Jury Theorem does not play

an important role within discussions around the epistemology of democracy. On the contrary – within

democratic theory, the Condorcet Jury Theorem is still an important theorem which serves to refute a

key objections which advocates of epistocracy have towards democratic practice: namely that

democratic principles such as universal suffrage may lead a generally uninformed majority to make

inconsistent if not generally poor voting decisions.[299] In showing how democratic processes do not lead

to worse (but perhaps better) epistemic results because of its inclusive voting principles, the Condorcet

---

[297] Anderson, "The Epistemology of Democracy," 12.

[298] Anderson, "The Epistemology of Democracy," 12.

[299] Knight and Johnson note that the developments on the Condorcet Jury Theorem have only strengthened the evidence for 'the wisdom of the multitude", claiming that "relatively recent research [into the Condorcet Jury Theorem] goes a considerable distance toward mitigating such deep-seated skepticism." In particular, Knight and Johnson point to List and Goodin's argument which goes beyond the initial conditions of the Condorcet Jury Theorem to consider a variety of other aggregation mechanisms which all seem to support the same conclusion as the original Condorcet Jury Theorem. See List and Goodin, "Epistemic Democracy".

Jury Theorem has a strong part to play within the justification and defence of democracy as a form of government.

Nevertheless, the point of Anderson's criticism has been to show how the Condorcet Jury Theorem, taken by itself, provides an epistemic account of democracy which is incomplete if not in tension with the actual epistemic features of democracy. Anderson's criticism has been supported by recent work on epistemic democracy: for example, Landemore's account of democratic reason sees democratic decision-making as comprising not only majority vote (as supported by the Condorcet Jury Theorem), but also *inclusive deliberation*.[300] Landemore rejects the claim made by aggregative democrats that democratic decision-making is achievable simply by majority rule, because "majority rule is constitutively unable to formulate the options that are voted on."[301] As such, Landemore argues that democratic decision-making also requires active deliberation within political organisations and parties, an epistemic process which is best captured by the DTA Theorem.

On this point, Anderson agrees with Landemore about the epistemic strengths of the DTA Theorem. For Anderson, the epistemic importance of active deliberation within democracy is clearly captured by the DTA Theorem, as seen by how the DTA Theorem accounts for two of the aforementioned constitutive features of democracy. Specifically, these are the need for epistemic diversity within participants, and the role of public discussion within democratic decision-making. With regards to the need for epistemic diversity, the basic result of the DTA Theorem shows how epistemic diversity (not

---

[300] Landemore states that these two features are the two "mechanisms" of democratic reason. See Landemore, *Democratic Reason*.

[301] Landemore, *Democratic Reason*, 146.

simply expert opinion) is required in order to find the correct solution in collective problem solving.[302] With regards to the role of public discussion, the DTA Theorem represents the role of active discussion and deliberation within its model of problem solving. In particular, Anderson notes that because the DTA Theorem was "initially constructed to model problem solving within firms…the DTA theorem models some of the epistemic functions of citizens' associations and political parties."[303] The DTA Theorem is therefore able to represent the role of active deliberation within democratic decision-making by explaining how political parties can "help diverse citizens work together in smaller groups to hammer out proposed solutions to problems [such that] discussion [is] epistemically productive, not merely as something that potentially interferes with the epistemic virtues of vote aggregation."[304]

However, despite its theoretical strengths, the DTA Theorem fails to provide an adequate epistemic account of democracy for two reasons.[305] First, one point which Anderson's argument does not discuss is how recent research has raised questions as to the scope of the validity of the DTA Theorem. In particular, Sakai's robustness analysis of the DTA Theorem suggests that the DTA phenomenon "is robust only for problem-solving cases about *unpredictable issues*."[306] This means that while "random

---

[302] Anderson, "The Epistemology of Democracy," 12.

[303] Anderson, "The Epistemology of Democracy," 12.

[304] Anderson, "The Epistemology of Democracy," 12–13.

[305] In Anderson's original argument, Anderson also notes how the DTA Theorem does not account for the noninstrumental importance of universal inclusion. However, because this point is more important for democratic ethics than for our purposes of social epistemology, I have omitted this point within my overview of Anderson's argument. See Anderson, "The Epistemology of Democracy," 13.

[306] Ryota Sakai, "Mathematical Models and Robustness Analysis in Epistemic Democracy: A Systematic Review of Diversity Trumps Ability Theorem Models," *Philosophy of the Social Sciences* 50, no. 3 (2020): 202. Italics is my own.

and diverse groups can perform quite well, and even optimally (e.g. when problem solving performance is a matter of luck)…they are not the best at identifying optimal solutions under all circumstances. When problems admit of true ability…expertise is plausibly important."[307] As such, the current literature suggests that the claim that diversity trumps ability is only valid in a limited amount of cases, and that "outside of the highly unpredictable range, some studies [support] a mixed population of experts and laypersons as the best composition of problem-solving groups."[308] The current research therefore shows that the DTA Theorem overdetermines the role which inclusive deliberation plays in effective collective decision-making, and that optimal collective decision-making actually a modest view on inclusive deliberation which respects the need to defer to expert judgment in certain situations. We discuss this problem of expertise, and the need for a more nuanced view on deliberative democracy in further detail in Appendix A of this thesis.

Secondly, one major part of democratic decision-making which Anderson argues is left unaccounted for by the DTA Theorem is the "epistemic functions of periodic elections and other feedback mechanisms designed to change the course of collective decisions in light of information about their consequences.[309] In representing democratic decision-making in terms of a single deliberative moment, the DTA Theorem suffers the same problem to the Condorcet Jury Theorem: namely, it fails to see how democratic decision-making is an *ongoing* process (what Anderson calls the"dynamic aspects of democracy") which requires feedback mechanisms so as to correct the fallible and developing process of policy-making.

---

[307] Holman et al., "Diversity and Democracy," 273–74.

[308] Sakai, "Mathematical Models and Robustness Analysis in Epistemic Democracy," 203.

[309] Anderson, "The Epistemology of Democracy," 13.

This brings us to Anderson's argument for why Dewey's experimentalist model of democracy is the best epistemic account of democracy. Part of the reason why Anderson privileges Dewey's account of democracy is its ability to account for the epistemic function of feedback mechanisms within democracy, something which we have seen is left unexplained by the Condorcet Jury Theorem and the DTA Theorem. However, for Anderson, Dewey's attention to these feedback mechanisms within democracy is not simply a slight bonus to Dewey's account of democracy: rather, it is reflective of how Dewey's account of democracy differs *in kind* to the epistemic accounts of democracy given by the Condorcet Jury Theorem and the DTA Theorem.

We can highlight the difference between Dewey's account of democracy and the other two theorems by invoking the familiar distinction between ideal and non-ideal theory. More specifically, whereas the Condorcet Jury Theorem and the DTA Theorem give *end-state* accounts of democracy as an epistemic institution, Dewey's model instead provides a *transition* account of democracy. The reason why both the Condorcet Jury Theorem and the DTA Theorem are end-state theories of democracy is because both theorems see democracy as acquiring certain epistemic *ends* which are achieved through certain means – be it the attaining of true belief through majority voting, or the discovery of the optimal solution to a social problem through inclusive deliberation. As a result, both the Condorcet Jury Theorem and the DTA Theorem provide a fixed conception of democracy as an epistemic institution, a conception which serves as a long-term goal for how democratic decision-making should ultimately be done.

While both theorems still play an important role in understanding and justifying democratic practice, Anderson's criticisms have shown how the Condorcet Jury Theorem and the DTA Theorem fail to provide adequate epistemic accounts of democracy. In particular, it is because both theorems are end-state theories about democracy which leads both theorems to be unable to account for actual political situations which do not meet their internal or external standards of validity. In terms of internal

validity, the Condorcet Jury Theorem does not justify democratic practice if we have reasons to believe that the group of voters are not competent, while the DTA Theorem does not justify democratic practice if we have reasons to believe that the particular political problem is not in fact an "unpredictable issue". In terms of external validity, the Condorcet Jury Theorem and the DTA Theorem do not account for actual political situations where majorities converge on an inefficient solution, or when a diverse group of problem-solvers do not find the optimal solution respectively. As such, in providing a theoretical epistemic result which is restricted to certain *ideal* situations, both the Condorcet Jury Theorem and the DTA Theorem fail in two regards: both theorems fail to provide an accurate account of the actual epistemic features of democracy, and both theorems fail to provide a normative account that we can use to evaluate or guide our actual political practice.

By contrast, what makes Dewey's model of democracy different from the other two theorems is the fact that Dewey's account of democracy is not an epistemic *theorem* about democracy, but rather an explanatory account of the epistemic *functions* of democracy. In other words, whereas the Condorcet Jury Theorem and the DTA Theorem predict certain theoretical results about the epistemic efficacy of democratic processes from the attributes of the democratic community (e.g. from the size or epistemic diversity of the participants in democracy), Dewey's epistemic model of democracy argue that it is the *structural* features of democratic procedure and their functions which allows democratic decision-making to revise and improve its epistemic practice.[310] In particular, Dewey's model of democracy is

---

[310] As a sidenote, while Dewey's model of democracy does see the epistemic virtues of democratic decision-making as resulting from the *procedure* of democracy, MacGilvray helpfully distinguishes Dewey's view on democracy from David Estlund's epistemic proceduralism. The main difference for MacGilvray is this: while MacGilvray "agree[s] with Estlund in thinking that an adequate defense of democracy must take both moral and epistemic considerations into account, [he believes] that [Estlund] errs in trying to define the relationship between them in general terms." Here, we can once again use the ideal/non-ideal distinction to explain the difference more clearly: Estlund's epistemic proceduralism is an end-state epistemic account insofar as it aims to provide a general theory of democratic legitimacy which explicates democracy as a

described as "experimentalist" because Dewey sees political deliberation as following the same *procedure* as the experimental method of scientific inquiry: for Dewey, democratic decision-making can be seen to follow the experimental method in how democratic procedure consists of a group of participants who collectively gather to propose potential solutions to problems; discuss the foreseeable consequences of each solution and decide which solution is most likely to work; test the solution by implementing a policy which we act in accordance with them; and then evaluate the result of our actions so as to improve our future policies.[311] It is this experimental method which allows Dewey's account to provide a *transition* epistemology of democracy – one which sees democracy as comprising a series of institutions and processes which collectively serve as the epistemic *means* by which we continually correct and therefore improve democratic decision-making.

To illustrate Dewey's conception of democracy as social experimentation, and how it provides a transition account of democracy, let us examine how Dewey's account explains the epistemic function of feedback mechanisms within democracy. For Anderson, it is in the explanation of the feedback

political end on the basis of democratic procedure and its epistemic value. By contrast, MacGilvray (following Dewey) sees democratic legitimacy as a political form of inquiry, such that democratic procedure is explicated in terms of the theory of inquiry. In seeing democratic procedure as a form of political inquiry not unlike scientific inquiry, MacGilvray argues that the "procedure" of political inquiry has no general form – since, as we have seen, the scientific "method" contains a series of contextual maxims rather than an universally applicable algorithm. As such, the pragmatist explanation of the epistemic strengths of democracy is reflective of how the pragmatist theory of inquiry is a *transition* account: namely, that the epistemic strengths of democracy can only be properly accounted for by assessing how democratic inquiry is functioning on a case-by-case basis. Or, in MacGilvray's own words: "This, then, is the answer that pragmatism gives to the question of democratic legitimacy: that there is no general answer, only particular doubts, that particular people have about particular political practices." See MacGilvray, "Democratic Doubts," 120.

[311] Anderson, "The Epistemology of Democracy," 13.

mechanisms of democracy in which Dewey's experimentalist model most clearly distinguishes itself

from the other two epistemic models. Consider the following passage:

> *Most importantly*, Dewey's experimentalist model of democracy helps us see the epistemic import of several democratic institutions that sustain its dynamism, its capacity for change: periodic elections, a free press skeptical of state power, petitions to government, public opinion polling, protests, public comment on proposed regulations of administrative agencies. In Dewey's model, these are mechanisms of feedback and accountability that function to institutionalize fallibilism and an experimental attitude with respect to state policies. They push governments to revise their policies in light of evidence—public complaints, as expressed in both votes and discussion—that they are not working, or expected not to work…Dewey stressed that for democracy to work, it was not enough simply to institute legal arrangements such as representation and periodic elections. Culture had to change too, so that citizens at large, interacting with one another in civil society, welcome diversity and discussion, and take an experimental attitude toward social arrangements…To realize the epistemic powers of democracy, citizens must follow norms that welcome or at least tolerate diversity and dissent, that recognize the equality of participants in discussion by giving all a respectful hearing, regardless of their social status, and that institute deliberation and reason-giving, rather than threats and insults, as the basis of their communication with one another. An epistemic analysis of democracy helps us see that it is not just a matter of legal arrangements. It is a way of life governed by cultural norms of equality, discussion, and tolerance of diversity…Diversity and disagreement are central features of democracy.[312]

There are a number of important points to unpack from this passage. First, Anderson explains how

Dewey's experimentalist model of democracy helps us to see the epistemic role which the feedback

mechanisms of democracy play within democratic decision-making. As we can see in this passage,

Anderson explains how the function of the feedback mechanisms of democracy is to contribute to the

"dynamism" of democracy: that is, in consisting of democratic institutions and processes which give

democratic decision-making "the capacity to change". While Anderson does not explicitly state what

this "capacity to change" entails, we can see that Anderson is alluding to how democratic decision-

making consists not only of one-time decisions (e.g. a single majority vote or a single act of collective

political deliberation), but an ongoing process where decisions change over time. We can therefore

---

[312] Anderson, "The Epistemology of Democracy," 14–15. Italics is my own.

clarify and add to Anderson's point by making a distinction between synchronic and diachronic collective decision-making.

