

THE EPISTEMOLOGY OF DISAGREEMENT: HUME, KANT, AND THE
CURRENT DEBATE

by

Robert Kyle Whitaker

A Dissertation submitted to the Faculty of the Graduate School, Marquette
University, in Partial Fulfillment of the Requirements for the Degree of Doctor of
Philosophy

Milwaukee, Wisconsin

May 2020

ABSTRACT
THE EPISTEMOLOGY OF DISAGREEMENT: HUME, KANT, AND THE
CURRENT DEBATE

Robert Kyle Whitaker

Marquette University, 2020

The epistemological issue of disagreement comprises several related problems which arise in relation to disagreeing with another person. The central questions at issue are: (1) Can a body of evidence confer rationality on opposed propositions? (2) What is the relevance of unshareable evidence to disagreement? (3) What are one's epistemic responsibilities in the context of disagreement?

I consider several arguments from the recent disagreement literature which suggest that reasonable disagreements between people who have shared their evidence and are epistemic *peers*--i.e., they are equally informed about the disputed issue, and are roughly equal with respect to intelligence, thoughtfulness, carefulness, alertness, and so on--are not possible. I also consider several arguments which suggest that one may rationally persist in one's views in spite of peer disagreement. In the course of considering these arguments, I discuss the role of "evidential seemings," seemings of the form "It seems to one that evidence *E* supports proposition *p*"; their analogous relationship to perceptual seemings; and how such seemings might be outweighed by higher-order evidence provided by peer disagreement.

I propose that the most difficult and interesting aspect of the peer disagreement problem is *prolonged* disagreement with a recognized epistemic peer, which may undermine one's evidential seemings in the form of Parity: in the context of peer disagreement, both parties may be aware that things would seem just the same to them as they do if in fact the other party were correct and they were mistaken. This presents a localized skeptical problem that affects what one should believe in the context of such a disagreement, because prolonged peer disagreement puts one's own evidential seemings into conflict.

I propose and discuss three possible solutions to the Parity Problem, one inspired by David Hume, and two inspired by Immanuel Kant: (1) Practical considerations stemming from the underdetermination of one's views by the available evidence as well as one's larger epistemic goals can justify maintaining one's view. Alternatively, judgments that someone else is an epistemic peer may be classified as either (2) empty regulative theoretical judgments (a Kantian category), or (3) non-theoretic judgments (analogous to judgments about taste). In either case, rational disagreement under Parity is possible.

ACKNOWLEDGEMENTS

Robert Kyle Whitaker

I would like to thank my dissertation director, Prof. Michael Wreen, for his patient guidance throughout this project. My ability to write and assess arguments is in large part due to him, as he was my professor for graduate logic and several courses on analytic philosophy. Any argumentative errors that may remain in this work are likely due to my failing to consider his advice more carefully.

I also want to thank Dr. Noel Adams and Dr. Stanley Harrison for serving on my dissertation committee, and for the positive influence--both professional and personal--they've had on me during my graduate studies. Dr. Adams introduced me to Søren Kierkegaard, and Dr. Harrison to C.S. Peirce, two philosophers from whom I will continue to learn for the rest of my life. I am indebted to Drs. Adams and Harrison for these introductions, and for representing those great philosophers so well.

I would like to extend special thanks to Dr. Paul Moser for serving as outside reader on this project. I have benefitted greatly from his work in both epistemology and philosophy of religion, as well as from the graduate seminar on Kierkegaard that I was fortunate to take with him. In many ways he serves as a model for me of Christian philosophy done well: with rigor, inclusion, and guided by a genuine faith. I am grateful for this example.

There are many others who were not engaged directly in this project, but whose influence has in one way or another contributed to its completion. Thanks are due to all of those who have taught me to think clearly, and shown me by example the difference that philosophy can make with respect to forming a good character. These people include Dr. Cynthia Gayman, Dr. Rory Goggins, Dr. Theresa Tobin, Dr. Nancy Snow, Dr. Bronwyn Finnigan, and many others.

In addition, I must thank my fellow Marquette graduate students, whose friendships have sustained me through this process, especially Nick Oschman, Brett Yardley, and Shaun Blanchard. I have learned as much from them about how to be a scholar as anyone, and they've all provided much needed levity and kindness.

Finally, to my family. My parents, George and Linda Whitaker, first told me that I was capable, and guided me to many of the ideas that still occupy my attention. They gave support of every kind at every stage of my studies, and never once was any condition attached. And to my wife, Emily, the biggest thanks of all. Her contributions are too many and too varied to name. Completing a dissertation while remaining mentally healthy is a difficult task, and my accomplishment of it is due almost solely to her.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	i
INTRODUCTION	1
CHAPTER 1: CAN EVIDENCE SUPPORT OPPOSED PROPOSITIONS?.....	10
1.1: Peter van Inwagen on Evidence and Exclusivity	10
1.2: Richard Feldman's Uniqueness Thesis	19
CHAPTER 2: PRIVATE EVIDENCE.....	33
2.1: Van Inwagen on the Use of Private Evidence	33
2.2: Feldman on the Relevance of Private Evidence.....	40
2.3: Christensen on the Symmetry of Private Evidence.....	46
CHAPTER 3: THE PERCEPTUAL ANALOGY	54
3.1: Perceptual Seeming in the Disagreement Debate	54
3.2: Implications of the Analogy	62
CHAPTER 4: THE EVIDENTIAL VALUE AND DEFEAT OF SEEMINGS	69
4.1: Rosen on the Contingency of Evidential Seemings	69
4.2: The Epistemic Parity of Seemings	81
CHAPTER 5: THE EVIDENTIAL VALUE OF DISAGREEMENT ITSELF	95
5.1: Idealization, Independence, and Bootstrapping	95
5.2: Christensen on Symmetry between Apparent Peers	103
CHAPTER 6: HIGHER-ORDER EVIDENCE AND TESTIMONY	110
6.1: The Meta-Evidence Principle.....	110
6.2: Disagreement and Testimony.....	122

6.3: The Problem of Possible Rational Disagreement.....	130
CHAPTER 7: PEERHOOD, CONFORMISM, AND NONCONFORMISM	140
7.1: Defining Peerhood.....	140
7.2: Responding to Parity: Conformism and Nonconformism.....	149
CHAPTER 8: PROPOSED SOLUTIONS I: HUME’S CONTRIBUTION	161
8.1: A Humean Approach to Disagreement	162
8.2: Moffett’s Millian Solution to Peer Disagreement	172
8.3: Hume Again	180
CHAPTER 9: PROPOSED SOLUTIONS II: KANT’S CONTRIBUTION	187
9.1: Kant’s Categories of Judgments and Disagreement.....	187
9.2: Conclusion.....	196
CONCLUSION	200
BIBLIOGRAPHY	206

Introduction

The literature on the epistemology of disagreement contains several questions that must be clearly distinguished. A survey of this literature yields many similar but distinct emphases with respect to what the debate over rational disagreement is fundamentally about. If we ask, “What is the guiding question in this debate?” the answer is: There is no single guiding question. Rather, there is an array of related questions. Five of the fundamental questions are:

(1) Can a single body of evidence support opposed conclusions drawn on the basis of that evidence?

For example, Richard Feldman asks: “...how exactly can there be reasonable disagreements? And how can there be reasonable disagreements when the parties to the disagreement have been confronted with a single body of evidence?”¹

(2) Can one’s evidence for a proposition always be shared, and if not, what effect does unshareable (“private”) evidence have on the justification of one’s belief in a proposition?

For example, Peter van Inwagen, one of the earliest framers of what has become the disagreement debate, discusses what he takes to be a relatively common form of evidence:

... there seem to be plausible examples of “having evidence” that do not conform to the courtroom-and-laboratory paradigm of evidence.... There are, therefore, arguments by example for the conclusion that... evidence is not always of the public sort, that some evidence is not *exportable*, that some evidence cannot be passed from one person to another.²

(3) When is it reasonable for one to hold onto a belief in a disputed proposition?

¹ Richard Feldman, “Reasonable Religious Disagreements,” in *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, ed. Louise Antony (Oxford University Press, 2007), 201.

² Peter van Inwagen, “We’re Right. They’re Wrong,” in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford University Press, 2010), 25-6.

For example, John Pittard writes: "...when can we reasonably maintain a confident belief in some proposition despite the fact that many people who seem as qualified as ourselves to assess the proposition's plausibility hold it to be false?"³ He continues:

It is obvious that the discovery that someone disagrees with my belief that p should at least sometimes cause me to be less confident in my disputed belief. The primary area of controversy in the epistemology of disagreement has to do with the conditions under which the prima facie threat posed by disagreement can be reasonably resisted... Can... partisan, dispute-dependent reasons be a good basis for doubting the reliability of my disputant? This question serves as the primary point of contestation in debates over the epistemic significance of disagreement.⁴

Along the same lines, David Christensen asks:

How should I react when I discover that my friend and I have very different beliefs on some topic? Thinking about belief in a quantitative or graded way, the question concerns cases in which my friend and I have very different degrees of confidence in some proposition P. Should my discovery of her differing degree of belief in P lead me to revise my own confidence in P?⁵

Thomas Kelly puts it similarly: "...once you and I learn that the other has arrived at a different conclusion despite having been exposed to the same evidence and arguments, how (if at all) should we revise our original views?"⁶

(4) How should one's beliefs be affected by the beliefs of others?

Adam Elga frames the question more generally: "How should you take into account the opinions of an advisor? ... How, exactly, should we be guided by outside

³ John Pittard, *Disagreement, Rationality, and Religious Belief*, Doctoral Dissertation, Yale University (2013), 3.

⁴ *Ibid.*, 9-10.

⁵ David Christensen, "Epistemology of Disagreement: The Good News," *Philosophical Review* 116, no. 2 (2007): 188.

⁶ Thomas Kelly, "Peer Disagreement and Higher-Order Evidence," in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford University Press, 2010), 112.

opinions?”⁷ This is a question about testimony rather than disagreement per se, although disagreement between epistemic peers is what is prompting the question in this debate. Similarly, Marc Moffett frames the issue as a matter of competing demands of rationality regarding one’s epistemic relationship to others. If critical scrutiny of one’s views by one’s peers is conducive to the justification of one’s beliefs, which seems plausible, he asks how one can avoid the problem of having one’s view swamped by the contrary views of one’s.⁸

(5) Is it reasonable to maintain both one’s own belief in the face of disagreement, and also the reasonableness of one’s disputant?

Feldman asks:

Can epistemic peers who have shared their evidence reasonably maintain their own belief yet also think that the other party to the disagreement is also reasonable?⁹

The concept of *peerhood* occupies a central place in the debate, and its primary importance is its role in the present question.¹⁰ For example, for Michael Bergmann, the central issue reflects both the sharing of a single evidence base and the peer-status of one’s disputant:

The Question: Can two people – who are, and realize they are, intellectually virtuous to about the same degree – both be rational in continuing knowingly to

⁷ Adam Elga, “Reflection and Disagreement,” *Noûs* 41, no. 3 (2007): 478–79.

⁸ Marc Moffett, “Reasonable Disagreement and Rational Group Inquiry,” *Episteme: A Journal of Social Epistemology* 4, no. 3 (2007): 352–67.

⁹ Feldman, “Reasonable Religious Disagreements,” 201.

¹⁰ As we will see, what exactly peerhood amounts to is a contentious question. For present purposes, we can say that to be a peer is a context-dependent (i.e., not global) relational property whereby two or more knowers are similarly qualified to assess the evidence relevant to some disputed question. “Similarly qualified” is of course vague, and various commentators will argue for different interpretations. The two main divisions here are between those who prefer to define peerhood from the “bottom up”--so that two people are peers if they are roughly equal with respect to various relevant epistemic factors such as intelligence, thoroughness, background knowledge, etc.--and those who prefer to define peerhood from the “top down,” so that A and B are peers iff A is equally likely as B to judge correctly about the truth value of p, where p is some disputed proposition. These notions of peerhood will be discussed in Chapter Seven.

disagree after full disclosure (by each to the other of all the relevant evidence they can think of) while at the same time thinking that the other may well be rational too?¹¹

Clearly, these questions differ in important ways, though they also overlap. Many commentators discuss most of these questions at some point, or at least they take their arguments to have implications for all five questions. In my view, these questions can be reduced to a central three Main Questions (MQs):

(MQ1) What expectations can one reasonably have of a body of evidence for a proposition?

(MQ2) What is the relevance of private evidence to peer disagreement?

(MQ3) What are one's epistemic responsibilities in the context of disagreement?

(MQ1) encapsulates the concerns of question (1). It will be the subject of Chapter One.

(MQ2) encapsulates the concerns of question (2). It will be the subject of Chapters Two through Six. (MQ3) encapsulates the concerns of questions (3) through (5). It will be the subject of Chapters Seven through Nine.

The question of whether a single body of evidence supports more than one proposition of a competing set of propositions is answered in various ways in the literature on peer disagreement. The most notable figures who discuss it are Peter van Inwagen and Richard Feldman. As they frame the issue, the central question is whether two people could reasonably believe opposed propositions on the basis of the same evidence. Both van Inwagen and Feldman deny this possibility. Their reasons for this are discussed and critiqued in Chapter One. Both hold that evidence, of metaphysical necessity, can support only one of a competing set of propositions, a principle known as

¹¹ Michael Bergmann, "Rational Disagreement After Full Disclosure," *Episteme* 6, no. 3 (2009): 336.

the Uniqueness Thesis. I argue that the Uniqueness Thesis is false, as it is far too strong, has counterintuitive consequences (even on weaker versions), and assumes things at issue in the disagreement debate.

Chapter Two considers the possibility of private evidence. Private evidence is evidence that cannot be shared, despite one's best efforts. It may include the felt truth of some proposition, or the seeming to one that some piece of evidence supports p rather than $\sim p$. I call these latter sorts of experiences "evidential seemings." I argue that private evidence in the form of evidential seemings may justify one in maintaining one's view in the face of disagreement, and whether it does depends on the confidence one has in one's own private seemings, on the evidential value of disagreement itself, and on the confidence one has in the peer status of one's interlocutor.

Evidential seemings are often compared in the disagreement literature to perceptual seemings, such as the sense that one is experiencing a green expanse when standing in front of the quad in the springtime.¹² This comparison is the subject of Chapter Three. There I argue that evidential seemings share broad structural similarity with perceptual seemings, including being non-doxastic, pre-theoretical, sui generis, basic, generative of prima facie justification, and incapable of non-circular justification themselves. I argue that such seemings have evidential significance in a similar way that perceptual seemings do, and that this significance does not disappear on the discovery of disagreement, though the seemings do of course remain defeasible.

¹² I borrow this example from Robert Audi, *Epistemology: A Contemporary Introduction to the Theory of Knowledge*, 3rd ed. (New York: Routledge, 2011), 1.

Chapters Four and Five tackle the related questions of (1) the evidential value of and the defeasibility of evidential seemings, and (2) the evidential import of disagreement itself. The goal of these chapters is not to determine what one should do in the context of disagreement, but rather to assess (a) what weight the mere fact of disagreement, or agreement, should have in one's total evidence for some proposition, and (b) whether it may serve as a defeater for one's felt sense of what a body of evidence supports. I consider and reject several ways in which disagreement may defeat one's evidential seemings, such as by providing evidence of the contingency of one's belief (à la Gideon Rosen), or by creating epistemic symmetry between peers so that a peer has no way to identify which peer's seemings are correct (à la David Christensen and Adam Elga). I argue that the fact of disagreement itself does not necessarily provide counterbalancing evidence regarding what one should believe about a proposition. Chapter Four also introduces the Parity Problem: in a case of prolonged disagreement with a recognized peer, each party may be aware that things would seem just the same as they do if in fact the other party were correct and he/she were mistaken. This presents a skeptical difficulty that affects what one should believe about p in the context of such a disagreement, and may in fact undermine one's justification for one's prior beliefs. This is because disagreement under Parity puts one's own evidential seemings into conflict.

Chapter Six considers the role of higher-order evidence in peer disagreement in two forms: (1) as meta-evidence (i.e., evidence of there being evidence for a proposition), and (2) as the *possibility* that there is evidence, even if that possibility is not actually instantiated. In conjunction with (1), I also consider in this chapter the relationship between disagreement and testimony. The first section considers the claim, advanced by

Richard Feldman, that “evidence of evidence is evidence.”¹³ The idea here is that the fact of disagreement itself gives one extra evidence for or against some proposition that one would not otherwise have, because disagreement gives one evidence *that there is evidence* with respect to p. This fact itself is claimed to have evidential force, even if one is not aware of the evidence possessed by the other party. I argue that meta-evidence may not in fact constitute evidence for a subject that is directly relevant to the disputed proposition. However, I do hold that there is something important about meta-evidence: insofar as one is dealing with a perspective that is constituted by certain features of epistemic excellence, the views of a person with that perspective may count as evidence for a disputed proposition.

The second section considers the relationship between disagreement and testimony. I suggest that a virtue theory of testimony can help to clarify whether negative testimony (disagreement), or any other form of testimony, carries evidential weight, and that such a theory helps to justify the critique of Feldman’s argument in the previous section. Finally, the third section discusses the problem of *possible* rational disagreement. There is some reason to think that if actual peer disagreement presents a problem for maintaining a view rationally because the disagreement provides higher-order evidence that one’s judgment of the first-order evidence is mistaken, then this problem can be generalized to include merely possible disagreement as well. If it can, then general skepticism looms as a possibility. I argue that this problem does not actually have the force that has been claimed for it. Thus, even if peer disagreement creates a localized skeptical problem, this does not entail generalized skepticism.

¹³ See Richard Feldman, “Reasonable Religious Disagreements,” 208.

Chapter Seven considers the notion of peerhood. An influential view of peerhood advanced by Adam Elga, according to which epistemic peers are those who are equally likely to be correct or mistaken, is critiqued and rejected. In its place, a more common-sense approach to defining peerhood is defended. The approach is based on similarities between people with respect to various factors that make one a good assessor of evidence for a proposition, such as intelligence, carefulness, time spent considering the evidence, sobriety, etc. The chapter then turns to both the Conformist and Nonconformist responses to the rational disagreement problem.

The Conformist view holds that one ought to *conform* one's view to one's peer's view. The Nonconformist view holds that one is typically justified in *maintaining* one's own view in the midst of peer disagreement. I argue that both views, as presented, have problems. Conformist views (e.g., David Christensen, Richard Feldman, Adam Elga) place too much emphasis on the evidential value of disagreement itself, mistakenly taking the fact of disagreement to have some overriding power that it does not always have. They also tend to assume (e.g., Christensen) that peer disagreement is epistemically similar to disagreement with a superior, which is unwarranted. Nonconformist views (e.g., Thomas Kelly, Marc Moffett, Michael Bergmann), on the other hand, tend to underestimate the significance of epistemic parity produced by long disagreement with someone one takes to be a peer. Thus, I argue that the strength of the Conformist view is parasitic on the power of the Parity Problem. Likewise, the burden of the Nonconformist view is to overcome that problem and provide a clear explanation of how one remains justified in one's view in cases of prolonged disagreement after full disclosure of evidence with a recognized epistemic peer.

In Chapters Eight and Nine, I offer three proposed solutions to the problem of rational peer disagreement under Parity. The first, inspired by David Hume and Marc Moffett, is presented in Chapter Eight. There I argue that (1) owing to the nature of judgments that someone is an epistemic peer, judgments that are at best inexact estimates, an objectively rational solution to the peer disagreement problem is not to be had, but that (2) nonetheless it can be reasonable to maintain one's view in a peer disagreement given one's overall epistemic goals.

The second and third solutions, inspired by Immanuel Kant and William Alston, are presented in Chapter Nine. These solutions concern the nature of peerhood judgments--i.e., judgments that someone is an epistemic peer. I argue that Kant's classifications of judgments provide two ways in which peerhood judgments may allow for the possibility of rational disagreement, disagreement such that each peer may justifiably believe that the other is responding rationally to the evidence, despite coming to an opposed conclusion. I argue that peerhood judgments may be classified as either (1) empty regulative theoretical judgments, or (2) non-theoretic judgments (analogous to judgments about taste). In either case, rational disagreement under Parity may be possible. I conclude with a critique and defense of these solutions, and some reflections on where they leave us with respect to the peer disagreement problem.

Chapter One – Can Evidence Support Opposed Propositions?

1.1 – Peter van Inwagen on Evidence and Exclusivity

In his paper “We’re Right. They’re Wrong,” Peter van Inwagen introduces the issue of peer disagreement by distinguishing between what he calls “weakly exclusivist” and “strongly exclusivist” positions on religious disagreement.¹ He then extends the discussion to philosophical, political, and other forms of disagreement. He calls any religion an “Ism,” and says that an Ism is “weakly exclusivist” if

(1) it’s logically inconsistent with all other Ism’s, and

(2) according to Ism, its adherents are rational to accept its teachings.

(1) and (2) together entail that it is rational for the adherent of an Ism to believe that the teachings of all other Ism’s are at least partly wrong. Note that this does *not* entail that one should believe that the adherents of other Ism’s are irrational. An Ism is “strongly exclusivist” if

(A) it is weakly exclusivist, and

(B) according to Ism, its adherents (or anyone in their epistemic situation) are not rational to accept the teachings of any other Ism.

Note that this still does not *require* that Ismists must believe that adherents of other Isms are irrational, though they may believe that.

¹ Peter van Inwagen, “We’re Right. They’re Wrong,” in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford University Press, 2010), 10-28.

Nearly everyone, van Inwagen argues, holds weakly exclusivist views about controversial issues in politics, art, science, philosophy, religion, etc. But what about strong exclusivism? If one were strongly exclusivist, one would hold that

- (1) There are positions inconsistent with my own,
- (2) I am rational to hold the view I do, and
- (3) Anyone else in my epistemic position would *not* be rational to hold any view inconsistent with my own.

Van Inwagen believes that there are lots of strong exclusivists too; he mentions Darwinists, critics of Holocaust deniers, *lots* of political activists, various religious people (regardless of whether their religion actually *teaches* that it is strongly exclusivist), etc.

But philosophers, too, often hold strongly exclusivist views. Consider this principle:

If it is rational for S to accept p, then it cannot also be rational (in the same circumstances) for S to accept $\sim p$.²

Similarly:

If S is deciding whether to accept p, to accept $\sim p$, or to suspend judgment about p, it cannot be true *both* that it would be rational for S to accept p *and* to accept $\sim p$.³

And another:

...if it is rational for a person whose total evidence is E to accept p, then it would be rational for anyone whose total evidence was E to accept p.⁴

Van Inwagen calls this last principle “obviously correct.” These are intuitively compelling principles, accepted by many philosophers. The principles have the following

² Ibid., 20. Paraphrase.

³ Ibid. Paraphrase. This is similar to--though less restrictive than--what Richard Feldman calls the “Uniqueness Thesis,” which is the view that a body of evidence can support at most one among a variety of inconsistent propositions. This thesis will be discussed below.

⁴ Ibid., 22.

implication: philosophers who hold them are committed to the view that all other philosophers with whom they disagree about some contentious proposition are irrational. More precisely, they entail that any philosopher who *shares their exact epistemic situation*, and who still disagrees, is irrational. Such colleagues would, in van Inwagen's words, be "comparable to ordinary, educated people of the present day who believe that cigarette smoking does not cause lung cancer or that the positions of the stars and planets at the moment of one's birth determine one's fate."⁵

Van Inwagen states the argument for these principles as follows:

...a body of evidence cannot have the power to confer rationality on both a proposition and its denial. And...it is rational for one to accept a proposition at a certain moment if and only if one's total evidence at that moment bears this impersonal 'confers rationality on' relation to that proposition.⁶

This argument appeals to a metaphysically necessary property that evidence has *of itself*. Hence, what it is rational to do in a case of disagreement depends on this property of evidence: it cannot "confer rationality on" inconsistent propositions. From this it follows that:

If two people have the same evidence, and if one of them accepts a certain proposition and the other accepts its denial, at least one of them is not rational...⁷

But this just is strong exclusivism. So it seems that weak exclusivism entails strong exclusivism. This is significant because almost everyone holds propositions (especially in philosophy, politics, morality, etc.) in a weakly exclusivist way.

This leaves anyone who holds a view regarding a controversial proposition (i.e., nearly everyone) with three options:

⁵ Ibid., 19.

⁶ Ibid., 20-21.

⁷ Ibid., 21.

- (I) Doubt or deny one's own rationality;
- (II) Claim that at least some others who deny one's own view are less intelligent, less rational, or in some other cognitive way deficient with respect to seeing the truth of p ; or
- (III) Claim that at least some others who deny one's own view are in inferior epistemic circumstances, i.e., lack some evidence that one possesses.

But of course in real-world cases none of these options are very attractive. Van Inwagen is fond of referring to his longstanding disagreement with David Lewis over the compatibility of free will and determinism. It's just obvious, he thinks, that both he and Lewis were rational, in similar epistemic circumstances, comparably intelligent, and so forth. And yet it also *seems* obvious that both he and Lewis were rational to hold the views they did, even after they disagreed and shared evidence. So the argument for the principles above must be mistaken in some way. Yet it is hard to see what is wrong with it. Van Inwagen thus concludes his essay in a state of uncertainty, unable to believe that either he or David Lewis is less rational than the other, and yet unable to see how they could be equally rational with respect to a single evidence base.

But there are questions to be raised in regard to van Inwagen's argument for his principles. His argument, in full, is:

Consider any person as he is at a certain moment. That person has available to him, at that time, a body of evidence, his total evidence at that moment. Call it E . Of metaphysical necessity, E has the following property (of itself, regardless of who may have it) with respect to any proposition p : either (i) it would be rational for anyone to accept p on the basis of E , or (ii) it would be rational for anyone to accept the denial of p on the basis of E , or (iii) it would be not be [*sic*] rational for anyone to accept either p or the denial of p on the basis of E . This 'or' is exclusive. It *cannot* be that it would be rational to accept p on the basis of E *and* rational to accept the denial of p on the basis of E . Loosely speaking, a body of evidence cannot have the power to confer rationality on both a proposition and its

denial. And, finally, it is rational for one to accept a proposition at a certain moment if and only if one's total evidence at that moment bears this impersonal "confers rationality on" relation to that proposition.⁸

In essence, van Inwagen is claiming that a person cannot be justified in believing *p* on the basis of *E* unless anyone else who possessed *E* would also be justified in believing *p* on its basis. The last sentence of his argument also commits him to evidentialism, which holds that rationality is solely a matter of believing what the evidence supports.

This argument has several problems. First, van Inwagen defines "E" as the total evidence that is *available to a person* at a moment. "Available" is vague. Presumably, "available" does not imply that one can *express* all of his evidence, either to himself or to another, since this would rule out private evidence--evidence that cannot be shared with another, such as one's felt sense that *E* supports *p*--which van Inwagen himself seems to think is possible and relevant to whether one is rational to accept *p*. Also, one's private evidence, if such there be, intuitively *would* be a part of one's total evidence pertaining to *p* at a time. So, if *E* does not need to be entirely expressible (and thus includes private evidence), then why think that it must be either rational or (exclusive) irrational for *anyone* to accept *p* on its basis? Presumably van Inwagen thinks that if some bit of private evidence supports *p* for one person, then it would support *p* for any other person who possessed it. This implies that someone else *could* possess another person's *private* evidence. But how could this be? By definition, private evidence cannot be shared. Perhaps van Inwagen may reply here that he is not suggesting that someone *could* possess the private evidence of another, but that *if they could*, then it would support *p* just the same. I'm not sure this is coherent, since it seems to me to be logically necessary that

⁸ Ibid., 20-21.

someone cannot possess the private evidence of another. But even if van Inwagen could somehow avoid incoherence here, his point ceases to be relevant to real-world cases of disagreement, since in such cases no one can in fact obtain anyone else's private evidence.

To put the point a bit differently, imagine that two parties to a dispute share all of their evidence, including the private "seemings" of the form "It seems to me that this evidence supports p." Obviously, they differ in at least two ways: (1) with respect to their positions on p, and (2) in the indexical sense that when one says "It seems to *me* that this evidence supports ~p," the "me" depends on who is saying it. Can they be said to have the *same* evidence in that case? Clearly not, since one has the private evidential seeming that the evidence supports p, and the other has the private evidential seeming that the evidence supports ~p. These are different evidence sets. Since there is no epistemological puzzle about the rationality of believing opposed propositions with different evidence sets, the disagreement problem does not arise here if we allow for private evidence. So insofar as van Inwagen accepts the possibility of private evidence, which he seems to do, he cannot consistently maintain that E must "confer rationality on" only one of the three stated assessments of E (that it supports p, supports ~p, or supports neither p nor ~p).⁹

However, let's assume for the sake of argument that E is the same for the two disputants. Would the claim that E must either support p, support ~p, or support neither p nor ~p follow in this case? This would follow only if one takes E to be *everything that is relevant* to deciding the rationality of one's doxastic attitude toward p. But this is too

⁹ As we will see in the next chapter, van Inwagen recognizes the importance of private evidence, and comes close to admitting the argument in these last two paragraphs. However, he ultimately rejects the ability of private evidence to allow for reasonable disagreement for reasons that do not withstand scrutiny.

strong--there are considerations other than E itself which are relevant to the rationality of one's assessment of p, including (1) factors that are "unavailable" to one in the moment of a disagreement (and so not a part of E by van Inwagen's definition of E), and (2) one's interpretation of E, which is not itself reducible to E. A word of explanation about both (1) and (2) follows.

(1) As van Inwagen defines E, it clearly does not include any consideration pertinent to p that is *unavailable* to one at a moment. But it seems that considerations that are *not* available to one can make a difference with respect to what it is rational for one to believe about p. Consider an atheist and a theist disagreeing over a particular miraculous claim. The theist recalls believing on good evidence in the past that there is a successful refutation of Hume's argument against the rationality of belief in miracle reports, even though he cannot remember what the refutation is or where he learned it. This information is unavailable to him at the moment of the disagreement, and yet it seems relevant to the rationality of his assessment of the miracle claim. Alternatively, one may be in an epistemically lucky situation without realizing it, say by asking for directions in a neighborhood where people tell the truth about such things, but in a city where most people lie. This would make a difference with respect to the rationality of one's believing the directions one is given. Or, perhaps someone could have a "neural quirk," to borrow a phrase of van Inwagen's, that makes them better at discerning the truth about p than other people, though they don't know this about themselves. These possibilities suggest that factors other than what is included in E could be relevant to whether it is rational for one to believe p on the basis of E.

(2) Van Inwagen states that the ability of a body of evidence to “confer rationality on” the doxastic attitude that one takes toward a proposition is a metaphysically necessary property of the evidence *itself*, “regardless of who may have it.” What could this mean? Presumably, van Inwagen has in mind the objectivity of justification mentioned above: to be justified in believing *p* on the basis of *E*, *E* must objectively support belief in the truth of *p*--i.e., anyone else who possessed *E* would also be justified in believing *p* on its basis. But is this the right way to view how evidence supports a proposition? Jason Decker has argued, convincingly in my view, that it is not.¹⁰

Decker describes this view as the “pan balance model of evidential support,” because it assumes that we determine whether an evidence set supports a proposition in the same way that we determine which of two objects weighs more by putting the objects on a scale. We drop some evidence into the pan on one side of the scale, which represents *p*, and some into the pan on the opposite side of the scale, which represents $\sim p$, and the lowest pan at the end of this process is what the total evidence supports.¹¹ However, as he says,

It is at best elliptical--and, at worst, a category mistake--to say that a body of evidence supports a proposition. A body of evidence isn't the right sort of *thing* to support a proposition.... only an *interpreted body of evidence* can serve such a role.¹²

Decker gives two examples to illustrate his point. The first involves someone having an identical total evidence set with his counterpart on Hilary Putnam's Twin Earth: “For

¹⁰ Jason Decker, “Disagreement, Evidence, and Agnosticism,” *Synthese* 187, no. 2 (2012): 753–83. In the portion of the paper that I am referencing here, Decker is discussing Roger White's and Richard Feldman's “Uniqueness Thesis,” which I discuss in the next section. However, what he says applies just as well to van Inwagen's claim here about evidential support.

¹¹ *Ibid.*, 777-778.

¹² *Ibid.*, 778.

every experience you have had, your twin has had a phenomenally identical experience.”¹³ Now suppose that you come to think that your experiences are caused by mind-independent “material” objects, while your Twin Earth counterpart comes to adopt a Berkeleyan perspective. Decker thinks it is obvious that both you and your counterpart could be *rational* in your views, even if only one of you may be correct. However, van Inwagen’s view of evidence would insist that at most one of you could be rational.

Alternatively, Decker offers another case: two people are looking, under the same conditions, at an image like the famous duck-rabbit, which may be seen in two ways.¹⁴ One forms the belief that what they are looking at is only a duck; the other forms the belief that what they are looking at is only a rabbit. According to van Inwagen, at least one of them is irrational. But surely neither is irrational. “It’s not a failure of rationality,” Decker says, if either of the people fails “to be lucky enough to see that the [image] can be carved up in a second way.”¹⁵ The reason is that evidence cannot tell us how to interpret it. Decker goes on to give a somewhat crude description of what he means by “interpretation,” but this need not concern us here. The upshot is that van Inwagen’s view of evidence ignores the fact that evidence must be interpreted in order to support a proposition, and that therefore we cannot say that evidence can have support-properties “of itself.” Further, it won’t help to say (as, for example, Richard Feldman has said) that one’s interpretation of the evidence should itself be supported and hence become part of E, because (a) this won’t always be possible, as in the duck-rabbit case, and (b) even if it

¹³ Ibid.

¹⁴ I’ve modified this case slightly for ease of exposition.

¹⁵ Ibid., 779.

were possible to find some extra evidence for one's interpretation of E, *that* evidence would also need to be interpreted.

So, as stated, van Inwagen's argument that evidence can confer rationality on only one view will not work. Given that he offers no further support for this view, I will here set it aside, and turn to a closely related view of Richard Feldman's.