To explain, recall how Bayesian epistemology categorises the norms of belief revision into synchronic and diachronic norms of belief. Whereas the synchronic norms of belief determine how one's doxastic attitudes should be at a given moment in time by restricting one's doxastic attitudes to conform to the classical probability axioms, the diachronic norms of belief determine how one's doxastic attitudes should change in light of new information according to the Principle of Conditionalisation. In this distinction, Bayesian epistemology therefore characterises rational belief revision as consisting of two stages: one which considers how one's doxastic attitudes ought to be revised at a given moment (according to the synchronic norms), and one which considers how one's doxastic attitudes ought to be revised over time (according to the diachronic norm).

In a similar way, Dewey's experimentalist model of democracy allows us to see how democratic decision-making also consists of two stages: in addition to deciding particular policies and solutions to a given social problem through majority voting or inclusive deliberation (i.e. synchronic decision-making on a given social problem at a moment in time), democratic decision-making also consists of processes which allow for the revision of existing policies upon receiving new information (i.e. diachronic decision-making on an ongoing social problem).

It is this attention to the diachronic (or "dynamic") aspects of democratic decision-making which exemplifies how Dewey's experimentalist model is a transition account of democracy. For example, compare the explanation which Dewey provides for the epistemic role of majority voting and inclusive deliberation to that which is given by the Condorcet Jury Theorem and the DTA Theorem. For Dewey, democratic processes such as majority voting and inclusive deliberation do not *guarantee* the optimal solution to a social problem - rather, they serve an *experimental* function within democratic

procedure.[313] Just as how the role of experiments within scientific inquiry is not to provide conclusive knowledge of physical phenomena but to obtain empirical evidence in favour of certain hypotheses, the purpose of democratic processes such as majority voting and inclusive deliberation is not to decisively find the optimal solution to a social problem, but to obtain and use the information which is socially dispersed amongst different groups of individuals to make informed policies.[314]

Furthermore, given that the experimental method is an *iterative* process, the use of experiments within scientific inquiry are not as one-time events that definitively confirm a scientific hypothesis, but a continual process where further experimentation can always provide new empirical evidence that defeats previous experimental data and undercuts the evidence for currently held hypotheses. In the

---

[313] Anderson, "The Epistemology of Democracy," 14.

[314] At this point, it is important to note that within the literature on epistemic democracy, there has been questions as to what information really is gained from aggregative mechanisms such as majority voting. The classical view stemming from Condorcet and Rousseau has been to take democratic processes such as majority voting as giving us knowledge of the "general will" of the people. However, in the twentieth century, developments in social choice theory raised various concerns about whether aggregative mechanisms such as majority voting do in fact the general will of the people. In particular, theoretical results such as Arrow's impossibility theorem or the Condorcet paradox suggested that the aggregation of individual preferences could lead to a collective preference ordering which is either ambiguous or inconsistent respectively. As such, one might question the reliability of majority voting to give us the information that we need to find the information needed to make informed decisions about policy making. However, while these debates about the epistemic reliability of majority voting are important debates within epistemic democracy, they are to some extent orthogonal to the topic of this thesis. This is because Anderson's focus is not simply on aggregative mechanisms on majority voting, but on the epistemic functions of democracy as a whole. As such, as we have already seen, one of the constitutive features of democracy according to Anderson is the interaction between voting with public discussion, such that the voting decisions of individual epistemic agents are not independent, but influenced by public fora and discussion within political parties. However, for a helpful overview on the question of popular will in epistemic democracy, see MacGilvray, "Democratic Doubts", 106ff. MacGilvray also points to William H. Riker as raising the most recent major objection against the notion of popular will, and Joshua Cohen as providing the primary response to Riker in his notion of the "epistemic" interpretation of voting. See William H. Riker, *Liberalism Against Populism* (San Francisco: W. H. Freeman, 1982); Joshua Cohen, "An Epistemic Conception of Democracy," *Ethics* 97, no. 1 (1986): 26–38. For more on the recent developments on the epistemic interpretation of voting, see Knight and Johnson, *The Priority of Democracy*, 152ff.

same way, Dewey's experimentalist model of democracy does not see democratic processes such as majority voting and inclusive deliberation as one-time events within democratic decision-making, but as ongoing *processes* which are regularly enacted so that our current policies remain informed.

This brings us back to Dewey's explanation of the feedback mechanisms of democracy. For Dewey, the purpose of feedback mechanisms is to aid democratic processes such as majority voting and inclusive deliberation in their experimental function.[315] In particular, feedback mechanisms play an important part in initiating the iterative process of the experimental method, by serving as indicators for when an existing policy is ineffective and therefore in need of revision. Feedback mechanisms signal to democracies that a given policy needs to be revised by facilitating public complaint – for example, in the form of public opinion polling, protests, or petitions to government officials – which provide evidence that a given policy is not working or expected not to work. This in turn restarts democratic processes such as majority voting and deliberation within political parties, which gathers and utilises this new information about the consequences of previous policies so as to devise a better solution to a given social problem. As such, Dewey's model reveals how feedback mechanisms contribute to the diachronic aspect of democratic decision-making by "[institutionalising] fallibilism and an experimental attitude with respect to state policies". Through the facilitation of public complaint, the feedback mechanisms of democracy ensure that state policies are subject to the iterative process of the experimental method, a process which repeats democratic processes such as majority voting and active

---

[315] To be more precise, we can say that feedback mechanisms and democratic processes such as majority voting and inclusive deliberation are *both* playing experimental functions within democratic decision-making. For Anderson, the feedback mechanisms are themselves sites for further experimentation, such that in protests, petitions, and public opinion polling, we receive empirical "evidence" as to the ineffectiveness of certain policies. Once we acknowledge that certain policies are ineffective or predicted to be ineffective, we require further voting and deliberation to obtain more evidence as to what policies might be effective. As such, feedback mechanisms act both as a source of evidence for ineffective policies, and as an indicator that further voting and deliberation (i.e. further experimentation) is needed to find more effective policies.

deliberation within political parties in light of new information about previous policies to improve its epistemic practice.

This brings us to the next point: for Anderson and Dewey, the feedback mechanisms of democracy do not simply play a formal part within democratic procedure, but also a *social* role in institutionalising certain norms and values within democratic decision-making. The distinction I am making between formal and social in this case is this: the role of feedback mechanisms within democracy is not simply as a formal platform for public complaint, but a social institution which promotes a certain way of *responding* to public complaint. This is an important distinction because the role which the feedback mechanisms of democracy play within the experimental method of democratic procedure would not be possible if the feedback mechanisms were purely formal platforms for public complaint. In order for public complaint to have any bearing on democratic decision-making, feedback mechanisms cannot simply express the objections of the public, but must also incorporate these objections within the stages of democratic decision-making in order to ensure that these objections are addressed and taken seriously within present and future political deliberation. As such, as social institutions which are *sanctioned* within democratic procedure, the feedback mechanisms of democracy play a part in promoting certain ethical norms regarding how objections are treated within political deliberation.[316]

---

[316] At this point, it is also important to note that this distinction between formal and social institutions can be used to defend Anderson's exposition of Dewey's experimentalist model of democracy from a particular objection: namely, one might argue that Anderson provides an *end-state* account of institutions, such that Dewey's experimentalist model of democracy is just as much of an end-state account of democracy as the Condorcet Jury Theorem or the DTA Theorem. The reason why this objection fails, however, is because this distinction between formal and social institutions reveals that Anderson (and therefore Dewey) does not consider any democratic institutions to take any fixed structure or form. Instead, institutions such as the feedback mechanisms of democracy are continually changing and evaluated according to their social *function* in ensuring that certain norms and values are embodied within democratic decision-making. More on this below. For more

It is here where we see how Anderson's exposition of Dewey's experimentalist model of democracy coincides with Putnam's account of the democratic structure of inquiry. In both accounts of the pragmatist theory of inquiry, Putnam and Anderson highlight the role of ethical norms within social cooperative inquiry, and explain how the implementation of these ethical norms are necessary for epistemic reasons. On the topic of ethical norms within inquiry, Anderson is quick to point out that Dewey's conception of democracy sees democracy as more than a "matter of legal arrangements", but as a social and cultural movement which contains certain norms and values (i.e. that culture has to "change" in order to adopt "an experimental attitude towards social arrangements"). In particular, the norms which Anderson sees as governing political deliberation are remarkably similar (if not the same) as the norms of discourse ethics: Anderson notes that the discussion within political deliberation should allow each participant to introduce new assertions or objections to previous assertions (that is, that political deliberation should "welcome or at least tolerate diversity and dissent), and allow each participant to express any attitudes towards a given proposition without fear from internal or external coercion (or in Anderson's words, to"[give] all a respectful hearing [and to] institute deliberation and reason-giving, rather than threats and insults, as the basis of their communication with one another").

Furthermore, Anderson insists that these norms are not enforced for moral but epistemic reasons – that is, in order "to realize the epistemic powers of democracy". However, it is at this point in which Anderson's focus on social institutions allows Anderson to provide a different perspective on the role of the ethical norms within inquiry to that given by Putnam. This is because Anderson's exposition reveals how the norms of discourse ethics are not personal norms that are affirmed individually by

---

information on how pragmatists can be non-ideal theorists and still institutionalists, see Shane J. Ralston, "Can Pragmatists Be Institutionalists? John Dewey Joins the Non-Ideal/Ideal Theory Debate," *Human Studies* 33, no. 1 (May 2010): 65–84.

citizens within democracy, but norms which are structurally embedded within democratic institutions as to form a kind of social community.

This brings us to our final point: namely, Anderson's pragmatist exposition of democracy explains how democratic inquiry can use social institutions to channel disagreement in an epistemically productive way. To begin, Anderson explains the epistemic role which disagreement plays within democratic procedure in how disagreement is *institutionalised* in every stage of democratic decision-making.[317] In the section after the given passage, a section that is aptly entitled "The Epistemic Import of Dissent", Anderson claims that disagreement is a "central [feature] of democracy", such that disagreement is used within "all stages of decision-making: during deliberation, at the point of decision (voting), and after a decision has been made."[318] In each stage of democratic decision-making, disagreement is incorporated within political deliberation through institutions which promote public complaint and instate political opposition. However, rather than being an obstacle to inquiry, Anderson explains how the legitimising of political opposition is structurally designed to be beneficial for democratic decision-making: "Without an opposition to remind the public of continuing objections to collective decisions, and to pose alternatives, accountability of decision makers is impossible. Nothing would force decision makers to reconsider their decisions. Only with such

---

[317] The incorporating of disagreement within the procedure of inquiry is also affirmed and emphasised by other contemporary pragmatists. For example, Knight and Johnson agree with Anderson in stating that "democratic politics is perhaps best seen less as a way of reaching consensus or agreement than as an effective way of structuring the terms of such persistent disagreement." MacGilvray makes a similar point in emphasising the Peircean insight that inquiry is motivated by doubt rather than belief. For MacGilvray, this insight reveals how democratic decision-making welcomes rather than rejects disagreement, since it is particular doubts in particular situations which legitimates democratic procedure. See Knight and Johnson, *The Priority of Democracy*, 150–51; MacGilvray, "Democratic Doubts", 113-14, 120.

[318] Anderson, "The Epistemology of Democracy," 15.

continuing opposition can fallibilism and the institutional capacity for experimentation—revising

one's decisions on the basis of experience with their consequences—be realized."[319]

How do democratic institutions structure political opposition in a way which aids democratic

decision-making? Anderson addresses this question in the given quote by pointing to the epistemic

function of political opposition (and thus, disagreement) in 'realising' fallibilism and experimentation

within democratic procedure. Notice that this is the same reason which Anderson previously gave for

the epistemic function of feedback mechanisms within democracy: recall how in the above passage,

Anderson claims that feedback mechanisms "function to institutionalize fallibilism and an

experimental attitude with respect to state policies". It is these dual notions of fallibilism and

experimentation which for Anderson acts as normative principles which shape how epistemic agents

respond to disagreement within inquiry. However, given that Anderson does not elaborate on these

two notions in this paper, a brief explanation is needed as context about how Peirce's notion of

fallibilism is used within Dewey's theory of inquiry.[320] Peirce's original doctrine of fallibilism was

concerned with explaining the ever-present possibility that any of our beliefs are in error due to the

inherent fallibility of our epistemic practice. For Peirce, this possibility of error in our beliefs and

assertions is inevitable due to the multitude of ways in which our epistemic practice can be limited:

---

[319] Anderson, "The Epistemology of Democracy," 16–17.

[320] Margolis questions whether Peirce's principle of fallibilism has been consistently used throughout the pragmatist literature, stating that it is "developed in rather different ways by a number of other authors, often (and surprisingly) employing the very term"fallibilism" even where Peirce's account is either not mentioned or not pursued." Margolis then claims that Peirce's notion of fallibilism is often mistaken for Dewey's use of fallibilism within his theory of inquiry: "My impression is that most American discussants really favor Dewey's sense of self-corrective inquiry rather than Peirce's." To discuss what interpretation of Peirce's notion of fallibilism is correct, or to evaluate whether Dewey's notion of fallibilism is better or worse than Peirce's notion will be outside of the scope of this essay. As such, by "fallibilism", I will simply be focusing on Dewey's notion of fallibilism as it is used within his theory of inquiry. Margolis, "Peirce's Fallibilism," 535–36.

throughout his writings, Peirce forcefully argues for this point through examples of the limitations of

human cognitive ability, the limitations of the epistemic methods we employ (for example, the use of

approximations or human oversight in measurement), and even the inherent limitations within the

content of our epistemic practice (for example, the inherent defeasibility of inductive evidence).[321]

Given that these various limitations are an ineliminable part of our epistemic practice, Peirce therefore

concludes that any account (descriptive or normative) of our actual epistemic practice must begin by

acknowledging that our epistemic practice is fallible, such that there is an "inevitability of experimental

error" and an "unavoidable element of uncertainty" and doubt in all of our beliefs and assertions.[322]

Within Peirce's theory of inquiry, however, the inevitability of experimental error does not lead to

skepticism about the scientific experimental method – in fact, fallibilism provides the best justification

for the use of the scientific experimental method. This is because, for Peirce, the best way for us to

address the possibility of error in our beliefs is to use the experimental method to continually test our

beliefs for error by conducting further experimentation.[323] As such, beginning from Peirce, the

pragmatists saw fallibilism and the experimental method as interconnected features of the pragmatist

theory of inquiry: just as the experimental method serves as the solution to the problems caused by

fallibilism, so does the doctrine of fallibilism justify the use of the experimental method as the best

---

[321] Susan Haack, "Fallibilism and Necessity," *Synthese* 41, no. 1 (May 1979): 43.

[322] Haack, "Fallibilism and Necessity," 42–43.

[323] Peirce's connection of the fallibility of belief with the need for the experimental method is developed across these three papers: see Charles S. Peirce, "Some Consequences of Four Incapacities," in *The Pragmatism Reader: From Peirce Through the Present*, ed. Robert B. Talisse and Scott F. Aikin (Princeton; Oxford: Princeton University Press, 2011), 12–36; Charles S. Peirce, "How to Make Our Ideas Clear," in *The Pragmatism Reader: From Peirce Through the Present*, ed. Robert B. Talisse and Scott F. Aikin (Princeton; Oxford: Princeton University Press, 2011), 50–65; and Charles S. Peirce, "The Fixation of Belief," in *The Pragmatism Reader: From Peirce Through the Present*, ed. Robert B. Talisse and Scott F. Aikin (Princeton; Oxford: Princeton University Press, 2011), 37–49.

form of epistemic practice. Dewey agrees with Peirce in seeing the empiricism of the scientific experimental method as the best solution to the possible error within our fallible beliefs and assertions. However, what Dewey adds to this aspect of the pragmatist theory of inquiry is the claim that experimentation requires not only the testing of our beliefs according to our experience of natural phenomena, but also the testing of our beliefs *discursively within the community of inquirers*.[324] For Dewey, the experimental method of scientific inquiry does not only consist of the physical experimentation which tests our belief according to our experience of the natural world, but also a *social experimentation* which tests our beliefs according to the objections and alternative explanations which are raised by interlocutors within inquiry.

We can now see how the two themes of fallibilism and experimentation act both as justification for the epistemic importance of disagreement, and as normative principles for our response to disagreement. In terms of the pragmatist doctrine of fallibilism, the pragmatists explain how fallibilism reveals the epistemic importance of disagreement as a consistent reminder of the fallibility of our beliefs and assertions. This in turn teaches us to take disagreement seriously by treating disagreement as an opportunity to check whether our beliefs are in error, rather than simply dismissing disagreement as

---

[324] In their explication of Dewey's theory of inquiry, Fesmire notes the connection between fallibilism, the empiricist norm of exposing beliefs to experience, and the dialectic need for discussion within the community of inquirers in the following way: "For the classical pragmatists, an empiricist habit of exposing beliefs to experience as test for their viability was regarded as the only possible corrective for overactive bells that toll with subjective certainty. Dewey agreed with Peirce that the tried-and-true method for arriving at seaworthy assertions is that followed by a community of rigorous albeit fallible inquirers putting beliefs to the test in the world. Philosophy or science cut loose from answering to anyone or anything is anemic and irrelevant. Such intellectual insularity is a luxury of nonempirical philosophies." However, once again, it is important to note that the dialectical need for discussion within the community of inquirers is a feature that is unique to Dewey's theory of inquiry, and which is not clearly found within Peirce's original theory of inquiry. Fesmire, *Dewey*, 102. For more on how Dewey added a dialectical element to his theory of inquiry, see Margolis, "Peirce's Fallibilism", 537ff.

being the result of the mistakes of other epistemic agents, or as epistemically irrelevant and

insoluble.[325]

As a sidenote, it is important to note that while this doctrine of fallibilism might sound very similar to

conciliatory views such as the Equal Weight View, there are important epistemic differences between

the two normative views. Unlike conciliationism within the epistemology of disagreement, the

pragmatist doctrine of fallibilism does not outrightly claim that the best response to disagreement is

always to seek reconciliation with an epistemic peer, or that one always has to lower one's level of

confidence in their beliefs when encountering disagreement (e.g. so as to give equal weight to your

opinion and that of the dissenting peer). One can be a fallibilist and acknowledge the general

possibility of error within one's beliefs and assertions, and still come to the conclusion that in a given

epistemic context, one is in fact not in error.[326] As such, the pragmatist doctrine of fallibilism is not a

general norm towards disagreement, but a norm which is specifically designed to safeguard against the

problem of subjectivity within disagreement: given that our subjective attitudes can distort our

perception of an epistemic situation such that we are excessively certain of our own opinions, the

---

[325] This is precisely the conclusion which Cheryl Misak argues for about disagreements in political inquiry: rather than simply dismissing political disagreements as intractable, our intuition that there does exist some non-negotiable views about certain political issues necessitate that we take disagreement on this issue seriously, and that we attempt to resolve these disagreements in a manner that is reason-giving rather than coercive. See Cheryl Misak, "Making Disagreement Matter: Pragmatism and Deliberative Democracy," in *The Pragmatism Reader: From Peirce Through the Present*, ed. Robert B. Talisse and Scott F. Aikin (Princeton; Oxford: Princeton University Press, 2011), 471–83.

[326] Jonathan Matheson and Brandon Carey come to a similar conclusion in arguing that the Equal Weight View does not entail outright skepticism, but should be interpreted as a moderate view which acknowledges our actual lack of evidence about our first-order evidence in certain contexts. I would agree with Matheson and Carey's conclusion, while noting that the pragmatist doctrine of fallibilism provides a helpful epistemic norm/maxim which allows us to take this conclusion into consideration within our epistemic practice. See Jonathan Matheson and Brandon Carey, "How Skeptical Is the Equal Weight View?" in *Disagreement and Skepticism*, ed. Diego E. Machuca (New York; London: Routledge/Taylor; Francis Group, 2013), 131–49.

doctrine of fallibilism reminds us to be open to the criticism of other epistemic agents, and to consider ways in which we can evaluate our own beliefs (as well as those of the dissenting peer) in light of a given disagreement.

This brings us to the pragmatist emphasis on the experimental attitude of scientific inquiry. By analogising from the physical experimentation of scientific inquiry, the pragmatists explain how disagreement is epistemically important as a site of social experimentation by which we test our beliefs dialectically within a community of inquirers.[327] In treating disagreement as a site for social experimentation, the pragmatists highlight the need to adopt an experimental attitude which sees disagreement as a source of new empirical data rather than as an obstruction to inquiry. To compare with the analytic epistemology of disagreement, the initial set-up of the epistemic problem of disagreement implicitly teaches us to treat epistemic disagreement as an impasse between two (or more) epistemic agents which can only be resolved through some external arbiter such as evidence. By contrast, the pragmatists argue that the best response to epistemic disagreement is to treat it as a source of empirical data, data which can and should be evaluated and interpreted through cooperative reasoning and deliberation. For Anderson, it is these normative attitudes towards disagreement which are institutionalised within democracy, and which therefore allow democratic decision-making to use

---

[327] In fact, Dewey highlights that the experimental method works in much the same way in studying physical phenomena within scientific inquiry as it does in tackling social problems within political inquiry: "The conclusion that agreement of activities and their consequences is a test and a moving force in scientific advance is in harmony with the position that the ultimate end and test of all inquiry is the transformation of a problematic situation (which involves confusion and conflict) into a unified one. That it is much more difficult to accomplish this end in social inquiry [that is, inquiry into social situations] than in the restricted field of physical inquiry [that is, inquiry into physical phenomena] is a fact. But it is not a fact which constitutes an inherent logical or theoretical difference between the two kinds of inquiry. On the contrary, the presence of practical difficulties should operate, as within physical inquiry itself, as an intellectual stimulus and challenge to further application." Dewey, *Logic*, 490–91.

disagreement as a resource by which we improve our epistemic practice. In the concluding remarks to this analysis, Anderson explains how the main upshot of Dewey's experimentalist model of democracy is to see how disagreement (as it is processed within democratic reform) is closely tied with epistemic improvement within inquiry: "Epistemic improvement and democratic reform go hand-in-hand, just as Dewey's experimentalist model of democracy predicts. As naturalized epistemology is the application of scientific inquiry to improve inquiry itself, democratic reform is the application of experimental social epistemology to improve collective inquiry into the definition and solution of public problems."[328]

5.4   *The Pragmatist Response to Disagreement and the Empirical Basis of the Pragmatist Theory of Inquiry*

Given the arguments that we examined in this chapter, how then should we characterise the pragmatist response to disagreement? In this section, I will use the upshot of the respective arguments given by Kelly, Putnam and Anderson to provide a pragmatist response to our non-ideal epistemic problem of disagreement.

Let us begin with our critical examination of Kelly's analysis of belief polarisation. The reason why we began this chapter with Kelly's analysis of belief polarisation is because of the way in which belief polarisation highlighted the ineliminable role which our subjective attitudes play within epistemic practice. While an epistemic agent's subjective attitudes do not affect the norms of rational belief revision – that is, how one ought to rationally revise one's belief upon a body of evidence does not change depending on one's initial doxastic attitudes – Kelly reveals that an epistemic agent's subjective attitudes can nevertheless have a *causal* impact on the body of evidence one has. This is because our

---

[328] Anderson, "The Epistemology of Democracy," 21.

epistemic practice of revising our beliefs upon a body of evidence – the epistemic practice which must once again be flagged as the sole epistemic practice which evidentialists see as constituting epistemic rationality – is in fact situated within a larger process of inquiry wherein one must also decide which lines of inquiry are most promising and therefore worthy of pursuing. As such, given that these decisions require psychological mechanisms which predispose certain hypotheses over others, Kelly's analysis of belief polarisation reveals the ineliminable role which subjective attitudes play within epistemic practice.

This brings us back to Putnam's exposition on inquiry, where Putnam argues that the problems which arise from subjective attitudes within individual epistemic practice are best resolved by *intersubjective* discussion within cooperative epistemic practice. For Putnam, the best way of addressing the problems which subjective attitudes can create within epistemic practice is not to find an agent-free algorithm for belief revision (like the one proposed by Carnap), but to outline a normative account of how epistemic agents can correct each other's subjective blindspots through cooperative inquiry. For Putnam, this kind of normative account is given within Habermas and Apel's norms of discourse ethics, a series of ethical norms which govern public discussion to ensure that each participant within the discussion is able to introduce new assertions, to raise objections to previous assertions, and to express attitudes about previous assertions and objections without fear of internal and external coercion.

These discursive norms allow inquiry to address epistemic disagreement in two ways: first, the norms of discourse ethics create a social environment which best allows a community of inquirers to exchange and evaluate information. As a result, the free exchange of information within inquiry ensures that a community of inquirers have access to a common body of evidence, which therefore safeguards against the possibility that epistemic disagreements may occur on the basis of incomplete and skewed bodies of evidence from different individual epistemic agents (i.e. like those in belief

polarisation). Second, the norms of discourse ethics ensure that each participant is able to raise objections without any interference to the discussion from internal and external coercion – which is important given how coercion can often derail a discussion, if not outright prevent further discussion from continuing. This in turn ensures that disagreement is properly considered and addressed within public discussion, so that a disagreement is resolved if not settled in a way that is acceptable to all participants within a community of inquirers.