1.2 – Richard Feldman's Uniqueness Thesis

in his widely influential 2007 paper "Reasonable Religious Disagreements," Richard Feldman begins by noting that many people--e.g., his own students--take a tolerant attitude toward those with whom they disagree.¹⁶ They seem to think they can "agree to disagree" about an issue, and that all parties can remain reasonable. But what does this attitude amount to?¹⁷ Feldman notes two ways of *failing to* reasonably disagree: intolerance (viewing one's interlocutors as opponents to be beaten or mocked), and relativism (viewing all sides of a dispute as somehow equally correct). Each fails to take seriously the views of their interlocutors. So what sense is there to "reasonable disagreement"?

Feldman distinguishes two questions:

Q1 Can epistemic peers who have shared their evidence have reasonable disagreements?

¹⁶ Richard Feldman, "Reasonable Religious Disagreements," in *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, ed. Louise Antony (Oxford University Press, 2007), 194–214.

¹⁷ For what it's worth, Feldman describes himself as a "complacent atheist" and admits that this issue "challenges that complacency." *Ibid.*, 195. The emphasis here on complacency is not unique to Feldman; Gideon Rosen also speaks in these terms. See below, especially Chapter Four.

Q2 Can epistemic peers who have shared their evidence reasonably maintain their own belief yet also think that the other party to the disagreement is also reasonable?¹⁸

Feldman thinks that affirmative answers to these questions cannot be defended. “As will become apparent,” he writes, “open and honest discussion seems to have the puzzling effect of making reasonable disagreement impossible.”¹⁹

In this context, the cases under consideration are restricted to those in which the interlocutors hold that each party’s specific beliefs are *reasonable*, and not merely of practical value, or tied to desirable behavior. With the case so restricted, Feldman distinguishes four ways that reasonable disagreement might be possible.

First, different conclusions can be drawn from the same evidence. For example, consider the case of Righty and Lefty. A detective has strong evidence incriminating Righty and (presumably equally) strong evidence incriminating Lefty of the same crime.²⁰ Further, the detective knows that only one suspect could be guilty (though, presumably, he does not know that it must be either Righty or Lefty). One may think that the detective would be reasonable to draw *either* the conclusion that Righty is guilty *or* the conclusion that Lefty is guilty. But, Feldman argues, this isn’t right: any decision by the detective in this case is arbitrary; the reasonable thing to do is to suspend judgment.

This suggests what Feldman calls the Uniqueness Thesis:

UT: ...a body of evidence justifies at most one proposition out of a competing set of propositions...and...it justifies at most one attitude toward any particular proposition.²¹

¹⁸ Ibid., 201.

¹⁹ Ibid., 202.

²⁰ Ibid., 204-205.

²¹ Ibid., 205. Feldman thinks that “our options with respect to any proposition are believing, disbelieving, and suspending judgment.”

Feldman takes this to rule out a “yes” answer to Q1. As he says, “If The Uniqueness Thesis is correct, then there cannot be any reasonable disagreements in cases in which two people have exactly the same evidence.”²²

Second, reasonable people may weigh evidence differently because of differing initial assumptions about evidence. Imagine, for example--this is my case, not Feldman’s--that two jurors are analyzing four pieces of evidence in a murder trial: (1) the murder weapon with the defendant’s fingerprints, (2) the known fact that the defendant hated the victim, (3) the known fact that the victim was depressed, and (4) a suicide note apparently in the victim’s handwriting. It seems as if this body of evidence is consistent with both the proposition “The defendant is guilty” and the proposition “The defendant is innocent,” and further that a juror could be justified in believing *p* (guilt) or in believing not-*p* (innocence) on the basis of the same evidence. This would be because the evidence admits of multiple consistent relative weights, depending on how the jurors weigh handwriting analysis versus fingerprint analysis. Thus, rational disagreement with a single evidence base seems plausible.

However, Feldman argues that this only pushes the disagreement back a step: now the issue is which initial assumptions are true--in this case, which actually *is* more reliable, fingerprint analysis or handwriting analysis? Feldman says that the parties should share their evidence about *that* and come to a conclusion:

These different starting points help support the existence of reasonable disagreements only if each side can reasonably maintain its starting point after they have been brought out into the open. And this idea can support the tolerant attitude...only if people can think that their own starting point is reasonable and that different and incompatible starting points are reasonable as well. I cannot

²² Ibid.

understand how that could be true. Once you see that there are these alternative starting points, you need a reason to prefer one over the other.²³

The third way that reasonable disagreement might be possible is that both parties to a disagreement can have a bit of unshared, or “private,” evidence that makes their view reasonable. Perhaps, Feldman suggests, this is what we suspect is going on when we take the tolerant attitude. As he says,

If we take (Q1) and (Q2) to be about cases in which all the evidence is shared, then the answer to both questions is ‘no.’ But if we take the questions to be about cases in which the evidence is shared as fully as is realistically possible, then the answers are ‘yes.’ We might say that the reasonable disagreements are possible in those cases in which each side has private evidence supporting its view.²⁴

But Feldman argues that this won’t work either. There is an analogy here with perception: imagine that two people are facing a quad; one sees the dean standing there and the other doesn’t. The reasonable thing here, Feldman thinks, is to suspend judgment once the parties realize that something strange is going on, not to conclude that they’re having a reasonable disagreement owing to private evidence.²⁵ As he says:

Each may have his or her own special insight or sense of obviousness. But each knows about the other’s insight. Each knows that this insight has evidential force. And now I see no basis for either of them justifying his own belief simply because the one insight happens to occur inside of him. A point about evidence that plays a role here is this: evidence of evidence is evidence. More carefully, evidence that there is evidence for P is evidence for P. Knowing that the other has an insight provides each of them with evidence.²⁶

²³ Ibid., 206. One may reasonably wonder here what the evidence for “starting points”--like the principles about the reliability of handwriting and fingerprint analysis--would be like. Eventually, it may be that all we have to go on in preferring our own starting points is an intuitive “feel” or “seeming.” This will be discussed below.

²⁴ Ibid., 207. Private evidence will be discussed in more detail in the next chapter.

²⁵ For what it’s worth, my own intuition about this case is that, in the absence of further evidence, I should not consider my interlocutor an epistemic peer with respect to whether the dean is in the quad--something has clearly gone wrong with one of our perceptions, and *ex hypothesi*, I have no reason to assume the problem is with me.

²⁶ Ibid., 208. These points will be explored further in Chapter Six. See Chapter Two for a critique of Feldman’s argument against private evidence.

Fourth and finally, perhaps it is possible to *have* a reasonable disagreement (so, an affirmative answer to Q1), but not to *think* that one is having one (so, a negative answer to Q2). As Feldman says:

One way or another, then, the conclusion drawn is that the other person does not have a reasonable or justified belief. And the idea is that *both* parties to the disagreement can reasonably draw this conclusion. Thus, both parties have a reasonable belief, yet they reasonably think that the other side is not reasonable.²⁷

But, again, any reasons for thinking that one is cognitively superior can themselves be discussed and become part of the shared evidence. So here again, the best option would be to suspend judgment, not to insist on reasonable disagreement.

Apparently, then, there cannot be reasonable disagreements. So what should we say about situations where it seems like there are? Feldman says there are two options:

(1) The “hard line”: evidence can only support one view, and so whoever takes that view is reasonable, and everyone else is not. Feldman thinks the hard line view is probably appropriate in many cases (e.g., astrology, climate change, evolution, racism, etc.), but it is apparently inappropriate in many others (e.g., most complex moral, political, scientific, and religious disputes). Take, for example, the question of the existence of God: the theist and atheist may both take the hard line view and conclude that the evidence supports only her view, and so the other is mistaken and unreasonable. But is there any reason to suppose that one is more correct to do so than the other? Feldman thinks not.

(2) The skeptical alternative: we should suspend judgment in cases of peer disagreement, i.e., become agnostics about a great many things. This is Feldman’s own

²⁷ Ibid., 209.

stated view in this article (in some of his later work, he backs away from this conclusion).²⁸ Feldman insists that this is “a modest view, in the sense that it argues for a kind of compromise with those with whom one disagrees.” But it also:

implies that one should give up one’s beliefs in the light of the sort of disagreement under discussion. This is a kind of modesty in response to disagreement from one’s peers. This is also a skeptical view, in the limited sense that it denies the existence of reasonable beliefs in a significant range of cases.²⁹

Because Feldman here advocates the “skeptical” response to peer disagreement, he is usually classed as a “Conformist,” one who thinks that the rational response to peer disagreement, once evidence has been shared, is to suspend judgment about *p*. Several points Feldman makes here will be discussed in later chapters; here I want to focus on the Uniqueness Thesis (UT), which is very similar to van Inwagen’s view of evidence noted above.

First, what can be said in support of UT? In Thomas Kelly’s discussion of it in “Peer Disagreement and Higher-Order Evidence,” he mentions UT while discussing the “Equal Weight View.”³⁰ The Equal Weight View recommends that in cases of peer disagreement with shared evidence, one should accord *equal weight* to one’s own view and the view of one’s peer--i.e., one should split the difference. Kelly describes a case in which two parties consider evidence *E* and how well it supports hypothesis *H*. Assume that we can assign a numeric value to our confidence in a proposition such that 1 represents certainty that the proposition is true, 0 represents certainty that it is false, and

²⁸ See Richard Feldman, “Evidentialism, Higher-Order Evidence, and Disagreement,” *Episteme* 6, no. 3 (2009): 294–312.

²⁹ Feldman, “Reasonable Religious Disagreements,” 213.

³⁰ Thomas Kelly, “Peer Disagreement and Higher-Order Evidence,” in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford University Press, 2010), 111-74.

.5 represents suspension of judgment about the proposition--i.e., one takes the evidence to be 50/50 for and against that proposition. Party A is reasonable, let's say, to accord H a credence of 0.7, and moreover, A knows that it would also be reasonable to be slightly less confident in H on the basis of E, say, a credence as low as 0.6. Party B accords H a credence of 0.6 on the basis of E, but knows that a slightly higher credence would also be reasonable, say, as high as 0.7. Now suppose the parties realize their differing credence levels. It seems that the Equal Weight View would recommend that they split the difference and opt for a shared credence of 0.65. But, says Kelly, there's no reason at all to do this, since they both recognize that the other's credence is a reasonable response to the evidence. In order to rule out the possibility of cases like this and thereby save her view, he argues, the Equal Weight Theorist will need to adopt UT.³¹ As he says:

Suppose that the Uniqueness Thesis is true. Then, if it is in fact reasonable for me to give credence 0.7 to the hypothesis, it follows that you are guilty of unreasonable diffidence for being even slightly less confident. On the other hand, if you are reasonable in being slightly less confident than I am, then I am guilty of being unreasonably overconfident. Hence, the description of [the case] offered above is incoherent; [that] is not in fact a possible case.³²

But, Kelly says, the case just described makes UT highly counterintuitive, since if one has already allowed for some wiggle room in what she takes the reasonable credence for H to be, then there is no reason for one to alter her view, UT notwithstanding. In order to maintain her view, then, the Equal Weight Theorist must adopt a highly counterintuitive principle (UT), which seems ad hoc. But some of Kelly's phrasing here suggests a way to alter UT to make it less counterintuitive. For example, he says:

³¹ Whether or not this is a fair assessment of the Equal Weight View is not at issue here. My present purpose is merely to consider the merits of the Uniqueness Thesis itself.

³² *Ibid.*, 119-120.

Clearly, the Uniqueness Thesis is an extremely strong claim: for any given batch of evidence, there is some one correct way of responding to that evidence, any slight departure from which already constitutes a departure from perfect rationality.³³

Later, he says:

Let us suppose for the sake of argument, then, that the Uniqueness Thesis is correct: for a given batch of evidence, there is some one way of responding to that evidence that is the maximally rational way.³⁴

The emphasis on “perfect” or “maximal” rationality here is suggestive. If we view rationality as a degree concept, then the door is open to notions such as “acceptably rational, though not perfectly or maximally rational,” as responses to evidence. Along these lines, we can adjust UT in the following way:

UT*: For any batch of evidence E, there is at most one maximally rational assessment of that evidence.

This reformulation of the Thesis doesn’t entail a single rational view and a host of other equally irrational views. Rather, there is a single most rational view, and a variety of other less rational views. One might also build into this some constraint regarding the natural limitations on human cognition to determine exact, fine-grained distinctions of warranted credence. Perhaps, for example, it would be inappropriate to expect even careful thinkers to correctly ascertain evidential assignments of warranted credence between 0.65 and 0.7. Thus, the proponent of a graded, differential version of the Uniqueness Thesis could build in allowances for human limitations, perhaps by rephrasing the Thesis as:

³³ Ibid., 120.

³⁴ Ibid., 122. The “for the sake of argument” here is worth noting, as Kelly in fact rejects the Uniqueness Thesis, though he does not explicitly argue against it in this paper. His focus in this portion of “Peer Disagreement and Higher-Order Evidence” is refuting the Equal Weight View. See his comments to this effect on p. 121.

UT**: For any batch of evidence E, there is at most one maximally rational range of assessments of that evidence, consistent with the reasoning capacities of human thinkers.

Such a modified UT seems to avoid the sort of problem raised by Kelly's case.

Unfortunately, it is not available to Feldman. Note that while UT** avoids the counterintuitive implication that all views but one are automatically irrational, it still has the consequence that all assessments of any evidence set that fall outside of a certain range of credence levels regarding p are automatically suboptimal. The reasonableness of this view depends on how narrow this range is. Unfortunately for Feldman, the range would have to be quite narrow in order to justify his view that one should suspend judgment about all issues disputed by epistemic peers. UT** simply isn't restrictive enough to support this conclusion, since it is plausible that two careful thinkers could come to different conclusions about a set of evidence that are both within the maximally rational range of interpretations of that evidence. This is what Kelly's case shows.

What about the problems raised above for van Inwagen's closely related view?

Does UT** avoid those? Recall that the problems with that view are:

- (1) Considerations that are "unavailable" to one in the moment of a disagreement can be relevant to the rationality of one's assessment of p.
- (2) Evidence must be interpreted in order to support a proposition; it has no metaphysically necessary *intrinsic* support-properties.

With respect to (2), UT** does not include or depend on the claim that evidence has intrinsic support-properties. Rather, the question of *why* there is at most one maximally rational range of assessments of an evidence set is left open.

With respect to (1), if UT** were to be used in the way that Feldman and van Inwagen use UT--i.e., as a way to rule out the possibility of rational disagreement--then it would fare no better than UT, since in order to support that conclusion, one must also assume evidentialism, which both Feldman and van Inwagen do. But it is not necessary to use UT** in this way, and on its own it says nothing about whether or not considerations *other than* one's total evidence may play a role in the rationality of one's doxastic state regarding p.

But there is still something counterintuitive about UT**. Say that it seems to me that argument X supports p, and I assign p a credence of 0.7. However, I might also recognize that owing to the limitations on how much weight I should accord my "seemings" about such things, I may be off by as much as 0.1. It seems to you also that X supports p, but you assign it a credence of only 0.6. But you also realize the same limitations about your own "seemings," and so you accept a range of 0.1 higher as still consistent with the evidence, which in this case is exhausted by X. According to UT** we may both be maximally rational in this case. This would not be allowed on UT or UT*. However, if either of us has a private sense that is not extremely close to the private sense of the other (within, say, a 0.1 credence assignment range), then UT** requires that at most one of us can be maximally rational, even if it seems to both of us that we are *both* doing everything that we should be doing in assessing the evidence that we have. Applied to real-life disagreement cases, this is problematic, since often neither of us will be in a position to know which of our private seemings is more accurate. If neither of us can see how either of us could be doing any better than we are at assessing our evidence, it seems odd to be forced to conclude that one or both of us is doing worse than we could

with the evidence we have. Instead, we should conclude that we're both within the maximally rational range. But if this is possible, then UT** cannot do the work that Conformists like Feldman need it to do, since it would allow for the possibility of more than one maximally rational conclusion on the basis of a single set of evidence.

The Conformist may reply here that the distinctions we're talking about are so fine (0.7 vs. 0.6) that for most real-life issues, we'd likely consider that an *agreement*, and so UT** would be irrelevant. But of course we could easily think of situations where that fine of a difference would *make a difference*. It often happens that the divergence of expert opinion on some issue is very slight, and yet it constitutes real disagreement between them. It seems in such cases that we should allow that the experts who differ by a slight amount between themselves may all have beliefs within the maximally rational range allowed by UT**. The only motivation for saying otherwise would seem to be the requirement that to be maximally rational in assessing our evidence, our assessment must be *accurate*. But of course in a real case of peer disagreement, we won't be in a position to know this. So UT** cannot be used to decide between the private seemings of two disagreeing peers, which was the purpose of UT from the beginning.

For his part, van Inwagen does offer an argument for an evidentialist interpretation of rationality, which he borrows from W.K. Clifford. In his essay "Listening to Clifford's Ghost," van Inwagen discusses Clifford's famous essay "The Ethics of Belief."³⁵ He asks, "Why does Clifford think that it is wrong to believe things upon insufficient evidence?" And he answers that:

³⁵ Peter van Inwagen, "Listening to Clifford's Ghost," *Royal Institute of Philosophy Supplement* 65 (2009): 15–35. Also see W.K. Clifford, "The Ethics of Belief," in *The Ethics of Belief and Other Essays*, ed. Timothy Madigan (Amherst, MA: Prometheus, 1999), 70–96.

The central nerve of Clifford's reason for supposing this is...: Of all the forces in the human psyche that direct us toward and away from assent to propositions, only rational attention to relevant evidence *tracks the truth*. Believing things *only* upon sufficient evidence is, therefore, the only device we have for minimizing the extent of our false beliefs, or at least the only such device that has any prospect of providing us with a useful set of true beliefs... If we form our beliefs on any other basis – if we allow them to be formed by some factor that does not track the truth –, we are, in effect, believing things at random... Since there are a lot more ways to be wrong than there are to be right, beliefs formed by a method that does not track the truth will, to a high probability, be false.... A person who believes things upon insufficient evidence, therefore, is not taking care to minimize the extent of his false beliefs.³⁶

Although Clifford is less concerned with being rational than with being ethical, since his goal is to show that it is morally wrong to believe anything on insufficient evidence, not merely that it is irrational to do so, his argument applies to rationality as well as morality. This is because for Clifford, the only alternative to truth-tracking is randomness. It thus follows that rationality entails truth-tracking, since believing things at random is clearly not truth-tracking. To avoid the philosophical baggage associated with the term “truth-tracking” associated with Robert Nozick, I'll instead use the phrase “based on evidence.”³⁷ So a reformulated, Clifford-inspired maxim for rationality would be: Rationality requires holding beliefs which are based on evidence. But is it right that one is forced to choose between forming her beliefs in accordance with evidence or else at random? I think not. The literature on Clifford is vast, and full discussion of it would take us too far afield, but I'll briefly point out a few ways that one may respond to Clifford's argument.³⁸

³⁶ Van Inwagen, “Listening to Clifford's Ghost,” 33.

³⁷ See Robert Nozick, *Philosophical Explanations* (Cambridge, MA: Harvard University Press, 1981).

³⁸ For a nice overview of the topic, see Andrew Chignell, “The Ethics of Belief,” in *The Stanford Encyclopedia of Philosophy*, Spring 2018, <https://plato.stanford.edu/archives/spr2018/entries/ethics-belief/>.

First, it can be argued that it may be rational to believe not on the basis of evidence. The basic idea here is that sometimes evidence may not point decisively in one direction or another, and that when this is the case it may be rational--if certain conditions are met, and contra various commentators in the disagreement debate³⁹--to *choose* what to believe based on desires or practical concerns other than knowing the truth. This is a straightforward denial of the maxim above. William James famously responded to Clifford in just this way. He argued that it is neither immoral nor irrational to choose what to believe on such a basis.⁴⁰ If we go further and argue that it can in fact be positively rational to choose on such a basis, given one's overall epistemic goals, then we can reject Clifford's dilemma: evidence-based or irrational.⁴¹

Another avenue of response is to argue that non-evidential motivations may be rational in the truth-capturing sense.⁴² For example, it may be that the desire for goodness is itself productive of true beliefs more often than not, so that when one believes with the end of the good in view, one is *eo ipso* believing in a truth-capturing way, even though one's goal in believing is not truth *per se*.⁴³

Thirdly, one may argue that "evidence" is broader than Clifford imagined, so that aligning one's beliefs strictly in proportion to the evidence in his narrow sense is

³⁹ See, e.g., Feldman, "Reasonable Religious Disagreements," 204-205.

⁴⁰ See William James, *The Will to Believe and Other Essays in Popular Philosophy* (New York: Dover Publications, 1956), 1-31.

⁴¹ For a contemporary example of someone who argues this way in the context of the disagreement debate, see Marc Moffett, "Reasonable Disagreement and Rational Group Inquiry" *Episteme: A Journal of Social Epistemology* 4 (2007): 352-367. I will discuss Moffett's view at length in Chapter Eight.

⁴² James sometimes speaks this way as well, e.g.: "to say... 'Do not decide, but leave the question open,' is itself a passional decision,--just like deciding yes or no,--and is attended with the same risk of losing the truth." See James, "The Will to Believe," 11.

⁴³ This would be an example of a non-evidentialist, reliabilist approach to rationality. For an overview of reliabilism, see Alvin Goldman and Bob Beddor, "Reliabilist Epistemology," in *The Stanford Encyclopedia of Philosophy*, Winter 2015, <https://plato.stanford.edu/entries/reliabilism/>.

epistemically wrong-headed. On such a view, contra much of the modernist tradition, desires and natural inclinations would be evidential, and therefore beliefs formed on their basis may be rational.⁴⁴ This would be a way of affirming the maxim stated above, while insisting that it should be interpreted more broadly than Clifford does.

Fourth, one may hold that evidence is not required *at all* for at least some beliefs to be held rationally. This would be an even stronger denial of Clifford's maxim than the one James tendered, since his reply at least assumed that the evidence had been assessed and found indecisive. On this view, exemplified by so-called "reformed epistemology," sometimes no evidence is needed for rationality. Primary examples of such rationally held, non-evidential beliefs include belief in other minds, belief in an external world, basic perceptual beliefs, and more controversially, belief in God.⁴⁵

All of these proposals are reasonable responses to the Clifford-inspired argument offered by van Inwagen, and on any of them, the line between rationality and arbitrary belief is not where Clifford takes it to be.

Thus, for now it seems that no version of the Uniqueness Thesis will do what van Inwagen and Feldman want it to do: rule out the possibility of opposed rational responses to the same body of evidence.

But what does it mean to have the "same body of evidence" in the first place? Is full disclosure of one's total evidence possible, or is there some sense in which some evidence can remain private? This will be the topic of the next chapter.

⁴⁴ In certain places, Kierkegaard seems to take a view like this.

⁴⁵ For helpful overviews of reformed epistemology, see Peter Forrest, "The Epistemology of Religion," in *The Stanford Encyclopedia of Philosophy*, Summer 2017, <https://plato.stanford.edu/archives/sum2017/entries/religion-epistemology/>, Section 7; and Anthony Bolos and Kyle Scott, "Reformed Epistemology," in *Internet Encyclopedia of Philosophy*, n.d., <https://www.iep.utm.edu/ref-epis/>.

Chapter Two – Private Evidence

2.1 – Van Inwagen on the Use of Private Evidence

Can evidence always be fully disclosed, or is there such a thing as private evidence? Van Inwagen suggests some examples of such private evidence that he says seem “plausible”:

I sometimes know that my wife is angry when no one else does, for example, and I cannot explain to anyone how I know this--I cannot give what Plato would call an “account” of what underlies my conviction that she is angry. It seems to me to be plausible to say that in such cases my belief that my wife is angry is grounded in some body of evidence, evidence that lies entirely within my mind and that I cannot put into words. A second example is provided by the case of the chicken sexer, beloved of epistemologists in the far-off days of my graduate studies... Mathematicians are often intuitively certain that some mathematical proposition is true, although they are unable to prove it. (Gödel, I understand, was convinced that the power of the continuum was \aleph_2 .) Since they often later do discover proofs of these propositions, it seems likely that, prior to their discovery of the proofs, they had some sort of evidence for the truth of those propositions.¹

Van Inwagen considers private evidence in the context of his disagreement over free will with David Lewis. Given that Lewis was obviously rational, and that he had all the same public evidence (arguments, distinctions, and the like) as van Inwagen, it follows, given van Inwagen’s assumptions about evidence, that if the evidence did not point Lewis toward incompatibilism, then it didn’t point van Inwagen toward it either. That is, of course, unless van Inwagen had some evidence that Lewis lacked, of a sort that could not be disclosed. Van Inwagen here admits that if this is even possible, then his earlier attempt to demonstrate that strong exclusivism follows from weak exclusivism fails,

¹ Peter van Inwagen, “We’re Right. They’re Wrong,” in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford University Press, 2010), 25.

since if one has private evidence, then one needn't hold those with whom one disagrees to be irrational--they simply, through no fault of their own, lack some evidence.²

It is worth pausing for a moment here over the connection between private evidence and intuition. As described by van Inwagen, private evidence is a matter of its seeming to one that something is the case, in a way that one is unable to express other than with a locution of the form "It seems to me that..." Is this an intuition? I think it is, in the sense of a pre-reflective judgment that something is the case. In that way, seeming-intuitions are *intellectual* in that they are judgments with propositional content. In Chapter Three, I will argue that, in the disagreement literature, seemings that some proposition is true or false are treated as analogous to perceptual seemings, such as the seeming that there is a green expanse in front of me. One may argue that this makes seemings different from intuitions. For example, Joel Pust objects to the following definition of intuition on the grounds that it is "indiscriminate":

I: S has the intuition that p if and only if it seems to S that p .³

This is indiscriminate in the sense that it does not distinguish between intellectual and perceptual seemings, and therefore implies that there are perceptual intuitions. According to Pust, "Clarity is served by stipulating that such states are not intuitions even if they are all species of some common genus."⁴ Consequently, he prefers this definition of intuition:

² One may reasonably wonder whether it is possible for two people to have the same private evidence, in the sense that two competent chicken sexers could both have a private "seeming" that the chicken is a male. As I understand it, private evidence is indexical, so while two people may have similar private evidence, they cannot have the *same* private evidence. See Chapter One, pp. 14-15. It remains open, however, that one's private evidence may be *epistemically* equivalent to another's. As we will see below, private evidence in the form of propositional "seemings" bears a strong resemblance in this sense to perceptual seemings.

³ Joel Pust, "Intuition," in *The Stanford Encyclopedia of Philosophy*, Summer 2017, <https://plato.stanford.edu/archives/sum2017/entries/intuition/>, Section 1.3.

⁴ *Ibid.*

I*: *S* has the intuition that *p* if and only if it intellectually seems to *S* that *p*.⁵

While I have no philosophical objection to distinguishing between perceptual and intellectual seemings in this way, for the purposes of clarity in what follows, I will use the term “seeming” in the indiscriminate sense identified above. Nothing in my argument will turn on whether perceptual seemings are intellectual. Consequently, I will also generally avoid use of the word “intuition,” since it is not my aim here to wade into the murky waters of whether intuition is doxastic, dispositional, or *sui generis*, or whether intuitions should include perceptual states. In my usage, “seeming” can bypass such issues.

In addition, intuitions/seemings are often treated as evidence by philosophers, van Inwagen being a case in point. As Pust says:

... the view that intuitions are treated as evidence in philosophy is best thought of as the view that, with respect to many core questions of philosophy, our justification for believing an answer consists (at least substantially) in our having suitable intuitions.... Put simply, the view is that the occurrence of an *intuiting* is taken to provide the person in whom it occurs with *prima facie* justification for believing the *intuited*. Alternatively, it holds that *S*'s having an intuition that *p* *prima facie* justifies *S* in believing that *p*. Given a capacious view of evidence, then, both the *intuiting* and the *intuited* are treated as evidence, the occurrence of the former being treated as *prima facie* non-inferential justification for accepting the latter and the latter as a potential inferential justification for further propositions.⁶

I will follow this common usage here, treating both the having and the content of a seeming as evidential in the sense that they provide *prima facie* justification for the intuited proposition.

⁵ Ibid.

⁶ Ibid., Section 2.4.

So perhaps van Inwagen has some private evidence for the truth of incompatibilism--the sort of evidence that is exemplified by knowing when one's spouse is angry, or by "chicken sexers," or by mathematicians who intuit mathematical propositions prior to proving them, etc. Maybe so, but this is ultimately an unsatisfying response.

Consider the implications: surely not *all* philosophical (or other) disagreements are due to differing private evidence. As van Inwagen says:

Am I to believe that in every case in which I believe something many other philosophers deny (and this comes down to: in every case in which I accept some substantive philosophical thesis), I am right and they are wrong, and that, in every such case, my epistemic circumstances are superior to theirs? Am I to believe that in every such case this is because some neural quirk has provided me with evidence that is inaccessible to them? If I do believe this, I must ask myself, is it the same neural quirk in each case or a different one? If it is the same one, it begins to look more a case of "my superior cognitive architecture" than a case of "accidental feature of my cognitive architecture." If it is a different one in each case--well, that is quite a coincidence, isn't it?⁷

I take it that "neural quirk" here is a stand-in for all forms of private evidence that one might have, even though, technically speaking, a neural quirk would be the a *cause* of private evidence, not private evidence itself. Ultimately, van Inwagen cannot believe that he has such private evidence in every case in which he disagrees with a peer.⁸ Since he cannot otherwise justify his superiority to Lewis, or Lewis's to him, he is left with a dilemma: either he and Lewis are both irrational, or both rational. It's unlikely to be the former, but he also thinks it cannot be the latter, owing to the arguments given above about the nature of evidence. The implication here is that we should become agnostics

⁷ Peter van Inwagen, "We're Right. They're Wrong," 27.

⁸ Van Inwagen does not use "peer" language, though it is implied that Lewis is his epistemic equal with respect to the question of free will.

regarding a topic or proposition about which there's a dispute among the informed. But we can't do that--literally, are unable to--and so we're stuck believing that we're rational and that our disputants often are as well, although we have no way to justify this.

Fortunately, given the critique of the Uniqueness Thesis advanced above, we need not follow van Inwagen into this dilemma, as it remains an open question whether a body of evidence may support more than one conclusion.

But what about the possibility of private evidence? Is it reasonable to believe that one's view is correct while the view of one's peer is mistaken, because one has some private evidence that this is so? Two questions arise here:

- (1) Just what *is* this private evidence? What is its nature? Can we say anything more about it that will help us understand why it is not communicable, and/or how it provides justification?
- (2) Is it ever reasonable to rely on private evidence in a dispute with an epistemic peer?⁹

To my knowledge, (1) has not been explicitly discussed at length in the disagreement literature.¹⁰ It is often set aside, as in van Inwagen, in favor of focusing on (2). This is understandable if, whatever private evidence turns out to be, it is never reasonable to rely on it when disagreeing with a peer. In that case the concept of private evidence is of no use to the peer disagreement debate. But the "if" of two sentences back is a big one. So I think that (1) is worth considering in its own right. I will do so in the next chapter via a comparison with perceptual seeming, which is used throughout the professional literature

⁹ Related to this is the question whether non-arbitrary criteria can be given for the *application* of private evidence as a way to break the apparent epistemic symmetry between peers.

¹⁰ Even describing its relation to intuitions, as we did above, goes beyond what the major players in the disagreement debate have said, with the exception of Gideon Rosen, whose views we will discuss below. By and large, they are more focused on what one may reasonably do with one's seemings than on the seemings themselves. I think this is an oversight.

on disagreement as an analogy with what I call *evidential seeming*--the seeming to one that evidence E supports proposition p.

For the present, however, let's consider (2) more closely. Here is van Inwagen's argument for a negative answer to (2), paraphrased:

Private evidence cannot be sufficient to justify a preference for one's own view in *every* case of belief in a disputed proposition.¹¹ This is because one would be left with a dilemma if this were the case: either the cause of one's private evidence is the *same* in each dispute, or the cause is *different* in each dispute. If it is the same, then this is equivalent to saying that one is simply cognitively superior; i.e., it is equivalent to denying that *any* disputant is one's peer. If it is different in each dispute, then this is too unlikely to be believed, since one would have to be supremely lucky to have been born with access to all the private evidence one would need to form all the correct beliefs about everything.¹²

One obvious response to van Inwagen here is to say that private evidence need not justify a preference for one's own view in *every* case, but rather only in those cases in which one has done all she can to share her evidence and to consider that of her peer, but to no avail, and yet it still *really seems* to her that she is correctly assessing the total evidence. Perhaps van Inwagen's and David Lewis's disagreement over free will was just such a situation. But how often do philosophers (or anyone else) take that kind of time to really get to know one another and hash out an issue for a period of years? I suspect, though I have no hard evidence on the matter, that this is rare, and that therefore full disclosure even of the evidence that is *shareable* with respect to some dispute, is rare. If this is right, then perhaps van Inwagen's dilemma is not so bad, since a person would

¹¹ In the context of van Inwagen's discussion, the propositions in question are substantive philosophical theses, and the disputants are other philosophers, but we can safely extend his argument to other sorts of disputed propositions. Also, for the time being, I will limit the class of disputants to *experts*, since in those cases, differences of opinion cannot so easily be chalked up to a lack of sharable information.

¹² One could, I suppose, bite this bullet. I once heard Richard Swinburne, known both for his uses of Bayesian probability theory in the examination of the existence of God and religious belief and for the remarkable stability of his opinions, observe with apparent seriousness how unlikely it must have been for him to have begun with all the correct beliefs.

only rarely, if ever, find herself in a situation where she would need to fall back on the likelihood of her having private evidence for some proposition.

But this is not a completely satisfying response. Even if it is rare to share all of the evidence that one can, it surely happens, and in such a case it may still feel out of place to chalk the disagreement up to having one's own private evidence. This is because in the imagined cases, one considers the interlocutor to be her intellectual *equal*, or in the jargon of the disagreement debate, her *peer*. The situation is really this: person S has shared all of the evidence that she can with a peer, and the peer has clearly understood all of that evidence and the arguments offered in support of a view. However, S thinks that the peer must still be mistaken because S has some evidence that she can't articulate. To van Inwagen, this sounds like unjustified confidence. He claims that this is because S would not be able to use this same justification in *every* case of disagreement. It is for this reason that he is hesitant to accept that private evidence can account for both his and Lewis's rationality, although they disagree. However, it seems to me that van Inwagen's hesitancy to rely on private evidence is really rooted in the fact that his disagreement is with *David Lewis*, a man clearly at least his intellectual equal.¹³ This suggests that considering someone a peer has some evidential force. Indeed, the more confident I am in the peer-status of my interlocutor, the less comfortable I am likely to be with the notion

¹³ Elsewhere, van Inwagen says:

Look, it's *David Lewis* we're talking about here. I can remember a talented young philosopher saying to me in the 1970s, following his first encounter with Lewis, 'Lewis is so smart it's *scary!*'; and that has been more or less the response of all philosophers who have measured themselves against that formidable mind. Nor could anyone suggest with a straight face that Lewis was lacking in philosophical ability – not unless *all* human beings are lacking in philosophical ability. And he was scrupulously honest: he may have believed one or two odd things, but he did *believe* them, and believed them because he thought that they were straightforward objective truths.

See Peter van Inwagen, "Listening to Clifford's Ghost," *Royal Institute of Philosophy Supplement* 65 (2009): 25-26.

that I am correct by virtue of some private evidence that I have.¹⁴ It is peerhood itself that raises the question of when it is appropriate to rely on private evidence, because when I recognize someone as my epistemic peer, I cannot so easily discount their private evidence as less evidentially weighty than my own. This leads van Inwagen and others to answer “no” to (2) above. Whether this response is warranted will depend on the evidential value we assign to both private seemings (discussed in Chapter Four), and to disagreement itself (discussed in Chapters Five and Six). For the present it is sufficient to note that many commentators in the disagreement debate claim that there is a tension between the reasonableness of relying on private evidence and considering an interlocutor to be an epistemic peer. However, for all we’ve said so far, this tension does not yet provide sufficient reason to reject the reasonableness of relying on private evidence. So our answer to (2) for now is *maybe*, contingent on the answer to (1), which is just to say that, contra van Inwagen, there is no prima facie reason to reject the reasonableness of relying on private evidence in a dispute without a fuller account of the nature of that evidence.

2.2 – Feldman on the Relevance of Private Evidence

¹⁴ And yet I *am* still likely to consider it. This suggests that our intuitions about peerhood are compatible with the possibility of private evidence, since if the two were really incompatible we would presumably not feel *any* temptation to consider private evidence when disagreeing with a peer. This is a point in favor of one way of defining peerhood--as equality with respect to various “disagreement factors”--and against another way of defining it--as equal likelihood to be correct. For if I took you to be equally likely to be correct as me, then any private evidence either of us may have would already be accounted for in calculating our likelihoods, and would therefore be irrelevant to our present disagreement. For further discussion of these competing notions of peerhood, see Chapter Seven.

Richard Feldman also discusses, and dismisses, private evidence.¹⁵ Following van Inwagen and Gideon Rosen, whose view will be discussed below, Feldman understands private evidence to be a kind of seeming or intuition. As he says, “The idea... is that the seeming obviousness, or the intuitive correctness, of one’s position counts as evidence.”¹⁶ Recall here van Inwagen’s knowledge of his wife’s mood, or the intuition of the chicken sexer. The defender of reasonable disagreement might argue that this sort of evidence gives one sufficient reason to justify her view, even if a peer disagrees. But Feldman does not think this will work either, for the simple reason that each party to the disagreement knows--or can be brought to know--about the private insight of the other party.¹⁷ Here Feldman gives an argument against the *relevance* of private evidence that echoes van Inwagen’s argument with respect to (2) above. The argument is as follows:

- (1) If the parties to a disagreement each have private evidence with respect to a proposition *p*, then each may be brought to know that the other also has private evidence.

¹⁵ Richard Feldman, “Reasonable Religious Disagreements,” in *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, ed. Louise Antony (Oxford University Press, 2007), 206-209. In this context, Feldman notes that “if *all* the evidence were shared, there could not be a reasonable disagreement. This is the consequence of The Uniqueness Thesis” (207). But this is too strong: the defender of reasonable disagreement need not concede this point. She needs only maintain the possibility of justifying evidence that cannot be communicated. She need not grant the further claim that if it were communicated, then it would also be justifying in the same way for the other party. In other words, private evidence supports the possibility of reasonable disagreement whether or not the Uniqueness Thesis is true.

¹⁶ *Ibid.*, 207. This way of framing the issue makes it seem as though one’s seeming is *additional* evidence besides whatever arguments or observations led one to believe a proposition in the first place. This, however, is not the only way to think about private evidence. Here is another: I have a low confidence level in my ability to communicate certain elements of my experience, which suggests that they could in principle be communicated better than I can manage, and that they certainly could be understood better if the other person could simply share the experience. This difference--between my experience and their understanding of my experience--is itself evidence to me of the following proposition: “I possibly possess private evidence of the truth of *p*.” I cannot be sure that the other person ever really understands my evidence, or I hers (indeed, the lower my confidence level in my ability to communicate my experience, the stronger my evidence for private evidence). Nonetheless, my private evidence, if it exists, is a proper part of my already accounted-for total evidence for *p*, prior to recognizing that it may be private. I’ll return to this point below when I discuss David Christensen’s dismissal of private evidence.

¹⁷ If the reader has the strong reaction here that knowing *about* someone’s insight is obviously different from *having* that insight, keep reading.

- (2) Knowing that there is evidence for a proposition itself constitutes evidence for that proposition.
- (3) If one's disputant in a disagreement is taken to be an epistemic peer, then the peer's private evidence should be considered to be on par with one's own private evidence, unless one has some further reason to prefer one's own private evidence.
- (4) Therefore, in the absence of a reason to prefer one's own private evidence over a peer's, one should suspend judgment about the disputed proposition.¹⁸

Premise (1) is not intended to be controversial, and I do not wish to argue this point, so I'll stipulate that in most realistic cases of disagreement, each party can in fact be made aware that the other claims private evidence. Premise (2) constitutes what I'll call the "Meta-Evidence Principle," and will be discussed in Chapter Six. For the present, I will grant this premise. Premise (3) is a ban on what Ralph Wedgwood has called an "egocentric epistemic bias," i.e., a tendency to prefer one's own view or to weigh one's own insight more heavily simply because it is one's own.¹⁹ This will be discussed below, but it is on this point that I wish to focus my critique of Feldman's argument against the relevance of private evidence.

As I will show in the next chapter, private evidence is generally thought of in terms of an analogy with perception: evidential seemings are taken to be analogous in some sense to perceptual seemings. Feldman, for example, employs the perceptual analogy in the case of the dean in the quad, mentioned above. Indeed, the case does much of the argumentative work for Feldman, since it is designed to support the view that in cases of straightforward perceptual disagreement, it is unreasonable to assume that it is the *other* person whose eyes are malfunctioning. When I see the dean and you don't, it is

¹⁸ This is a reconstruction of Feldman's argument on pp. 208-209.

¹⁹ See Ralph Wedgwood, "Disagreement and the A Priori," in *The Nature of Normativity* (Oxford University Press, 2007): 261.

inappropriate for me to conclude, owing to my perceptual seeming which I cannot fully communicate, that you are in the wrong rather than me. As Feldman says,

I cannot reasonably say, “Well, it’s really seeming to me like the dean is there. So, even though you are justified in your belief, your appearance is deceptive.” I need some reason to think you, rather than me, are the one with the problem.... To think otherwise, it seems to me, is to think something like this: “You have an insight according to which *P* is not true. I have one according to which *P* is true. It’s reasonable for me to believe *P* in light of all this because, gosh darn it, *my* insight supports *P*.” If one’s conviction survives the “confrontation with the other,” to use Rosen’s phrase, this seems more a sign of tenacity and stubbornness than anything else.²⁰

The perceptual case is especially powerful here because in that sort of case, the intuition of equality between oneself and the other is very strong.

Even so, there are serious problems with the claim that a peer’s private evidence should be considered as evidentially equivalent to one’s own [premise (3)]. First, we should be clear about what each party knows. Feldman says: “Each may have his or her own special insight or sense of obviousness. But each knows about the other’s insight. Each knows that this insight has evidential force.”²¹ “About” is ambiguous here. It could mean that each is aware of there being evidence, or it could also mean that each is acquainted with that evidence. Clearly Feldman intends the former, and yet his conclusion seems to follow only from the latter. Each knows *that* the other has an insight or sense of obviousness, and each may assume or grant that that insight has evidential force *for the other person*, but neither fact provides sufficient justification for either party to conclude that the other person’s insight is *on par* with one’s own. How could it, when neither knows--is acquainted with, or *has* in a strong sense--that evidence? In other

²⁰ Feldman, “Reasonable Religious Disagreements,” 208.

²¹ Ibid.

words, the actual evidence here (whatever it is) remains private, even when the other party is made aware of its existence. So here we have a clear dissimilarity: one's own seeming or insight or sense of obviousness on the one hand, and the peer's *report* of a seeming, insight, sense of obviousness, etc., on the other. The decision to accord the peer's private evidence, to which one does not have access, equal weight as one's own is itself an irrational decision, since it cannot be motivated by the evidence one actually possesses. We do not need to deny that the peer's claim of private evidence carries evidential force (though we will complicate this claim later); but it does not follow that it carries sufficient force to overcome the prima facie presumption in favor of one's own seeming.

This is made clearer in Feldman's example of accused criminals and Lefty and Righty, which he modifies to illustrate his point:

Consider again the example involving the two suspects in a criminal case, Lefty and Righty. Suppose now that there are two detectives investigating the case, one who has the evidence about Lefty and one who has the evidence incriminating Righty. They each justifiably believe in their man's guilt. And then each finds out that the other detective has evidence incriminating the other suspect. If things are on a par, then suspension of judgment is called for. If one detective has no reason at all to think that the other's evidence is inferior to hers, yet she continues to believe that Lefty is guilty, she would be unjustified. She is giving special status to her own evidence with no reason to do so, and this is an epistemic error, a failure to treat like cases alike. She knows that there are two bodies of equally good evidence for incompatible propositions, and she is favoring the one that happens to have been hers originally.²²

Note the inclusion of "If things are on a par..." This is a very large *if*. How could either detective possibly know whether the other's evidence is comparable *until she knows what*

²² Ibid.

it is? But she cannot know that if it's private.²³ Granted, she may lack a reason for thinking that it's *inferior* to her own, but she also lacks any reason for thinking that it's *superior* or even *equal to* her own. All she has evidence for is that it's sufficient to motivate conviction in her peer. It may be, for all we have said so far, that this alone is reason enough to lead her to lower her credence in her own view, but if so that is due to the epistemic relevance of peerhood, not to private evidence per se. Given that peerhood is the subject of Chapter Seven, I will reserve discussion of this point.

Note also that Feldman claims the detective “knows that there are two bodies of equally good evidence for incompatible propositions...” But this isn't right--she knows only that there are two bodies of evidence sufficient to warrant belief in opposed propositions. If the evidence is really private, she can discern nothing about its objective value.

As we will see below, peerhood implies, among other things, that we are *generally* equal with respect to whatever faculties contribute to assessing evidence well. However, it does not imply anything in the *present* case, which concerns private evidence. If I believe that you and I are epistemic peers, then I will accord your view some evidential weight, even if I don't know the specifics of your evidence, but I must also weigh your view against my own seemings, and it is not yet clear that your view should counterbalance mine. Indeed, given that most instances of peer disagreement are unlikely to move far beyond the initial recognition of disagreement, we can already say

²³ This quote comes from the section where Feldman is considering private evidence. He is here using a case which includes in-principle shareable evidence to help him make the point that knowing that someone else has comparable evidence to your own implies that you should not prefer your own to theirs simply because it is your own. This comparison (between shareable evidence and private evidence) is problematic for the reasons I point out in the text.

that one is *usually* justified in maintaining one's view after recognition of peer disagreement, owing to one's possession of private evidential seemings which are not overcome by the report of another's contrary seemings. Nothing Feldman has said about private evidence should lead us to reject this conclusion. Indeed, it will be strengthened when we consider the relationship between evidential and perceptual seemings in the next chapter. So Feldman, like van Inwagen, has not given us sufficient reason to regard the possibility of private evidence as irrelevant or unreasonable.

2.3 – Christensen on the Symmetry of Private Evidence

David Christensen, another Conformist, has also argued against the relevance of private evidence.²⁴ Christensen argues that once I discover that a peer disagrees with me about *p*, this gives me two sorts of evidence: evidence that the peer has made a mistake--because she has reached what I take to be a wrong conclusion--and evidence that *I* have made a mistake, because an epistemic peer disagrees with me.²⁵ This results in an evidentially symmetrical situation, according to Christensen, and so the reasonable thing to do is to lower one's credence in *p*. He illustrates this with an example: I check my wristwatch for the time, and my friend does the same. The watches are of comparable quality and, so far as I know, equal reliability.²⁶ But on this occasion they disagree by about 10 minutes. Clearly, says Christensen, I should not prefer the time given by my watch simply because it is on my wrist.

²⁴ David Christensen, "Epistemology of Disagreement: The Good News," *Philosophical Review* 116, no. 2 (2007): 199-205.

²⁵ *Ibid.*, 196.

²⁶ Christensen assumes they are of the same make, age, have similar track records of recording the correct time, etc.

This analogy is crude: watches are external objects and can be neutrally compared. I, on the other hand, cannot get out of my epistemic perspective to compare it with a peer's. But Christensen says the watch analogy is accurate in the following way: I have no reason, outside of the disagreement itself, to prefer my view to my peer's. Indeed, explaining our disagreement by saying that my peer must be the mistaken one because she disagrees with me, is question-begging. Further, Christensen argues, stipulating that it seems to me that my peer is the incorrect one from within a first-person perspective won't change the epistemic situation. Though Christensen does not use the phrase "private evidence" here, I take it that this is what he has in mind with the "first-person perspective" language. As he says, "it seems to me that taking a first-person perspective on the situation does not license me in thinking that disagreement with my friend is better explained by her error than by mine."²⁷ So the challenge is to explain the disagreement from my first-person perspective "in a way that does not just beg the question in favor of the opinion I currently hold."²⁸

Christensen's solution is to require that any explanation of the disagreement that favors one party over the other must be based on reasoning that is *independent* of the matter under dispute.²⁹ The question is then whether private evidence is independent of the disputed matter.³⁰ Taking it to be so is anomalous, for private evidence is typically thought of in terms of a seeming or intuition that such-and-such evidence really does support p.

²⁷ Ibid., 197.

²⁸ Ibid., 198.

²⁹ Ibid., 199.

³⁰ See footnote 16, this chapter.

Consider, in this regard, another of Christensen's cases: two meteorologists study the same weather data and come to slightly different conclusions about the chance of rain tomorrow: meteorologist A puts it at 55%, and meteorologist B puts it at 45%.³¹ Assume that prior to examining the weather data, both A and B take one another to be peers with respect to discerning the likelihood of rain based on this kind of data. We can even stipulate that if we asked them prior to examining the data who they thought would be correct in the event of this kind of disagreement, they both would have said that they had an equal chance of being correct. When A learns about B's differing assessment of the evidence, she should probably be a bit less confident in her own assessment than she had been a moment before, given that she considers B a peer and had even predicted that in just this sort of situation, B would be as likely as herself to be correct. So let's say she goes back through the weather data and examines it even more carefully than before, and she still comes away with the view that there is a 55% chance of rain. Moreover, after having spent this extra time with the data, she is now more aware than before just how apparent this result is, and she has a hard time imagining how B could come away with a lower result. In other words, she has a seeming or a felt sense of obviousness that 55% is the chance that the data supports.

Is this seeming *new* evidence independent of the issue which is already in dispute? It seems not. Granted, she *noticed* the seeming after becoming aware of the disagreement, and we can even say that she would not have noticed it without the disagreement, and we can grant that the seeming is *strengthened* by reviewing the data, but it *always* seemed to her that the chance of rain was 55%. If she were asked why she

³¹ Ibid., 193-194.

thinks that is the right prediction, her answer would certainly involve the weather data that she and B examined, the application of various meteorological models, her calculation, and so on. It would certainly *not* involve her own seeming. If this is right--if seemings are not independent evidence--then Christensen's condition rules them out, and thereby prohibits private evidence in the form of seemings. Indeed, Christensen later implies that private evidence would *not* be independent: he references van Inwagen mentioning "incommunicable insight" and concludes that since "van Inwagen does not claim to be able to point to any reason, *independent of the disagreement itself*, for thinking that those who disagree with him lack some special insight he has," he therefore ought to conciliate.³²

Is Christensen's independent evidence condition warranted? First, note some of its counterintuitive consequences: the theist who has had what she takes to be obvious and compelling mystical experiences must consider these private seemings to be evidentially empty when she encounters an atheist peer who has had no such experiences. Similarly, a savant who can "see" the answer to a calculation immediately and clearly--and who has never been wrong--must consider this seeming to be evidentially irrelevant when confronted with a peer who gets a different answer.³³ Van Inwagen, who can sense when his wife is angry, must set that seeming aside when her sister (say) insists that she

³² Ibid., 205. Italics in original.

³³ Christensen mentions just such a case and concludes that the savant should in fact move her belief in the direction of her peer--presumably another savant (202). I disagree, both because my intuition does not suggest that the savant should discount the clear evidence of her seeming, and because I think that the disagreement of a *peer*--rather than a superior--does not give one positive reason to alter her view.

isn't. And so on.³⁴ This does not by itself show that the condition is unwarranted, but it does cast some doubt on it.

There are, in addition, more serious problems with it. The condition is open to counterexample in cases in which either: (i) there is no independent evidence but it is clear whether one should conciliate or not, or (ii) there is independent evidence, and it is clear whether, *ceteris paribus*, one should conciliate or not, but, all things considered, one should not do what is *ceteris paribus* required *because of* the independent evidence, or (iii) gathering independent evidence is unrealistic or practically costly.

Let us consider these cases in order:

(i) Consider Christensen's case of the bill calculation. One person arrives at an obviously false answer, say, a \$450 share. In this case, it is apparent that one should not move one's belief in the direction of one's peer at all, since the peer's result is plainly wrong. Christensen agrees but says that this is because there is independent evidence here: the peer's violation of a norm of common sense. After all, common sense dictates that a share of a bill cannot be more than the total bill. So once one realizes that her peer has violated this norm, she has evidence independent of the original dispute that justifies her in maintaining her own view. But Christensen acknowledges, rightly, that this may not be convincing because the person is still experiencing the violation of the norm as a felt sense of obviousness. He, however, rejects this point. But he shouldn't. Christensen himself is simply redescribing the basis for the act of non-conciliation so that it sounds

³⁴ These last two examples are complicated by the fact that the persons in question--the savant and van Inwagen--both also possess inductive evidence for the truth of their conclusion based on past success. Nonetheless, in both cases the independent evidence condition entails that their private seemings are evidentially irrelevant, which is counterintuitive, especially considering that their past success should be attributed to these very seemings.

independent of the dispute at issue, when it in fact is not. One's sense that her peer's answer is obviously wrong is part of the original evidence that constitutes the disagreement.³⁵

Christensen attempts to dodge this bullet by switching cases. He brings up a different case in which it is clearer that there is no independent evidence. He does this in order to show that in such cases, maintaining one's view is arbitrary and therefore unreasonable. The example he uses is the mathematical savant case. Imagined two peer savants who both "see" the answer to a math problem and yet disagree. He thinks that the savants should both conciliate, since neither can appeal to an outside piece of evidence to support the view that it is his peer who is mistaken.

That is Christensen's intuition. I have a different one. Mine is that if it seems to one in the clear and vivid way that the "seeing" metaphor suggests that *p* is true, then that is sufficient reason to maintain belief in *p*, even if a peer disagrees and also claims a clear and vivid seeming. This is due largely to my view of peerhood and its significance, which will be discussed in Chapter Seven. For now, it is enough to point out a problem with this case. Christensen apparently intends it to mollify the objector in the previous case, about which he acknowledges that one may not share his intuition that the evidence there is truly independent. But this case is no improvement, since it does not in fact show that in peer situations with no independent evidence, one should conciliate. Rather, it shows that in cases of perceptual disagreement between peers (for that is what the savant case amounts to), there is no good reason to break the stalemate. It is true, but unremarkable,

³⁵ Consider: the calculation would not violate common sense except in the context of *this particular bill*. Thanks to Michael Wreen for this observation.

that in such a situation one cannot be sure that one's seeming is veridical. Nonetheless, it may still be reasonable to stick to one's vivid seeming, rather than to cast about for some "independent evidence," find none, and on that basis decide to lower one's credence.

(ii) In the bill calculation case with an obviously false answer, a person is reasonable to "stick to her guns" because she has a clear and vivid seeming that the peer must have made a mistake. This is why in fact most people in such a situation would refuse to alter their view, and this is the proximal explanation of the reasonability of refusing to alter one's view, even if there is, as Christensen argues, a possible *distal* explanation having to do with violations of norms of common sense. If Christensen is right about the independent evidence condition and about the norm violation being independent evidence, this would entail that most people in such circumstances are *not* in fact acting reasonably when they maintain their view, as he admits they should, since they would be doing so for the wrong reason: a seeming. Even if we agree with Christensen that in the bill calculation case one's appeal to a violation of a norm of common sense is in fact an appeal to independent evidence, if one is not conciliating *because of* that independent evidence, then one is not in fact reasonable to conciliate. I take it this consequence is sufficiently counterintuitive--counter to Christensen's *own* stated intuitions--to reject the independent evidence condition.

(iii) It is also possible that there would be practical costs to seeking independent evidence that would outweigh the potential benefit of finding such evidence, even if there is independent evidence to be found.³⁶ Take, for example, the case of two epistemic peers

³⁶ Marc Moffett considers something like this; see Marc Moffett, "Reasonable Disagreement and Rational Group Inquiry" *Episteme: A Journal of Social Epistemology* 4 (2007): 362.

disagreeing over abortion.³⁷ Neither has any particular expertise in the matter, but both have a felt sense of obviousness about the morality of abortion. They've shared all of the arguments and evidence in the form of data, statistics, moral principles, and so on, that they can think of, and each remains convinced that the other is mistaken. As things stand, on Christensen's account, neither has any reason to prefer her own view, and so both should conciliate, unless there is independent evidence to appeal to that would make it reasonable to conclude that the other person is mistaken. Presumably, either party could, if she chose, go out and begin interviewing people who have direct experience with abortion, physicians, policy-makers, judges, ethicists, and the like. On the basis of these interviews, one may find some additional support for or against one's own view that was not included in the previously shared evidence. But given the practical difficulty of undertaking such an enterprise, can we really consider someone irrational for not conciliating? Some commentators have allowed for a rational resistance to this sort of effort under the label of "epistemic conservatism," the view that one may be rational in resisting belief revision if the practical cost is very high, or more modestly, that there is an epistemic presumption in favor of one's standing beliefs, even when one is aware of the existence of equally justified conflicting beliefs.³⁸ If some version of epistemic conservatism is plausible, then one may be justified in avoiding gathering independent evidence on practical grounds.

Thus, Christensen's independent evidence condition is open to counterarguments.

It therefore cannot be used as a prohibition on private evidence.

³⁷ I borrow and modify this case from Adam Elga.

³⁸ See Gideon Rosen, "Nominalism, Naturalism, Epistemic Relativism," *Noûs* 35, no. 15 (2001): 85-86; and Moffett, "Reasonable Disagreement and Rational Group Inquiry," 361-363.

Chapter Three – The Perceptual Analogy

3.1 – Perceptual Seeming in the Disagreement Debate

What is a seeming? As we saw above, there is a tendency in the literature to understand *seeming to be the case that p*--call this *evidential seeming*--as analogous to perceptual seeming--seeming to see that *p*.¹

Alvin Plantinga gives a nice characterization of what he calls “doxastic experience,” based primarily on memorial and a priori sources of belief. It can, I think, be applied to what I’m calling evidential seeming:²

...experience comes in several varieties. First, there is *sensuous imagery*, the kind of experience you have most prominently in vision but in hearing, smelling, tasting, and touching as well. To use Roderick Chisholm's terminology, in this kind of experience one is *appeared to* in such and such a way. Sensuous imagery plays an enormously significant role in perception; perceptual beliefs are formed *in response to* sensuous imagery and *on the basis of* such imagery. Still, this isn't the only kind of experience that goes with belief formation...the formation of *memory* beliefs is often unaccompanied by phenomenal experience, or else accompanied only by fragmentary, fleeting, indistinct, hard-to-focus sensuous imagery. You remember that you went to a party in Novosibirsk; there is a bit of imagery, all right, although it is fleeting, partial, indistinct, and such that when you try to focus your attention on it, it disappears. But there is another kind of experience present: the belief that it was *Novosibirsk* (and not, say, Cleveland) seems *right, acceptable, natural*; it forces itself upon you; it seems somehow inevitable (the right words are hard to find). The belief *feels* right, acceptable, and natural; it feels different from what you think is a false belief. The same goes for

¹ One may wonder here if such evidential and/or perceptual seemings are essentially *conceptual*, in the sense that they essentially have conceptual items, such as propositions, as objects. I intend my arguments here to remain neutral between conceptual and nonconceptual understandings of evidential and perceptual seemings. I want only to illustrate the widely assumed view in the disagreement literature that whatever can be said about the justifying role of perceptual seemings should apply, *mutatis mutandis*, to the justifying role of evidential seemings. For an account of the epistemological role of conceptual/nonconceptual and perceptual/nonperceptual experiences, including an extended argument regarding how nonconceptual perceptual contents play an evidential, justifying role, see Paul K. Moser, *Knowledge and Evidence* (Cambridge University Press, 1989), especially Section 2.3.2 (esp. pp. 84-6), Section 3.4 (esp. pp. 156-8), Section 4.1.1. (esp. pp. 181-3), and Section 4.2 (esp. p. 192). See also pp. 34-35 above.

² One difference is that evidential seeming need not yield belief--it may yield a variety of other cognitive attitudes such as acceptance, assumption, or deferral of judgment. In this sense it is not properly “doxastic” according to Plantinga’s usage.

a priori belief. You believe that no dogs are sets. This belief, too, involves little by way of sensuous imagery. When you consider that proposition, perhaps it is as if you catch a momentary and fleeting glimpse of part of a sentence expressing the proposition, or perhaps a fragmentary glimpse of a dog, or perhaps of a dog enclosed within braces; this imagery seems unimportant, however, more like mere decoration than something on the basis of which the belief in question is formed. And here, too, there is also this other sort of experience: it's just seeming *true* and indeed *necessarily* true that no dogs are sets. Thinking about this proposition *feels* different from thinking about the proposition that some dogs...*are* sets. Still a third kind of example...: the knowledge that it is *you* (as opposed to someone else) who is now perceiving the page in front of you. This too is not a matter of sensuous imagery: it is not on the basis of sensuous imagery that you believe it is *you* who are perceiving that page, rather than your cousin in Cleveland. Here too there is that *other* sort of phenomenal experience, that feeling that the proposition in question is the *right* one. Suppose we call this second kind of phenomenal experience *doxastic* experience because it always goes with the formation of belief...³

This sense that an epistemic state of affairs is “right, acceptable, natural” also applies, I believe, to evidential situations. It seems to a person that her evidence supports *p*, that concluding that *p* is the case on the basis of this evidence is the “right, acceptable, natural” thing to do. As Plantinga says, there need be no imagery that accompanies this felt sense, but it nonetheless feels different from the thought that the evidence does not support *p*. *That* thought feels *wrong, unacceptable, unnatural*. Perhaps the closest we can get to a fuller account of this feeling is by comparing it with the more imagistic experience of sense perception. Note that Plantinga implies that the paradigmatic case of experience or seeming is perceptual, and that what goes for it epistemically also typically goes for other forms of experience.

This comparison is also common in the disagreement literature. For example, van Inwagen, speaking of private evidence, writes:

Owing to some neural accident (I might say) I have a kind of insight into the, oh, I don't know, entailment relations among various of the propositions that figure in

³ Alvin Plantinga, *Warranted Christian Belief* (New York: Oxford University Press, 2000), 110-111.

the Ism/Nism debate that is denied to the Nismists. I *see*, perhaps, that p entails q (although I am unable to formulate this insight verbally) and they are unable to see that p entails q .⁴

Here the sense that some proposition follows from another is analogous to the sense that perceptually something is the case. Similarly, Feldman's dean-in-the-quad analogy suggests that he takes perceptual and evidential seemings to be analogous, and so what goes for one goes for the other. For example, he disallows the evidential relevance of private evidence analogous to van Inwagen's "insight" or "sense of obviousness" by citing the case of perceptual disagreement:

...compare a more straightforward case of regular sight, rather than insight. Suppose you and I are standing by the window looking out on the quad. We think we have comparable vision and we know each other to be honest. I seem to see what looks to me like the dean standing out in the middle of the quad.... I believe that the dean is standing on the quad. Meanwhile, you seem to see nothing of the kind there. You think that no one, and thus not the dean, is standing in the middle of the quad. We disagree.... Either I am "seeing things," or you are missing something. I would not be reasonable in thinking that the problem is in your head, nor would you be reasonable in thinking that the problem is in mine.⁵

This case would not be a sufficient response to the epistemic relevance of private evidential seemings unless Feldman took such seemings to be analogous to perceptual seemings.

Gideon Rosen has perhaps been most explicit in drawing this analogy, and it is worth quoting him at length. In the context of discussing how the realization of the contingency of one's beliefs may affect one's confidence in those beliefs, Rosen airs the

⁴ Peter van Inwagen, "We're Right. They're Wrong," in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford University Press, 2010), 27.

⁵ Richard Feldman, "Reasonable Religious Disagreements," in *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, ed. Louise Antony (Oxford University Press, 2007), 207-208.

possibility that perception-like seemings could play a justificatory role in non-perceptual contexts:

It is a familiar thought that perceptual knowledge rests on a principle to the effect that certain non-doxastic ‘perceptual seemings’ constitute grounds for belief. If it seems to me that there is a cat on the mat, then I have reason to believe that there is a cat on the mat. Its seeming to me that P is not a matter of my believing anything. It is a sui generis propositional attitude that does not stand in need of justification but which is capable of supporting a perceptual judgment that P. The justificatory connection between seeming and believing is immediate. It does not depend on the reliability of the transition... Nor does it depend on the subject’s believing that perceptual seemings are reliable.⁶

After providing this characterization, Rosen discusses three cases in which such seemings--such “non-doxastic appearing-true” as he puts it--are justificatory in an analogous way to perceptual seemings.⁷ The first is desire:

Mark Johnston has argued that a certain sort of desire is best understood as a matter of being “struck by the appeal of things.” This “being struck by the appeal” is a quasi-perceptual matter--it is phenomenologically vivid, for example; but since it is typically directed at hypothetical states of affairs, it is not literally a matter of perception.⁸

The second is moral evaluation:

Like perceptual seeming, evaluative seeming is not doxastic. It is a matter of an object’s seeming to one to possess one or another highly determinate evaluative property. There may be other routes to evaluative knowledge. But it is plausible that insofar as perceptual judgment is canonically grounded in perceptual appearances, evaluative judgment is canonically grounded in evaluative appearances. That one finds a certain contemplated course of action appalling is a reason for taking it to be appalling. And once again, it is plausible that this is a basic epistemic norm, ungrounded in further facts (or thoughts) about the reliability of evaluative seeming.⁹

The third is modal judgment:

⁶ Gideon Rosen, “Nominalism, Naturalism, Epistemic Relativism,” *Noûs* 35, no. 15 (2001): 86-87.

⁷ *Ibid.*, 88.

⁸ *Ibid.*, 87. Rosen here cites Mark Johnston, “The Authority of Affect,” *Philosophy and Phenomenological Research* 63, no. 1 (2001): 181–214.

⁹ *Ibid.*

Steve Yablo has argued that judgments of metaphysical possibility are typically grounded in intuitions of conceivability. On Yablo's view, to conceive that P in the relevant sense is a compound act: One first imagines a P-world, and that world then *strikes one* as possible. The model is as above. The non-doxastic modal seeming provides a sufficient though defeasible ground for the modal judgment. And in this case in particular there is no thought that the seeming need be a matter of sensitivity to the facts in question.¹⁰

He concludes:

These cases suggest a pattern. It may be that as a general matter, one's finding it obvious upon reflection that P is the case--its striking one that P; it seeming clearly to be the case that P--amounts to a ground for believing that P.¹¹

Seemings are non-doxastic, pre-theoretical, sui generis, basic, and generative of prima facie justification.¹² This is the fullest description one gets in the disagreement literature of this phenomenon. And it is noteworthy that this description is assumed by almost everyone else in the literature, along with the perceptual analogy on which it is based.

For example, Adam Elga uses the perceptual analogy to highlight what he takes to be the strengths of his "Equal Weight View," which recommends according equal weight to one's own assessment of the evidence and to that of one's epistemic peer. Considering a case of perceptual disagreement in which two peers differently judge the outcome of a horse race, he says:

¹⁰ Ibid. Rosen here cites Stephen Yablo, "Is Conceivability a Guide to Possibility?," *Philosophy and Phenomenological Research* 53, no. 1 (1993): 1–42. Cf. Kripke on modal intuitions: "Of course, some philosophers think that something's having intuitive content is very inconclusive evidence in favor of it. I think it is very heavy evidence in favor of anything, myself. I really don't know, in a way, what more conclusive evidence one can have about anything, ultimately speaking." See Saul Kripke, *Naming and Necessity* (Cambridge: Harvard University Press, 1972), 42.