But how is disagreement properly considered and addressed over the course of discussion within inquiry? It is here where Anderson's epistemic analysis of Dewey's experimentalist account of democracy provides an illuminating answer. For Anderson, the example of democratic procedure reveals how disagreement must be institutionalised through various systems and mechanisms within inquiry, to ensure that disagreement is taken seriously within collective decision-making. In particular, Anderson uses the example of the feedback mechanisms of democracy to explain how social institutions play two important epistemic functions within the process of inquiry. First, the feedback mechanisms of democracy ensure that democratic procedure is *dynamic* in its decision-making: in facilitating disagreement in the form of public complaint, the feedback mechanisms of democracy serve as an indicator within democratic procedure for current policies which are shown or expected to be ineffective. This in turn allows democratic procedure to restart processes such as majority voting and inclusive deliberation in order to make better policies in light of this new information. As such, the feedback mechanisms of democracy allows democratic procedure to follow the iterative process of scientific experimentation, in instating a stage within democratic decision-making after political deliberation where policies are revisited in light of their consequences, and revised or changed accordingly.

Second, the feedback mechanisms of democracy not only serve as indicators for ineffective policies, but also as social institutions which hold decision-makers accountable for their decisions. Anderson

therefore reveals that the appropriate response to disagreement within social inquiry involves not only

formal procedures which continually examine disagreement through further social experimentation,

but social attitudes which shape the response by the community of inquirers to the consequences of

inquiry. By holding decision-makers accountable to their decisions, Anderson reveals how the

feedback mechanisms of democracy serve as an example of social institutions which promote a

fallibilistic and experimental attitude within inquiry: one which teaches us to see disagreement as a

sign of possible errors within our beliefs and assertions (following fallibilism), and as a reason for

conducting further social experimentation (following the experimental method).

As such, the best way to summarise the upshot of both Putnam and Anderson's account of inquiry is

this: if the evidentialist response to disagreement is summarised in the norm "follow your evidence",

then the norm which best captures the pragmatist response to disagreement is to *continue inquiring*

cooperatively.[329] This is because, as we have seen from Putnam and Anderson, the pragmatist response

to disagreement sees the best solution to epistemic disagreement as continual inquiry within a

community of inquirers: inquiry where new hypotheses, objections to these hypotheses, and responses

to these objections are continually raised; and inquiry which assesses these hypotheses and objections

through new and better forms of experimentation. It is this process of constant inquiry which Putnam

and Anderson are aiming to explicate in their respective epistemic accounts. As we have seen, for

---

[329] Cheryl Misak comes to much the same conclusion in her Peircean defence of a pragmatist moral epistemology, where she concludes her work like so: "It is fitting to return, at the close of this examination of pragmatism and morals, to the principle which Peirce thought should 'be inscribed upon every wall of the city of philosophy: Do not block the path of inquiry' (CP 1.135)." As a personal remark, I encountered Misak's work towards the very end of my research, and was shocked to find how similar her conclusion was, despite the notable differences between Peirce and Dewey's theory of inquiry. This suggests that the pragmatist response to disagreement should not only be seen as representative of Deweyan pragmatists, but pragmatism as a whole. Misak, *Truth, Politics, Morality*, 155. For more information about the differences between Peirce and Dewey's theory of inquiry, see fnn. 52, 320, and 324 of this thesis.

Putnam, the norms of discourse ethics are important precisely because they allow others to cooperate within public discussion, so as to continue inquiring. Similarly, for Anderson, the main purpose of social institutions of feedback and accountability is to dynamically drive collective decision-making so as to update social policy with further social experimentation. As such, for both Putnam and Anderson, the process of inquiry is seen to be the best solution to epistemic disagreement because it is the main method by which we can improve our epistemic practice so as to arrive at a position where we know what the rational response to a particular disagreement is.

For what reasons do the pragmatists believe and justify the claim that the process of inquiry can and will consistently improve our epistemic practice? For the pragmatists, the only reasons we can give to this question are answers which themselves have been derived from previous inquiries. This is why Hilary Putnam and Ruth-Anna Putnam both highlight the fact that *the pragmatist theory of inquiry is itself the product of previous inquiries*: "Logic as the theory of inquiry is itself the result of an inquiry. Thus what Dewey says about inquiry in general is true of inquiry into inquiry."[330] In other words, rather than being an *a priori* analysis of epistemic practice, the pragmatist theory of inquiry is formulated and continually developed as an empirical account about inquiry which draws from the successes of previous inquiries.

This is what we've seen in the epistemic accounts of the structural features of inquiry given by Putnam and Anderson: without the examples of science and democracy as empirical examples of successful social epistemic practice, it would not have been possible to identify the norms of discourse ethics and social institutions for feedback and accountability as crucial epistemic features in addressing epistemic disagreement within inquiry. However, by making careful empirical observations of inquiry as it is

---

[330] Putnam and Putnam, "Dewey's Logic," 199.

conducted within science and democracy, Putnam and Anderson have played a part in developing the

pragmatist theory of inquiry from a potentially vacuous norm (given that the norm to continue

inquiring does not say anything about *how* to inquire) to a substantive epistemic account of social

epistemic practice, and the structural features which inquiry need in order to exchange and evaluate

information in the most efficient manner. While the epistemic features which Putnam and Anderson

pick out in their respective accounts do not guarantee the epistemic success of an inquiry, they are

justified in their use because of their consistent success within previous inquiries.

As a quick sidenote, the development of the pragmatist theory of inquiry therefore parallels recent

developments within inductive inference: while previous attempts to theorise about inductive

inference were impeded because of the inability to provide an *a priori* justification or formal system for

the use of induction, these theoretical concerns are no longer an issue within current research into

inductive inference within Bayesian epistemology, machine learning and artificial intelligence, and

statistics.[331] There are two reasons for why the lack of a formal inductive system does not deter current

research into inductive inference: first, the sheer utility and reliability of inductive inference is such an

indispensable part of our epistemic practice – both within everyday decision-making, and in

---

[331] Gilbert Harman and Sanjeev Kulkarni make this point in their analysis of the traditional problem of induction from the
perspective of statistical machine learning theory. For Harman and Kulkarni, while the problem of the reliability of inductive
methods does not admit a formal or *a priori* proof as the solution, this problem has been fruitfully investigated within
empirical research: "that the problem of induction as we have described it—the problem of finding reliable inductive
methods—can be and is being fruitfully investigated in statistical learning theory". Harman and Kulkarni therefore conclude
that the best method of studying inductive inference is an empirical approach: one which "stud[ies] empirically how well
various inductive methods do in real life." Harman and Kulkarni, *Reliable Reasoning*, 20, 26. For more on the recent research
into inductive inference within Bayesian epistemology, see the citations in fn. 91. For a classic paper on inductive inference
within statistics, which also argues for piece-meal research such that one investigates inductive inference within statistics on a
case-by-case basis, see Deborah G. Mayo and D. R. Cox, "Frequentist Statistics as a Theory of Inductive Inference," in
Institute of Mathematical Statistics Lecture Notes - Monograph Series, ed. Javier Rojo, vol. 49 (Institute of Mathematical
Statistics, 2006), 77–97.

specialised research – that the use of inductive inference is seen as justifiable despite the possibility that one draws unsound inferences from induction in practice. We discussed this briefly when we discussed how traditional sceptical arguments fail to bear any practical relevance in §1.2.2. Second, actual research into inductive inference is not deterred by the inability to find an *a priori* account of inductive inference, as different forms of inductive inference are continually developed and refined through *empirical inquiry*: that is, through the active use and testing of certain methods of inductive inferences, and through the assessment of the consequences of these methods (be it success or failure) in order to find better forms of inductive inference.[332]

As such, it is this empirical basis which allows the pragmatist theory of inquiry to provide a *transition* epistemic account of disagreement. As we discussed in the previous chapter, the reason why the evidentialist response to disagreement fails to resolve the non-ideal epistemic problem of disagreement is because the evidentialist norm to "follow your evidence" can only be satisfied when one is in a position to know what one's evidence is, or what doxastic justification one receives from one's evidence. This condition on the satisfiability of the evidentialist norm therefore highlights how evidentialism is an *end-state* account of epistemic rationality: one which only applies when one is in a position to attain rational belief as an epistemic end. By contrast, the pragmatist theory of inquiry is centrally concerned with the question of how we can use cooperative discussion and experimentation in order to arrive at a position where we do have sufficient evidence and understanding to know what the rational response to a particular disagreement is.

---

[332] As we have seen earlier, one example of this within machine learning is the problem of algorithmic bias: even if there is no formal or algorithmic solution to the problem of algorithmic bias, the problem of algorithmic bias is effectively addressed through oversight by epistemic agents who make evaluations about biases within a machine learning model on a case-by-case basis. See fn. 273 for more details.

As a sidenote, the pragmatist theory of inquiry highlights how cooperative social inquiry provides the underlying conditions which allow us to satisfy the end-state norms of rationality in the first place. More specifically, we can see that once a particular line of inquiry has conclusively settled a particular question or problem, the solution of that particular problem becomes an end-state epistemic norm: in other words, once a particular scientific or engineering problem has been resolved (e.g. how one should find the chemical composition to a substance, or how one should test the durability of some construction material), then that problem is one which we as an epistemic community can and do resolve through an end-state norm (e.g. a norm about what test one should conduct, or what technological equipment one should use).

However, where a disagreement places us in a situation where we do not have sufficient evidence or understanding about what rational belief entails, the pragmatist response to disagreement states that one can only resort to further inquiry to resolve disagreement, because only in collective social inquiry do we find the means by which we can improve our individual and collective practice, so as to arrive at a position to settle a particular disagreement. In this way, the pragmatist response to disagreement allows us to provide a *transition* account of our epistemic practice, one which explains how social inquiry can allow us to apply the lessons we've learnt from previous inquiries to settle new disagreements which we encounter in our everyday life, and then to learn from our inquiries these new disagreements to find better ways of responding to other disagreements in the future.

Two final remarks about the pragmatist response to disagreement. First, now that we see how collective social inquiry functions as the source of the evidence and epistemic norms which we use and have available in our social-epistemic context, we can also see how the various norms, concepts, and methods which are discussed within the analytic epistemology of disagreement connects to the pragmatist theory of inquiry. To re-iterate, it is important to note that the main critique which we raised towards the evidentialist response to disagreement in Chapter 4 was not against the norm to

"follow your evidence" as a valid epistemic norm, but towards the reductionist view that saw all of our epistemic practice as reducible to this norm. It is this reductionist project which led strong evidentialists such as Richard Feldman to distinguish between methodological rationality and current-state rationality, and then to argue that current-state rationality is the only kind of reasoning which determines whether our beliefs are doxastically justified. Once we reject this reductionist claim, however, we can see that it is possible to re-contextualise the norms, concepts, and methods within the analytic epistemology of disagreement within the wider epistemic context of inquiry. More specifically, we can see that the various norms we have examined within the analytic epistemology of disagreement – norms about remaining steadfast in our belief, seeking conciliation, taking into consideration the epistemic peerhood of a dissenting peer, and making sure that one is responding according to one's *total* body of evidence – are norms which can still be used in certain situations to assess what the rational response is to a given disagreement. However, given that these norms can always be interpreted or applied in an incorrect way, the pragmatist norm to continue inquiring acts as the fundamental norm for epistemic practice: one which is applicable in every epistemic context, and which allows us to re-assess our beliefs within a community of inquirers in order to improve our epistemic practice.

Second, it is important to clarify that while the pragmatist response to disagreement claims that social inquiry can lead an epistemic community to find the rational response to a given disagreement, this does not mean that the rational response to disagreement always require dissenting parties to come to the *same conclusion* about a given proposition. To explain, recall how in §2.3.1, we discussed how the kind of *epistemic* disagreement which are focused on in this thesis are disagreement over matters of fact. The pragmatist response to disagreement provides a comprehensive answer to this particular problem of *epistemic* disagreement, insofar as it provides an account of how we as a community of inquirers can settle the *facts* about a given topic through discussion and experimentation. Nevertheless,

it is important to see how disagreements can also be *non-epistemic*, in the sense that a disagreement is not always over a matter of fact, but can also be about the values that different dissenting parties hold, or different ways in which people can conceptualise a situation. In these disagreements, it is sometimes the case that the best that social inquiry can do is to reveal how the rational response is to "agree to disagree": that is, to show that a given disagreement might be based on the fact that the dissenting parties are simply different in their biological or psychological makeup, in their social or cultural identity, or otherwise. This means that while the pragmatist response to disagreement does still allow us to rationally resolve a disagreement (insofar as the dissenting parties can all agree that the facts show that the cause of the disagreement is in the respective differences between the dissenting parties), this does not mean that dissenting parties have to come to the same conclusion (insofar as a given disagreement is one in which dissenting parties cannot or choose not to come to the same conclusion). Agreeing on what is rational does not require agreeing to be the same person.

That being said, it is important not to overstate this point about non-epistemic disagreements, given that the pragmatists reject a strong fact-value distinction which would claim that there are no facts about matters of value, or that there are no value judgments in matters of fact. In other words, the pragmatists claim that the question of whether a particular disagreement is an epistemic disagreement or not is one where the answer must itself be discovered in *inquiry*: rather than broadly categorising certain topics (such as the traditionally "value-laden" areas of ethics, aesthetics, and politics) as being *non-epistemic*, the pragmatists argue that our very understanding of factivity or what the epistemic is is something which can be revised in inquiry, and therefore that any topic is one which we can at least *attempt* to collectively inquire into. For more on the pragmatist argument against a strong fact/value distinction, and how that opens up the possibility of inquiring into ethical and political disagreement, see Appendix A of this thesis.