¹¹ Ibid.

¹² On the non-doxastic character of seemings, see footnote 2, this chapter. See also pp. 34-35 above. Essentially, while all seemings may be "intellectual" in the sense of having propositional content, they do not necessarily yield or justify *beliefs*--they may instead justify weaker cognitive states, like acceptance. Going forward, I will gloss this distinction and speak of "belief" in a loose way to connote any of the class of propositional attitudes which may be justified by a seeming.

Here is the bottom line. When you find out that you and your friend have come to opposite conclusions about a race, you should think that the two of you are equally likely to be correct. The same goes for other sorts of disagreements.¹³

Elga then adapts an argument brought against reliabilism about perceptual knowledge, and uses it against those who would argue that it is reasonable to maintain one's view in the context of peer disagreement, regardless of the type of disagreement.¹⁴ The details of these arguments don't matter in this context. The important point is that Elga takes for granted that an objection to a view about perceptual disagreement can be applied to all other sorts of disagreement.

Similarly, Thomas Kelly, in his critique of the Equal Weight View, does not deny the perceptual analogy--he implicitly accepts it--and then takes pains to show that his own "Total Evidence View" yields the same verdict as the Equal Weight View in perceptual cases.¹⁵

Elsewhere, Michael Bergmann adopts the perceptual analogy in defending the claim that disagreement with an epistemic peer after full disclosure of evidence need not lead one to revise her view. To the conformist, Bergmann says:

I think that the rational response for S1 to E is to continue believing p. However, I don't have an argument for that conclusion, just as I don't have an argument for the conclusion that the rational response for us to a tactile experience like the one we typically have when grabbing a billiard ball is to believe something like "that's a small hard spherical object". I can see, in the case of the billiard ball, that that belief *is* a rational response to the tactile experience in question. But I don't have an argument for why that is so. Likewise, I don't have an argument for the view that the rational response for S1 to E is to continue believing p.¹⁶

¹³ Adam Elga, "Reflection and Disagreement," *Noûs* 41, no. 3 (2007): 487.

¹⁴ *Ibid.*, 488.

¹⁵ Thomas Kelly, "Peer Disagreement and Higher-Order Evidence," in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford University Press, 2010), 150-152.

¹⁶ Michael Bergmann, "Rational Disagreement After Full Disclosure," *Episteme* 6, no. 3 (2009): 345.

Here the analogy with perception is especially clear, as is the assumption that like perceptual seemings, evidential seemings are basic and prima facie justified (i.e., without need of argument or further evidence). Bergmann goes on to give examples of other cases in which one need not modify one's view, despite knowing that a peer disagrees for reasons similar to those one uses to reject the peer's view. The examples need not detain us here, but it is noteworthy that they are inspired by an epistemological issue in which perceptual beliefs are typically implicated: skepticism about the external world. "We realize," Bergmann says, "that it's possible for things to seem to us perceptually just as they do now even though we are mistaken."¹⁷ Nonetheless, it is rational to hold onto our belief in an external world, even if we can't prove it. The same applies, *mutatis mutandis*, to cases of peer disagreement in realms as varied as morality, religion, and politics.

Other commentators also assume the perceptual analogy. Marc Moffett argues that it may be reasonable to use disagreement as the basis for reassessing an interlocutor's status as a peer. To illustrate this, he uses the now-familiar perceptual case of looking out the window and seeing someone in the quad. Interestingly, however, his conclusion about the case differs from those of other commentators: Moffett argues that in such a case, "my visual evidence is so strong, that it simply makes more sense to jettison my antecedent belief that you are an epistemic peer than to accept that I am mistaken."¹⁸ This is because the changes to my "overall theory of the world" are less radical on the supposition that you are not my peer than on the supposition that my senses are presently

¹⁷ Ibid.

¹⁸ Marc Moffett, "Reasonable Disagreement and Rational Group Inquiry," *Episteme: A Journal of Social Epistemology* 4, no. 3 (2007): 358.

deceiving me.¹⁹ Moffett then takes this point to apply broadly to all types of peer disagreement, including non-perceptual cases, implying that what holds for visual evidence also holds for other sorts of evidence.

Lastly, Jennifer Lackey suggests that conclusions drawn on the basis of perceptual cases can and should be applied more broadly. For example, while she rejects the Nonconformist view, she admits that the Nonconformist position yields the correct verdict in cases of perceptual disagreement, and that this helps to explain what is attractive about the Nonconformist position in general.²⁰ Moreover, she notes as much in the course of observing that Nonconformists (and everyone else in the debate) hold to “Uniformity,” the view that peer disagreement functions the same--and therefore has the same type of solution--in all epistemic contexts.²¹ She eventually rejects Uniformity, since some types of disagreement have apparently Nonconformist solutions, such as perceptual disagreement, while other types have apparently Conformist solutions. However, she does not question the basic idea that perceptual disagreement is analogous to other sorts of disagreement. She merely attributes the connection between disagreement-types to the fact that many cases of disagreement, perception among them, involve highly justified underwriting beliefs.²² Insofar as a case of disagreement does not involve highly justified underwriting beliefs--i.e., beliefs which resemble the justification

¹⁹ Ibid., 365 (footnote 4).

²⁰ Jennifer Lackey, “A Justificationist View of Disagreement’s Epistemic Significance,” in *Social Epistemology*, ed. Adrian Haddock, Alan Millar, and Duncan Pritchard (Oxford University Press, 2010), 305-14.

²¹ Ibid., 302.

²² Ibid., 323.

structure of perceptual beliefs--it is more likely to be susceptible to a Conformist-style solution. So on Lackey's justificationist account, the perceptual analogy is preserved.

3.2 – Implications of the Analogy

I take it that this survey of the disagreement literature is adequate to establish that the analogy between perceptual seemings and evidential seemings is widely assumed and, indeed, apparently uncontroversial. This is a telling feature of the debate. For one thing, it suggests broad agreement among commentators on the basic justificational structure of evidential seemings, namely, that they are like perceptual seemings, and that this entails that they are *prima facie* justified. It also has the interesting consequence that attempts to argue for the epistemic legitimacy of one's evidential seemings, as with one's perceptual seemings, will be inevitably circular (hence Bergmann's comment above about the proper epistemic response to his perception of the billiard ball). This is because any argument for rational reliance on one's perceptual seemings will necessarily involve premises which appeal to one's perceptual experience. As Rosen says:

...insofar as we are justified in believing that perceptual seemings are generally veridical, it is because we are independently justified in accepting a detailed account of the perceptual mechanisms which can only be grounded in particular instances of the transition from 'seems' to 'is'.²³

Similarly, Ralph Wedgwood says, "...the only primitively rational way of coming to know facts about the material world is by relying on the very experiences that are in question."²⁴ Continuing, Wedgwood explains his justification for holding that there is an

²³ Rosen, "Nominalism, Naturalism, Epistemic Relativism," 87.

²⁴ Ralph Wedgwood, *The Nature of Normativity* (Oxford University Press, 2007), 246.

“egocentric epistemic bias” by claiming a dispositional capacity shared by all humans capable of perceptual experiences. As he says:

I propose that the practice of forming beliefs by taking one’s sensory experiences at face value has the following essential connection with the truth. It may be that it is essential to sensory experiences that any subject who has such experiences at all has some *disposition* to have experiences that veridically represent certain aspects of her environment.²⁵

Notice that this is equivalent to answering the question “How is sensory belief connected with the truth?” by saying: “To have a sensory experience is to be disposed to have an *accurate* sensory experience.” And this, in turn, is equivalent to answering: “Sensory belief is connected to the truth by virtue of being a sensory belief.” Wedgwood continues: “Whenever one’s experience does consist in the manifestation of this essential disposition, then the content of the experience will be true.”²⁶

This, however, is just unquestioning and absolute reliance on the assumed accuracy of perceptual seemings, and is certainly too strong a claim, even if such seemings are basic and therefore *prima facie*--not absolutely--justified and fundamental. Wedgwood is essentially building truth into mechanisms that aim at truth, which is question-begging.

The epistemic fundamentality, or “basicity,” of perceptual seemings, and the epistemic circularity engendered by it, is widely acknowledged. For example, William Alston has argued that owing to the social nature of our epistemic practices, justifying our basic sources of information inferentially, or giving a non-circular argument for them, is impossible. At the end of “Perceiving God,” Alston writes:

²⁵ *Ibid.*, 231.

²⁶ *Ibid.*

Apart from reliance on doxastic tendencies with which we find ourselves, we literally have nothing to go on. Indeed, what Descartes did, as Thomas Reid trenchantly pointed out, was arbitrarily to pick one doxastic practice he found himself engaged in--accepting propositions that seem self-evident--and set that as a judge over all the others, with what results we are all too familiar. This is not to say that we must acquiesce in our prereflective doxastic tendencies in every respect. We can tidy things up, modify our established practices so as to make each more internally consistent and more consistent with the others. But, on the whole and for the most part, we have no choice but to continue to form beliefs in accordance with these practices and to take these ways of forming beliefs as paradigmatically conferring epistemic justification. And this is the way that epistemology has in fact gone, except for some arbitrary partiality. Of course it would be satisfying to economize our basic commitments by taking one or a few of these practices as basic and using them to validate the others; but we have made little progress in this enterprise over the centuries. It is not self-evident that sense perception is reliable, nor can we establish its reliability if we restrict ourselves to premises drawn from introspection; we cannot show that deductive reasoning is valid without using deductive reasoning to do so; and so on. We are endowed with strong tendencies to engage in a number of distinct doxastic practices, none of which can be warranted on the basis of others. It is clearly the better part of wisdom to recognize beliefs that emerge from these practices to be rational and justified, at least once they are properly sifted and refined.²⁷

Alvin Plantinga concurs with Alston, saying:

I believe he succeeds in establishing the important conclusion that it is not possible to show in a noncircular fashion that SP [sense perception] is reliable--at any rate he gets as close to establishing this conclusion as philosophers ever get to establishing any important conclusion.²⁸

Elsewhere, Plantinga himself argues for a similar point in connection with his defense of the idea that belief in religious propositions can be justified, rational, and “warranted” (i.e., can amount to knowledge) *without* any argument or evidence whatsoever to support them. He draws upon the perceptual analogy we’ve noted in the disagreement literature, and says:

We have perceptual beliefs: in general, there won’t be ‘appropriately cogent evidence’ for these beliefs by way of beliefs that are themselves evidentially independent of perceptual beliefs. The same goes for the beliefs about the past,

²⁷ William P. Alston, “Perceiving God,” *The Journal of Philosophy* 83, no. 11 (1986): 664.

²⁸ Plantinga, *Warranted Christian Belief*, 119.

inductive beliefs, beliefs about other minds, and the like. In each of these cases, there is a source of warranted belief; and beliefs produced by the source in question can't in general be provided with that "appropriately cogent evidence" from beliefs emanating from some other source. But of course that is nothing against them.²⁹

Similarly, Robert Audi has said that "To acquire justified beliefs about the reliability of perception, moreover, we would need to use perception, say in looking at objects again to see if our initial color perceptions were accurate."³⁰ To the response that this justification may be a priori, Audi replies:

On the most plausible kind of reliability theory of justification... a belief is justified by virtue of being grounded in reliable belief-producing processes such as perceptual ones; and it is apparently not an a priori matter what processes are reliable, that is, actually produce a suitably large proportion of true beliefs. This is the sort of thing that must be determined largely by observation.³¹

So epistemic circularity, in the case of perceptual experience, is held by many philosophers to be unproblematic.³² So if the perceptual analogy holds, and perceptual seemings as justificatory cannot be non-circularly justified, then neither can evidential seemings of the form "It seems to me that E supports p." Demanding a further argument for the justificatory force of this seeming, or expecting one for comparison and assessment,³³ is out of place. One cannot justify a seeming that E supports p without relying on the same faculty of judgment that produced the initial seeming, just as one

²⁹ Alvin Plantinga, "On 'Proper Basicity,'" *Philosophy and Phenomenological Research* 75, no. 3 (2007): 619–20.

³⁰ Robert Audi, *Epistemology: A Contemporary Introduction to the Theory of Knowledge*, 3rd ed. (New York: Routledge, 2011), 371.

³¹ *Ibid.*

³² For more on this, see William P. Alston, "Epistemic Circularity," *Philosophy and Phenomenological Research* 47, no. 1 (1986): 1–30; and Ernest Sosa and Barry Stroud, "Philosophical Scepticism," *Proceedings of the Aristotelian Society, Supplementary Volumes* 68 (1994): 263–307.

³³ Recall Feldman's demand in the last chapter.

cannot justify perceptual seemings without relying on the same faculty of perception which entails that perceptual seemings have evidential value.

It is clear that if the perceptual analogy we've been discussing holds, then *if* perceptual seemings yield basic and prima facie justified beliefs, then so do evidential seemings of the sort exemplified by van Inwagen's intuition about his wife or the chicken sexer. This has the apparent consequence that cases of peer disagreement involving conflicting conclusions based on the same shared evidence base, where the evidence *seems* to one to support *p*, and the evidence *seems* to the other to support *not-p*, ought to be judged on the basis of analogous cases of perceptual disagreement. That is to say, when I find myself in disagreement with a peer about what evidence supports, it should be useful to ask myself what I would do in a case of disagreeing with a peer about what I am currently seeing. If in such a case I would be justified in holding onto the perceptual belief, then I would, *ceteris paribus*, be similarly justified in the evidential case (assuming the seemings involved are similarly clear, vivid, stable, and so forth). But recall that opinions on what one should do in the perceptual case are divided: some (e.g., Feldman) think one ought to suspend judgment, while others (e.g., Moffett) think that one ought to maintain one's view. Thus, the perceptual analogy will unfortunately not satisfy all of the commentators in the disagreement debate.

It does, however, help to constrain potential reasonable solutions to the problem. For example, it suggests that one should not simply disregard one's evidential intuitions (e.g., that the preponderance of the evidence supports *p*), counting them as irrelevant in the case of peer disagreement, unless it would be reasonable to do so in the case of perceptual disagreement. I believe this is an important consequence, since this is

precisely what is recommended by Conformists such as Christensen and Elga who insist that only evidence *independent* of the disputed issue can be relevant to breaking the stalemate created by peer disagreement. If this is correct, then such Conformists are committed to the same policy in cases of perceptual disagreement, or else they must admit inconsistency. Should they reject the analogy, they would bear the burden of showing what is special about perceptual seemings that does not hold for evidential seemings, such that the former are evidentially relevant under dispute, while the latter are not.

But if they bite the bullet and maintain that disputed perceptual seemings *are* in fact evidentially irrelevant to what one should believe, then their position becomes implausible. For if perceptual seemings are evidentially *irrelevant* in the context of perceptual disagreement, then any perceptual seeming anyone is having at any time could be evidentially neutralized simply by someone else having a different seeming, or, at least, by a *peer* having a different seeming. In the case of perception, peers are remarkably easy to come by: I assume that almost any stranger I meet has perceptual capacities which are roughly equal to my own. It therefore follows on the Conformist “independent evidence only” view that I actually have *no evidence* that is relevant to the truth of any proposition about the world that things are the way they seem to me to be. But this is absurd: I of course *do* have perceptual evidence that things are some way or other, and this evidence is highly relevant to my beliefs about the world, even if it can be outweighed by enough competing evidence (e.g., finding out that I was given a hallucinogen).

So the Conformist independent evidence requirement cannot be sensibly applied to cases of perceptual disagreement, and if not to those, then not to other cases of non-perceptual disagreement either, at least where one has a seeming that some evidence supports a proposition. Thus, the perceptual analogy, apparently assumed by everyone in the disagreement debate, has consequences that have been missed by those in the debate. But as we said above, this will not convince everyone in the disagreement debate, since there is dispute about what one should do in cases of *perceptual* disagreement between peers. It should be clear that one should not accord one's perceptual seemings *no* evidential weight, but it does not follow that they should be trusted blindly either. And there are cases in which it is apparent that even a clear, vivid, and stable perceptual seeming can be overridden by considerations of the views of others. We must, then, determine the evidential weight of seemings--of the perceptual or evidential type--within the context of disagreement as such. This is the subject of the next two chapters.

Chapter Four – The Evidential Value and Defeat of Seemings

4.1 – Rosen on the Contingency of Evidential Seemings

As we saw above, seemings of either the perceptual or evidential variety are widely held to be *basic*, i.e., they are not justified by virtue of an inferential relationship with other cognitive states. They are not *based on* anything further; rather, they are themselves that on which other reasonably held cognitive states (beliefs primarily, but also desires, volitions, and others) are based, if not themselves basic. We also saw that owing to their basic nature, and perhaps to their tendency to produce true beliefs, seemings are also typically taken to be productive of prima facie justification.¹ We concluded in the last chapter that evidential seemings--seemings of the form “It seems to me that E supports p”--have evidential weight, and that this weight does not disappear on the discovery of disagreement. However, it is in the nature of prima facie justified beliefs that they may be overridden, or *defeated*, in the parlance common among epistemologists. Classic examples of perceptual defeat include learning that one is not in a sound state of mind, learning that one is in a simulation, learning that there are lots of barn facades around, etc. Similar examples could be adduced for evidential seemings: learning that one has missed some crucial evidence, learning that one has been slipped a drug with the side effect of making 99% of people misjudge evidence sets, and so on.²

¹ The “perhaps” here alludes to the ongoing disagreement between epistemologists about internalism and externalism, and the role of reliability in both justification and knowledge.

² This second example comes from David Christensen, “Epistemology of Disagreement: The Good News,” *Philosophical Review* 116, no. 2 (2007): 207-208.

Our question in this and the next chapter is whether disagreement serves as a defeater for evidential seemings, and if so, to what extent.³

Gideon Rosen has argued that disagreement can be a defeater. In his 2001 paper “Nominalism, Naturalism, Epistemic Relativism,” Rosen takes as his point of departure the experiences of nominalists versus anti-nominalists regarding the existence of abstract objects. He uses this to prod the reader’s intuitions about cases of disagreement. He begins by noting that belief in abstract objects is rationally permissible but not obligatory. One may rationally believe in them, suspend judgment about them, or even reject them. He lists some reasons to think that it’s permissible to believe in them (he calls these his “dogmatic assumptions”):

1. We take them for granted in nearly everything we do.

As he says, “...if our considered worldview is roughly accurate, abstract entities exist.”⁴

2. By ordinary standards, we are justified in accepting this worldview--i.e., we are not normally obliged to establish the existence of, say, numbers, before using them.⁵

This shows “that any case for compulsory nominalism must be a skeptical case: a case for revising a pervasive (and by ordinary standards, unproblematic) commitment of common sense and established science.”⁶ The nominalist can make such a case in two ways: she can argue either that abstract objects are inadmissible from the point of view of science, or from that of philosophy. Unfortunately,

³ The language of “defeat” is preserved here because of its use in the literature. Defeat, however, has the connotation of the complete removal the evidential value, which is not always appropriate for disagreement contexts. For this reason, it would be more precise to speak of disagreement as providing countervailing considerations which may or may not override or outweigh one’s first-order evidential seemings.

⁴ Gideon Rosen, “Nominalism, Naturalism, Epistemic Relativism,” *Noûs* 35, no. 15 (2001): 70.

⁵ This is analogous to the lack of an obligation to rule out radical skepticism about the external world.

⁶ *Ibid.*, 71.

3. No such arguments against abstract objects are compelling.

If the above assumptions are correct, Rosen claims, then belief in abstract objects (antinominalism) is rationally permissible.

But can it also be rationally permissible to take a contrary view? Can it be rationally permissible, for example, to suspend judgment on the existence of abstract objects, especially considering that, per assumption 1, this entails the rejection of much of our usual view of the world? If it can, then the person who suspends judgment--or the nominalist who advocates suspension of judgment--must still be able to take "science seriously as a source of information about concrete nature."⁷ That is, straightforward skepticism about the deliverances of science is *not* rationally permissible. But it seems that there is a way for the suspend-judgment-nominalist to do this--to take science seriously and still not be committed to the existence of abstract objects.

There is, Rosen thinks, a "trick" for "nominalizing at a stroke any theory whatsoever."⁸ Rosen asks us to imagine a place called Bedrock, where instead of taking Platonism for granted as we do (at least as Rosen assumes we do), people believe, and can articulately support, that platonistic language is a useful fiction. When they "assert" some mathematical truth, for instance, they are doing something different from what we are doing. We presuppose that we're aiming for truth, but they presuppose only that they're aiming for "nominalistic adequacy." All they are presupposing is that the aggregate of concrete objects in the world in which the assertion is made is identical to the aggregate of concrete objects in a world in which the assertion would be true. That is,

⁷ Ibid., 73.

⁸ Ibid.

“things [in the actual world] are in all concrete respects *as if* [the assertion] were true.”⁹

And nominalistic adequacy has all the same practical effects as platonic realism: “To the extent that practically relevant conclusions are all claims about the concrete array, it follows that there is no practical difference between believing a theory to be true and believing it to be nominalistically adequate.”¹⁰

Now we ask: “Are the Bedrockers rationally entitled to persist in their fictionalism? And is it a rational option for us?”¹¹ If the answer to either question is yes, then rational people can disagree; if the answer to the second question is yes, then *we* can be rational both in accepting *p* and in accepting $\sim p$, even when we have no compelling case against *p* (where *p* is the proposition that abstract objects exist).¹² With respect to the first question: if the Bedrockers are not rationally entitled to their fictionalism, then we should be able to show them their error, but it’s hard to see how we could, especially since they’re as lucid and well-informed about the relevant scientific evidence as we are.

Does this mean we should become fictionalists too, since we cannot show that it would be unreasonable to do so? Or is it rational to hold onto our complacent platonism? One thing to note is that our platonism is *contingent*--if things had gone differently

⁹ Ibid., 75.

¹⁰ Ibid., 77.

¹¹ Ibid.

¹² Note that acceptance here does not entail belief. The Bedrockers, strictly speaking, are agnostic about the existence of abstract objects, preferring to speak in terms of nominalistic adequacy. Rosen notes a similarity between this position and Bas van Fraassen’s empiricism:

The task is analogous to the challenge that confronts van Fraassen’s constructive empiricism. Van Fraassen’s main claim is that agnosticism about unobservables is consistent with immersed participation in science. The claim in the present section is a weaker claim to the same effect, viz. that immersed participation in science is compatible with agnosticism about the abstract.

Ibid., fn. 13, p. 90. For further discussion, see Gideon Rosen, “What Is Constructive Empiricism?,” *Philosophical Studies* 74, no. 2 (1994): 143–78; and Bas van Fraassen, “Gideon Rosen on Constructive Empiricism,” *Philosophical Studies* 74, no. 2 (1994): 179–92.

historically, we might have been Bedrockers (indeed, there is some reason to suspect that this is close to our natural attitude).¹³ Does this recognition of contingency mean we should modify our belief? Not necessarily: this will depend on the case. It is true both that it is “psychologically possible for us to maintain a belief in full awareness of its strong contingency,” and that in “some cases there is absolutely nothing wrong with doing so.”¹⁴ For example, if one were raised by Genghis Khan, one would likely believe different things about cruelty than one in fact does; if one were raised by a certain kind of empiricist, one would reject the existence of atoms. Moreover, these positions might withstand reflection. Nonetheless, one need not *now* revise her view simply because it is contingent in this way. On the other hand, in some cases--and Rosen thinks complacent theism and complacent platonism are such cases--he does recommend that one revise one’s view, though he thinks it is not obligatory to do so. As he says:

When I reflect on the fact that I might have found it natural to suspend judgment on the existence of mathematical objects and that I have nothing to say to someone who does find this natural, I find my complacency somewhat shaken, and rightly so, in my opinion. I grant that I am entitled to persist on grounds of conservatism, or on the ground that it takes effort and ingenuity to sustain the fictionalist *epoché*. But belief on such a basis strikes me as hollow: as a lamentable concession to necessity or to laziness. It may be permissible, but it is hard to see the virtue in it.¹⁵

So what is the epistemic difference between the cases in which one *should* revise a belief upon recognizing its contingency and those cases in which one should not? Rosen

¹³ There is a sense in which *every* belief is contingent: there is a possible world where I do not have that belief. This is not the sense of contingent that Rosen has in mind. Instead, he seems to think of contingent in the sense of the Islamic concept of *taqlid*, discussed below, which he takes to imply “‘derivative’ or ‘second-hand’” belief. (Ibid., 84.) That is, it is belief which is held only because it is the accepted belief of one’s group. Using the example of theism, he says, “I am imagining... a theist who believes simply because he has been immersed in a culture in which God’s existence is taken for granted.” (Ibid., 85.) For more on *taqlid*, see note 28 below.

¹⁴ Ibid., 85.

¹⁵ Ibid., 85-86.

does not say--or rather, he says that different cases just *seem*, phenomenally, different to him. Revising in the cases of the alternate morality or science feels or seems unnecessary, whereas revising in the case of theism feels or seems right.¹⁶ But why should we care about individual psychology? The question, after all, is whether we *should* revise our beliefs upon recognition of coherent, apparently equally motivated alternatives, not whether we are *inclined* to do so.

Rosen responds by drawing the analogy discussed in the last chapter between perceptual seemings and evidential seemings. For perception, seemings serve as the justifying basis of beliefs, and the beliefs need no other ground. Similarly, there are non-perceptual seemings (e.g., evaluative seemings) that can ground their analogous beliefs. And so, if an encounter with an alternate view removes this natural tendency to ascribe prima facie justificatory force, then it may be recommended--or even "virtuous"--to revise one's belief, though still not *obligatory*, owing to the potential costs involved to do so. "But if one does have the time, and one can see one's way through to a more coherent overall view, then, as I say, there would seem to be virtue in rethinking."¹⁷ Rosen classes platonism along with theism as topics that he personally finds to be affected by the realization of contingency, i.e., the force of the initial seemings are altered for him. This results in a form of *relativism* based on one's non-doxastic seemings. Rosen concludes by reiterating that if one can avoid agnosticism on some issue only by appealing to the

¹⁶ Ibid., 86.

¹⁷ Ibid., 88.

practical difficulty of belief revision (he calls this “conservatism”), then one may still be “rational,” but is nonetheless epistemically defective.¹⁸

So, to recapitulate: Rosen understands the potential for disagreement to defeat one’s evidential seemings to be a matter of the recognition of the *contingency* of one’s belief. But he also admits that not *all* contingent beliefs are defeated upon the recognition of their contingency. This happens only when this recognition is able to overcome the seeming that one has--and thus the *prima facie* justification--that the belief is true. When does this happen? His answer seems to be that there is no way to tell, except by experiencing it, hence his conclusion of epistemic relativism. That is, whether or not one’s evidential seemings are defeated is relative to the felt sensibility of the force of recognizing their contingency. As such, Rosen’s is a very strong relativism, for it makes the justification of one’s belief entirely a matter of the felt sensibility of the individual.¹⁹

Is this strong relativist conclusion correct? I think not. Indeed, I do not think that contingency is the right place to focus in order to locate the evidential significance of seemings, or how they might be defeated. To see this, consider more closely the examples of belief contingency that Rosen himself discusses.

His first and most prominent example is “complacent theism”:

When the complacent theist first asks himself whether or not there is a God, and then notices that there is nothing intrinsically unreasonable in agnosticism, it

¹⁸ One may legitimately wonder here what use the word “rational” has if it does not imply lack of epistemic defect. On p. 85, Rosen says, “even if such persistence is rationally permissible, it seems to me that there is something admirable in the choice to reconsider. Reconsideration manifests a virtue for which we have no standard name: non-complacency, a concern to be responsive to reasons.” But being responsive to reasons *just is* being rational. Analogously to moral norms, if it is a virtue, then to knowingly fail to practice it is a vice. Presumably, by “defective,” Rosen means something like “suboptimal.” On this reading, one is rational to persist, in the sense of not having violated any epistemic norms, even though one is not doing the best one could.

¹⁹ While Rosen does not use the word “justification” here, he does strongly imply that reconsidering due to the recognition of contingency is the epistemically virtuous state of affairs.

would be bizarre to insist that he is nonetheless obliged to remain a theist simply because that is where he started out, and he has been given no ‘positive grounds’ for thinking that he is mistaken. In a case of this sort, retrenchment is obviously an option. The case of nominalism is no different in principle in my opinion.²⁰

It is certainly true that the complacent theist is not *obliged* to remain a theist, but I think that it is at best misleading to say that “retrenchment” is simply an “option.” Surely it is in some sense *unreasonable* for him to become an agnostic just on the basis of coming to believe that agnosticism is not “intrinsically unreasonable” (whatever that may mean).

The complacent theist may have *some* reason to think he’s wrong, but so has the agnostic--or atheist--who meets a reasonable theist. It would be at best misleading to say that theism is not obviously an option because of such meetings. The complacent theist merely sees that it is possible to be a reasonable agnostic--but these do not become *his* reasons just because he now recognizes that they exist.²¹

Rosen continues:

It can be disconcerting to realize that one’s views are strongly contingent in this way. When it first dawns upon the complacent theist that his belief in God is not forced upon him by compelling reasons, but is rather a matter of what he has been raised to take for granted, one response is to say, “There but for the grace of God go I. It’s a matter of luck that I wound up believing. But thank goodness for it. If I had been raised differently, I would have been mistaken.”²² But another possible response is to wonder, “With what right do I suppose that I am one of the lucky ones?” The structure of this familiar predicament is as follows. You have always believed that P for no reason in particular. P is part of the worldview you take for granted: neither inculcated by argument nor supported by argument after the

²⁰ Ibid., 83-84.

²¹ In what sense the recognition that someone else has evidence counts as *one’s own* new evidence will be considered in Chapter Six, which concerns Richard Feldman’s “Meta-Evidence” principle.

²² Interestingly, David Enoch has argued that this sort of reflection is an indicator of the *objectivity* of the facts about which one is speaking. Specifically, the fact that this sort of statement can be made without any humor implied is an indicator that the referent of one’s beliefs are independent of one’s cognitive dispositions. To see this, replace the word “believing” in this response with “disliking spinach,” and note that it becomes humorous, presumably because liking spinach is not independent of one’s cognitive dispositions. See David Enoch, “Why I Am an Objectivist about Ethics (And Why You Are, Too),” in *The Ethical Life: Fundamental Readings in Ethics and Moral Problems*, ed. Russ Shafer-Landau, 4th ed. (New York: Oxford University Press, 2018), 208-211.

fact.²³ You then encounter someone else who finds it natural to doubt it. Neither of you can offer arguments to sway the other. Nor is it plausible to suppose that you have special access to the truth in the relevant domain (e.g., a special sense for detecting the truths in question). In such circumstances the only thing to think is that the differences between you and them [*sic*] are due to historical contingencies which in themselves favor neither option.... How should one respond to this sort of predicament?²⁴

Let's leave aside for the moment the unmotivated claim that the "historical contingencies... favor neither option," a claim which neither party is likely to be in a position to know. Why should the recognition of historical contingency be disconcerting in the first place? Presumably, it is because one may think of "contingent" in the sense of "arbitrary," or "lacking supporting reasons." But this is a confusion to be avoided. It is a matter of contingent history that we now take general relativity for granted. Things may have gone differently if, say, Eddington and his team had failed to detect the bending of light around the sun during the 1919 solar eclipse.²⁵ Nonetheless, belief in general relativity is not *arbitrary*, even if one holds it only because it is an assumption of one's culture. This distinction helps explain why Rosen thinks that two of his cases--the Genghis Khan case and the atomist case--are different from his complacent theist and complacent platonist cases. In the former, one's views--that Khan was morally bad and that atomism is true--are not *arbitrary* even though they are historically contingent. In the

²³ Of course, it does not follow that a belief's being unsupported by argument means that there is no reason to hold it, as Plantinga and others have famously argued. For example, the belief that the world was not created five minutes ago with the appearance of age is a belief for which I can provide no argument, and yet I think I have very good reason to hold it. If they are right about this (and I think they are), then there is no *prima facie* reason to doubt one's belief upon realizing that it is assumed as part of one's worldview, nor is there any *prima facie* reason to criticize another's rationality for not so doubting.

²⁴ Rosen, "Nominalism, Naturalism, Epistemic Relativism," 84.

²⁵ When this possibility was suggested to Einstein, he reportedly replied: "The theory is correct all the same." See Ilse Rosenthal-Schneider, *Reality and Scientific Truth: Discussions With Einstein, Von Laue, and Planck*, ed. Thomas Braun (Detroit: Wayne State University Press, 1980), 74.

latter cases, Rosen thinks, one's beliefs are both historically contingent and arbitrary.²⁶

So when Rosen says, "I can acknowledge the fact that my commitments are contingent without finding myself with anything like a compelling ground for retreating from them," this is because in these cases he is admitting contingency, but denying arbitrariness.²⁷

Call cases like this cases of *mere* contingency. In such cases--Khan and atomism--one has reasons to prefer one's view to others. For example, one has the strong sense that Khan's lifestyle was morally repugnant, and some reason to think that such a view is not culturally relative.²⁸ Also, one knows that scientists from various cultures were apparently motivated by the discovery of truth, and have converged on the idea that some version of atomism is true. One is also not likely to be complacent about these issues--in the sense that one would be unbothered by the suggestion that one's beliefs about them are false--which would constitute a relevant difference from the complacent theist and platonist cases. But even if one were complacent about them, one's view would not be arbitrary. It would be undergirded by lots of reasons, reasons that would presumably be persuasive even to people of different eras if they could be brought to understand them. The only way I see to make the Khan case comparable to the complacent platonist case is if in fact one thinks that one has no better reasons for one's moral views than a reflective Genghis Khan--in other words, that there could be a community of Khanians analogous

²⁶ Note that this reflects an apparent bias on Rosen's part against theism. In the cases of Khanian morality and atomism, the beliefs are not arbitrary because there is reason for one view over others, even if I do not *possess* those reasons. To contrast these views with theism is to assume that it lacks such reasons.

²⁷ Rosen, "Nominalism, Naturalism, Epistemic Relativism," 85.

²⁸ See note 20. If one is tempted to object here that one's moral sense should not count as a reason, note that Rosen himself maintains that it does: "The recognition that my moral view is an historical artifact does not destroy the palpable obviousness that attaches to certain claims about right and wrong. I can dwell indefinitely on how I would have felt if I had been raised by Ghengis [*sic*] Khan. I can acknowledge that I would have reveled in deliberate cruelty, and that this stance would been [*sic*] stable upon reflection. And yet this exercise does not lead me to find deliberate cruelty any the less repellent." *Ibid.*, 86.

to the community of Bedrockers who recognize all the same data about concrete actions but draw wildly different moral conclusions. Presumably Rosen doesn't think this. If he did, then based on what he says about the complacent platonist, it would be a virtue to reconsider his view in light of the reflective Khan community.

Rosen's notion of contingency is made clearer in what he says immediately following the quotation above. There, he references al-Ghazali's use of *taqlid*, a concept in Islamic thought that signifies "'derivative' or 'second-hand' belief."²⁹ Rosen apparently takes this to mean belief that is held "for no reason in particular"; it is simply the belief of one's culture or tradition.³⁰ He then notes that al-Ghazali thinks that abandoning *taqlid* when one becomes aware of it is both the epistemically appropriate and psychologically necessary course of action. Rosen then disagrees with al-Ghazali, and introduces the Khan and atomism cases as examples where it is both psychologically possible and epistemically permissible to maintain one's *taqlid* belief. But if I am right that Rosen is confusing contingency with arbitrariness, then perhaps there is room to interpret al-Ghazali more charitably here. This is because abandoning one's belief when one discovers that it is arbitrary *does* seem to be the appropriate response. On the other

²⁹ Ibid., 84.

³⁰ Note, however, that *taqlid* also has the connotation of "authoritative," in the sense of deference to past interpreters of religious texts whose views are held to be unsurpassed. This would not be an arbitrary belief. Notably for our discussion of disagreement, *taqlid* also had a classical sense in which one was unjustifiably conforming one's view to another. This happened when one who was an authority (a *mujtahid*) conformed his view to one who was not, or to the view of another *mujtahid*, who was considered a peer. This is somewhat ironic considering Rosen's position that there is virtue in rethinking a view when one is presented with evidence of the view's contingency in the form of a peer's disagreement. Here the classical Islamic view just referenced would seem to suggest that the opinion of a peer should have no bearing on one's own. See Bernard G. Weiss, "Taqlid," in *The Oxford Encyclopedia of the Islamic World* (Oxford University Press, 2009), <http://www.oxfordreference.com/view/10.1093/acref/9780195305135.001.0001/acref-9780195305135-e-0785>.

hand, abandoning it when it is contingent but not arbitrary--as with many moral and scientific beliefs--is obviously not appropriate.

The next question to ask is whether disagreement provides one with evidence that one's belief is arbitrary, rather than simply contingent. And the answer, of course, is no. This is the case whether one is an expert about the disputed issue, an informed non-expert, or someone who has no reasons for the belief other than that it is the received view of his tradition. In any of these cases, disagreement highlights the contingency of the belief, but in no way suggests that it is arbitrary in the strong sense of lacking any reasons. Granted, disagreement may serve as an *occasion* for the discovery of arbitrariness, but it is not itself evidence of it. Hence, disagreement is not a defeater of evidential seemings, insofar as it is not an indicator of arbitrariness. But it may still be a defeater in other ways, which will be discussed in the next chapter.

Another thing to note about Rosen's case is that there is an element of tautology in his view, in the sense that the thing that actually motivates his relativism is trivially true for all views of belief. As we've seen, he thinks that one should reconsider a view when one recognizes that his seeming is contingent. As he says,

When the encounter with an alien sensibility destroys this felt obviousness--or when it makes it plain that it was never there to begin with--it knocks the ground out from under what one has previously taken for granted.³¹

The language of "never there to begin with" is telling: it suggests that the virtue of reconsideration accompanies beliefs that are held on the basis of *nothing*. In other words, if one's seeming is removed ("destroyed")--if one has nothing to go on--then it is recommended to revise one's view. But surely this holds for all of the cases Rosen cites.

³¹ Rosen, "Nominalism, Naturalism, Epistemic Relativism," 87.

Who *wouldn't* say that one should revise his belief when he discovers that he has nothing to support it? Notice also that belief revision here is not due to anything peculiar to disagreement, or even to contingency. It is based solely on the loss of one's seeming. This effectively trivializes Rosen's position.

So far, then, we have no reason to think that disagreement provides a defeater for one's evidential seemings, except insofar as it provides evidence of arbitrariness, which it usually does not.

4.2 - *The Epistemic Parity of Seemings*

Some have construed seemings in such a way to make it very unlikely that they *could* be defeated by disagreement. For example, Thomas Kelly has argued that if one "recognizes" that some evidence supports a proposition, then one *knows* that it does, or at any rate, one is justified in thinking that it does, and is also justified in believing the higher-order proposition that one is reasonable to hold that belief.³² If he is right, then disagreement is unlikely to defeat one's recognition in cases where one is in fact correct,³³ and guaranteed not to defeat it if recognition entails knowledge.

The problem here is that "recognize" is ambiguous. It can mean *see*, where one *perceives* that p, and is thereby prima facie justified in p in the normal way that perceptual beliefs are prima facie justified. This would be the sense of the word in usages such as "I recognize that bird as a cardinal." But it can also mean "know," as in "I

³² Thomas Kelly, "Peer Disagreement and Higher-Order Evidence," in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford University Press, 2010), 139.

³³ This is the thrust of the passage: Kelly wants to defend the idea that when one is in fact correct, then one is better justified in thinking that one is correct than when one is not correct. The reason comes from one's recognition.

recognized my friend in the lineup,” but in that case Kelly’s argument would be clearly question-begging. If recognition entails knowing, then the word “recognize” is simply inappropriate for what happens in many cases of determining what a body of evidence supports, which is much more akin to the “seeing” sense of recognition. Kelly is surely right to say that “in some cases, one’s recognition that one’s evidence supports a given belief is based on an unmediated appreciation of that evidence itself.”³⁴ But this is just what we’ve been calling a seeming. It would be more precise to think of the relevant sense of “recognition” in the following way: the claim “My evidence supports p” is prima facie justified when based on the experience of E supporting p, but can be defeated by, for example, further higher-order evidence that one’s initial evidence really did not support p. The question, again, is whether disagreement provides such higher-order defeating evidence.

Kelly raises an interesting point here: it seems that if one person has done a better job assessing the evidence than another, then the conclusion of the former should not be defeasible in the same way as the conclusion of the latter. For example, if you and I disagree about, say, Davidson’s theory of indirect discourse, and you have assessed his argument correctly, and I have not--assume that we’ve shared all the reasons we can think to share and remain deadlocked, each finding the other’s assessment weak--it seems wrong that your view should be overcome just as easily as mine. The problem is that neither of us are in a position to know which of us has made the incorrect assessment. If each of us has a seeming (or “recognition”) that the evidence supports his view, then each of us is prima facie justified in believing that the evidence supports his view. Now, for all

³⁴ Ibid., 159.

we've said about the Uniqueness Thesis, it may be possible that the evidence supports both our views, but let's say in this case it does not. How is either of us to decide which of our beliefs ought to be defeated by the evidence presented by the other's disagreement? "From the inside," as Kelly says, "a case in which you fail to appreciate the genuine merits of what I say on behalf of my view because of dogmatic commitment on your part might seem just like a case in which my defense is indeed without merit."³⁵

I'll call this the Parity Problem:

Parity: From the perspective of either party in a peer disagreement, the epistemic situation is indistinguishable from one in which the strength of the peers' evidence assessments are inverted. Things would appear just as they do to the correct party if she were incorrect, and vice versa.

For his part, Kelly says that "the fact that it might be difficult to tell which of these is the case does not mean that it makes no difference."³⁶ But it may mean that it can make no difference *to me*, even if it is true that it ought to make a difference in some all-things-considered, epistemically idealized way.

Another philosopher who holds that seemings are unlikely to be defeated by disagreement is Michael Bergmann. In essence, he concludes that there simply is no way to solve the Parity Problem--the best one can do is to maintain one's view based on one's own seemings, no matter how dialectically unsatisfying this may be. It will be helpful to consider his view here.

In "Rational Disagreement after Full Disclosure," Bergmann defends a form of Nonconformism, according to which reasonable disagreement with an epistemic peer is

³⁵ Ibid., 165.

³⁶ Ibid.

possible.³⁷ Bergmann understands peerhood to be a matter of two (or more, presumably) peers who are, and take themselves to be, equally “intellectually virtuous,” where intellectual virtue implies “intelligent, thoughtful, and sincerely seeking the truth.”³⁸ “Full disclosure” means that each party has taken the time to lay out all the evidence for his position that he can think of. Suppose that disagreement persists after such disclosure, and that both parties are still convinced of their equal intellectual virtue. Are they rational to be convinced of this?

Bergmann distinguishes between “internal” and “external” rationality. Consider a person who comes to believe *p* (where *p* is a noninferential belief) on the basis of experience *e*, and suppose it is the case that the reasonable thing to do when confronted with experiences of “the same phenomenal type as *e*” is to believe *p*. Suppose further that this person believes *p* as a result of brain damage which caused the experience. Is this belief rational? Bergmann here borrows Alvin Plantinga’s distinction between internal and external rationality.³⁹ For Bergmann, “a belief is *internally rational* if and only if it is an epistemically appropriate response to the subject’s mental states,” and it is “*externally rational* if and only if the believer’s cognitive processing mechanisms are working as they epistemically should be in producing the belief (including where their working well is not in response to the subject’s mental states).”⁴⁰ So external rationality rules out rational beliefs formed due to cognitive malfunction, while internal rationality does not. With this distinction in mind, Bergmann provides some answers to the question of

³⁷ Michael Bergmann, “Rational Disagreement After Full Disclosure,” *Episteme* 6, no. 3 (2009): 336-53.

³⁸ *Ibid.*, 336.

³⁹ See Alvin Plantinga, *Warranted Christian Belief* (New York: Oxford University Press, 2000), 110-112.

⁴⁰ Bergmann, “Rational Disagreement after Full Disclosure,” 337.

whether the parties to the disagreement are rational to hold each other to be equally intellectually virtuous.

With respect to internal rationality: imagine a case of peer disagreement in which peers S1 and S2 are such that S1 believes p and S2 believes $\sim p$.⁴¹ Further, each has a broader outlook, O1 and O2, of which their beliefs about p are a part. Part of S1's outlook O1 is an error theory about why O2 seems true but isn't, and vice versa for S2's outlook. Further, these beliefs of S1 and S2 in p and $\sim p$, and in the error theories of O1 and O2, respectively, are based partially on "apparent insight" into the truth of these propositions, where an "insight" is a special case of "seeing that" a proposition is true, in other words, evidential seemings.⁴² As Bergmann says, "If there are no defeaters, the rational response to a strong apparent insight that p --like the rational response to a strong apparent memory that p --is to believe p ."⁴³ Now, note that S1 and S2 *do not have the same evidence*, even after full disclosure. After all, *having* an insight (even an apparent one) and *hearing* about someone else's insight are not the same thing. Note also that S2's apparent insight that $\sim p$ does not count as a defeater for S1's belief that p because S1 has an error theory that anticipates that belief (even from someone as intellectually virtuous as oneself) and explains why it's wrong. The same thing is of course also true for S2. So, it seems that when internal rationality is intended, it *is* possible for both parties to rationally hold themselves to be equally intellectually virtuous after full disclosure of evidence.

With respect to external rationality, Bergmann argues as follows. If it is possible for two disagreeing parties to be internally rational while taking themselves to be equally

⁴¹ Bergmann does not use the word "peer," but it is convenient here.

⁴² See Bergmann's endnote 5, p. 350.

⁴³ *Ibid.*, 339.

intellectually virtuous, then this suggests that they may be externally rational while doing so as well. This is because all that is required in order to be wrong and yet still externally rational is a little epistemic bad luck.⁴⁴ One of the parties is generally reliable but in this case (against the odds) wrong. Nonetheless, both parties remain externally rational.

Bergmann says it is also possible that “both parties are externally rational in continuing to disagree *and* in thinking that the other may well be externally rational in continuing to disagree.”⁴⁵ This is because “may well be” implies that, from the vantage point of one of the parties, the other party may be--for all one knows--a victim of said epistemic bad luck, even though everything else is working as it should. Nonetheless, Bergmann holds that in such cases--in which one thinks that another is internally rational but mistaken--one will *likely* conclude that the other is in fact externally irrational, especially if the disagreement is a persistent one (more on this below).

So it turns out that “recognized disagreement about p with someone you view as roughly equal to you in intellectual virtue” *need not* provide a defeater for your belief that p. But Bergmann also thinks that it *may*, in the following kind of case:

If in response to recognizing that S disagrees with you about p (which you believe), you either do or epistemically should disbelieve or seriously question or doubt the claim that you are, on this occasion, more trustworthy than S with respect to p, then your belief that p is defeated by this recognition; otherwise, not.⁴⁶

⁴⁴ As Bergmann says, “...although the other person’s belief is externally rational and formed in an epistemically appropriate environment, that guarantees only that it is *likely* to be true, not that it is true. Suppose you believe you own a car and that this belief is externally rational and formed in an epistemically appropriate environment. If, a few seconds ago, unbeknownst to you, a stray meteorite destroyed your only car, then your belief is false. Here the explanation is that the reliability of externally rational beliefs formed in an epistemically appropriate environment is combined with unlikely falsehood.” Ibid., 341.

⁴⁵ Ibid.

⁴⁶ Ibid., 343.

So when is it the case that one “epistemically should” question whether one is more trustworthy on this occasion than a peer? Unfortunately, says Bergmann, there is no uncontroversial answer to this question, even if we look at specific cases:

It’s unfortunate that, even after full disclosure of evidence among people seemingly equal in intellectual virtue, disagreement persists about precisely when disagreement provides a defeater. But I think that’s what we’ll find.⁴⁷

A Conformist like Elga or Christensen might here raise the following objection: what reason does either party have to assume that *she* is the one who has avoided error on this occasion? Bergmann’s response to this is already quoted above in another context:

I think that the rational response... is to continue believing *p*. However, I don’t have an argument for that conclusion, just as I don’t have an argument for the conclusion that the rational response for us to a tactile experience like the one we typically have when grabbing a billiard ball is to believe something like ‘that’s a small hard spherical object’.⁴⁸

Bergmann is here appealing to the basic nature of evidential seemings in order to reject the demand for an argument supporting a bias in one’s own favor. He provides two cases to help make his point:

Case 1: We may not be able to satisfactorily answer the radical skeptic, but this does not thereby provide a defeater for our perceptual beliefs.

Likewise, the fact that we think someone is mistaken (in a moral or religious or political disagreement with us) despite her being confident in ways similar to the ways in which we are confident doesn’t automatically mean that we should seriously question or doubt the reliability of our beliefs in which we’re so confident.⁴⁹

⁴⁷ Ibid.

⁴⁸ Ibid., 345.

⁴⁹ Ibid.

Case 2: Imagine an instance of obviously morally wrong behavior,⁵⁰ and imagine further that two other parties disagree with you about its wrongness--one because he is an ethical egoist, and the other because he is a moral nihilist. Assume that you're peers with respect to intellectual virtue and other relevant factors, and that you discuss it and reach full disclosure. Should you "have significant doubts about the reliability of your *own* apparent insight that [the] behavior is morally wrong? Is this what internal rationality requires?"⁵¹ Of course Bergmann thinks not. Your own apparent insights are based on seemings which provide great confidence in your belief and which you have no reason to give up just because someone else apparently lacks those seemings or has different ones.

Unfortunately, Bergmann underestimates the force of the seeming that one should modify her view in light of the disagreement of an informed peer about issues regarding which one's seemings are not as clear and vivid as those in the cases of current perceptual experiences or grossly immoral behavior. When informed peers disagree about, say, the causes of the gender pay gap, and they reach full disclosure, it is much less clear that either of them is entitled to say, "Well, my insight supports p, and even though I have no argument for preferring my insight to yours, I am reasonable to believe p." If asked what makes one reasonable to do so in this case, one will simply not be able to say. And as Rosen says, if one has nothing to say, then one should be less confident in one's view.⁵² Since we are imagining a situation of parity, neither party to the disagreement will be able to determine that her own assessment is the correct one, outside of her own

⁵⁰ Bergmann here falls prey to the regrettable habit among philosophers of using examples involving depraved cruelty.

⁵¹ *Ibid.*, 346.

⁵² Rosen, "Nominalism, Naturalism, Epistemic Relativism," 83.

evidential seemings, and both parties will know this about their own and their peer's situation.

So even if Bergmann has successfully shown that it is *possible* for peer disagreements to be rational,⁵³ as in the case of epistemic bad luck, he has not established that one is reasonable to think that one's *own* peer disagreement is rational, especially if, as he admits, it is a persistent disagreement. In his own words:

I think the most natural outcome of a case where both parties to a disagreement sensibly believe the other is internally rational but mistaken is that both parties will also sensibly think the other's belief is *probably* externally irrational--i.e., the other probably has some cognitive processing problem that has resulted in her having misleading experiential evidence for a false proposition. ... But at most one of the parties to the dispute will be basing her belief in the disputed proposition (and in the probable external irrationality of the other) on *nonmisleading* evidence.⁵⁴

But of course neither party can know if *she* is the one with the nonmisleading evidence, so we're back to the Parity Problem.

Bergmann avoids this problem only by defining "intellectual virtue," which is his stand-in for "peerhood," in an idiosyncratic and unmotivated way, saying that "in thinking that one's friend is roughly on a par with oneself in terms of intellectual virtue, one needn't think that the apparent insights of one's friend are produced in a cognitively healthy and fitting way." But this is precisely what most commentators on peerhood have thought, and what strict peerhood would seem to require. If I do not think that your seemings are produced in a cognitively healthy way, then there is no interesting epistemic

⁵³ Technically speaking, he has argued only that disagreements between equally epistemically virtuous people are rational, with epistemic virtue being taken to be a matter of "intelligence, thoughtfulness, and sincerity in seeking truth." Were he to take certain other views of peerhood, such as that of Elga, for example, his argument would not work. See Bergmann, "Rational Disagreement after Full Disclosure," 342. Competing views of peerhood will be discussed in Chapter Seven. See especially footnote 8 in that chapter.

⁵⁴ *Ibid.*, 342.

problem with respect to our disagreement. So Bergmann avoids the primary issue by instead answering a more trivial problem.

Another possible avenue for dealing with the Parity Problem is to claim that there is a presumption in favor of one's own view just by virtue of it being one's own. This has been called an "egocentric epistemic bias," most notably defended by Ralph Wedgwood.⁵⁵ According to Wedgwood, it is reasonable to assume that the *other* party in a disagreement is mistaken, and not oneself. This is a rational stance because it is rational to have a primitive trust in one's own intuitions--i.e., to regard them as reliable, take them at face value, even if one has no independent or antecedent reason to do so. But, Wedgwood argues, it is more rational to have this sort of primitive trust only in one's *own* intuitions, rather than in someone else's, because this minimizes risk of error, which is already quite high. He explains:

In general, it seems to be a general requirement of rationality that one must in some non-arbitrary way minimize the number of sources of information that one has this sort of primitive trust in. (The non-arbitrariness condition is crucial here: it would, for example, be arbitrary to have a primitive trust in one's vision but not in one's hearing or one's sense of touch.) If I have a primitive trust only in my own intuitions, then clearly there is a smaller number of sources of information that I have such a primitive trust in than there would be if I had such a primitive trust in everyone's intuitions. Moreover, it seems to me that it would not be arbitrary for me to have this sort of primitive disposition to trust only my own intuitions: after all, my own intuitions by their very nature involve my having an inclination to believe the corresponding proposition, whereas other people's intuitions do not involve *my* having any such inclination.⁵⁶

In other words, there is an asymmetry between the epistemic advisability of relying on one's own intuitions and relying on those of other people. Hence, owing to this rational

⁵⁵ Ralph Wedgwood, "Disagreement and the A Priori," in *The Nature of Normativity* (Oxford University Press, 2007): 261.

⁵⁶ *Ibid.*

egocentric bias, one may reasonably conclude in the face of peer disagreement (even if one is forced to lower one's confidence in one's view) that it is the *other* party who is wrong.

Is the egocentric epistemic bias a solution to the Parity Problem? I think not, though I do think it gets at something important that is worth taking into account. It is true that there is an initial asymmetry between the trust one should place in one's own seemings and the trust one should place in the reported seemings of a peer. And it may be true that this asymmetry is justified in much the way that Wedgwood describes. This is not an insignificant point. It entails that there is a *prima facie* presumption in favor of one's standing view (assuming it is accompanied by the right sorts of evidential seemings) in any dispute. This puts the burden of proof squarely on the disagreeing party to overcome this presumption. Given that many disputes never move past the initial stage of recognition of disagreement, this entails that people are often within their epistemic rights to go on believing as they have after becoming aware of disagreement with those they take to be peers. I am, after all, epistemically responsible for *me*, not for my peer, and all I have to go on are my own seemings.

Nonetheless, the egocentric bias cannot solve the Parity Problem because there the relevant comparison is not between one's own seeming and the report of someone else's--it is between the *subject's own conflicting seemings*. In a sustained disagreement with an epistemic peer and shared evidence, the Parity Problem stipulates that I will eventually reach a point where I cannot tell, *phenomenally*, whether I am in the right or not, because I recognize that things would seem just the same to me in either situation. So I have two seemings:

- (1) the evidential seeming that suggests (and produces justified belief that) E supports p, and
- (2) the higher-order seeming that I have reason to believe that (1) is not accurate.

An egocentric epistemic bias, even if true, will not help here.

Now, it is of course true that I *always* have some reason to think that I might be mistaken. But this worry only becomes salient if I am confronted in some way with the need to worry about it. The peer disagreement problem is akin in this way to the broader epistemological problem of skepticism, and it has been recognized as such by prominent commentators.⁵⁷ Bryan Frances calls this “disagreement skepticism,” and defines it as follows:

The disagreement skeptic focuses on beliefs that satisfy the following recognition-of-controversy conditions. You know that the belief *B* in question has been investigated and debated (i) for a very long time by (ii) a great many (iii) very smart people who (iv) are your epistemic peers and superiors on the matter and (v) have worked very hard (vi) under optimal circumstances to figure out if *B* is true. But you also know that (vii) these experts have not come to any significant agreement on *B* and (viii) those who agree with you are not, as a group, in an appreciably better position to judge *B* than those who disagree with you.⁵⁸

Thus, disagreement skepticism is significantly different from traditional global skepticism, as its effects are limited to beliefs that satisfy conditions (i)-(viii). As Frances says, this form of skepticism is limited in three ways:

⁵⁷ See Bryan Frances, “Scepticism and Disagreement,” in *Skepticism: From Antiquity to the Present*, ed. Diego Machuca and Baron Reed (New York: Bloomsbury Academic, 2018), 581–91; David Christensen, “Disagreement as Evidence: The Epistemology of Controversy,” *Philosophy Compass* 4, no. 5 (2009): 760; and the articles assembled in Diego Machuca, ed., *Disagreement and Skepticism* (New York: Routledge, 2013).

⁵⁸ Bryan Frances and Jonathan Matheson, “Disagreement,” in *The Stanford Encyclopedia of Philosophy*, 2018, <https://plato.stanford.edu/archives/spr2018/entries/disagreement/>, Section 7.

- (1) It is *contained*: “Only beliefs that meet something like the recognition-of-controversy conditions are subject to this skeptical threat.”⁵⁹
- (2) It is *contingent*: “The nature and extent of disagreements are both contingent matters, so since disagreement skepticism relies on these factors, the skeptical consequences of disagreement are also contingent.”⁶⁰
- (3) It is exclusively *epistemic*: it does not imply anything about “what moral, prudential, or even religious reasons you may have for holding a controversial belief. The skeptical threat from disagreement only concerns our epistemic reasons.”⁶¹

The effects of the Parity Problem would also be similarly limited, since there are various conditions one must meet to be in it, analogous to Frances’s recognition-of-controversy conditions. Nonetheless, these effects are serious: they include the fact that many controversial beliefs, in areas as important as morality, religion, politics, and philosophy, can be defeated—even when they start out as justified—by prolonged peer disagreement. This includes beliefs for which one has clear, vivid, and stable evidential seemings. Thus, there is something to the Conformist stance: via the Parity Problem, there are many cases in which a person is seemingly epistemically obligated to alter her view because of peer disagreement. Potential solutions to the Parity Problem will be discussed in Chapters Eight and Nine.

⁵⁹ Ibid.

⁶⁰ Ibid. Of course, one could have a contingent disagreement that entails contingent skepticism about a necessary proposition, such as (arguably) the proposition that God exists.

⁶¹ Ibid.

Are there other ways that peer disagreement may lead to the defeat of one's evidential seemings? To answer this, we need to understand what evidential value agreement or disagreement with another carries on its own. This is the subject of the next chapter.

Chapter Five – The Evidential Value of Disagreement Itself

5.1 - Idealization, Independence, and Bootstrapping

The clearest explanation in the disagreement literature of the epistemic significance of disagreement--and the clearest justification for the discussion of peerhood--is due to Jennifer Lackey. She notes that many disagreements are due to differences in access to evidence, cognitive inequalities between the parties, or other epistemic defects such as bias. Such cases would not easily yield what is epistemically interesting about disagreement *per se*. This is because, as she says,

...adjusting our doxastic states in all of these sorts of cases does not reveal anything significant about disagreement itself, since the fact that you and I disagree drops out of the explanation of this adjustment; it is, for instance, the difference in our familiarity with the relevant evidence or the asymmetry in our cognitive capacities that does the explanatory work.¹

To examine the effect of disagreement itself, we must control for these sorts of inequality, which means we must assume some notion of *peerhood*. No other method captures the intrinsic significance of disagreement as well as this, since it allows us to isolate disagreement as an epistemic issue.²

Defining peerhood is not my aim here--that is the subject of Chapter Seven--I want merely to note that the idealization of peer disagreement cases in the literature is primarily motivated by the reason Lackey explains.³ Idealizing cases somewhat (e.g., by

¹ Jennifer Lackey, "A Justificationist View of Disagreement's Epistemic Significance," in *Social Epistemology*, ed. Adrian Haddock, Alan Millar, and Duncan Pritchard (Oxford University Press, 2010), 298-299. Note that Lackey's point here follows only if factors other than the disagreement account *completely* for the belief adjustment.

² *Ibid.*

³ For her part, Lackey finds this idealization a bit overwrought. She says, "...it is very common for philosophers writing on this topic to motivate interest in it by citing debates in history, philosophy, politics, religion, and other areas where disagreement is widespread and impassioned. But these debates bear very

considering cases of epistemic peers) helps to focus on the epistemic nature of disagreement itself. Our question here is what contribution the disagreement with another makes to one's total evidence base. One way of approaching this question is to begin with the related question: What evidential difference does *agreement* make?

Thomas Kelly approaches this question via a consideration of the Common Consent Argument (CCA) for the existence of God. He first notes that when a large number of peers *independently* converge on some opinion, that affords that opinion significant evidential weight.⁴ He then notes that:

As arguments go, the Common Consent Argument for the Existence of God is not exactly an overwhelming one, possessing as it does the twin defects of transparent invalidity and the having of an obviously false claim as its sole premise. Nevertheless, even though *God exists* does not follow from *Everyone believes that God exists*, we can ask: if it were true that everyone, or almost everyone, believed that God exists, how much support would that lend (if any) to the proposition that God exists?⁵

The answer, he says, depends on whether those people came to their belief independently or not. If I believe *p* merely because you believe *p*, then the fact that two people now believe *p* is no more epistemically significant (i.e., it affords no greater evidence for *p*)

little resemblance to the hyper-idealized scenarios under consideration here.” (Ibid., 311) I sympathize with Lackey's view here, even while disagreeing with other statements she makes. For example, she continues: “It would rarely, if ever, happen that two people continue to disagree with one another about, say, gay marriage where there is evidential and cognitive equality with known epistemic symmetry in situations of full disclosure.” I disagree. Even idealized symmetry seems plausible in cases where the parties have known each other for a long time, and/or have had sufficient interaction to become acquainted with the personal characteristics of the other. Lackey claims that “...conclusions drawn from hyper-idealized situations involving disagreement ultimately tend to have very little connection to the disagreements we face every day.... I think that, at least for the most part, ordinary disagreement ought to be the focus of this debate.” (Ibid.) I certainly agree that the philosophical debate ought to be sensitive to--even constrained by--what disagreement is like in real life. However, I deny that the cases Lackey describes are all unrealistic. The seminal examples in the disagreement debate (e.g., the disagreement we discussed above between van Inwagen & Lewis on free will, or the older debates in religious pluralism) are very real and yet exhibit a high degree of what Lackey calls idealism.

⁴ Thomas Kelly, “Peer Disagreement and Higher-Order Evidence,” in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford University Press, 2010), 146.

⁵ Ibid., 147.

than the fact that you believed *p* all by yourself before convincing me. “In principle,” he says, “the fact that a small handful of people arrive at the same belief independently of one another might be better evidence that that belief is true than if many millions of people arrive at the same belief non-independently.”⁶ This entails that if one is in disagreement with a whole group of people about some issue, the scales are not obviously shifted in favor of the group, if the members of the group have not arrived at their view independently of one another.⁷ Kelly borrows an example here from G.A. Cohen, who recounts a division among philosophers from his days as a student at Oxford:

Now people of my generation who studied philosophy at Harvard rather than at Oxford for the most part *reject* the analytic/synthetic distinction. And I can’t believe that this is an accident. That is, I can’t believe that Harvard just *happened* to be a place where both its leading thinker rejected that distinction and its graduate students, for independent reasons--merely, for example, in the independent light of reason itself--also came to reject it. And vice versa, of course, for Oxford. I believe, rather, that in each case students were especially impressed by the reasons respectively for and against believing in the distinction, because in each case the reasons came with all the added persuasiveness of personal presentation, personal relationship, and so forth.⁸

So, as someone trying to decide what to think about the analytic/synthetic distinction, one need not be overwhelmed by the number of people who reject (or accept) it, since in actuality, there is only one independent argument, or set of arguments, represented by that group.⁹

⁶ Ibid.

⁷ A danger lurks here with respect to testimony. Kelly’s view could be read as Lockean to the extent that beliefs rooted in testimony are seen as non-independent, and therefore epistemically less valuable, than beliefs rooted in some other knowledge source, such as direct perception. If in fact this is an entailment of Kelly’s view, I take that to be a further weakness of his account, in addition to the one I here discuss. I discuss the relationship between testimony and disagreement below.

⁸ G.A. Cohen, *If You’re an Egalitarian, How Come You’re So Rich?* (Cambridge: Harvard University Press, 2000), 18. Quoted in Kelly, “Peer Disagreement and Higher-Order Evidence,” 147-148.

⁹ Note again the implications for testimony: nearly *every* belief the average person holds about the world is susceptible to Cohen’s critique, since almost no one is in a position to independently assess most of the things they believe. So if being affected by the “persuasiveness of personal presentation, personal

Now, this is true to an extent, but there is a danger here: we do not want to go so far as to say that a substantial number of people who believe *p*, and who trace that belief to a common source, carries *no* epistemic weight. That would ignore the fact that it's well-nigh impossible for every individual to do all of her own investigating, given how many propositions most individuals need to believe or disbelieve. It is perfectly rational to accept the well-reasoned views of one's community, with or without investigating them for oneself, and if the views are sufficiently detailed and accurate to have survived in a large population for a long time (a population in which it is the job of *some* members to investigate, refine, and if necessary, alter these views), then there is no reason the average person shouldn't take the commonality of a view as evidence of its truth.¹⁰

Kelly clarifies his view as follows: "... insofar as one believes as one does because this is what one's teacher believes, the fact that one believes as one does is not an *additional* piece of psychological evidence, over and above the psychological evidence

relationship, and so forth" is a mark against the epistemic value of a belief, then it is a mark against nearly all beliefs.

¹⁰ J.L. Austin says something similar, not regarding ordinary belief but ordinary language. Still, what Austin says seems to apply, *mutatis mutandis*, to ordinary belief:

Certainly ordinary language has no claim to be the last word, if there is such a thing. It embodies, indeed, something better than the metaphysics of the Stone Age, namely, as was said, the inherited experience and acumen of many generations of men. But then, that acumen has been concentrated primarily upon the practical business of life. If a distinction works well for practical purposes in ordinary life (no mean feat, for even ordinary life is full of hard cases), then there is sure to be something in it, it will not mark nothing: yet this is likely enough to be not the best way of arranging things if our interests are more extensive or intellectual than the ordinary. And again, that experience has been derived only from the sources available to ordinary men throughout most of civilised history: it has not been fed from the resources of the microscope and its successors. And it must be added too, that superstition and error and fantasy of all kinds do become incorporated in ordinary language and even sometimes stand up to the survival test (only, when they do, why should we not detect it?). Certainly, then, ordinary language is *not* the last word: in principle it can everywhere be supplemented and improved upon and superseded. Only remember, it *is* the *first* word.