5.5   Conclusion: On the Task of Creative Democracy

As such, the pragmatist response to disagreement therefore reveals how the only thing that limits what

we can inquire to as an epistemic community is *us*. This is the key idea in Dewey's most powerful

exhortation of his view of democracy – namely, in his 1938 address entitled "Creative Democracy –

The Task before Us". In it, Dewey explains how he has "been accused more than once and from

opposed quarters of an undue, a *Utopian*, faith in the possibilities of intelligence and in education as a

correlate of intelligence."[333] Yet, Dewey continues by explaining that he

> did not invent this faith [but] acquired it from [his] surroundings as far as those surroundings are
> animated by the democratic spirit. For what is the faith of democracy in the role of consultation, of
> conference, of persuasion, of discussion, in formation of public opinion, which in the long run is self-
> corrective, except faith in the capacity of the intelligence of the common man to respond with
> commonsense to the free play of facts and ideas which are secured by effective guarantees of free
> inquiry, free assembly, and free communication? [334]

It is for this reason which Dewey states that his view of democracy is not only a utopian, but "creative"

one: one which sees the pragmatist notion of democratic inquiry as something which is not realistic

only insofar as it has not yet been created anew by real epistemic agents (Recall Estlund and Dewey's

notion of the aspirational function of utopian theory in §1.2.1). Or, in Dewey's words, "to get rid of the

habit of thinking of democracy as something institutional and external…is to realize that democracy is

a moral ideal and so far as it becomes a fact is a moral fact. It is to realize that democracy is a reality

only as it is indeed a commonplace of living."[335] The "task" of creative democracy which Dewey sets

before us in this address is therefore one that is at once both epistemic and ethical: epistemic, because

---

[333] John Dewey, "Creative Democracy – the Task Before Us," in *The Pragmatism Reader: From Peirce Through the Present*, ed.
Robert B. Talisse and Scott F. Aikin (Princeton; Oxford: Princeton University Press, 2011), 152. Italics is my own.

[334] Dewey, "Creative Democracy – the Task Before Us," 152.

[335] Dewey, "Creative Democracy – the Task Before Us," 153.

it requires us to improve our epistemic practice through "inventive effort and creative activity"; and

ethical, because it requires us "to take as far as possible every conflict which arises…[and] to cooperate

by giving differences a chance to show themselves because of the belief that the expression of

difference is not only a right of the other persons but is a means of enriching one's own life-

experience."[336]

The aim of this thesis has been to consider the epistemic problem of disagreement from this Deweyan

perspective. In particular, in this thesis, I have attempted to provide the starting point for how further

research into the non-ideal epistemology of disagreement, or non-ideal epistemology more generally,

can be done. I have argued that the project of non-ideal epistemology differs from the typical 'ideal'

epistemology done within contemporary analytic epistemology, because of its focus on the epistemic

practice of epistemic agents, rather than on the epistemic *status* of their beliefs. I have elaborated on

this difference through the notion of epistemic context or position, by showing how ideal epistemology

(such as that done within Bayesian evidentialism) focuses on how one's belief can attain the status of

justified or rational belief in certain *contexts*, while non-ideal epistemology (such as the pragmatist

theory of inquiry) focuses on how one can arrive at a *better* position through social inquiry. And I have

shown that a truly non-ideal account of our epistemic practice must extend beyond simply

epistemology, because the efficacy of our social epistemic practice depends on, at the very least, a series

of ethical and political factors which ensure that information is exchanged and evaluated in the most

efficient way.

In fact, on this last point, it is clear that our exposition of Putnam and Anderson's arguments has only

scratched the surface when it comes to explaining the factors which play into how we address

_____

[336] Dewey, "Creative Democracy – the Task Before Us", 151, 153.

disagreement in social epistemic practice, since the complexity of our social epistemic practice cannot be captured simply by accounts about the norms of discourse ethics or mechanisms for formal complaint. Instead, our social epistemic practice can be influenced by a whole host of other factors – factors such as the role of social media algorithms and search engines in facilitating the information we have available to us, the role of economic factors in funding political campaigns or news companies, and the role of rhetoric (and its psychological, or even aesthetic(?), effects) in influencing our perception of the reliability of testimony.

As such, once we see how the pragmatist theory of inquiry incorporates ideas from ethics and political philosophy within its explanation of effective epistemic practice, the pragmatist theory of inquiry naturally opens up new avenues of research by revealing how it is possible to further our understanding of effective social inquiry by considering still other areas of philosophy, or areas *outside* of philosophy. To take the example of the literature on trust and trustworthiness, I mentioned in §1.2 that the recent work done on the non-ideal epistemology of testimony has been heavily influenced by feminism, and how feminist philosophers have contributed to our understanding of the use of testimony in actual epistemic practice – not only in terms of the ways in which we can fail to receive doxastic justification from certain sources of testimony due to epistemic prejudice, but also in the ways in which we can *improve* our reception and transmission of testimony.[337] As such, just as how the pragmatists use empirical observations about our conduct of science and democracy to explain effective social epistemic practice, feminists have also used empirical observations from gendered experience in order to explain the barriers to cooperative epistemic practice, and the ways to address these barriers.

---

[337] For more information, see fnn. 18, and 37 of this thesis.

In this way, the non-ideal epistemic account of testimony which arises from the recent feminist literature reinforces the pragmatist claim that the means to better epistemic practice does not only come from certain fields of specialised research, but can come from the experience and epistemic practice of *any* individual epistemic agent, or any epistemic community.[338] It is for this reason which the pragmatist response invites us to listen and seriously test what other dissenting parties are saying, because what we might learn from inquiry with dissenting parties may not be what we expected, but something even better.

---

[338] To give another example, in his historical account of the origins of probability, Ian Hacking explains how the origins of probability theory surprisingly came from the study of probable statements (or *opinio*) in the "low" sciences of astrology, alchemy, and medicine – studies which were (rightfully) scorned by the epistemically revered "high" sciences of optics, astronomy, and mechanics. As such, the development of probability theory is an extremely surprising development, since it would have been impossible for pre-Renaissance scientists to predict that the probability – that which was the study of alchemists and astrologists – would become a central aspect in practically every field of science and engineering. For more information, see Chapters 4 and 5 of Hacking, *The Emergence of Probability*.

B<small>IBLIOGRAPHY</small>

Adler, Mortimer J., and Charles Van Doren. *How to Read a Book*. New York: MJF Books, 1972.

Alston, William P. "Perceptual Knowledge." In *The Blackwell Guide to Epistemology*, edited by John Greco and Ernest Sosa, 221–42. Oxford, UK: Blackwell, 1999.

———. "The Deontological Conception of Epistemic Justification." *Philosophical Perspectives* 2 (1988): 257–99.

Anderson, Elizabeth. "The Epistemology of Democracy." *Episteme* 3, no. 1-2 (2006): 8–22.

Apel, Karl-Otto. *C. S. Peirce: From Pragmatism to Pragmaticism*. Amherst: University of Massachusetts Press, 1981.

———. *Diskurs Und Verantwortung: Das Problem Des Übergangs Zur Postkoventionellen Moral*. Frankfurt am Main: Suhrkamp, 1985.

Audi, Robert. "Theoretical Rationality: Its Sources, Structures, and Scope." In *The Oxford Handbook of Rationality*, edited by Alfred R. Mele and Piers Rawling, 17–44. Oxford: Oxford University Press, 2004.

Baier, Annette C. "Trust and Antitrust." *Ethics* 96, no. 2 (1986): 231–60.

Bajaj, Sameer. "Review of Democratic Reason: Pollitics, Collective Intelligence, and the Rule of the Many, by hélène Landemore." *Ethics* 124, no. 2 (2014): 426–31.

Baker, Judith. "Trust and Rationality." *Pacific Philosophical Quarterly* 68, no. 1 (1987): 1–13.

Ballantyne, Nathan, and E. J. Coffman. "Uniqueness, Evidence, and Rationality." *Philosophers' Imprint* 11, no. 18 (2011): 1–13.

Belot, Gordon. "Bayesian Orgulity." *Philosophy of Science* 80, no. 4 (2013): 483–503.

Berend, Daniel, and Jacob Paroush. "When Is Condorcet's Jury Theorem Valid?" *Social Choice and Welfare* 15, no. 4 (1998): 418–88.

"Bias, n., Adj., And Adv." In OED Online. Oxford University Press, December 2022. https://www.oed.com/view/Entry/18564. (accessed December 26, 2022.)

Black, Duncan. *The Theory of Committees and Elections*. Second. Cambridge: Cambridge University Press, 1963.

Blackburn, Simon. "Correspondence Theory of Truth." In *The Oxford Dictionary of Philosophy*, Second., 81. Oxford: Oxford University Press, 2005.

Blackwell, David, and Lester Dubins. "Merging of Opinions with Increasing Information." *The Annals of Mathematical Statistics* 33, no. 3 (1962): 882–86.

Bolstad, William M., and James M. Curran. *Introduction to Bayesian Statistics*. Third. Wiley, 2016.

Bolton, Martha. "Primary and Secondary Qualities in Early Modern Philosophy." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2022. Metaphysics Research Lab, Stanford University, 2022.

BonJour, Laurence. *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press, 1985.

Bradley, Seamus. "A Counterexample to Three Imprecise Decision Theories." *Theoria* 85 (2019): 18–30.

———. "Constraints on Rational Theory Choice." *British Journal for the Philosophy of Science* 68, no. 3 (2017): 639–61.

———. "Imprecise Probabilities." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2019. Metaphysics Research Lab, Stanford University, 2019.

Buchak, Lara. "Normative Theories of Rational Choice: Rivals to Expected Utility." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2022. Metaphysics Research Lab, Stanford University, 2022.

Carnap, Rudolf. "Carnap's Intellectual Autobiography." In *The Philosophy of Rudolf Carnap*, edited by Paul Arthur Schilpp, 3–84. LaSalle, IL: Open Court, 1963.

———. *Logical Foundations of Probability*. Chicago: University of Chicago Press, 1950.

———. *The Continuum of Inductive Methods*. Chicago: University of Chicago Press, 1952.

———. "The Two Concepts of Probability: The Problem of Probability." *Philosophy and Phenomenological Research* 5, no. 4 (1945): 513–32.

———. *The Unity of Science*. London: Kegan Paul, Trench, Hubner, 1934.

Carr, Jennifer Rose. "Why Ideal Epistemology?" *Mind* 131, no. 524 (2021): 1131–62.

Castelvecchi, Davide. "Mathematical Proof That Rocked Number Theory Will Be Published." *Nature* 580 (2020): 177.

Chang, Hasok. "Compositionism as a Dominant Way of Knowing in Modern Chemistry." *History of Science* 49, no. 3 (2011): 247–68.

———. "The Hidden History of Phlogiston: How Philosophical Failure Can Generate Historiographical Refinement." *Hyle: An International Journal for the Philosophy of Chemistry* 16, no. 2 (2010): 47–79.

Christensen, David. "Disagreement as Evidence: The Epistemology of Controversy." *Philosophy Compass* 4, no. 5 (2009): 756–67.

———. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116, no. 2 (2007): 187–217.

Christensen, David, and Jennifer Lackey. "Introduction." In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey, 1–3. Oxford: Oxford University Press, 2013.

Churchland, Paul M. "Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor." *Philosophy of Science* 55, no. 2 (1988): 167–87.

Coady, Cecil AJ. "Testimony and Observation." *American Philosophical Quarterly* 10, no. 2 (1973): 149–55.

———. *Testimony: A Philosophical Study*. Oxford: Clarendon Press, 1992.

Cohen, G. A. "Facts and Principles." *Philosophy & Public Affairs* 31, no. 3 (2003): 211–45.

———. *Rescuing Justice and Equality*. Cambridge, Massachusetts: Harvard University Press, 2008.

Cohen, Joshua. "An Epistemic Conception of Democracy." *Ethics* 97, no. 1 (1986): 26–38.

Cohen, L. Jonathan. *The Probable and the Provable*. Oxford: Clarendon Press, 1977.

Cohen, Stewart. "A Defense of the (Almost) Equal Weight View." In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey, 98–117. Oxford: Oxford University Press, 2013.

Collin, Finn. "The Twin Roots and Branches of Social Epistemology." In *The Routledge Handbook of Social Epistemology*, edited by Miranda Fricker, Peter J. Graham, David Henderson, and Nikolaj J. L. L. Pedersen, 21–30. New York; London: Routledge/Taylor; Francis Group, 2020.

Conee, Earl, and Richard Feldman. "Evidence." In *Epistemology: New Essays*, edited by Quentin Smith. Oxford: Oxford University Press, 2008.

———. "Evidentialism." *Philosophical Studies* 48, no. 1 (1985): 15–34.

———. *Evidentialism*. Oxford: Oxford University Press, 2004.

———. "Internalism Defended." In *Evidentialism*, 53–82. Oxford: Oxford University Press, 2004.

Cross, Richard. "Testimony, Error, and Reasonable Belief in Medieval Religious Epistemology." In *Knowledge, Belief, and God: New Insights in Religious Epistemology*, edited by Matthew A. Benton, John Hawthorne, and Dani Rabinowitz, 29–53. Oxford: Oxford University Press, 2018.