See J.L. Austin, "A Plea for Excuses: The Presidential Address," *Proceedings of the Aristotelian Society, New Series* 57 (1956-1957): 11.

afforded by the teacher's belief."¹¹ This seems correct: one person believing *p* merely because someone else believes *p* does not provide additional evidence for *p*. But Kelly continues this line of thought in a footnote:

... the need to discount the numbers is not limited to cases in which there is causal dependence present, as in the examples considered above. If I know that two individuals will respond to given evidence in the same manner, then I should treat their having arrived at some particular answer as *one* piece of evidence, and not two pieces of evidence, in favor of that answer (even if their both having arrived at that answer is in no way underwritten by some causal link).¹²

This is surely mistaken. I think here of William P. Alston, who reportedly defended his dissertation in front of Charles Hartshorne, Rudolf Carnap, and W.V.O. Quine, the first two of whom were on his committee.¹³ Let's say Alston agreed with something his committee members also believed, and for the same reasons--surely it is possible that the assent of that prodigious mind could have contributed *additional* support for the view besides what it already enjoyed. Similarly, imagine two mathematicians examine the same proof and decide that it is successful. If I believe that it is successful on the basis of both of their testimony, then my belief is better justified than if it were based on the testimony of only one of them, even though I know that they responded to the evidence "in the same manner."¹⁴ I am inclined to think that the perspective of certain thinkers, who exhibit certain characteristics like incisiveness, care, thoroughness, creativity, etc., count for something, even when they agree with the perspective of others,

¹¹ Kelly, "Peer Disagreement and Higher-Order Evidence," 148.

¹² *Ibid.*, 149, footnote 28.

¹³ This anecdote is related in a memoriam written by Tom Senor, archived here: <http://web.archive.org/web/20130510142744/http://prosblogion.ektopos.com/archives/2009/09/memories-of-bil.html>.

¹⁴ Thanks to Michael Wreen for this example.

and *for the very same reasons*. One might argue that this is still a matter of independence, with the difference that here the independence is a trait of personality, rather than of reasons. Nonetheless, this serves as a counterexample to Kelly's claim that two individuals responding to the evidence in the same way is only ever one piece of evidence. I believe that it can be additional evidence, and whether or not it is depends on the epistemic characteristics of the individuals, and not merely on how similar the contents of the arguments are between them.

So we can discern two senses in which agreement has epistemic weight over and above what a testimonial report would typically carry on its own: when the view of others is either (i) supported by multiple independent lines of evidence (represented by individuals arriving at the view independently of one another), or (ii) supported by the assent of individuals who are independent and especially careful reasoners.

Another way to approach this question is via a disagreement over the so-called Equal Weight View, championed by Adam Elga. This view is a species of the Conformist view, which holds that peer disagreement provides one with reason to adjust one's own view, and adds that it is most rational to accord *equal* weight to the view of one's peer and to one's own view.¹⁵ Elga contends that views which reject this conclusion--i.e., views which hold that it is reasonable to accord more weight to one's own view than to one's peer's--incur a bootstrapping problem. As he says:

...suppose that it was legitimate to give your own evaluations more weight than those of a friend who you initially count as a peer. Then it could be legitimate for you to "bootstrap"--to come to be confident that you are a better evaluator than the friend merely by noting cases of disagreement, and taking it that the friend made most of the errors. But that is absurd. So it is not legitimate to give your own evaluations more weight than those who you count as peers. A similar

¹⁵ Adam Elga, "Reflection and Disagreement," *Noûs* 41, no. 3 (2007): 478-479.

argument shows that it is not legitimate to give your own evaluations less weight than those who you count as peers. So the equal weight view is correct.¹⁶

If this is correct--if Nonconformist views have a bootstrapping problem--then it follows that disagreement itself is evidentially significant. This is because if it is illegitimate to give one's own view more weight than a peer's based on the same evidence, this can only be because the fact of disagreement itself acts as a counterweight against the first-order evidence one had prior to the disagreement. It is not that one assumes the peer has evidence one lacks--this is ruled out by stipulating shared evidence. And it is not that one assumes that the peer's reasoning abilities indicate evidence one has missed--Elga explicitly disallows this by insisting on the "independent evidence only" view discussed above.¹⁷ What makes the Equal Weight View the correct view, Elga thinks, is that it recognizes that disagreement itself carries epistemic weight sufficient to defeat the evidential seemings one possesses for her view that *p*. The bootstrapping objection is intended to bring this out by showing that denying disagreement's epistemic significance commits one to inventing justification out of thin air by saying, in effect, "We disagree; you're wrong; therefore, we're not peers on this issue." If disagreement provides its own counterbalancing evidence, one would not be able to dismiss the view of an assumed peer.

¹⁶ Ibid., 488. Because Elga, like most commentators in the disagreement debate, limits peerhood to local disagreements over specific propositions, it would not follow immediately that non-Equal Weight views entail that a person is justified in thinking that he is a better evaluator *full stop* than anyone who disagrees with him. It follows only that he can bootstrap his way into thinking he is a better evaluator *of p*. Nonetheless, Elga would likely think that enough such disagreements *would* legitimate the belief that one is a better evaluator than one's peer entirely.

¹⁷ That is, if we are peers, then for Elga, we are equally likely to be correct, regardless of what we take our evidence about *p* to be. He therefore thinks that the only thing that can break a stalemate between us is other evidence that is independent of our dispute about *p* but that bears on our likelihood to be correct in the dispute. This has the effect of rendering the specific propositional content of our evidence, *and* the specific characteristics that make us peers, evidentially irrelevant.

However, the bootstrapping objection does not work. Or, to put it differently, it works *too well*, since it cuts against the Equal Weight View itself. Elga contends that when someone considers her own view to be superior to her opponent's, she is bootstrapping--even if she is *correct*--since she has no evidence that she has in fact evaluated the evidence correctly. This much is true. However, for Elga, the person also has no evidence that she is *mistaken*. This counterintuitive result is a consequence of Elga's independent evidence only requirement. Notice: intuitively, the way that peer disagreement justifies counterbalancing one's confidence in a standing belief with that of a peer is by providing higher-order evidence that one has misjudged the first-order evidence for p. Why does peer disagreement provide this meta-evidence? Because a *peer* is as capable as oneself to correctly judge the evidence for p. How do we know if someone is a peer with respect to p? According to Elga, it is by observing how that person judges the evidence pertaining to allied issues related to p. Considering a disagreement about abortion, for example, Elga says: "setting aside [another's] reasoning about abortion does not set aside her reasoning about allied issues.... the accuracy of an advisor's views on these allied issues indicates how accurate the advisor is likely to be, when it comes to abortion."¹⁸ Since in "messy real-world cases, the disputed issues are tangled in clusters of controversy," one cannot help but determine the peerhood of an interlocutor based on issues related to the one under dispute.¹⁹

The upshot of this, says Elga, is that one will only count as peers those who "have views that are similar to one's own."²⁰ Hence, in cases like the abortion case, I *cannot*

¹⁸ Ibid., 493.

¹⁹ Ibid.

²⁰ Ibid., 494.

appeal to independent evidence to break a stalemate with a peer because a peer is simply someone who probably already agrees with me in general, so there shouldn't be a peer disagreement in the first place. But let's say that, against the odds, there is. Given Elga's view of how judgments of peerhood work, one would not be justified in *lowering* her confidence in p in response to peer disagreement, because there is no independent evidence to appeal to that would justify this change. Any such evidence has already been accounted for in deciding that one is a peer with respect to p, which in Elga's view means that one is *equally likely to be correct* about p as oneself.²¹ But if one has no evidence independent of the dispute to appeal to in order to justify thinking that one is more likely to be correct about p, then for the same reasons one has no evidence independent of the dispute which justifies thinking that one is *less* likely to be correct about p. Elga thinks this implies that one should "split the difference" and assign equal weight to both views. But if I am right, any change in confidence, up or down, would be unjustified on Elga's account. This is based on the intuitively plausible principle that a change in credence for p (in either direction) requires justification, and owing to the independent evidence only restriction on judgments about peerhood, no such justification exists in these cases. So *decreasing* one's confidence in p (say, by according it "equal weight" as the opponent's view) as a result of disagreement is a kind of "reverse" bootstrapping.²²

5.2 - David Christensen on Symmetry between Apparent Peers

²¹ This conception of peerhood will be critiqued in Chapter Seven.

²² Kelly also makes a compelling bootstrapping argument against the Equal Weight View itself. See Kelly, "Peer Disagreement and Higher-Order Evidence," Section 3.2, 125-128.

David Christensen has also taken peer disagreement to constitute defeating evidence which counterbalances one's evidential seeming. Consider his restaurant bill case: two friends at a moderately priced restaurant mentally calculate their split bill and come to slightly different totals for their share, say, \$43 and \$45. Christensen describes his initial intuition about this case as follows:

If we set up the case in this way, it seems quite clear that I should lower my confidence that my share is \$43 and raise my confidence that it's \$45. In fact, I think (though this is perhaps less obvious) that I should now accord these two hypotheses roughly equal credence.²³

What could motivate the idea that the hypotheses should be accorded equal credence?

Christensen admits that this is not obvious, and I suspect it is not obvious because Christensen's view assumes that disagreement itself provides *equal* evidence rather than merely being an indicator that *there is* evidence against one's view. My own intuition about the bill calculation case is that one should lower one's credence about the amount pending recalculation, and form no beliefs *at all* about the peer's answer. Assigning the answers equal credence is unwarranted because I lack sufficient evidence to overrule my seeming until I know more about why the peer takes a different view.

Alternatively, consider Christensen's discussion of a case he borrows from Kelly, in which two meteorologist peers interpret a set of weather data differently. Kelly says that in a case like this--in which each takes the other to be a peer initially and then discovers their disagreement--from the perspective of one of them, the disagreement provides evidence that they're not peers after all. About this, Christensen says:

It seems to me that there is clearly something right about this line. As applied to the weather case, my discovering that my friend has reached what seems to me to

²³ David Christensen, "Epistemology of Disagreement: The Good News," *Philosophical Review* 116, no. 2 (2007): 193.

be the wrong conclusion does constitute evidence that she has made a mistake, and thus does give me reason to trust her opinion less than I ordinarily would. However, another point needs equal emphasis: the fact that she disagrees with my prediction also constitutes evidence that *I* have made a mistake. So it's not clear so far that any asymmetry has developed.²⁴

So Christensen thinks the evidential situation here is a wash, because the fact that she disagrees constitutes evidence that I have made a mistake, and that evidence counts the same as the evidence I have that she made a mistake.

But is this right? It seems to me that Christensen (and Kelly) are correct that learning of peer disagreement in a case like this does provide *some* evidence that one's peer is mistaken, and therefore (Kelly, contra Christensen) contributes to breaking the symmetry between the peers, i.e., making them not peers in that case. However, the symmetry is not broken because of the discovery that the other party interprets the evidence differently from me, in the sense that because one disagrees with *me*, therefore she must be mistaken. After all, it's not as though *being mine* makes a belief more likely to be justified. Rather, the symmetry is broken because of what the disagreement is *based on*--namely, my own seeming that the evidence points my way.²⁵ But *I* do not have my peer's seeming--that the evidence points *her* way--and so it does not follow that her disagreement gives me evidence that counterbalances my own seemings. I can simultaneously acknowledge that she is usually my peer about this sort of evidence, but that she is apparently making a mistake in this case, and also acknowledge that her view counts *somewhat* against mine (owing to her peer status), but not enough to outweigh my first-order evidential seemings.

²⁴ Ibid., 195-196.

²⁵ So the inference is: "We disagree; the evidence supports my view and not yours; therefore, we are not peers in this case." It is not: "We disagree; I am usually right; therefore, we are not peers in this case."

Further, in some cases the discovery of disagreement will not make any difference with respect to the total evidence one has for or against *p*. For example, if you predict a 30% chance of rain on the basis of evidence which, to me, clearly supports a much higher percentage (say, 70%), then this is evidence to me that you may not be my peer after all with respect to predicting rain in circumstances like this. It is not clear that my evidence for the proposition “There is a high chance of rain” has changed at all because of our disagreement, even if I thought you my peer prior to disagreeing.²⁶ Compare disagreeing with a child: intuitively, the disagreement alone typically provides *no* evidence against *p*. In the case of disagreeing with a presumed peer, on the other hand, the effect of the disagreement on one’s total evidence for *p* depends on the strength of one’s judgment that the other is in fact a peer in this case. But we’ve just said that the discovery of disagreement provides some evidence against peerhood. So, because I have only *my* evidential seeming to go on, not my peer’s, the fact of disagreement will often only count against *her* view, not mine, unless of course my seeming is that the evidence could go either way. The exception to this will be cases where I am *very* confident that she is my peer on issues like this one and I cannot bring myself to think that she is suffering some kind of unusual defect in this case, such as in a situation of Parity. Thus, a typical (non-Parity) disagreement does produce an asymmetry between us, assuming that I take the evidence to support my view.

Now, the same will of course be true from my friend’s perspective: she will also think (rightly) that an asymmetry has been created by the disagreement, and that I am no

²⁶ Some, e.g., Feldman, would object here that disagreement is evidence in the following way: it is evidence *that there is evidence*, which in turn is evidence for or against *p*. This claim will be considered in the next chapter.

longer *her* peer with respect to this question. Perhaps this is what Christensen is getting at. But he is wrong to conclude that “it’s not clear so far that any asymmetry has developed.” It is more accurate to say that an asymmetry has developed *from both perspectives*, but in opposite directions. This suggests that both parties are within their epistemic rights to regard the other as temporarily less than a peer, and that therefore rational disagreement between *apparent* peers (viewed from a third party perspective) is possible. The case also suggests that from the perspective of either party, one is justified in maintaining her view, and there is nothing here to suggest that it would be somehow more epistemically virtuous to conciliate (contra Rosen).

Christensen also takes his wristwatch analogy, mentioned above, to indicate symmetry between peers:

To focus in on the symmetry question, let me begin with an admittedly crude analogy: I look at my watch, a one-year-old Acme that has worked fine so far, and see that it says 4:10. Simultaneously, however, my friend consults her watch--also a one-year-old Acme with a fine track record--and it reads 4:20. When she tells me this, it clearly gives me new evidence that her watch is fast: I should not trust her watch as much as I would have before finding out that it disagreed with mine. But just as clearly, I’ve just gotten new evidence that my watch is slow, and this should diminish my trust in it. In this case, it’s obvious that the fact that one of the watches is on my wrist does not introduce an epistemically relevant asymmetry.²⁷

But here, as above, the symmetry is confused. It is true that the watch’s being *on my wrist* is not a relevant asymmetry. However, whatever reasons I have for trusting my watch *are* relevant, and if they are no more than as stated, I should pretty clearly suspend judgment. This is not simply because the other person disagrees with me, but because I have two objective, conflicting pieces of evidence (the two watches), and no way to tell which is

²⁷ Ibid., 196.

more reliable. Christensen himself comes close to acknowledging something like my points here when he says:

... it is certainly true that when I consider how to regard my friend's disagreement, I must do so from within the first-person perspective--that is, using my own beliefs. Might this give my beliefs--which include my belief about the matter on which my friend and I disagree--a kind of privileged position that the watch on my wrist doesn't share?²⁸

Of course Christensen goes on to argue that it does not, for reasons we saw above in Chapter Two. But again, one's beliefs should be privileged only if things really do appear to be a certain way after investigation (checking for defeaters, etc.), *not* merely because they are one's own beliefs.

So Christensen has not provided a compelling argument for the view that peer disagreement necessarily produces epistemic symmetry between the peers which entails according the views of a peer equal weight as one's own. While peer disagreement may certainly (in fact, typically does) provide some evidence against one's view--though I will argue below that it does not do so necessarily--it does not follow that it provides counterbalancing or defeating evidence. However, Christensen is right about one thing: it can happen in a case of peer disagreement that I am justified in counting the present disagreement against you, and therefore considering you temporarily not my peer, and I am also aware that you are justified in doing the same with me, and I cannot say for sure which of us is right. And as we saw above, the longer the disagreement lasts and the better we know one another's reasons--i.e., the more confident we are that we are peers--the more pernicious this problem becomes. Hence the Parity Problem returns.

²⁸ Ibid.

Chapter Six – Higher-Order Evidence and Testimony

6.1 – The Meta-Evidence Principle

As we saw in Chapter One, Richard Feldman understands the epistemic significance of disagreement to be a matter of higher-order evidence that there is evidence against one's view. This provides another way in which disagreement may defeat one's private evidential seemings, namely by providing evidence that one's initial evidence assessment was mistaken. This differs from David Christensen's proposed symmetry between peers, discussed in the previous chapter. With symmetry, there are two equally good evidence assessors, neither of which has any reason (allegedly) to prefer her view over the other. From a first-person perspective, Christensen thinks, one should weigh the peer's contrary testimony (disagreement) the same as one's own evidential seemings. In other words, peer disagreement defeats one's evidential seemings by providing counterbalancing evidence. On the other hand, with higher-order evidence, one obtains evidence about one's own evidence via disagreement which suggests that one's first-order evidential assessment was sub-par. The difference between these two justifications for conciliating in a case of peer disagreement roughly tracks the distinction between "undercutting" and "rebutting" defeaters. As Thomas Kelly says:

Following Pollock (1986), we can distinguish between *undercutting* and *rebutting* defeaters. Intuitively, where *E* is evidence for *H*, an *undercutting* defeater is evidence which undermines the evidential connection between *E* and *H*... In contrast, a *rebutting* defeater is evidence which prevents *E* from justifying belief in *H* by supporting not-*H* in a more direct

way.... It is something of an open question how deeply the distinction between ‘undermining’ and ‘rebutting’ defeaters cuts.¹

Alternatively, it may be better to think of Christensen’s argument for epistemic symmetry between peers not as a matter of gaining a rebutting defeater, but rather as a simple case of gaining new first-order evidence which renders one’s total evidence set inconclusive. Christensen describes the difference between first-order and higher-order evidence as follows:

Any new evidence... that rationalizes a change of belief state... is, in a trivial way, evidence against one’s previous belief state.... But while this sort of evidence does indicate that my previous beliefs are, in a certain sense, suboptimal, it does not indicate that I’ve been anything less than a perfectly rational believer. The evidence that there’s something suboptimal about my pre-change beliefs is merely a byproduct of the evidence bearing directly on the subject matter of the beliefs. Sometimes, however, evidence rationalizes a change of belief precisely because it indicates that my former beliefs were rationally sub-par. This is evidence of my own rational failure.... in this case, the indication that my former beliefs are suboptimal is no mere byproduct of my reasoning...²

The focus of this chapter is on the evidential effect of disagreement considered as higher-order evidence, which provides, as Christensen says, “evidence of my own rational failure.”

One of the most prominent commentators about the role of higher-order evidence in peer disagreement is Richard Feldman. He describes the views of a peer as providing “evidence that there is evidence for p,” which in a case of disagreement can sometimes defeat my own first-order evidence about p. Discussing private evidence, Feldman says:

Each may have his or her own special insight or sense of obviousness. But each knows about the other’s insight. Each knows that this insight has evidential force. And now I see no basis for either of them justifying his own belief simply because

¹ Thomas Kelly, “Evidence,” in *The Stanford Encyclopedia of Philosophy*, 2014, <https://plato.stanford.edu/entries/evidence/>, Section 1. Kelly is here citing John Pollock, *Contemporary Theories of Knowledge*, 1st ed. (Towota, NJ: Rowman and Littlefield Publishers, 1986).

² David Christensen, “Higher-Order Evidence,” *Philosophy and Phenomenological Research* LXXXI, no. 1 (2010): 185.

the one insight happens to occur inside of him. A point about evidence that plays a role here is this: evidence of evidence is evidence. More carefully, evidence that there is evidence for P is evidence for P . Knowing that the other has an insight provides each of them with evidence.³

This section will explore the claim that “evidence of evidence is evidence,” which we earlier labeled the Meta-Evidence Principle (MEP). In two later papers, Feldman clarifies his intent in advancing this principle and responds to some objections.⁴ The idea, he says, is not that one peer acquires the other’s experiential evidence--“When you tell me that you have a headache, I don’t thereby get your headache,” he says.⁵ Rather, it is that your believing that p gives me reason to believe that there is some reason for p , which in turn gives me reason for p .

The function of MEP in the context of peer disagreement is to explain why one is not justified in maintaining one’s prior belief after discovering the disagreement. The idea is that encountering a disagreeing peer gives one higher-order evidence that there is first-order evidence against one’s view that one has either missed, misinterpreted, or weighed incorrectly. Further, this holds even if one does not know what the peer’s first-order evidence is. It is enough that a peer disagrees to undercut one’s first-order evidence assessment. Or so Feldman claims. But is MEP defensible?

Several other commentators in the disagreement debate have had some (mostly negative) things to say about this principle. Much of this discussion centers on the relationship between evidence and justification. Plausibly, evidence typically produces

³ Richard Feldman, “Reasonable Religious Disagreements,” in *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, ed. Louise Antony (Oxford University Press, 2007), 208.

⁴ Richard Feldman, “Evidentialism, Higher-Order Evidence, and Disagreement,” *Episteme* 6, no. 3 (2009): 308-310; Richard Feldman, “Evidence of Evidence Is Evidence,” in *The Ethics of Belief*, ed. Jonathan Matheson and Rico Vitz (Oxford University Press, 2014), 284–300.

⁵ Feldman, “Evidentialism, Higher-Order Evidence, and Disagreement,” 309.

justification for believing a proposition. If one is an evidentialist, like Feldman, then in fact evidence is *the only thing that can* produce justification. So on first glance, the principle seems to imply that evidence that someone else has evidence for p gives one *justification* for the belief that p. But several commentators have wondered about this implication.

For example, Earl Conee notes that the principle does *not* entail justification because justification is a “summary evaluation” of a belief.⁶ Meta-evidence could easily be defeated by other considerations. For example, if I learn that you, normally my epistemic peer about such things, believe that it will snow today (despite having been quite warm yesterday), then by the principle, I thereby gain some evidence that it will snow today. However, say I also find out that the weather app you’ve been using has a known malfunction that results in false predictions of snow. Then the evidence I gained from your evidence is canceled out, and so I am not justified in believing that it will snow.

Nonetheless, both Conee and Feldman suggest that evidence of evidence provides *some* defeasible evidence for a person--even if that evidence is *always* defeated. Conee considers a case where someone has some evidence that there is evidence for some proposition, but that person also knows that her belief in the proposition is a guess, with no better than a 50/50 chance of being true. This knowledge that it is a guess plausibly

⁶ Earl Conee, “Rational Disagreement Defended,” in *Disagreement*, ed. Richard Feldman and Ted A Warfield (Oxford University Press, 2010), 76-78, and Appendix 1, 84-89. Conee is specifically discussing the principle:

EE: Evidence for Evidence. If S has evidence for the proposition that evidence exists in support of X, then S has evidence for X. (76)

defeats the meta-evidence, but it--the meta-evidence--nonetheless counts as evidence.⁷ Similarly, Conee considers another case in which one gains evidence that someone else has evidence for a proposition, in the context of being told about a hoax being perpetrated on that person. So Smith tells you that he made it seem to Jones as though Black was in the room today, when in fact Black was not. By the Meta-Evidence Principle, you now have evidence that Black was in the room--Jones's evidence--but that evidence is immediately defeated by the fact that the source of your evidence--Smith's testimony--is also defeating evidence for the proposition that Black was in the room. Feldman has a similar case: you read a retraction in the newspaper of a story from yesterday that you had not read. You now, says Feldman, get simultaneous evidence for and against what was reported yesterday, because the retraction gives you meta-evidence (and therefore evidence) for the content of yesterday's report. But the retraction also gives you *defeating* evidence. As he says, "if you were asked whether you had any evidence at all in support of P, it is clear that the answer is 'Yes'--you have evidence about yesterday's newspaper report. To ignore this would be wrong."⁸ In other words, "defeated evidence is evidence."⁹

Another commentator, Alvin Goldman, interprets Feldman's principle to mean that if one person "truthfully reports evidence Q concerning P, this evidence is also acquired by the hearer."¹⁰ Now, this is not what Feldman means--and Goldman admits

⁷ Ibid., 78.

⁸ Feldman, "Evidence of Evidence Is Evidence," 297.

⁹ Ibid., 298.

¹⁰ Alvin Goldman, "Epistemic Relativism and Reasonable Disagreement," in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford University Press, 2010), 211.

that it isn't--but nonetheless Goldman makes a point that is relevant to Feldman's intended usage:

Hearing such testimony may give the hearer default justification for believing *Q*, but such default justification can be defeated by other information in the hearer's possession. *In that case, Q does not qualify as an item of evidence for the hearer.*"¹¹

Goldman seems to imply here that evidence that is defeated--i.e., evidence that cannot justify a belief--is *not really* evidence.¹²

A related objection to the Meta-Evidence Principle comes from Hud Hudson:

Suppose that at some APA conference we're having our yearly beer together and I say, "Hey Rich, something kind of cool...it's my birthday today!" And I do it in that winning and trustworthy way you've come to trust over the years. But I'm lying and I know I'm lying. I now have evidence for the proposition that evidence exists in support of the claim that it's my birthday today, but I don't have evidence (not even a little bit) for the claim that it's my birthday today.¹³

Hudson's point here is that if MEP is correct, then Hudson *would* have evidence that it's his own birthday, when he knows that it isn't, which seems absurd.

¹¹ Ibid. Italics added.

¹² Other critiques of Feldman's principle include: John Biro and Fabio Lampert, "What Is Evidence of Evidence Evidence Of?," *Logos & Episteme* 8, no. 2 (2017): 195–206, John Biro and Fabio Lampert, "'Peer Disagreement' and Evidence of Evidence," *Logos & Episteme* 9, no. 4 (2018): 379–402; Juan Comesaña and Eyal Tal, "Evidence of Evidence Is Evidence (Trivially)," *Analysis* 75, no. 4 (2015): 557–59; Juan Comesaña and Eyal Tal, "Is Evidence of Evidence Evidence?," *Nous* 51, no. 4 (2017): 95–112; Jason Decker, "Disagreement, Evidence, and Agnosticism," *Synthese* 187, no. 2 (2012): 771; Anna-Maria A. Eder and Peter Brössel, "Evidence of Evidence as Higher Order Evidence," in *Higher-Order Evidence: New Essays*, ed. Mattias Skipper and Asbjørn Steglich-Petersen (Oxford University Press, n.d.); Branden Fitelson, "Evidence of Evidence Is Not (Necessarily) Evidence," *Analysis* 72, no. 1 (2012): 85–88; Andrew Reisner, "Peer Disagreement, Rational Requirements, and Evidence of Evidence as Evidence Against," in *Epistemic Reasons, Norms and Goals*, ed. Pedro Schmechtig and Martin Grajner (De Gruyter, 2016), 95–114; William Roche, "Evidence of Evidence Is Evidence Under Screening-Off," *Episteme* 11, no. 1 (2014): 119–24; William Roche, "Is Evidence of Evidence Evidence? Screening-Off Vs. No-Defeaters," *Episteme* 15, no. 4 (2018): 451–62. For further discussion, see my "Disagreement and Meta-Evidence," in preparation.

¹³ Feldman, "Evidentialism, Higher-Order Evidence, and Disagreement," 309. To be more precise, Feldman would have to *indicate* to Hudson in some way that he believes it is Hudson's birthday in order for Hudson to gain the evidence.

It seems to me that the way to respond to this objection is as follows: MEP is not a norm governing one's abilities to create evidence by lying. It is a principle about what the opinions of *others* give one reason to believe. This is due to the obvious fact (at least it seems obvious to me), which Hudson is here exploiting, that my *own* claims cannot create evidence *for me* for the content of those claims.¹⁴

But this is not how Feldman responds to Hudson's objection, and indeed it seems that he cannot. For if he were to acknowledge that the evidence possessed by one person may not be evidence for another, then the principle would fail to do the work he initially set for it, which is to show that evidence that a peer has evidence for *p* is evidence for me that has defeating power for my own first-order evidence about *p*. Goldman makes a similar point:

... even if [a report of there being some evidence is not defeated]..., the hearer does not necessarily acquire the same evidence possessed by Smith. Smith's saying that he had a certain visual experience, for example, does not reproduce in the hearer the same visual experience, with its full evidential load. So the hearer does not acquire the same evidence for *P* as Smith has.¹⁵

Clearly if evidence for one person cannot be assumed to be evidence for another person, then it does not follow that Feldman's evidence that it is Hudson's birthday is also evidence for Hudson that it is his birthday.

Instead of making this (I think) sensible reply, Feldman replies by biting the bullet: "I think that in his example Hud does get some evidence for the proposition that it is his birthday today. But this is not problematic."¹⁶ It isn't problematic, supposedly, because the evidence that Hudson receives through his lie is minimal and is easily

¹⁴ This is barring various unlikely scenarios involving amnesia and the like.

¹⁵ Goldman, "Epistemic Relativism and Reasonable Disagreement," 211.

¹⁶ Feldman, "Evidentialism, Higher-Order Evidence, and Disagreement," 309.

swamped by the other evidence that he has that it is not his birthday. Feldman reminds the reader, rightly, that to have some evidence for *p* is not to be justified in believing *p*, since justification has to do with the balance of one's total evidence. So, since Hudson does not obtain *justification* through his lie, there is no problem.

This seems to me to miss the force of the objection. The objection is not that Hudson obtains a *justified belief* that it is his birthday by lying about it being his birthday, but rather that he obtains *evidence* for the proposition that it is his birthday by lying about it being his birthday.¹⁷ Shifting attention away from this odd result and toward an even stranger result (justification) does not remove the original oddity. Surely evidence cannot be manufactured so easily. If it can, then all I need to do to obtain evidence for some view is to make a convincing show of evidence to someone else. But evidence, intuitively, has to do with having a reason to believe that something is true. As Kelly says:

... the concept of evidence is closely related to other fundamental normative concepts such as the concept of *a reason*. Indeed, it is natural to think that 'reason to believe' and 'evidence' are more or less synonymous, being distinguished chiefly by the fact that the former functions grammatically as a count noun while the latter functions as a mass term.¹⁸

And surely the person who lies to another doesn't thereby gain a reason to believe anything about the content of the lie. If, after Feldman forms the justified belief that it is Hudson's birthday, we ask Hudson, "Do you have any more reason to believe that it is your birthday than you did a moment ago?" the intuitive answer is no. But if evidence

¹⁷ Of course, it is possible--even common--that people obtain evidence through lies. If a stranger tells me he's from Detroit, I have evidence that he is, even if he's in fact from Seattle. The point is that *he* does not thereby gain evidence that he's from Detroit. The evidence he has for *that* proposition--which is presumably nil--remains exactly as it was before the lie.

¹⁸ Thomas Kelly, "Evidence," Section 1.

gives one a reason to believe something, then MEP entails that the answer is yes. I find this highly counterintuitive, *even if* such reasons are always defeated by other competing reasons. But as I say, Feldman cannot respond differently, for if he concedes that Hudson's claim creates evidence, but not evidence *for Hudson*, then he must admit that evidence of evidence is *not always* evidence *for* just anyone. And this would undercut the usefulness of the principle for his position, which is to establish that in a case of peer disagreement with private evidence, a peer's reported disagreement can defeat one's own private seemings.

It is also important to be clear about what peer disagreement is actually evidence *of*. Consider Christensen's mention of Feldman's principle:

Suppose I have good reason to believe that my friend's evidence, though different from mine, is just as good (for example, suppose my friend and I have done similar polls of distinct but comparable sample populations and have reached conflicting conclusions). Here, the reasons for our disagreement need not be cognitive--it could just be that one of the samples was not representative. But absent some special reason for treating her evidence or reasoning differently from mine, it seems clear that in evaluating explanations for our disagreement, I should regard our opinions as equally likely to be accurate, and thus I should alter my opinion toward hers.¹⁹

Assuming that one has not seen the peer's evidence directly, why should it follow that the peer's opinion is "equally likely to be accurate"?²⁰ To see the point a bit more clearly, consider another case. Returning to the courtroom scenario discussed above, imagine that you are a defense attorney. Your client has claimed his innocence, and given the evidence

¹⁹ David Christensen, "Epistemology of Disagreement: The Good News," *Philosophical Review* 116, no. 2 (2007): 211-212. He adds in a footnote here that "This is an example instance of Feldman's principle that 'evidence of evidence is evidence.'"

²⁰ I take it that this is what Christensen intends. If one *had* seen the evidence directly, and found it to be "just as good" as one's own, then the uncontroversially correct response would be to suspend judgment, but not because of anything relating to the peer or her disagreement; it would simply be a case of inconclusive evidence.

you've seen, you're inclined to believe him. However, the prosecution is not behaving as if the defendant is innocent. In fact, it is behaving in a way that suggests it believes something you do not. You infer this because you know how *you* would behave as a prosecutor, given the evidence you have, and the prosecution is going well beyond that. So here you have some evidence in the form of the prosecution's behavior. But what exactly is this evidence *of*? The proposition in question here--the proposition in dispute between you and the prosecution--is: "The defendant is guilty." Call that p. You deny p, while the prosecution affirms it (or at any rate, they're tasked with defending it, whether or not they actually believe it). You each have evidence relevant to p, some shared and some, presumably, not (such as your own interviews with your client). Does the prosecution's behavior give you any evidence for or against p itself? As things stand, without knowing more about the prosecution's evidence, I think that it does not. It gives you, rather, evidence for another proposition: "The prosecution has information I lack." Call this q. Note that even your evidence for q is weak, because many other things besides extra information could explain the behavior, and some alternative explanations are plausible. For example, the prosecution could be posturing, behaving this way to get you to think they have more information so that you'll be more likely to make a plea deal. Owing to these alternative explanations, q does not entail that p is either more or less likely than you took it to be before observing the prosecution's behavior. For all you know, the prosecution has indeed obtained damning evidence against your client. On the other hand, they may just as easily be working with evidence that would be compelling if not for a strong defeater that you possess and they lack. Or they could be posturing. Whether q turns out to be evidence for p cannot be predicted before you know the content

of the prosecution's information. So it seems that here evidence that there is evidence for p (for someone else) turns out *not* to be evidence *for p* (for you). In this case, evidence of evidence *is* evidence, but not necessarily evidence *of* the disputed proposition, and not necessarily *for* both parties to the dispute.

But there is, it seems to me, something right about MEP. I think what this is comes through in a comment van Inwagen makes about David Lewis:

Consider... the body of public evidence that I can appeal to in support of incompatibilism (arguments and other philosophical considerations that can be expressed in sentences or diagrams on a blackboard or other objects of intersubjective awareness). David Lewis "had" the same evidence (he had seen and he remembered and understood these objects) and was, nevertheless, a compatibilist. If I know, as I do, that David had these features (and this feature, too: he was a brilliant philosopher), that he had these features is itself evidence that is (or so it would seem to me) relevant to the truth of incompatibilism.²¹

Notice that what is relevant to what van Inwagen should believe about p (in this case, incompatibilism) is *that Lewis had these features*--namely, he had seen, remembered, and understood the evidence, and was brilliant.²² And those features are themselves evidence pertinent to p. If none of these obtained, the question of the evidential value of Lewis's opinions would be moot. So Feldman's principle is right in one sense: insofar as one is dealing with a perspective that is constituted by certain features of epistemic excellence,

²¹ Peter van Inwagen, "We're Right. They're Wrong," in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford University Press, 2010), 26.

²² One may note here that it is also relevant that Lewis testified to the contrary to van Inwagen; otherwise, van Inwagen would not have obtained any additional evidence pertinent to p. This is of course true. But I want to point out van Inwagen's emphasis here on Lewis's cognitive virtues as *constitutive* of the newly-acquired evidence relevant to p. Van Inwagen takes Lewis's testimony to be evidence which can actually challenge his own evidence because he is his *peer*, not merely because it is contrary testimony. The relationship between the evidential relevance of disagreement and the actual reporting of contrary testimony will be examined in the next two sections.

then the views of a person with that perspective may count as evidence for a disputed proposition which challenges one's own evidential seemings.²³

Kelly makes a similar point when he says:

Let us set aside, for the moment, the special case of disagreement among peers, and reflect on a much more general question: in what circumstances does it make sense for me to treat the fact that someone else believes as she does as evidence for the truth of that which she believes? A true (although perhaps not especially informative) answer: exactly when I take her belief to be a reliable indication of how things stand in the relevant part of reality. Thus, suppose that I know, on the basis of extensive past experience, that, when my weather forecaster judges that it will rain the next day, it tends to rain 80 percent of the time. In that case, I will treat her judgments to the effect that it will rain as evidence that it will rain, inasmuch as I take there to be a positive correlation between the two. Notice that, in this respect, there is absolutely nothing special about the way in which the judgments of another person come to count as evidence.²⁴

This implies that it is the reliability of a person that makes her opinion evidentially valuable, which in turn depends on the person's reasoning habits. This is not to say that testimony from someone with poor reasoning habits has *no* evidential weight--such people can still effectively transmit information, and perhaps even knowledge.²⁵ It is just to say that I am not likely to consider the testimony of another person capable of undercutting my own evidential assessment unless I take her to be an especially careful

²³ Note the implications here for Hudson's counterexample to MEP: if Feldman takes Hudson to have these features of epistemic excellence, then Hudson's testimony that it is his birthday does give Feldman evidence that it is Hudson's birthday. However, it does not give Hudson *any* evidence that it is his own birthday, because he knows himself to be lying, and therefore to be contravening the conditions for his own epistemic excellence, one of which is accurately representing the available evidence.

²⁴ Thomas Kelly, "Peer Disagreement and Higher-Order Evidence," in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford University Press, 2010), 132. Note that Kelly seems to imply here a reductionist view of testimony, since the forecaster's report acquires evidential weight only via an observed correlation between the report and the world. Notice also his emphasis on "taking there to be a positive correlation" between a peer's judgment and some "part of reality." It is this, he says, that grounds the evidential value of the judgment. It follows that judgments of epistemic excellence--and consequently, judgments that one is a peer--are based on observations of this sort of correlation. This strikes me as a Humean point about the empirical basis of judgments regarding the beliefs and reasoning habits of others. How to understand such judgments along Humean (and Kantian) lines will be the theme of Chapters Eight and Nine.

²⁵ See the discussion of testimony in the next section.

reasoner. In other words, I must take her testimony regarding *p* to be *trustworthy* if it is to be capable of defeating or compelling me to alter my first-order belief. So we can modify Feldman's principle to yield a more plausible version as follows:

MEP*: For peers S1 and S2, evidence of S2's trustworthy judgment that *p* is evidence for S1 that *p*.

This takes into account the perspectival nature of evidence (that it's *for* S1, though not for S2); it specifies the object of the evidence (the proposition *p*); and it emphasizes that evidence rooted in a trustworthy judgment, which implies various features of epistemic excellence, is evidence that is pertinent to *p* for S1.²⁶ It thus avoids the objections to Feldman's original principle that we've mentioned.

As stated, MEP* leaves open the question of whether the reported judgment of another that is *not* marked by such features may count as evidence for S1 that *p*. In fact, it is pretty clear that it can: most of us would accept directions in an unfamiliar city from a perfect stranger, and accept that we thereby obtain evidence for propositions about where things are located. But the focus of this section has been on when the disagreement of another provides evidence which can defeat (undercut) one's first-order evidential seemings. MEP* allows for that to happen via higher-order evidence obtained through disagreement, and it specifies what is required for that meta-evidence to be capable of defeating one's own prior evidence: namely, a contrary judgment from someone one takes to be trustworthy in the relevant domain.

6.2 - Disagreement and Testimony

²⁶ A list of such features will be provided in Chapter Seven when I discuss competing notions of peerhood.

It is worth noting a connection here between peer disagreement and testimony. The epistemological problem of peer disagreement may be thought of as a special version of the problem of testimony--i.e., the problem of whether the reports of others can be a basic source of knowledge or justified belief.²⁷ This is because the evidence that one has in a case of peer disagreement is testimonial evidence: the report of a peer about the truth of some proposition. Thus, one is weighing testimonial evidence against the evidence of one's own seemings (including one's reasoning, perceptions, memory, etc.). Thinking of the problem in this way does not necessarily get us closer to a solution, since the issue of testimonial justification itself is still controversial. But there is at least widespread agreement that testimony *is* a source of evidence (the debate is mostly about whether it is a *fundamental* or *autonomous* source of evidence), so if peer disagreement is merely a problem of weighing testimony, then the subject of this chapter can be clarified: insofar as testimonial reports provide potentially defeating evidence, so do the disagreements of a peer.

Note, however, that in idealized cases of peer disagreement (especially those that give rise to the Parity Problem), we are assuming that the peers have done their best to share their evidence. The means of sharing is testimonial, but once shared, it is thought that the evidence can become one's own, so to speak. If you write your argument down for me, and I read it, whatever inferential justification went into your testimonial report is now part of my own inferential evidence base (barring, of course, private evidence). As Kelly says, "In cases of peer disagreement, one gets to go underneath the hood, as it

²⁷ For a nice overview of the epistemological issues surrounding testimony, see Jonathan Adler, "Epistemological Problems of Testimony," in *The Stanford Encyclopedia of Philosophy*, Winter 2017, <https://plato.stanford.edu/archives/win2017/entries/testimony-episprob/>.

were.”²⁸ So while the epistemic problem of belief adjustment arises because of testimony, whether or not there is additional reason to adjust one’s belief beyond the content of the shared evidence depends on the nature of the peer’s reasoning habits, and not merely on the instance of testimony that gave rise to the disagreement.

In other words, recognizing peer disagreement as a problem about testimony leaves open the question of when the testimony of others should provide defeating (undercutting or rebutting) evidence for or against some proposition. I suggested in the last section that in order to defeat one’s first-order evidence for *p*, such testimony needs to be rooted in a trustworthy judgment, which is to say that I take your reasoning to be marked by certain features of epistemic excellence, features which I take into account.

One way of understanding how this works, which I want to raise here as a suggestion, without pursuing it in depth, is to think of the disagreement problem in the context of a virtue ethical theory of testimony.²⁹ Such a theory has been proposed by Paul Faulkner, who discusses the epistemic role of the virtue of trust, utilizing Miranda Fricker’s work on epistemic injustice.³⁰ He considers his own version of a “problem of disagreement,” which arises from the possibility of conflicting sets of ethical concepts, which make it impossible, he argues, to give a purely epistemological account of the virtue of trust. This is because it is not reasonable, for various reasons, to *presume* that a speaker is trustworthy in the absence of a relationship that makes ethical demands.³¹ His

²⁸ Kelly, “Peer Disagreement and Higher-Order Evidence,” 154.

²⁹ See my and Brett Yardley’s “Epistemic Trust, Testimonial Evidence, and Autonomy,” in preparation.

³⁰ Paul Faulkner, “A Virtue Theory of Testimony,” *Proceedings of the Aristotelian Society* CXIV, no. 2 (2014): 189-211. See also Miranda Fricker, “Group Testimony? The Making of A Collective Good Informant,” *Philosophy and Phenomenological Research* LXXXIV, no. 2 (2012): 249-76.

³¹ Paul Faulkner, “A Virtue Theory of Testimony,” 194-195, 202-207.

solution to this problem involves an explanation of the way that trust as an ethical virtue creates the epistemic tools needed to justify trusting the testimony of others.

Faulkner notes that testimony does not occur in a vacuum--the partners in a conversation always have some interest in the outcome of the conversation, and thus “there are a multitude of potential explanations of any given bit of testimony, where each explanation starts from the interest the speaker has in the conversation.”³² Note that while “an audience’s basic reason for entering into a testimonial exchange is to find things out,” it is not necessarily the case that a speaker’s interest is to help the audience accomplish this goal.³³ As Faulkner says:

... from the multitude of potential explanations of any given bit of testimony, there is no reason to single out ‘satisfying the audience’s epistemic interest’ as the default explanation. And this is to say that a presumption of trustworthiness cannot be established as the epistemic default, because testimony does not have the proper function of servicing an audience’s epistemic interests.³⁴

It follows from this that

... what an audience needs, in every case, in order to epistemically rationalize testimonial uptake is some judgement that this explanation applies, that the speaker’s purpose in communicating is indeed informative, and that the speaker is thereby trustworthy. More generally, what is thereby needed is some particular reason for thinking that a given bit of testimony is true.³⁵

Now, we must be careful here in interpreting Faulkner’s claim that “a presumption of trustworthiness cannot be established as the epistemic default.” It is plausible, and widely acknowledged, that under the usual conditions in which one receives and depends on testimony, one is reasonable to assume a norm of truthfulness.

³² Ibid., 202.

³³ Ibid., 203.

³⁴ Ibid.

³⁵ Ibid.

For example, Jonathan Adler describes a class of “core cases” of testimonial transmission which meet certain conditions, including that the purpose of the speaker is primarily to inform. Regarding the core cases, he observes:

These are ordinary contexts where the norm of truthfulness holds, the purpose is primarily to inform, the hearer’s information about the speaker is minimal and there is little or no motivation to deceive e.g., the time, the weather, driving directions, the location of notable places, prominent historical facts, sports scores, the whereabouts of acquaintances, explaining why you are going to the shopping mall. In these core cases, hearers generally have no special reason to doubt the speaker’s word, as they would if the speaker’s assertion is controversial or self-serving.³⁶

Further, some epistemologists have offered arguments for the a priori justification of such a norm of truthfulness. For example, Tyler Burge says:

We are *a priori* entitled to accept something that is prima facie intelligible and presented as true.... Intelligible affirmation is the face of reason; reason is a guide to truth. We are *a priori prima facie* entitled to take intelligible affirmation at face value.³⁷

Burge here represents a version of the *anti-reductionist* view of testimony, according to which testimony is an autonomous source of epistemic authority, not reducible to other sources such as perception and memory, even if it depends on them causally. On the other hand, *reductionists* hold that testimony derives its epistemic authority from these other sources. For example, Elizabeth Fricker says that the assumption of a norm of truthfulness is “an epistemic charter for the gullible and indiscriminating,” and that “a hearer should always engage in some assessment of the speaker for trustworthiness. To

³⁶ Adler, “Epistemological Problems of Testimony,” Section 1.

³⁷ Tyler Burge, “Content Preservation,” *The Philosophical Review* 102 (1993): 472–3. Quoted in Adler, “Epistemological Problems of Testimony,” Section 1.

believe what is asserted without doing so is to believe blindly, uncritically. This is gullibility.”³⁸

Faulkner defends a theory that is midway between these views, though it is more accurately described as anti-reductionist than reductionist. His is a version of the *Assurance View of Testimony*, which holds that “The hearer’s entitlement to belief [on the basis of testimony] resides in the speaker’s standing behind his word, giving his assurance.”³⁹ Proponents of the Assurance View are suspicious of interpreting the epistemic significance of testimony in an evidential sense, because the epistemic import of evidence is independent of the presenter of the evidence, which is not typically the case in testimonial exchanges. Compare the difference between being told and overhearing: the one who is told has a claim against the speaker if the testimony turns out to be false, whereas the overhearer has no such claim because she is not entitled to expect an intention to inform from what is, to her, mere third-person evidence for the truth of what is said. As Adler says, “The overhearer has no... social-ethical duty to the speaker, and the speaker feels no ethical obligation of veracity toward the overhearer. Yet the overhearer has, it would seem, the same evidence as the hearer.”⁴⁰

On Faulkner’s version of the Assurance View, a hearer is not entitled to assume without reason that a speaker’s intention in the testimonial exchange is to inform.

³⁸ Elizabeth Fricker, “Against Gullibility,” in *Knowing from Words*, ed. B.K. Matilal and A. Chakrabarti (Boston: Kluwer, 1994), 126, 145. Quoted in Adler, “Epistemological Problems of Testimony,” Section 6.2. Adler qualifies Fricker’s position as follows: “... if the critical monitoring [of the speaker by the hearer] required is only the ‘counterfactual sensitivity’ that ‘if there were any signs of untrustworthiness, she would pick them up’... so that no special efforts are required of hearers, [Fricker’s] differences with those who uphold... a minimal anti-reductionist thesis seems to vanish.” Ibid.

³⁹ Adler, “Epistemological Problems of Testimony,” Section 7.3.

⁴⁰ Ibid. Compare the utterance of an AI like Amazon’s Alexa: I have no claim against a speaker here because there is no speaker; nonetheless, the utterance contains the same evidence as a syntactically equivalent utterance overheard between people.

Nonetheless, this view is anti-reductionist in the sense that one may be justified in relying on the speaker's testimony without further investigating it for oneself, insofar as the speaker presented the testimony to the hearer as true. The hearer is *trusting* the speaker, and this is sufficient for forming a justified belief via testimony. On the other hand, confirming the content of a speaker's testimony with some other faculty, like perception, implies that the speaker's word is insufficient on its own to justify belief.⁴¹ Assurance views treat testimony as an autonomous source of knowledge. They also hold, however, that trust in the truthfulness of a speaker is required for the justification of one's testimony-based beliefs. But this does not entail that one's testimonial beliefs are typically unjustified, because one generally *does* have such a basis for trust in the sorts of core cases Adler notes. For example, if I get the time from someone on the street, I can trust that person's testimony because there are implicit social norms governing trusting one another in such circumstances. I know that *people don't normally lie about that*, they're *not normally mistaken about that*, they *don't normally have a rational reason to be deceptive about that*, and so on. Such interactions are still ultimately trust-based, even though the relational connection between speaker and hearer is weak. This explains why I would feel like I had been wronged if a stranger asserted to me some directions which turned out to be misleading.

So when Faulkner says that "a presumption of trustworthiness cannot be established as the epistemic default," the keyword is *epistemic*. Treating testimony as evidence, disconnected from the assurance to the truth of the testimony by the speaker, does not yield a presumption of truthfulness. However, one *may* presume trustworthiness

⁴¹ Note that this would also make any knowledge obtained largely non-testimonial.

as an *ethical* default, which in turn can yield an epistemic norm. The moral reason to believe a peer can become an epistemic reason because, as Faulkner says, “the presumption of trustworthiness in effect articulates a truth-based explanation of utterance: the speaker tells one that *p* because one visibly needs to know whether *p*.”⁴² He continues:

... in the good case, there is an explanatory connection between one’s trusting and a speaker’s being trustworthy that runs via the existence of... social norms of trust. This connection ensures that it is no accident that the truth-based explanation of utterance holds, given one’s presumption that it does.⁴³

This way of thinking about the epistemic significance of testimony helps to clarify the problem with Feldman’s assumption, discussed in the previous section, that testimony that *p* is *automatically always* evidence for *p*. Feldman would likely reply that testimony always provides some (defeasible) evidence because of the presumption that it is intended to be informative, and that this evidence is outweighed when one finds out that it wasn’t so intended. But it is, I think, both *less epistemologically loaded* and *closer to real-world experience* to say that testimony provides a reason to believe something--a reason which could play a defeating role for one’s prior beliefs--only if one is justified in thinking that the speaker’s interest is to “satisfy the audience’s epistemic interest.” This is less epistemologically loaded in the sense that it does not presume that the proper function of testimony is to be informative, nor does it require an evidentialist interpretation of testimony’s epistemic significance. It is closer to real-world experience because we are all well-acquainted with the experience of discovering that being informative was in fact

⁴² Paul Faulkner, “A Virtue Theory of Testimony,” 208.

⁴³ Ibid.

not a speaker's intention in delivering some testimony, and feeling as though the speaker violated an obligation toward us as a result.

This view also implies a solution to Hudson's case: I do not gain evidence by giving false testimony, because--even though the other person testifies with the intention to inform--that testimony relies on other testimony (mine) that lacks this intention, and I am in a position to know that. The whole testimonial chain must preserve the intention to inform.

When is it reasonable, then, to believe that a speaker intends to be informative, and so is trustworthy? Faulkner's answer is that "trust, morally understood, is central to the epistemology of testimony in that it can be our 'reason' for testimonial uptake."⁴⁴ In other words, we need a virtue theory of testimony. The details of such a theory need not concern us here; I want merely to suggest that one plausible way of understanding when the testimony of a peer should count as evidence for oneself that could defeat one's own evidential seemings is that this depends in part on whether it is morally appropriate to trust that peer.

6.3 – The Problem of Possible Rational Disagreement

In this section I consider the argument that *possible* peer disagreement presents a similar epistemic problem to actual peer disagreement, and that therefore real-life peer disagreements are unnecessary to give rise to the epistemological issues at play in the debate.

⁴⁴ Ibid., 204.

This issue is raised by Kelly and considered by Christensen, who responds to a thought experiment of Kelly's.⁴⁵ Kelly considers whether possible disagreement should be "intellectually threatening." He thinks that the answer turns on how rational such disagreement would be. That in turn depends on how strong one takes the evidence and arguments for a possible contrary view to be. As he says, "the extent to which merely possible dissent should be seen as intellectually threatening effectively reduces to questions about the strength of the reasons that might be put forward on behalf of such dissent."⁴⁶

Kelly offers a thought experiment involving Newcomb's Problem to support his case.⁴⁷ Imagine a student, call her S, who is considering Newcomb's problem, weighing all of the arguments for both One-Boxing and Two-Boxing, and attempting to make up her mind about which strategy is best. Imagine two possible worlds that differ only with respect to opinions on this problem. In World A, opinion is divided evenly: about half the people who consider the matter are One-Boxers, and about half are Two-Boxers.⁴⁸ In World B, as a matter of contingent fact, everyone is convinced by the arguments for One-Boxing. Assume that the arguments and evidence pertinent to Newcomb's problem are

⁴⁵ See Thomas Kelly, "The Epistemic Significance of Disagreement," in *Oxford Studies in Epistemology, Volume 1*, ed. John Hawthorne and Tamar Gendler (New York: Oxford University Press, 2005), Section 5, 181-185; and Christensen, "Epistemology of Disagreement: The Good News," Section 5, 505-10. The issue is also present in Rosen's discussion of the nominalist "Bedrockers" discussed above.

⁴⁶ Kelly, "The Epistemic Significance of Disagreement," 181-2.

⁴⁷ This problem was made famous by Robert Nozick in his "Newcomb's Problem and Two Principles of Choice," in *Essays in Honor of Carl G. Hempel*, ed. Nicholas Rescher (Dordrecht: Reidel Publishing Company, 1969), 114-146. The problem is as follows: A player must choose either one opaque box or two boxes--the same opaque one and a transparent one. The transparent box contains \$1,000. The opaque box contains either \$1,000,000 (if a predictor has predicted that the player will choose only that box), or nothing (if the predictor has predicted that the player will choose both boxes).

⁴⁸ This was, apparently, the state of things in the actual world according to Nozick's first reporting. See Kelly, "The Epistemic Significance of Disagreement," 19-23, esp. footnote 17.

identical in both worlds. Should S differ with respect to her view of the problem depending on which world she is in? Specifically, assuming S has considered all of the arguments and evidence pertinent to the problem, and is leaning toward One-Boxing, should she be less confident in this opinion in World A than she would be in World B? Kelly thinks not, since the arguments in the two worlds are identical, and since, even in World B, she should realize that she has epistemic peers who are Two-Boxers *in World A*. In other words, the existence of these possible peers who are convinced by the same arguments that S already possesses should count the same epistemically as the existence of actual peers who were convinced by those arguments.⁴⁹ As Christensen says, “The lesson, if I understand Kelly correctly, is that it’s the *arguments* that determine what’s rational to believe about Newcomb’s problem, not contingent sociological facts about what other people think.”⁵⁰

Note that given what we said about Kelly in the last chapter, this interpretation is probably not *quite* fair: arguments are only as good as the evidence supporting them, and part of the evidence here is purportedly what other people think, which *can* matter. I take it this is why Kelly apparently accepts that multiple *independent* lines of evidence are weightier than a single independent line or multiple non-independent lines, even when the arguments in each are the same. So the views of others *do* seem to have evidential

⁴⁹ Kelly goes on to argue that the case of radical skepticism in epistemology is a real-life example of his point, since almost no philosopher professes skepticism, and yet almost all epistemologists have thought it a serious problem worth consideration. Kelly here approvingly quotes Laurence Bonjour: “...the need to consider skepticism does not depend in any crucial way... on whether or not serious proponents of skepticism are actually to be found; if skeptics did not exist, one might reasonably say, the serious epistemologist would have to invent them.” See Laurence Bonjour, *The Structure of Empirical Knowledge* (Cambridge: Harvard University Press, 1985), 14-15, quoted in Kelly, “The Epistemic Significance of Disagreement,” 185.

⁵⁰ Christensen, “Epistemology of Disagreement: The Good News,” 206.

weight for Kelly--he just does not think that this is a special kind of evidence that automatically has overriding power when the other people in question are peers. Nonetheless, Kelly does suggest with the Newcomb case that possible peers are epistemically equivalent to actual peers. This is important for Kelly because he wants to maintain that since the arguments in Worlds A and B are the same, if it is the case that S shouldn't lose confidence in World B when no one disagrees with her, then she also shouldn't lose confidence in World A, just because her peers there are actual. What really matters in both worlds is what the arguments support.

Christensen responds by returning to his restaurant bill case, though modifying it slightly. Imagine that instead of one peer, there are seventeen *superiors*, all expert calculators, say, and they all independently arrive at the same answer that differs from mine. It would obviously be irrational for me to continue to insist on my answer in this case, and Christensen says this remains true “even if, in this particular case, [my] answer was actually the correct one.”⁵¹ So perhaps what the arguments support is *not* all that matters. He applies his conclusion to Kelly's case, saying that if S:

...reasonably sees herself as being among others who... are her epistemic peers, then the denizen of the divided World A *should* be less confident of the truth of One-Boxing than the denizen of World B, where all the other smart and knowledgeable people take the evidence and arguments to support One-Boxing. The One-Boxing denizen of World A has real evidence that she may have made some mistake in evaluating the arguments.⁵²

Note the phrase “real evidence” in the last sentence. Christensen is implying that the evidence provided by the consideration of merely possible peers is not real, or more charitably, that it is considerably less weighty than that provided by actual peers.

⁵¹ Ibid., 207.

⁵² Ibid., 209.

For what it's worth, I am not sure that Christensen is right about the bill calculation case with seventeen superiors. I grant that my own intuitions are that I should weaken my confidence in my belief. Nevertheless, there are some cases in which being the "lone ranger" is rational. Consider, for example, Plantinga's case of the person on trial for a crime she knows she did not commit--surely here one should *not* revise her belief.⁵³ Or consider Einstein's claiming, before wide acceptance, that space is curved. Surely such a stance--even *resilience* about such a stance--can be rational. What is the difference between these cases and the restaurant bill case?

It is tempting to think that the difference is making the seventeen people in the restaurant case expert calculators, so that we are no longer peers. One might argue that, were we all peers, then perhaps there being seventeen of them would not make it less likely that I am the one that is right, because it could be that adding peers is like flipping coins--their probabilities are not additive. But this is incorrect. After all, the testimonial reports of eyewitnesses presumably *are* additive: seventeen of them gives us better reason to believe what they report than one of them. And if what I argued against Kelly above is correct--if the views of a particularly excellent reasoner can provide evidence in addition to the views of that reasoner's teacher, even when they agree and have the same evidence--then the reports of several such excellent peers would count for more than the report of only one. In that case, seventeen peers would have a similar effect on my evidence for *p* as seventeen superiors. So Christensen is probably correct that what the arguments support is not all that matters with respect to the justification of my own view:

⁵³ Alvin Plantinga, *Warranted Christian Belief* (New York: Oxford University Press, 2000), 450-1.

it also matters what other excellent reasoners *think* the arguments support. This, too, is evidence.⁵⁴

Where does this leave us with respect to the question of merely possible disagreement? Christensen has one more thing to say about this. As we saw above, he thinks that disagreement is a special version of the problem of possible error, and this is why it's epistemically significant. But this implies, he thinks, that actual disagreement and possible disagreement are different in the same way that actual error and possible error are different:

Actual disagreement with peers is informative because it provides evidence that a certain possibility--the possibility of our having made an epistemic error--has been actualized. It makes what we already know possible more probable. Could we get this same sort of evidence simply by asking ourselves whether merely possible peers might disagree with us rationally? It seems not--for in those cases where we have made a mistake in assessing the evidence about P, we are overwhelmingly likely to make the same mistake in assessing how a rational peer would assess the same evidence about the same proposition.⁵⁵

So the reason we could not get the same kind of evidence from possible disagreement that we get from actual disagreement is that when considering possible disagreement, we are likely to make the same mistakes in imagining possible disagreeing peers as we make in assessing the evidence for ourselves. In other words, our epistemic frailty extends as

⁵⁴ Christensen has another case: a doctor who finds out she's been given a drug that causes 99 percent of the people who take it to make mistakes on tasks like dosage determination. The doctor, says Christensen, is morally obligated to have her dosage determinations checked in these circumstances, *even if* it turns out that she isn't affected by the drug. But in this case one is not rational to insist on her view because to do so she must assume that she got *lucky*. So the new information about having been drugged serves as a defeater for her belief that she prescribed the correct dosage, because she has no defeater-defeaters, such as evidence that she is immune to that kind of drug. So this case is sufficiently unlike the cases of peer disagreement we've been considering as to be unhelpful. See Christensen, "Epistemology of Disagreement: The Good News," 207-208.

⁵⁵ *Ibid.*, 208-209. Note here that in the case of possible disagreement as imagined, one need not *assess* how a rational peer would assess the evidence for p; one need only realize that a contrary assessment by a rational peer is possible.

much to our ability to conceive of the reasons of rational disagreeing peers as it does to our assessment of the evidence for p.

The problem here is that Christensen misses the force of the possible disagreement problem. For the defeater provided by disagreement to kick in, we do not need to accurately predict (or predict at all) how a possible peer might assess the evidence differently. We need only see that in fact a different assessment is possible. We can even build in the assumption, to avoid the sort of worry Christensen has, that whatever possible assessment a possible peer would have *would be* as coherent, detailed, fruitful, and so on as our own. The problem Christensen mentions is avoidable if care is taken about what views are really possible, consistent with the evidence. And of course, ideally, a plurality of people will help us to figure this out, but this is not necessary. It may be that the cogency of our ideas of what alternate views are possible is likely to vary with our experience of actual peer disagreement. This may not be the case, however, for a few excellent reasoners, but it's sufficient to note that *any* alternative view consistent with the evidence raises the possible disagreement problem. One need not be a philosopher to think of an alternate, consistent conclusion drawn from most evidence sets, especially given the underdetermination of theory by evidence.⁵⁶

Christensen here cites Alston, who argues that possible competing epistemic practices “would pose as great a threat to our practice as would actual competing practices.”⁵⁷ Alston takes the upshot of this to be that if we don't think possible

⁵⁶ The underdetermination issue will be discussed in Chapter Eight.

⁵⁷ *Ibid.*, 208, footnote 23. See William P. Alston, *Perceiving God: The Epistemology of Religious Experience* (Ithaca: Cornell University Press, 1991), 266-278. In context, Alston is considering Aristotelian vs. Cartesian vs. Whiteheadian forms of possible epistemic practice. He says:

Suppose that there were a diversity of sense perceptual doxastic practices as diverse as forms of [mystical practice] are in fact. Suppose that in certain cultures there were a well established

competing epistemic practices are worrisome, then we shouldn't think actual competing epistemic practices are worrisome either.

For what it's worth, Alston seems right here: if it's really possible to have fundamentally different "sense perceptual doxastic practices" (see note 46) that all give rise to similar practical results, then that should cause just as much concern as if those forms were actualized. We may not be able to foresee the details of how a rational peer would assess the same evidence about the same proposition from one of these fundamentally different perspectives (though we can surely try), but nonetheless the possibility that they might be able to do so rationally should give us just as much pause as if we actually encountered such a person. This is because, in both cases, "we have no reason to suppose [our perspective] to be more reliable than these other possibilities."⁵⁸ But Alston thinks that in both cases, we need not be worried about this situation. The reason he gives is telling:

Yet it is difficult to work ourselves up into worrying about these possibilities, and even difficult to convince ourselves that it is our intellectual duty to do so. Insofar as this has a respectable basis, it is that it is not really clear that the alleged possibilities are indeed possible. At least they may not be real possibilities for us,

"Cartesian" practice of seeing what is visually perceived as an indefinitely extended medium that is more or less concentrated at various points, rather than, as in our "Aristotelian" practice, as made up of more or less discrete objects scattered about in space. In other cultures we find a "Whiteheadian" SP to be equally socially established; here the visual field is seen as made up of momentary events growing out of each other in a continuous process. Let's further suppose that each of these practices serves its practitioners equally well in their dealings with the environment. We may even suppose that each group has developed physical science, in its own terms, to about as high a pitch as the others. But suppose further that, in this imagined situation, we are as firmly wedded to our "Aristotelian" form of [sense perception] as we are in fact. The Cartesian and Whiteheadian *ausländer* seem utterly outlandish to us, and we find it difficult to take seriously the idea that they may be telling it like it is. Nevertheless, we can find no neutral grounds on which to argue effectively for the greater accuracy of our way of doing it. In such a situation would it be clear that it is irrational or us to continue to form perceptual beliefs in our "Aristotelian" way, given that the practice is proving itself by its fruits? It seems to me that quite the opposite is clear. In the absence of any external reason for supposing that one of the competing practices is more accurate than my own, the only rational course for me is to sit tight with the practice of which I am a master and which serves me so well in guiding my activity in the world. (273-274)

⁵⁸ Alston, *Perceiving God*, 274.

given our actual constitution and the lawful structure of the world. It may be that we are innately programmed to perceive the environment in Aristotelian terms, and that this program does not allow for social reprogramming into a “Cartesian” or a “Whiteheadian” mode. Nevertheless, it seems clear to me that *if* these alternatives really are possible, then they give rise to difficulties for the rationality of engaging in SP that are quite parallel to those arising from the actual diversity of religions for the rationality of engaging in CP. And so, in that case, if it is rational to engage in SP despite these difficulties, the same conclusion follows for CP.⁵⁹

Alston’s point is that it’s not clear that these alternate viewpoints really are possible, and it’s hard to care about something that you’re not sure is possible. This gives us a way to respond to Kelly’s Newcomb argument more fully than Christensen does. Possible disagreements should not be treated the same as actual disagreements, but this is not because--contra Christensen--we’re likely to make mistakes in imagining them. Rather, it is because until the possibility is actualized in our experience, we don’t have good reason to think that the alternative viewpoints really are nomologically possible for us--possible, as Alston says, “given our actual constitution and the lawful structure of the world.” So S *should* be less sure of her One-Boxing in World A, where opinion is divided, than in World B, where everyone agrees with her, because in World A the possibility of an alternative consistent view is actualized, thereby confirming that it is a real possibility for people like S.

This response allows us to avoid an objection that Christensen apparently does not appreciate but that people like Wedgwood, with his “moral evil demon,” do:⁶⁰ if disagreement itself provides defeating evidence (as Christensen maintains), and if possible disagreement is as evidential as actual disagreement (which Christensen denies,

⁵⁹ Ibid. Here “SP” is sense perceptual practice, and “CP” is Christian mystical practice.

⁶⁰ See Ralph Wedgwood, “Disagreement and the A Priori,” in *The Nature of Normativity* (Oxford University Press, 2007): 259.

but Kelly affirms), then a general skepticism ensues. This is because it is *nearly always* possible to think of an alternative, consistent belief for any issue, and for any actually controversial issue, it is possible to think of one that seems rational. It is fortunate, then, that we need not grant merely possible disagreement equal evidential weight as actual disagreement.⁶¹

So the problem of merely possible rational disagreement, since it cannot motivate genuine epistemic concern, need not detain us any further. It is best to keep the peer disagreement issue rooted firmly in everyday experiences of actual disagreement. Such experiences motivate interest in the abstract version of the problem, and the epistemological discussion of that problem has value only insofar as it helps us to frame and clarify our epistemic obligations when we disagree.