Dancy, Jonathan. *An Introduction to Contemporary Epistemology*. Oxford: Blackwell, 1985.

Davidson, Donald. *Essays on Actions and Events*. Oxford: Clarendon Press, 1960.

De Finetti, Bruno. "Foresight: Its Logical Laws, Its Subjective Sources." In *Breakthroughs in Statistics*, 134–74. Springer, 1992.

———. *Theory of Probability: A Critical Introductory Treatment*. Chichester: John Wiley & Sons Ltd, 2017.

Dempster, A. P. "Upper and Lower Probabilities Induced by a Multivalued Mapping." *The Annals of Mathematical Statistics* 38, no. 2 (April 1967): 325–39.

Dewey, John. "Creative Democracy – the Task Before Us." In *The Pragmatism Reader: From Peirce Through the Present*, edited by Robert B. Talisse and Scott F. Aikin, 150–54. Princeton; Oxford: Princeton University Press, 2011.

———. *How We Think*. Boston, New York, Chicago: D. C. Heath & Co., 1910.

———. *Logic: The Theory of Inquiry*. New York: Henry Holt; Company, 1938.

———. "Nature, Communication and Meaning." In *Experience and Nature*. Mineola, N.Y.: Dover, 1958.

———. *Reconstruction in Philosophy*. New York: Henry Holt, 1920.

———. *The Public and Its Problems: An Essay in Political Inquiry*. Edited by Melvin Rogers. Athens, Ohio: Swallow Press, 2016.

———. *The Question for Certainty: A Study of the Relation of Knowledge and Action*. Gifford Lectures. London: G.Allen & Unwin, 1929.

Dietrich, Franz, and Kai Spiekermann. "Jury Theorems." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2022. Metaphysics Research Lab, Stanford University, 2022.

Dougherty, Trent. "Introduction." In *Evidentialism and Its Discontents*, edited by Trent Dougherty, 1–14. New York: Oxford University Press, 2011.

———. "The Ethics of Belief Is Ethics (Period): Reassigning Responsibilism." In *The Ethics of Belief: Individual and Social*, edited by Jonathan Matheson and Rico Vitz, 146–68. Oxford: Oxford University Press, 2014.

Earman, John. *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press, 1992.

Elgin, Catherine Z. "Persistent Disagreement." In *Disagreement*, edited by Richard Feldman and Ted A. Warfield, 53–68. New York: Oxford University Press, 2010.

Endriss, Ulle. "Judgment Aggregation." In *Handbook of Computational Social Choice*, edited by Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, 399–426. Cambridge University Press, 2016.

Enoch, David. "Not Just a Truthometer: Taking Oneself Seriously (but Not Too Seriously) in Cases of Peer Disagreement." *Mind* 119 (2010): 953–97.

Estlund, David. "The Epistemic Dimension of Democratic Authority." *The Modern Schoolman* 74, no. 4 (1997): 259–76.

Estlund, David M. *Democratic Authority: A Philosophical Framework*. Princeton; Oxford: Princeton University Press, 2008.

———. "The Irrelevance of the Jury Theorem." In *Democratic Authority: A Philosophical Framework*, 223–36. Princeton; Oxford: Princeton University Press, 2008.

Fantl, Jeremy, and Matthew McGrath. *Knowledge in an Uncertain World*. Oxford, England: Oxford University Press, 2009.

Farrington, Benjamin. *The Philosophy of Francis Bacon*. Liverpool: Liverpool University Press, 1964.

Feldman, Richard. "Epistemic Obligations." *Philosophical Perspectives* 2 (1988): 235–56.

———. "Epistemological Duties." In *The Oxford Handbook of Epistemology*, edited by Paul K. Moser, 362–83. Oxford: Oxford University Press, 2002.

———. "Epistemological Puzzles about Disagreement." In *Epistemology Futures*, edited by Stephen Hetherington, 216–36. Oxford: Oxford University Press, 2006.

———. *Epistemology*. Upper Saddle River, NJ: Prentice Hall, 2003.

———. "Evidence of Evidence Is Evidence." In *The Ethics of Belief: Individual and Social*, edited by Jonathan Matheson and Rico Vitz, 284–300. Oxford: Oxford University Press, 2014.

———. "Having Evidence." In *Evidentialism*, 219–41. Oxford: Oxford University Press, 2004.

———. "Reasonable Religious Disagreements." In *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, edited by Louise Antony, 194–214. Oxford University Press, 2006.

Fesmire, Steven. *Dewey*. Edited by Brian Leiter. Routledge Philosophers. London; New York: Routledge, 2015.

Festenstein, Matthew. "Ideal and Actual in Dewey's Political Theory." In *Pragmatism and Justice*, edited by Susan Dieleman, David Rondel, and Christopher Voparil, 97–114. New York, NY: Oxford University Press, 2017.

———. "Inquiry as Critique: On the Legacy of Deweyan Pragmatism for Political Theory." *Political Studies* 49 (2001): 730–48.

Feyerabend, Paul Karl. "An Attempt at a Realistic Interpretation of Experience." In *Realism, Rationalism, and Scientific Method*, edited by Paul Karl Feyerabend. Cambridge: Cambridge University Press, 1985.

Fodor, Jerry A. "A Reply to Churchland's "Perceptual Plasticity and Theoretical Neutrality"." *Philosophy of Science* 55, no. 2 (1988): 188–98.

———. "Observation Reconsidered." *Philosophy of Science* 51, no. 1 (1984): 23–43.

Foley, Richard. "Egoism in Epistemology." In *Socializing Epistemology*, edited by Frederick F. Schmitt. Lanham, MD: Rowman; Littlefield, 1994.

———. *Intellectual Trust in Oneself and Others*. Cambridge: Cambridge University Press, 2001.

———. "Self-Trust." In *The Routledge Handbook of Trust and Philosophy*, edited by Judith Simon, First., 231–42. New York; London: Routledge/Taylor; Francis Group, 2020.

Frances, Bryan. *Disagreement*. Cambridge, UK: Polity Press, 2014.

Frances, Bryan, and Jonathan Matheson. "Disagreement." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2019. Metaphysics Research Lab, Stanford University, 2019.

Freeman, Samuel. *Rawls*. Edited by Brian Leiter. Routledge Philosophers. London; New York: Routledge, 2007.

Fricker, Miranda. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press, 2007.

Friedman, Jane. "The Epistemic and the Zetetic." *Philosophical Review* 129, no. 4 (2020): 501–36.

Friedman, Michael. "Explanation and Scientific Understanding." *Journal of Philosophy* 71, no. 1 (1974): 5–19.

Frost-Arnold, Karen. "Social Media, Trust, and the Epistemology of Prejudice." *Social Epistemology* 30, no. 5-6 (2016): 513–31.

———. "Trustworthiness and Truth: The Epistemic Pitfalls of Internet Accountability." *Episteme* 11, no. 1 (2014): 63–81.

Fuller, Steve. *Social Epistemology*. 2nd ed. Bloomington: Indiana University Press, 2002.

Gaifman, Haim, and Marc Snir. "Probabilities over Rich Languages, Testing and Randomness." *Journal of Symbolic Logic* 47, no. 3 (1982): 495–548.

Gigerenzer, Gerd, and Wolfgang Gaissmaier. "Heuristic Decision Making." *Annual Review of Psychology* 62, no. 1 (2011): 451–82.

Gigerenzer, Gerd, and Daniel G. Goldstein. "Reasoning the Fast and Frugal Way: Models of Bounded Rationality." *Psychological Review* 103, no. 4 (1996): 650–69.

Gigerenzer, Gerd, and Reinhard Selten, eds. *Bounded Rationality: The Adaptive Toolbox*. Cambridge, Massachusetts; London, England: MIT Press, 2002.

Goldman, Alvin I. "A Guide to Social Epistemology." In *Social Epistemology: Essential Readings*, edited by Alvin I. Goldman and Dennis Whitcomb, 11–37. Oxford University Press, 2011.

———. *Epistemology and Cognition*. Cambridge: Harvard University Press, 1986.

———. "Experts: Which Ones Should You Trust?" In *Social Epistemology: Essential Readings*, edited by Alvin I. Goldman and Dennis Whitcomb, 109–33. Oxford University Press, 2011.

———. "The Internalist Conception of Justification." *Midwest Studies in Philosophy* 5, no. 1 (1980): 27–52.

———. "The What, Why, and How of Social Epistemology." In *The Routledge Handbook of Social Epistemology*, edited by Miranda Fricker, Peter J. Graham, David Henderson, and Nikolaj J. L. L. Pedersen, 10–20. New York; London: Routledge/Taylor; Francis Group, 2020.

Goldman, Alvin I., and Cailin O'Connor. "Social Epistemology." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2021. Metaphysics Research Lab, Stanford University, 2021.

Greco, Daniel. "Significance Testing in Theory and Practice." *The British Journal for the Philosophy of Science* 62, no. 3 (2011): 607–37.

Grim, Patrick, Daniel J. Singer, Aaron Bramson, Bennett Holman, Sean McGeehan, and William J. Berger. "Diversity, Ability, and Expertise in Epistemic Communities." *Philosophy of Science* 86, no. 1 (2018): 98–123.

Grossi, Davide, and Gabriella Pigozzi. *Judgment Aggregation: A Primer*. San Rafael, CA: Morgan & Claypool, 2014.

Haack, Susan. "Fallibilism and Necessity." *Synthese* 41, no. 1 (May 1979): 37–63.

Habermas, Jurgen. "Discourse Ethics: Notes on a Program of Philosophical Justification." In *Moral Consciousness and Communicative Action*, translated by Christian Lenhart and Shierry Weber Nicholson, 43–115. Cambridge: MIT Press, 1990.

———. *Justification and Application: Remarks on Discourse Ethics*. Translated by Ciaran Cronin. Cambridge, Massachusetts; London, England: MIT Press, 1993.

———. *Moral Consciousness and Communicative Action*. Translated by Christian Lenhart and Shierry Weber Nicholson. Cambridge: MIT Press, 1990.

———. *The Theory of Communicative Action*. 2 vols. Boston: Beacon Press, 1985/1989.

Hacking, Ian. *The Emergence of Probability: A Philosophical Study of Early Ideas about Probability, Induction and Statistical Inference*. 2nd ed. New York: Cambridge University Press, 2006.

Hamlin, Alan, and Zofia Stemplowska. "Theory, Ideal Theory, and the Theory of Ideals." *Political Studies Review* 10, no. 1 (2012): 48–62.

Hanson, Norwood Russell. *Patterns of Discovery*. Cambridge: Cambridge University Press, 1958.

Hansson, Sven Ove, and Till Grüne-Yanoff. "Preferences." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2022. Metaphysics Research Lab, Stanford University, 2022.

Hardwig, John. "Epistemic Dependence." *Journal of Philosophy* 82, no. 7 (1985): 335–49.

———. "The Role of Trust in Knowledge." *The Journal of Philosophy* 88, no. 12 (1991): 693–708.

Harman, Gilbert. *Change in View: Principles of Reasoning*. Cambridge, Massachusetts; London, England: The MIT Press, 1986.

———. "Practical Aspects of Theoretical Reasoning." In *The Oxford Handbook of Rationality*, edited by Alfred R. Mele and Piers Rawling, 45–56. Oxford: Oxford University Press, 2004.

———. *Thought*. Princeton University Press, 1973.

Harman, Gilbert, and Sanjeev Kulkarni. *Reliable Reasoning: Induction and Statistical Learning Theory*. Edited by François Recanati. The Jean Nicod Lectures. Cambridge, Massachusetts; London, England: MIT Press, 2007.

Hatfield, Gary. "Sense Data." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2021. Metaphysics Research Lab, Stanford University, 2021.

Hawthorne, John, and Amia Srinivasan. "Disagreement Without Transparence: Some Bleak Thoughts." In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey, 9–30. Oxford: Oxford University Press, 2013.

Hayek, F. A. "Economics and Knowledge." In *L. S. E. Essays on Cost*, edited by James M. Buchanan and G. F. Thirlby. New York: New York University Press, 1981.

———. "The Use of Knowledge in Society." *The American Economic Review* 35, no. 4 (1945): 519–30.

Hedden, Brian. "A Defense of Objectivism about Evidential Support." *Canadian Journal of Philosophy* 45, no. 5 (2015): 716–43.

———. "Options and Diachronic Tragedy." *Philosophy and Phenomenological Research* 90, no. 2 (2015): 423–51.

———. *Reasons Without Persons: Rationality, Identity, and Time*. Oxford University Press UK, 2015.

———. "Time-Slice Rationality." *Mind* 124 (2015): 449–91.

Henderson, David. "On the Background of Social Epistemology." In *The Routledge Handbook of Social Epistemology*, edited by Miranda Fricker, Peter J. Graham, David Henderson, and Nikolaj J. L. L. Pedersen, 3–9. New York; London: Routledge/Taylor; Francis Group, 2020.

Henderson, David, Terrance Horgan, Matjaz Potrc, and Hannah Tierney. "Nonconciliation in Peer Disagreement: Its Phenomenology and Its Rationality." *Grazer Philosophische Studien* 94 (2017): 194–225.

Herzog, Lisa. "Ideal and Non-Ideal Theory and the Problem of Knowledge." *Journal of Applied Philosophy* 29, no. 4 (2012): 271–88.

Hieronymi, Pamela. "The Reasons of Trust." *Australasian Journal of Philosophy* 86, no. 2 (2008): 213–36.

Holman, Bennett, William J. Berger, Daniel J. Singer, Patrick Grim, and Aaron Bramson. "Diversity and Democracy: Agent-Based Modeling in Political Philosophy." *Historical Social Research* 43, no. 1 (2018): 259–84.

Honneth, Axel, and Jacques Rancière. *Recognition or Disagreement: A Critical Encounter on the Politics of Freedom, Equality, and Identity*. Edited by Katia Genel and Jean-Philippe Deranty. New York: Columbia University Press, 2016.

Hricko, Jonathon. "Scientific Rationality: Phlogiston as a Case Study." In *Rationality: Constraints and Contexts*, edited by Tzu-Wei Hung and Timothy Joseph Lane, London., 37–59. Elsevier Academic Press, 2017.

Hughes, Nick. "Epistemology Without Guidance." *Philosophy Studies*, 2021.

Hume, David. *An Enquiry Concerning Human Understanding*. Edited by Tom L. Beauchamp. Oxford/New York: Oxford University Press, 1999.

Huttegger, Simon M. "Merging of Opinions and Probability Kinematics." *The Review of Symbolic Logic* 8, no. 4 (2015): 611–48.

Jacobs, R. A., and Kruschke J. K. "Bayesian Learning Theory Applied to Human Cognition." *Wiley Interdisciplinary Reviews: Cognitive Science* 2, no. 1 (2011): 8–21.

James, William. "The Moral Philosopher and the Moral Life." In *The Will to Believe and Other Essays in Popular Philosophy*. Cambridge: Harvard University Press, 1979.

Jaynes, E. T. *Papers on Probability, Statistics, and Statistical Physics*. Edited by Roger D. Rosenkrantz. Dordrecht, Holland: D. Reidel Publishing Co., 1983.

———. *Probability Theory: The Logic of Science*. Edited by G. Larry Bretthorst. Cambridge: Cambridge University Press, 2003.

Jeffrey, Richard C. *The Logic of Decision*. 2nd ed. University of Chicago Press, 1990.

Jehle, David, and Branden Fitelson. "What Is the "Equal Weight View"?" *Episteme* 6, no. 3 (2009): 280–93.

Jenkin, Zoe, and Susanna Siegel. "Cognitive Penetrability: Modularity, Epistemology, and Ethics." *Review of Philosophy and Psychology* 6, no. 4 (2015): 531–45.

Johnson, Gabbrielle M. "Algorithmic Bias: On the Implicit Biases of Social Technology." *Synthese* 198, no. 10 (2020): 9941–61.

Jones, Karen. "'But i Was Counting on You!'." In *The Philosophy of Trust*, edited by Paul Faulkner and Thomas W. Simpson, 161–76. Oxford: Oxford University Press, 2017.

———. "The Politics of Intellectual Self-Trust." *Social Epistemology* 26, no. 2 (2012): 237–51.

———. "Trustworthiness." *Ethics* 123, no. 1 (2012): 61–85.

Joyce, James M. "Bayesianism." In *The Oxford Handbook of Rationality*, edited by Alfred R. Mele and Piers Rawling, 132–55. Oxford: Oxford University Press, 2004.

Jubb, Robert. "Tragedies of Non-Ideal Theory." *European Journal of Political Theory* 11, no. 3 (2012): 229–46.

Kahneman, Daniel, and Amos Tversky. "Judgment Under Uncertainty: Heuristics and Biases." *Science* 185, no. 4157 (1974): 1124–31.

———. "On the Reality of Cognitive Illusions." *Psychological Review* 103, no. 3 (1996): 582–91.

Kelly, Thomas. *Bias: A Philosophical Study*. Oxford: Oxford University Press, 2022.

———. "Disagreement and the Burdens of Judgment." In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey, 31–53. Oxford: Oxford University Press, 2013.

———. "Disagreement, Dogmatism, and Belief Polarization." *Journal of Philosophy* 105, no. 10 (2008): 611–33.

———. "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66, no. 3 (May 2003): 612–40.

———. "Evidence." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2016. Metaphysics Research Lab, Stanford University, 2016.

———. "Evidence: Fundamental Concepts and the Phenomenal Conception." *Philosophy Compass* 3, no. 5 (2008): 933–55.

———. "Peer Disagreement and Higher Order Evidence." In *Social Epistemology: Essential Readings*, edited by Alvin I. Goldman and Dennis Whitcomb, 183–217. Oxford University Press, 2011.

———. "The Epistemic Significance of Disagreement." In *Oxford Studies in Epistemology*, edited by Tamar Szabó Gendler and John Hawthorne, 1:167–96. Oxford: Oxford University Press, 2005.

———. "The Rationality of Belief and Some Other Propositional Attitudes." *Philosophical Studies* 110 (2002): 163–96.

Kemeny, J. "Fair Bets and Inductive Probabilities." *Journal of Symbolic Logic* 20 (1955): 263–73.

Kenny, Anthony. *A New History of Western Philosophy*. Oxford: Clarendon Press, 2010.

Keynes, John Maynard. *A Treatise on Probability*. London: Macmillan; Co., 1921.

Kim, Brian. "Pragmatic Encroachment in Epistemology." *Philosophy Compass* 12, no. 5 (2017).

Kim, Brian, and Matthew McGrath, eds. *Pragmatic Encroachment in Epistemology*. Routledge Studies in Epistemology. New York; London: Routledge, 2019.

King, Nathan L. "Disagreement: What's the Problem? Or a Good Peer Is Hard to Find." *Philosophy and Phenomenological Research* 85, no. 2 (2012): 249–72.

Kitcher, Philip. "Explanatory Unification and the Causal Structure of the World." In *Scientific Explanation*, edited by Philip Kitcher and Wesley Salmon, 410–505. Minneapolis: University of Minnesota Press, 1989.

———. "The Naturalists Return." *The Philosophical Review* 101, no. 1 (1992).

Klein, Jürgen, and Guido Giglioni. "Francis Bacon." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2020. Metaphysics Research Lab, Stanford University, 2020.

Knight, Jack, and James Johnson. *The Priority of Democracy: The Political Consequences of Pragmatism*. Princeton, NJ: Princeton University Press, 2011.

Kopec, Matthew, and Michael G. Titelbaum. "The Uniqueness Thesis." *Philosophy Compass* 11, no. 4 (2016): 189–200.

Kordzadeh, Nima, and Maryam Ghasemaghaei. "Algorithmic Bias: Review, Synthesis and Future Research Directions." *European Journal of Information Systems* 31, no. 3 (2022): 388–409.

Kuhn, Thomas S. "Objectivity, Value Judgment, and Theory Choice." In *"The Essential Tension: Selected Studies in Scientific Tradition and Change"*, 320–29. Chicago; London: University of Chicago Press, 1977.

———. *The Structure of Scientific Revolutions*. Fourth. Chicago; London: University of Chicago Press, 2012.

Kvanvig, Jonathan L. "Perspectivalism and Reflective Ascent." In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey, 223–42. Oxford, UK: Oxford University Press, 2013.

Lackey, Jennifer. "It Takes Two to Tango: Beyond Reductionism and Non-Reductionism in the Epistemology of Testimony." In *The Epistemology of Testimony*, edited by Ernest Sosa and Jennifer Lackey, 160–89. Oxford University Press, 2006.

———. *Learning from Words: Testimony as a Source of Knowledge*. Oxford University Press, 2008.

———. "Testimony: Acquiring Knowledge from Others." In *Social Epistemology: Essential Readings*, edited by Alvin I. Goldman and Dennis Whitcomb, 71–91. Oxford University Press, 2011.

———. "What Should We Do When We Disagree?" In *Oxford Studies in Epistemology*, edited by Tamar Szabó Gendler and John Hawthorne. Oxford: Oxford University Press, 2010.

Ladha, Krishna K. "The Condorcet Jury Theorem, Free Speech, and Correlated Votes." *American Journal of Political Science* 36, no. 3 (1992): 617–34.

Landemore, Hélène. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton: Princeton University Press, 2013.

Laudan, Larry. *Science and Values: An Essay on the Aims of Science and Their Role in Scientific Debate*. Berkeley/Los Angeles: University of California Press, 1984.

Lee, Matthew. "Conciliationism Without Uniqueness." *Grazer Philosophische Studien* 88 (2013): 161–88.

Lee, Peter M. *Bayesian Statistics: An Introduction*. Fourth. Wiley, 2012.

Leitgeb, Hannes, and André Carus. "Rudolf Carnap." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2022. Metaphysics Research Lab, Stanford University, 2022.

Leonard, Nick. "Epistemological Problems of Testimony." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2023. Metaphysics Research Lab, Stanford University, 2023.

Levenson, Laurie L. "Good Faith Defenses: Reshaping Strict Liability Crimes." *Cornell Law Review* 78, no. 3 (2017): 401–69.

Levinstein, Benjamin. "Permissive Rationality and Sensitivity." *Philosophy and Phenomenological Research* 94, no. 2 (2017): 342–70.

Lewis, David. "A Subjectivist's Guide to Objective Chance." In *Studies in Inductive Logic and Probability*, edited by Richard C. Jeffrey, Vol. 2. Berkeley: University of California Press, 1980.

———. "Elusive Knowledge." *Australasian Journal of Philosophy* 74, no. 4 (1996): 549–67.

———. "Humean Supervenience Debugged." *Mind* 103 (1994): 473–90.

Li, Han. "How Supererogation Can Save Intrapersonal Permissivism." *American Philosophical Quarterly* 56, no. 2 (April 2019): 171–86.

List, Christian. "Group Knowledge and Group Rationality: A Judgment Aggregation Perspective." In *Social Epistemology: Essential Readings*, edited by Alvin I. Goldman and Dennis Whitcomb, 221–41. Oxford University Press, 2011.

List, Christian, and Robert E. Goodin. "Epistemic Democracy: Generalizing the Condorcet Jury Theorem." *Journal of Political Philosophy* 9, no. 3 (2001): 277–306.

Lord, Charles, Lee Ross, and Lepper Mark. "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence." *Journal of Personality and Social Psychology* 37, no. 11 (1979): 2098–2109.

Macarthur, David. "A Kant-Inspired Vision of Pragmatism as Democratic Experimentalism." In *Pragmatism, Kant and Transcendental Philosophy*, edited by Gabrielle Gava and Robert Stern, 67–84. Routledge, 2015.

MacGilvray, Eric. "Democratic Doubts: Pragmatism and the Epistemic Defense of Democracy." *Journal of Political Philosophy* 22 (2014): 104–23.

Malherbe, Michel. "Bacon's Method of Science." In *The Cambridge Companion to Bacon*, edited by Markku Peltonen, 75–98. Cambridge: Cambridge University Press, 1996.

Margolis, Joseph. "Peirce's Fallibilism." *Transactions of the Charles S. Peirce Society* 34, no. 3 (1998): 535–69.

Matheson, Jonathan. "Disagreement: Idealized and Everyday." In *The Ethics of Belief: Individual and Social*, edited by Rico Vitz and Jonathan Matheson, 315–30. Oxford University Press, 2014.

Matheson, Jonathan, and Brandon Carey. "How Skeptical Is the Equal Weight View?" In *Disagreement and Skepticism*, edited by Diego E. Machuca, 131–49. New York; London: Routledge/Taylor; Francis Group, 2013.

Mayo, Deborah G., and D. R. Cox. "Frequentist Statistics as a Theory of Inductive Inference." In *Institute of Mathematical Statistics Lecture Notes - Monograph Series*, edited by Javier Rojo, 49:77–97. Institute of Mathematical Statistics, 2006.

McCain, Kevin. *Evidentialism and Epistemic Justification*. New York. Routledge, 2014.

———. "Evidentialism: A Primer." In *Believing in Accordance with the Evidence: New Essays on Evidentialism*, edited by Kevin McCain, 1–4. Synthese Library. Springer Cham, 2018.

McKenna, Robin. *Non-Ideal Epistemology*. Oxford: Oxford University Press, forthcoming.

McMullin, Ernan. "The Virtues of a Good Theory." In *The Routledge Companion to Philosophy of Science*, edited by Martin Curd and Stathis Psillos. London: Routledge, 2009.

———. "Values in Science." *PSA: Proceeedings of the Biennial Meeting of the Philosophy of Science Association 1982*, 1982, 3–28.

Mellor, D. H. *Probability: A Philosophical Introduction*. London; New York: Routledge, 2005.

Millgram, Elijah. *The Great Endarkenment: Philosophy for an Age of Hyperspecialization*. Oxford University Press USA, 2015.

Misak, Cheryl. "Making Disagreement Matter: Pragmatism and Deliberative Democracy." In *The Pragmatism Reader: From Peirce Through the Present*, edited by Robert B. Talisse and Scott F. Aikin, 471–83. Princeton; Oxford: Princeton University Press, 2011.

———. *Truth and the End of Inquiry: A Peircean Account of Truth*. Oxford: Clarendon Press, 1991.

———. *Truth, Politics, Morality: Pragmatism and Deliberation*. New York: Routledge, 2000.

Misak, Cheryl, and Robert B. Talisse. "Debate: Pragmatist Epsitemology and Democratic Theory: A Reply to Eric MacGilvray." *The Journal of Political Philosophy* 22, no. 3 (2014): 366–76.

Morreau, Michael. "Theory Choice and Social Choice: Kuhn Vindicated." *Mind* 124, no. 493 (2015): 239–62.

Moss, Sarah. "Credal Dilemmas." *Noûs* 48, no. 3 (2014): 665–83.

Nguyen, C. Thi. "Echo Chambers and Epistemic Bubbles." *Episteme*, 2018, 1–21.

O'Neill, Onora. *A Question of Trust*. Cambridge: Cambridge University Press, 2002.

———. *Autonomy and Trust in Bioethics*. Cambridge: Cambridge University Press, 2002.

Okasha, Samir. "On Arrow's Theorem and Scientific Rationality: Reply to Morreau and Stegenga." *Mind* 124, no. 493 (2015): 279–94.

———. "Theory Choice and Social Choice: Kuhn Versus Arrow." *Mind* 120, no. 477 (2011): 83–115.

Oreskes, Naomi, and Erik M. Conway. *Merchants of Doubt*. New York: Bloomsbury Press, 2010.

Pasnau, Robert. "Disagreement and the Value of Self-Trust." *Philosophical Studies* 172, no. 9 (2015): 2315–39.

Pearl, Judea. *Causality: Models, Reasoning, and Inferences*. Second. Cambridge: Cambridge University Press, 2009.

Peels, Rik, and Anthony Booth. "Why Responsible Belief Is Permissible Belief." *Analytic Philosophy* 55 (2014): 75–88.

Peirce, Charles S. "How to Make Our Ideas Clear." In *The Pragmatism Reader: From Peirce Through the Present*, edited by Robert B. Talisse and Scott F. Aikin, 50–65. Princeton; Oxford: Princeton University Press, 2011.

———. "Some Consequences of Four Incapacities." In *The Pragmatism Reader: From Peirce Through the Present*, edited by Robert B. Talisse and Scott F. Aikin, 12–36. Princeton; Oxford: Princeton University Press, 2011.

———. "The Fixation of Belief." In *The Pragmatism Reader: From Peirce Through the Present*, edited by Robert B. Talisse and Scott F. Aikin, 37–49. Princeton; Oxford: Princeton University Press, 2011.

Pettit, Philip. "Groups with Minds of Their Own." In *Social Epistemology: Essential Readings*, edited by Alvin I. Goldman and Dennis Whitcomb, 242–68. Oxford University Press, 2011.

Poslajko, Krzysztof. "The Lycan-Stich Argument and the Plasticity of "Belief"." *Erkenntnis* 87 (2022): 1257–73.

Price, Huw. *Naturalism Without Mirrors*. Oxford: Oxford University Press, 2011.

Putnam, Hilary. "A Reconsideration of Deweyan Democracy." In *The Pragmatism Reader: From Peirce Through the Present*, edited by Robert B. Talisse and Scott F. Aikin, 331–52. Princeton; Oxford: Princeton University Press, 2011.

———. "Language and Philosophy." In *Mind, Language and Reality*, 1–32. Cambridge: Cambridge University Press, 1975.

———. "Pragmatism and Moral Objectivity." In *Words and Life*, edited by James Conant, 151–81. Cambridge, Massachusetts; London, England: Harvard University Press, 1994.

———. *The Collapse of the Fact/Value Dichotomy and Other Essays*. Cambridge, Massachusetts; London, England: Harvard University Press, 2002.

———. *The Threefold Cord: Mind, Body, and World*. New York: Columbia University Press, 1999.

Putnam, Hilary, and Ruth Anna Putnam. "Dewey's Logic: Epistemology as Hypothesis." In *Words and Life*, edited by James Conant, 198–220. Cambridge, Massachusetts; London, England: Harvard University Press, 1994.

Quine, W. V. O. "Comment on Parsons." In *Perspectives on Quine*, edited by Robert Barrett and Roger Gibson. Oxford: Basil Blackwell, 1990.

Quirk, Paul J. "Making It up on Volume: Are Larger Groups Really Smarter?" *Critical Review* 26, no. 1-2 (2014): 129–50.

Ralston, Shane J. "Can Pragmatists Be Institutionalists? John Dewey Joins the Non-Ideal/Ideal Theory Debate." *Human Studies* 33, no. 1 (May 2010): 65–84.

Ramsey, Frank. "Truth and Probability." In *Philosophy of Probability: Contemporary Readings*, edited by Antony Eagle, 52–94. Routledge, 1926.

Rancière, Jacques. *Disagreement: Politics and Philosophy*. Translated by Julie Rose. Minneapolis; London: University of Minnesota Press, 1999.

Rawls, John. *A Theory of Justice*. Cambridge, Massachusetts: Harvard University Press, 1999.

———. "Justice as Fairness: Political Not Metaphysical." *Philosophy & Public Affairs* 14, no. 3 (1985): 223–51.

Reisner, Andrew. "The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem." *Philosophical Studies* 145, no. 2 (2009): 257–72.

———. "Weighing Pragmatic and Evidential Reasons for Belief." *Philosophical Studies* 138, no. 1 (2008): 17–27.

Resnik, Michael D. *Choices: An Introduction to Decision Theory*. Minneapolis: University of Minnesota Press, 1987.

Riker, William H. *Liberalism Against Populism*. San Francisco: W. H. Freeman, 1982.

Rinard, Susanna. "A Decision Theory for Imprecise Probabilities." *Philosophers' Imprint* 15, no. 7 (2015).

———. "Believing for Practical Reasons." *Noûs* 53, no. 4 (2018): 763–84.

Robert, Christian P. *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*. Second. New York: Springer, 2007.

Robeyns, Ingrid. "Ideal Theory in Theory and Practice." *Social Theory and Practice* 34, no. 3 (2008): 341–62.

Rorty, Richard. *Philosophy and the Mirror of Nature*. Princeton University Press, 1979.

Rosenkrantz, Roger D. *Foundations and Applications of Inductive Probability*. Atascadero, CA: Ridgeview Publishing, 1981.

Ross, W. D. *Aristotle's Metaphysics, a Revised Text with Introduction and Commentary*. Oxford University Press, 1924.

Roush, Sherrilyn. "Epistemic Self-Doubt." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2017. Metaphysics Research Lab, Stanford University, 2023.

Russell, Bertrand. *The Problems of Philosophy*. London: Williams & Norgate, 1912.

Saint-Croix, Catharine. "Non-Ideal Epistemology in a Social World." PhD thesis, University of Michigan, 2018.

Sakai, Ryota. "Mathematical Models and Robustness Analysis in Epistemic Democracy: A Systematic Review of Diversity Trumps Ability Theorem Models." *Philosophy of the Social Sciences* 50, no. 3 (2020): 195–214.

Schafer, Karl. "How Common Is Peer Disagreement? On Self-Trust and Rational Symmetry." *Philosophy and Phenomenological Research* 91, no. 1 (2015): 25–46.

Scheman, Naomi. "Trust and Trustworthiness." In *The Routledge Handbook of Trust and Philosophy*, edited by Judith Simon, First., 28–40. New York; London: Routledge/Taylor; Francis Group, 2020.

Schwitzgebel, Eric. "The Pragmatic Metaphysics of Belief." In *The Fragmented Mind*, edited by Cristina Borgoni, Dirk Kindermann, and Andrea Onofri, 350–75. Oxford: Oxford University Press, 2021.

Sen, Amartya. "What Do We Want from a Theory of Justice?" *The Journal of Philosophy* 103, no. 5 (2006): 215–38.

Shafer, Glenn. *A Mathematical Theory of Evidence*. Princeton; London: Princeton University Press, 1976.

Shafer, Glenn, and Judea Pearl. *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kaufmann, 1988.

Siegel, Susanna. "How Is Wishful Seeing Like Wishful Thinking?" *Philosophy and Phenomenological Research* 95, no. 2 (2017): 408–35.

———. "The Epistemic Impact of the Etiology of Experience." *Philosophical Studies* 162, no. 3 (2013): 697–722.

Siegfried, Robert. *From Elements to Atoms: A History of Chemical Composition*. Vol. 92. Transactions of the American Philosophical Society 4. Philadelphia: American Philosophical Society, 2002.

Simmons, A. John. "Ideal and Nonideal Theory." *Philosophy & Public Affairs* 38, no. 1 (2010): 5–36.

Singer, Daniel J. "Diversity, Not Randomness, Trumps Ability." *Philosophy of Science* 86 (2019): 178–91.

Skyrms, B. *Pragmatics and Empiricism*. New Haven: Yale University Press, 1984.

Smith, Nicholas J. J. "Problems of Precision in Fuzzy Theories of Vagueness and Bayesian Epistemology." In *Vagueness and Rationality in Language Use and Cognition*, edited by Richard Dietz, 5:31–48. Switzerland: Springer Nature Switzerland AG, 2019.

Soon, Valerie. "An Intrapersonal, Intertemporal Solution to an Interpersonal Dilemma." *Philosophical Studies* 178 (2021): 3353–70.

Stanley, Jason. *Knowledge and Practical Interests*. Oxford: Oxford University Press, 2005.

Stegenga, Jacob. "Theory Choice and Social Choice: Okasha Versus Sen." *Mind* 124, no. 493 (2015): 263–77.

Steglich-Petersen, Asbjørn, and Mattias Skipper. "An Instrumentalist Account of How to Weigh Epistemic and Practical Reasons for Belief." *Mind* 129, no. 516 (2020): 1071–94.

Stemplowska, Zofia. "What's Ideal about Ideal Theory?" *Social Theory and Practice* 34, no. 3 (2008): 319–40.

Steup, Matthias, ed. *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford: Oxford University Press, 2001.

Steup, Matthias, and Ram Neta. "Epistemology." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2020. Metaphysics Research Lab, Stanford University, 2020.

Strawson, P. F. *Scepticism and Naturalism: Some Varieties*. Routledge, 1987.

Stroud, Barry. *The Significance of Philosophical Scepticism*. Oxford University Press, 1984.

Sud, Rohan. "A Forward Looking Decision Rule for Imprecise Credences." *Philosophical Studies* 167, no. 1 (2014): 119–39.

Talisse, Robert B. *A Pragmatist Philosophy of Democracy*. New York: Routledge, 2007.

———. *Democracy After Liberalism: Pragmatism and Deliberative Politics*. New York: Routledge, 2005.

———. "Saving Pragmatist Democratic Theory (from Itself)." *Ethics & Politics* 12, no. 1 (2010): 12–27.

Talisse, Robert B., and Scott F. Aikin, eds. *The Pragmatism Reader: From Peirce Through the Present*. Princeton; Oxford: Princeton University Press, 2011.

Teller, Paul. "Conditionalization, Observation, and Change of Preference." In *Foundations of Probability Theory*, edited by W. Harper and C. A. Hooker. Dordrecht: D. Reidel, 1976.

Thompson, Abigail. "Does Diversity Trump Ability?" *Notices of the AMS* 61, no. 9 (2014): 1–24.

Thorstad, David. "There Are No Epistemic Norms of Inquiry." *Synthese* 200 (2022).

Troffaes, Matthias C. M. "Decision Making Under Uncertainty Using Imprecise Probabilities." *International Journal of Approximate Reasoning* 45 (2007): 17–29.

Urbach, Peter. *Francis Bacon's Philosophy of Science: An Account and a Reappraisal*. La Salle, IL: Open Court, 1987.

Valentini, Laura. "Ideal Vs. Non-Ideal Theory: A Conceptual Map." *Philosophy Compass* 7, no. 9 (2012): 654–64.

Vineberg, Susan. "Dutch Book Arguments." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2022. Metaphysics Research Lab, Stanford University, 2022.

Waldron, Jeremy. "The Wisdom of the Multitude: Some Reflections on Book 3, Chapter 11 of Aristotle's Politics." *Political Theory* 23, no. 4 (1995): 563–84.

Wedgwood, Ralph. "The Moral Evil Demons." In *Disagreement*, edited by Richard Feldman and Ted A. Warfield. Oxford: Oxford University Press, 2010.

White, Roger. "Epistemic Permissiveness." *Philosophical Perspectives* 19 (2005): 445–59.

———. "Evidence Cannot Be Permissive." In *Contemporary Debates in Epistemology*, edited by Matthias Steup, John Turri, and Ernest Sosa, Second., 312–23. UK: Wiley-Blackwell, 2014.

Williamson, Jon. *In Defence of Objective Bayesianism*. Oxford: Oxford University Press, 2010.

Williamson, Timothy. "Improbable Knowing." In *Evidentialism and Its Discontents*, edited by Trent Dougherty, 147–64. New York: Oxford University Press, 2011.

———. *Knowledge and Its Limits*. Oxford: Oxford University Press, 2000.

———. *Vagueness*. London: Routledge, 1994.

———. "Why Epistemology Cannot Be Operationalized." In *Epistemology: New Essays*, edited by Quentin Smith. Oxford: Oxford University Press, 2008.

Zimmerman, Aaron Z. *Belief: A Pragmatic Picture*. Oxford: Oxford University Press, 2018.