⁶¹ Note what seems like an agreement between Christensen and Alston: Christensen admits that “an ideally rational intellect that knew for certain that it was ideally rational would have no reason to defer to the opinions of others who were equally informed, no matter how numerous.” Christensen, “Epistemology of Disagreement: The Good News,” 208. Of course, Christensen does not think this amounts to much because we are not in fact ideal intellects and so “we would be irrational simply to dismiss the possibility of error.” Ibid., 210. But we must also weigh this possibility against the other evidence that we have, and often it just will not be enough to tip the scales. In addition, Christensen’s admission regarding the ideal intellect suggests that if one could *approximate* ideally rational conditions, then one should to that extent trust the judgments formed under those conditions, irrespective of the views of others. Given what was said about the perceptual analogy in Chapter Three, it seems that we *have* such approximately ideal conditions when we have clear, vivid, and stable evidential seemings.

Chapter Seven – Peerhood, Conformism, and Nonconformism

7.1 – Defining Peerhood

This chapter has three aims: (1) to define more clearly the notion of peerhood, (2) to critique an alternative conception, and (3) to demonstrate that both Conformist and Nonconformist solutions to the problem of rational disagreement are incorrect.

In *Disagreement*, Bryan Frances distinguishes three notions of peerhood:

1. Equality on “Disagreement Factors”
2. Equal likelihood to be correct
3. The conditional likelihood that *if* you turn out to disagree with me, then you’re as likely to be correct as I am (i.e., we are “C-peers”).¹

Option 1 is the common sense notion of peerhood: two people are peers if they are equivalent with respect to the things that make one a good assessor of the evidence for the truth of some proposition. Call these things “Disagreement Factors” and the view of peerhood which depends on them the “Factor Definition of Peerhood” (FP). Frances provides a list of such factors:

- cognitive ability had while answering the question
- evidence brought to bear in answering the question
- relevant background knowledge
- time devoted to answering the question
- distractions encountered in answering the question
- relevant biases
- attentiveness when answering the question
- intellectual virtues possessed²

¹ Bryan Frances, *Disagreement* (Cambridge: Polity Press, 2014), 188-190.

² Bryan Frances and Jonathan Matheson, “Disagreement,” in *The Stanford Encyclopedia of Philosophy*, 2018, <https://plato.stanford.edu/archives/spr2018/entries/disagreement/>, Section 4.

He does not specify the intellectual virtues required, but they would include things like conscientiousness, carefulness, creativity, thoroughness, fairness, honesty, etc. This is not intended to be an exhaustive list, and it should be apparent that one will rarely if ever be in a position to know with great precision how similar one is to someone else with respect to such factors. That is, the factors can provide only a *rough* approximation of the similarity between the epistemic positions of two or more people. But they do give one some idea of how seriously to take the views of others, especially in situations where one has had the time and the reason to measure oneself against someone else with respect to these factors.

Options 2 and 3 are advanced by Adam Elga, and have subsequently been influential in the disagreement debate. Option 3 is a development of option 2, so we can consider them together, as, in effect, a single view of peerhood. Call it the “Likelihood Definition of Peerhood” (LP). Elga begins by noting that taking someone to be an epistemic peer means thinking that that person is “as good as you at evaluating” the sort of claim you’re disagreeing about, adding the qualification that what matters “is that you *count* your friend as your epistemic peer,” not that your friend really *is* your epistemic peer.³ So far, a proponent of FP might very well agree: given the difficulty of knowing for sure when someone else is equivalent with respect to the various factors, what really matters for determining what to do in a case of disagreement is whether one *takes* another to be a peer in that sense. But then Elga elaborates: “In other words, you think that, conditional on a disagreement arising, the two of you are equally likely to be mistaken.”⁴

³ Adam Elga, “Reflection and Disagreement,” *Noûs* 41, no. 3 (2007): 484, 499 (fn. 14).

⁴ *Ibid.*, 487.

This is a departure from FP. Recognizing that this will not be obvious to everyone, Elga then explains in a footnote:

My use of the term “epistemic peer” is nonstandard. On my usage, you count your friend as an epistemic peer with respect to an about-to-be-judged claim if and only if you think that, conditional the two of you disagreeing about the claim, the two of you are equally likely to be mistaken. On more standard usages, an epistemic peer is defined to be an equal with respect to such factors as “intelligence, perspicacity, honesty, thoroughness, and other relevant epistemic virtues” (Gutting 1982, 83), “familiarity with the evidence and arguments which bear on [the relevant] question”, and “general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias” (Kelly 2005). In defense of my use, suppose that you think that conditional on the two of you disagreeing about a claim, your friend is more likely than you to be mistaken. Then however intelligent, perspicacious, honest, thorough, well-informed, and unbiased you may think your friend is, it would seem odd to count her as an epistemic peer with respect to that claim, at least on that occasion. You think that on the supposition that there is disagreement, she is more likely to get things wrong.⁵

So Elga thinks that “likelihood to be correct/mistaken” is the right concept to invoke, not various cognitive virtues, and therefore that LP is preferable to FP.

This, however, is a mistake. First, note that in a case like the one Elga describes, he apparently assumes that one must think *either* that one’s friend is more likely to be correct, equally likely to be correct/mistaken, or else more likely to be mistaken. But this is false. It is possible, and I suspect more common than not in real peer disagreement cases, that one simply has no view at all about who is more likely to be mistaken.⁶ What one *does* have are various past observations relevant to how good one’s friend is at

⁵ Ibid., 499 (fn. 21). The works that Elga cites here are: Gary Gutting, *Religious Belief and Religious Skepticism* (Notre Dame: University of Notre Dame Press, 1982); and Thomas Kelly, “The Epistemic Significance of Disagreement,” in *Oxford Studies in Epistemology, Volume 1*, ed. John Hawthorne and Tamar Gendler (New York: Oxford University Press, 2005).

⁶ Perhaps Elga means this: *if* one were to form a judgment about who is more likely to be correct, *then* one would think that the other is less likely. Nonetheless, there are cases, such as the one that follows in the main text, in which one would in fact not take a view about who is more likely to be correct. If one were asked in such a case who is more likely to be correct, the honest answer would be “I don’t know.” Nonetheless, there can be peers in such cases, and further, one can be reasonable in thinking that someone else is a peer in such cases, based on similarity with respect to the peer factors.

assessing *this* sort of evidence. In other words, one has some information about how similar both parties are with respect to the disagreement factors. For example, if we are meteorologists familiar with one another's predictive track records, and we are encountering a weather problem that is unfamiliar to both of us, I can have good reason to think you're my peer with respect to this question, while having no view of who is more likely to be correct. So one can think someone else is a peer in the FP sense while having no views about whether that person is a peer in the LP sense, which implies at the very least that LP will not suffice for all cases of peer disagreement.

Second, even if one does think that one's friend is more likely to be mistaken than oneself, it must be *because* of some dissimilarity in one or more of the disagreement factors (assuming the friend is rational). I must think that you are not as familiar with this kind of evidence as I am, or that you're distracted, or tired, or what have you. Judgments of peerhood depend on judgments about the disagreement factors, and therefore the intuitive strength of LP depends on FP.

Marc Moffett accepts Elga's notion of peerhood and offers further support for it:

Intuitively, we take an individual x to be an epistemic peer with respect to a given domain (if and) only if we regard x 's judgments concerning that domain to be as likely as our own to be correct given the same evidence. Something like this is a very plausible account of the notion of an epistemic peer. For suppose we use a different criterion which does not imply this one. For instance, we might try to define the concept of an epistemic peer in terms of general intelligence or general level of expertise or both, but nevertheless maintain that this criterion does not imply an equal likelihood of being correct given the same evidence. Then there must be some further truth-conducive factor F which we believe is relevant to judgments in this domain and we must believe that we are superior to x with respect to F . But this is counterintuitive. If we genuinely believe that F is a truth-conducive factor in judgment in a given domain and we genuinely believe that we

are superior to x with respect to F , then we do not regard x as an epistemic peer in the relevant domain.⁷

But Moffett is here making the same flawed assumption that Elga makes: that one party must take the other to be either more, less, or equally likely to be mistaken as oneself, and that if one does not, then one cannot consider the other to be an epistemic peer. But again, this misses the fact that I can be justified in believing someone is my epistemic peer even if I have no idea how likely either of us is to be right.⁸ For example: I and a friend each consider a case that is brand new to both of us, so that neither of us has any idea how likely either of us is to be right about it. Intuitively, I can still justifiedly believe that the other is my peer because I take her to be more or less the same as me with respect to the various cognitive and evidential factors. It is odd to insist, as Elga must, that I cannot hold someone to be a peer unless I have a belief about our comparative likelihood to be correct.

Moreover, I can also be *unjustified* in believing someone to be my epistemic peer even if I have good grounds for thinking that we *are* about equally likely to be right. If I have a friend who comes up with the correct answer to math problems exactly as

⁷ Marc Moffett, "Reasonable Disagreement and Rational Group Inquiry," *Episteme: A Journal of Social Epistemology* 4, no. 3 (2007): 356-357.

⁸ Michael Bergmann makes the same point:

It's in part because I think that many interesting cases of disagreement... don't involve viewing the other as equally reliable with respect to the disputed proposition that I won't be focusing on cases of disagreement between those who view each other as epistemic peers. Instead... I'll be focusing on cases of disagreement between people who view each other as roughly equal in intellectual virtue and who have tried to share all the relevant evidence they can think of with each other. (To be clear, I should note that it doesn't follow from the fact that the parties to the dispute don't view each other as being equally reliable on the topic in question that each views the other as being more reliable or as less reliable. It may be that each person holds no views at all about how reliable the other is on the issue.)

I take Bergmann's use of "reliable" to imply "likelihood to be correct." His mistake is to give over the notion of peerhood to Elga, and then set it aside altogether, rather than to insist on the common-sense FP understanding of it (which is what he does in practice; he merely rejects the label of "peer"). See Michael Bergmann, "Rational Disagreement After Full Disclosure," *Episteme* 6, no. 3 (2009): 351, fn. 10.

frequently as I do, but he uses a calculator and I do the problems in my head, then we are obviously not epistemic peers with respect to math, even though we are equally likely to be correct about it. Equally obviously, it is our presumed cognitive and evidential inequalities that make us not peers.

Or consider a case in which an apparent peer draws an absurd conclusion from some shared evidence. For example: I and a friend have the same kind of disagreement about what actor starred in a film that we've had a hundred times before. Based on past experience, we both believe that the other is about as likely to be right as ourselves conditional on a future disagreement about this topic (we can even stipulate that we've counted the times that each of us has been right). We then have a disagreement about who starred in a particular American Western--I think it was John Wayne, and you think it was Jane Seymour. Our answers are so drastically different that it seems insane to me that you could honestly think it was Jane Seymour. It seems that in this case our disagreement gives me evidence that you are in fact *not* my peer with respect to this question, *even though* prior to our disagreement I would have held us to be equally likely to be correct about this sort of question. Now, Elga does consider a case like this--in which I believe your answer to be obviously incorrect--and he says that "in order to apply the equal weight view, we must determine your prior probability that you would be right, conditional on these circumstances arising."⁹ In other words, I must ask myself what is the likelihood that I would be correct, *given* that when we disagree I find your answer to be utterly preposterous. If I think that in a case like that, I am more likely to be correct than you, then I should not consider you my peer.

⁹ Elga, "Reflection and Disagreement," 491.

One problem with this is that I can *always* imagine, prior to a disagreement occurring, that someone else might give an answer that I find absurd. I am then left in a situation where I must either disqualify that person from being a peer in Elga's sense because she *might* give an answer I find absurd--hence, no one would ever qualify as a peer--or else I must convince myself that *if* she gives an absurd answer, we are still equally likely to be correct. Again, this condition would seem to rule out the possibility of epistemic peers in almost all cases. The problem here is that Elga is smuggling in a *truth*-based notion of peerhood, so that peers are those who are *really* equally likely to be correct, when he himself admits, as we saw, that what matters for peerhood is that we *count* each other as peers. But, to the contrary, peerhood as an epistemic issue is about what one is *justified* in believing, not about whether one is correct or incorrect.¹⁰

Another problem is that it is possible that we *both* consider the other's answer to be ridiculous. Let's say you find my answer of "John Wayne" to be just as preposterous as I find your answer of "Jane Seymour" to be. In that case, Elga says, his view states that if, prior to our disagreement, we both agree that, conditional on a disagreement arising where we both consider the other's view to be utterly preposterous, we think we would be equally likely to be correct in that situation, then we are peers, and therefore (he thinks) when the disagreement does arise, we should accord equal weight to one another's view. Elga says that "that verdict is independently plausible."¹¹ I disagree. It is *not* reasonable to conclude that your friend is your peer with respect to the identity of the actor when she thinks it was Jane Seymour (and you clearly remember it being a man) and finds your

¹⁰ Thanks to Michael Wreen for drawing this point to my attention.

¹¹ Ibid.

answer of John Wayne to be absurd, *even if*, per improbable, your prior self would have concluded that in the instance where you both find the other's answer absurd, that you're equally likely to be correct. In this case, she's just obviously mistaken, despite what you might have predicted beforehand.

Jennifer Lackey offers her own critique of Elga's conception of peerhood.¹² Her own definition of peerhood includes evidential equality, cognitive equality, *and* full disclosure of evidence, making it a somewhat stronger notion than others in the disagreement debate.¹³ She defines cognitive equality as parties who are "equally competent, intelligent, and fair-minded in their assessment of the evidence and arguments that bear on the question whether *p*," and then notes how Elga's definition of peerhood deviates from this.¹⁴

There are... reasons to question Elga's non-standard use of this term. For on his account, two people could radically differ in both their evidential backgrounds and their cognitive abilities with respect to the question whether *p*, yet nonetheless turn out to be epistemic peers regarding this question. For instance, I may be a complete novice with respect to identifying birds of prey, and you may be an expert ornithologist. When I am sober and you are highly intoxicated, however, we may be equally likely to be mistaken about whether the bird flying overhead is an osprey. On Elga's account, then, you and I would be epistemic peers with respect to this question, but this strikes me as quite a counter-intuitive result.

I do not share Lackey's reaction to this case. I think that, given the sensible restriction of peerhood to the question whether *p*, it isn't hard to believe that we *could* be peers in circumstances in which I am sober and you, the ornithologist, are not. This is because, in

¹² Jennifer Lackey, "A Justificationist View of Disagreement's Epistemic Significance," in *Social Epistemology*, ed. Adrian Haddock, Alan Millar, and Duncan Pritchard (Oxford University Press, 2010), 302 (see fn. 17).

¹³ Note that this makes her notion of peerhood highly idealized, despite her concerns about such idealization discussed earlier.

¹⁴ *Ibid.*

those circumstances, the cognitive features relevant to identifying ospreys might be equal between us. So I do not think this is a successful counterexample to Elga's LP as stated.

The problem with his view of peerhood, rather, runs deeper. As we've seen, equal likelihood to be correct depends on equality of peer factors, not vice versa.¹⁵ This can be seen by asking why two people who disagree about something are *not* typically peers: it is nearly always because they differ on some peer factor or other. If we answer that they are not peers because they are not equally likely to be correct, while this may be true, it is not a satisfactory explanation of why they are not peers. In fact, it may not even be true: as Lackey hints here without adequately defending (owing to the weakness of her example), people who are *not* peers can be equally likely to be correct/mistaken. For example, let's say that you're an expert on the Crumple-Horned Snorkack, and I am a novice.¹⁶ You're convinced they don't exist, but I'm pretty sure I just saw one. We are not peers with respect to whether they exist (i.e., you are vastly superior to me on that question), and yet we may be equally likely to be mistaken about whether they exist.¹⁷ Or suppose that you're my defense attorney, but you're convinced I'm guilty. You know more about the evidence than I do, so we're not peers with respect to my guilt, and yet I clearly remember *not* committing the crime. I'd say that, in cases like these, we may be equally likely to be mistaken, even though we are not peers with respect to the factors

¹⁵ I do not mean to claim here that equal likelihood to be correct is conceptually parasitic on cognitive and evidential equality, in the sense that if one has no view about relative cognitive/evidential equality, then one can have no view about equal likelihood to be correct. That seems wrong. Consider again the case of a math "peer" in the equal likelihood sense who uses a calculator, or is fed the answers by someone else. Cases like this suggest that the two views of peerhood are conceptually distinct. What I am claiming in the main text is that if one believes that someone else is or isn't equally likely to be correct about p, that is typically because one has a prior belief about the other's cognitive/evidential equality with oneself.

¹⁶ See J.K. Rowling, *Harry Potter and the Order of the Phoenix* (Scholastic, 2004), ch. 13.

¹⁷ Indeed, I may be *more* likely to be correct than you about this question, even though you are the expert and I am the novice. But see qualification to follow in the main text.

involved in assessing the evidence for the relevant propositions. I say “may be” because it’s difficult to know how to weigh one type of evidence (experiential) against another type (theoretical). A similar point could be made about other sorts of claims, such as those involving religious experience, or claims that a miracle has occurred. To avoid this complication, consider a case in which one party has reasons the other lacks, but the second party is under some circumstantial fluke whereby they are as likely to be correct as the first party. Say, the second party is likely to be right about propositions of the p-type when they have recently encountered nice weather. If the first party has reasons pertaining to p which make him likely to be correct about p, but also knows that the weather has been nice recently, then on Elga’s definition of peerhood, the two parties are peers. This is highly counterintuitive.

So equal likelihood to be correct is neither necessary nor sufficient for peerhood, and its plausibility as a criterion of peerhood derives entirely from the primary criteria of equality in disagreement factors. Therefore, FP is preferable to LP. This is important because one’s definition of peerhood affects one’s view of what it is reasonable to *do* in a case of peer disagreement.

7.2 – Responding to Parity: Conformism and Nonconformism

Elga uses his notion of peerhood to support his Equal Weight View, which he defines as follows:

Upon finding out that an advisor disagrees, your probability that you are right should equal your prior conditional probability that you would be right. Prior to what? Prior to your thinking through the disputed issue, and finding out what the

advisor thinks of it. Conditional on what? On whatever you have learned about the circumstances of the disagreement.¹⁸

This view suffers from a host of problems, some of which have been canvassed in earlier chapters and so need not detain us here.¹⁹ For our purposes it is sufficient to note two things:

1. If we reject Elga's notion of conditional peerhood (LP), then the Equal Weight View is unmotivated. This is because if peerhood is not a matter of likelihood to be correct, then there is no reason to "split the difference"--i.e., assign the competing views equal credence probabilities--with an epistemic peer. Further, if evidence pertinent to the issue under dispute can count for or against assigning peer status to an interlocutor--which Elga's independent evidence condition discussed previously rules out--then whether someone counts as a peer does not depend on the conditional likelihood to be correct that one would have assigned prior to the disagreement occurring.²⁰

2. The Equal Weight View, even if correct, could not help with the Parity Problem.

This is both because (a) the Equal Weight View would *recommend* considering a

¹⁸ Elga, "Reflection and Disagreement," 490.

¹⁹ Thomas Kelly provides what is to my mind a satisfactory refutation of the Equal Weight View in "Peer Disagreement and Higher-Order Evidence," in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford University Press, 2010), 111-74.

²⁰ Ironically, Elga's view--which is a form of Conformism--entails that people are in fact seldom in a situation in which there is a need to conform their view to that of a peer. This is because he thinks that a person may rationally assign peer status based on how much agreement she thinks there is between herself and a potential peer. As he says, the Equal Weight View states what you should "defer to an advisor in proportion to your prior probability that the advisor would be correct in case of disagreement. In practice, this means deferring most to advisors whose views (on matters closely linked to the issue in question) are similar to one's own." Elga, "Reflection and Disagreement," 495. Of course, if only those who already agree with you are likely to count as peers, then Elga can avoid the objection that the Equal Weight View leads to widespread skepticism on controversial issues--but at the significant cost of recommending as a serious epistemic norm the consideration of only views which one already holds. Elga says this is not burying one's head in the sand, but it is hard to see how it could be anything else when it recommends choosing peers based on agreement about allied or prior issues. See *ibid.*, 492-493.

peer's view as equally good as one's own, and therefore would offer a straightforward "No" to the question whether reasonable disagreement in a situation of parity is possible; and because (b) the Parity Problem represents a conflict in one's *own* evidential seemings, and not merely between one's first-order seemings and the testimonial report of another. The Equal Weight View does not recognize such a possibility.

So, insofar as we are attempting to find a solution to the Parity Problem that allows for the possibility of rational peer disagreement, the Equal Weight View must be set aside.

Likewise, Christensen's version of Conformism cannot help with the Parity Problem, since it too helps to *create* the problem, and does not consider the need to suspend judgment in all cases of legitimate peer disagreement to be a significant cost. Christensen, Elga, and even Kelly all mention that most people--philosophers not excepted--tend to hold controversial views too confidently, and that therefore the Conformist principle of belief revision up to and including suspension of judgment or even outright disbelief in one's prior view, is not so drastic as it might at first appear.²¹ Nonetheless, the intuition that began this debate remains strong: people of equal epistemic virtue can have long and careful disagreements and remain rational, including rationally believing that the disagreeing party is also rational. We want to know if this can be a justified view, given a situation of parity.

If Conformist views cannot assist us in this task, can Nonconformism do any better? Kelly's Nonconformist view cannot help. This is because Kelly's "Total Evidence

²¹ See Elga, "Reflection and Disagreement," 497; David Christensen, "Epistemology of Disagreement: The Good News," *Philosophical Review* 116, no. 2 (2007): 214-216; Kelly, "Peer Disagreement and Higher-Order Evidence," 128.

View” stipulates that the preferable view is the one that *actually* accords best—objectively—with the *total* evidence, which includes both first-order evidence for p, and higher-order evidence *about one’s first-order evidence* provided by peer disagreement. The problem is that in a situation of parity, one is not in a position to know what the total evidence objectively supports. As Elga says, “Even if in fact you have done a much better job than your friend at evaluating the claims, simply comparing your verdicts to those of your friend gives you no evidence that this is so.”²²

Kelly considers this sort of objection, and he admits that there is something right about it, namely that “what it is reasonable to believe about the world on the basis of one’s evidence is constrained by what it is reasonable for one to believe *about one’s* evidence.”²³ This insight, he says, can naturally lead to the phenomenon of “downward epistemic push”:

even in a case in which [evidence set] E genuinely supports [hypothesis] H, one’s justification for believing H on the basis of E will tend to be undermined by evidence against the epistemic proposition that *E is good evidence for H*. And one’s justification for believing this epistemic proposition would seem to be hostage to what one’s peers think.²⁴

However, Kelly is quick to point out that there is also the complementary phenomenon of “upward epistemic push.” This happens if one’s correct assessment of first-order evidence provides justification for one’s second-order epistemic beliefs about one’s evidence.²⁵ In other words,

²² Elga, “Reflection and Disagreement,” 487.

²³ Kelly, “Peer Disagreement and Higher-Order Evidence,” 158.

²⁴ *Ibid.*, 159.

²⁵ Kelly provides an argument for this involving the ability to “recognize” that E supports H, which entails justification. This argument is similar to the one I advanced above regarding the *prima facie* justification of evidential seemings, though with some complications. See above, Chapter Four, pp. 81-83.

It is not only that one's higher-order evidence typically makes a difference to what one is justified in believing about the world; it is also the case that one's first-order evidence makes a difference to what one is justified in believing about higher-level epistemic matters.²⁶

For this reason, he thinks, the Total Evidence View is preferable to the Equal Weight View, which does not allow for the phenomenon of upward epistemic push.

However, he admits that there is a difficulty with the Total Evidence View's insistence that reasonableness in disagreement situations depends largely on who has more accurately evaluated the first-order evidence. As he says:

Of course, there is no magic red light that illuminates when one responds to the evidence correctly, no warning bell that sounds when one does not. Indeed, as a phenomenological matter, there might be no introspectible difference between how things seem when one is responding correctly and how things seem when one is not.²⁷

I have argued that it is just this possibility--that there may be no introspectible difference between being right and being wrong--that gives rise to the Parity Problem, and that becomes epistemically problematic when one is in the circumstances of prolonged disagreement with a peer. Kelly does not consider such circumstances; instead, he ends the article by arguing that evidence should not be judged according to what it is *dialectically* appropriate to cite as evidence.²⁸ But the issue in contention between the Conformist and the Nonconformist is just this: *what should I do* when I find myself disagreeing with a peer? While Kelly is surely correct that one may be justified in her view without knowing that she is, this is hollow succor when the person one is in a prolonged disagreement with is obviously an epistemic peer.²⁹ While I may be

²⁶ Ibid.

²⁷ Ibid., 169.

²⁸ Ibid., 169-172.

²⁹ Recall van Inwagen's disagreement with David Lewis.

reasonable, in Kelly's objective sense, to continue believing that my evidence is adequate to justify *p*, when my admitted peer is convinced--with conviction, after thorough consideration--that it justifies not-*p*, such an attitude does not seem commendable. Or, as Rosen says, "it is hard to see the virtue in it."³⁰ So again we see that prolonged disagreement with someone marked by the features of epistemic excellence that we have labeled "disagreement factors" gives rise to a situation of epistemic parity, in which it is difficult to see how the maintenance of one's view could be commendable or epistemically virtuous.

Before leaving the topic of what to do in a situation of peer disagreement, it will be useful to consider in some detail the view of one more author. Richard Feldman's Conformist stance has already been discussed in earlier chapters. However, after the article discussed above, he expressed a somewhat different view of the disagreement debate which, to my mind, is much more compelling.³¹ There, he turns his sights onto the various general principles that epistemologists have defended as responses to the problem of disagreement, and finds them all wanting, including those in defense of his own Conformist view.³² He then adopts and defends *evidentialism*, the view that (1) one is justified in believing a proposition if and only if one's evidence on balance supports that

³⁰ See Gideon Rosen, "Nominalism, Naturalism, Epistemic Relativism," *Noûs* 35, no. 15 (2001): 85-86; quoted above in Chapter Four, p. 73.

³¹ See Richard Feldman, "Evidentialism, Higher-Order Evidence, and Disagreement," *Episteme* 6, no. 3 (2009): 294-312.

³² Jennifer Lackey argues for a similar view in her "A Justificationist View of Disagreement's Epistemic Significance." Instead of evidentialism, Lackey focuses on justification. For her, instead of affecting one's overall evidence, as for Feldman, peer disagreement affects one's degree of justified confidence in a proposition. Thus, Lackey rejects, as does Feldman, both general Conformist and Nonconformist principles, and maintains that her justificationist account can explain the insights of both of those views.

proposition, and that (2) one's degree of justification is determined by how strongly one's evidence supports the proposition.³³ Feldman's thesis is that

there are no true special principles about justified responses to disagreement. That is, there are no general and widely applicable principles specifying what attitudes are justified for someone who encounters a disagreeing peer.³⁴

All we need, Feldman thinks, is evidentialism. Nonetheless, disagreement does often have an evidential impact, often in the form of *higher-order* evidence: evidence about the significance of one's first-order evidence.

Note that principles about conciliating or conforming in the face of disagreement--principles that Feldman himself came close to endorsing--do not count against evidentialism. Rather, they can be seen as applications of the general evidentialist principle in cases of peer disagreement. Nonetheless, Feldman now rejects such principles, because they entail that the evidence one obtains from the fact of disagreement must overwhelm whatever other evidence one had for the disputed proposition, and this is often not the case. For example, consider the case of disagreeing with a peer who is a lone dissenter among a larger group of agreeing peers, or the case of disagreeing with a wishy-washy or unsure peer when one is well-justified oneself, or the case of a peer who is known to be performing below par in this particular instance. In such cases, one's first-order evidence for *p* may outweigh the higher-order evidence provided by the peer disagreement. In other words, what's justified depends on one's *overall* or *total*

³³ Feldman, "Evidentialism, Higher-Order Evidence, and Disagreement," 294.

³⁴ *Ibid.*, 295.

evidence,³⁵ and there is no guarantee that the higher-order evidence of the disagreement will offset the rest of one's relevant evidence.³⁶

Feldman notes that there are several ways of modifying conciliatory principles (i.e., those that recommend suspending judgment in light of peer disagreement), none of which are successful or helpful for understanding what is behind the conciliatory intuition. That intuition can be stated simply without mentioning one's own belief state: "The proposition that S's peer who shares S's evidence concerning P disbelieves P is evidence against P."³⁷ In Feldman's view, nothing more than this needs to be said about disagreement generally, and trying to say more gets one into trouble.

For example, conciliationists might try to strengthen their principle by building in more descriptive features of the disagreement scenario that rule out Feldman's counterexamples. That is, one disagrees with a peer, *and* it is not the case that there is a preponderance of agreeing peers, *and* the peer is not unsure, *and* the peer is performing normally, etc. But Feldman claims that this list of *and*'s will run to infinity: "There will always be the possibility of some other factor being present that provides the subject with evidence that she ought not defer to the peer this time."³⁸ If one tries to counteract this by adding to the principle that no additional counteracting evidence is acquired, then one is still left with the possibility of one's *initial* evidence outweighing one's new evidence

³⁵ Cf. Kelly, "Peer Disagreement and Higher-Order Evidence."

³⁶ Note that this is an apparent departure from Feldman's earlier view. Though, as he says, "I continue to endorse the view that peer disagreement is epistemically significant and threatens to undermine the justification of many beliefs in areas of controversy." Feldman, "Evidentialism, Higher-Order Evidence, and Disagreement," 304.

³⁷ *Ibid.*, 298.

³⁸ *Ibid.*, 298-299.

gained from the disagreement. If one tries to allow for *this* in one's principle, then one will be left with something like this:

If S justifiably believes P immediately prior to t, and S learns at t that an epistemic peer with respect to P who has evidence comparable to S's concerning P disbelieves P, and S does not learn anything that counteracts the evidence provided by this, then S becomes less justified in believing P at t.³⁹

This principle or one like it may well be true, says Feldman, but it is no more than the general evidentialist principle combined with the insight that peer disagreement is evidence against one's view.

Another strategy to salvage conciliationist principles is to modify the definition of "peer." But here again, the only options are to build in requirements on peerhood that rule out potential flaws--and this will never be accomplished adequately--or to simply define peerhood so that it is true by definition that disagreement with such a person renders one's belief unjustified.⁴⁰ But of course this latter option is simply defining oneself into victory, and leaves the original issue untouched. We can either eliminate the vagueness of peerhood by making its requirements so strict that no one will ever actually be peers, or we can allow for looser "similarity" criteria of peerhood, in which case we must consider each case individually. In the latter case, we are back to our basic evidentialist principles again. The same can be said about the vagueness inherent in the concept of "sharing evidence." At the end of the day, Feldman thinks that "there are no simple facts about the epistemological significance of evidence of peer disagreement except for the fact (and its implications)...that evidence that a peer disagrees has evidential impact."⁴¹

³⁹ Ibid., 299.

⁴⁰ This, I think, is the result of Elga's LP.

⁴¹ Ibid., 303.

Feldman next turns to some reflections on higher-order evidence, which is evidence “about the existence, merits, or significance of a body of evidence.”⁴² There are at least three different sorts of disagreements involving higher-order evidence: (1) Disagreements in which peers disagree about the *significance* or *impact* or *weight* of a body of evidence, (2) Disagreements in which one learns that a peer has some unspecified evidence supporting the opposing view, (3) Disagreements in which one learns that a peer has some specific evidence that competes with one’s own view, but one does not acquire that evidence oneself (as in the case of differing perceptions). Feldman argues that (2) and (3) are simply cases of how to weigh competing evidence, and nothing more than the general evidentialist principles are needed. The first case, however, is more difficult. In a case like that--in which a peer and I agree about what the first-order evidence is, but disagree about its evidential weight--there are three possible positions: (i) I am justified in maintaining my view, and justified in believing that the first-order evidence supports P (contra my peer); (ii) I am justified either in maintaining my view, or (inclusive) justified in suspending judgment, or (inclusive) justified in disbelieving that the first-order evidence supports P (in agreement with my peer); (iii) I am not justified in maintaining my view, and I am justified in either suspending judgment or disbelieving that the first-order evidence supports P (in agreement with my peer). Feldman argues that (i) is unreasonable because its defense would involve arguing that one’s view of the evidential support relation in question cannot be mistaken. He argues that (ii) is also unreasonable because it implies the odd view that evidential support relations have no impact on the justification of one’s first-order beliefs. That leaves only (iii), which

⁴² Ibid., 304.

Feldman believes is very often the reasonable response, hence his conciliationist stance-- though he admits that suspense of judgment or disbelief may not *always* be recommended owing to the possibility of the overriding evidential weight of one's first-order views, or even competing second-order views. In this way, higher-order evidence *may* serve as a defeater for first-order evidence. As he says, "It is, in some ways, like what happens to the belief that an object is red when one learns that a red light is shining [on] it."⁴³

But again, there is no universal principle to be had here. Feldman concludes: "...the main thing to say is that [the evidence from one's disagreeing peers] should be added to the mix. This may seem disappointing--you want some crisper principles that you can object to. But, as I said, we shouldn't expect them."⁴⁴

I agree with Feldman that there are probably no general epistemological principles to be had concerning what it is reasonable to do in a case of peer disagreement.⁴⁵ For me, this implies that there is no absolutist, exceptionless solution to the Parity Problem, nor any guarantee of a solution that everyone would agree to in any specific case of disagreement. Fortunately, finding such a principle is not my aim. In Chapters Eight and Nine, I will sketch three possible solutions to the Parity Problem that may allow one to reasonably maintain one's view in the face of prolonged disagreement with a recognized epistemic peer. However, I do not expect these solutions to convince everyone, nor do I believe that they will settle every case of peer disagreement.

⁴³ Ibid., 308.

⁴⁴ Ibid., 310.

⁴⁵ Lackey notes something similar when she rejects what she calls the "Uniformity" thesis, which says that peer disagreement functions the same epistemically in all circumstances, and therefore is susceptible to a single, one-size-fits-all solution. See Lackey, "A Justificationist View of Disagreement's Epistemic Significance," 302.

Nonetheless, I do think that they can help with many cases, especially the cases that motivated the disagreement debate to begin with: prolonged philosophical and religious disagreement.

Chapter Eight – Proposed Solutions I: Hume’s Contribution

In this chapter and the next, I propose two solutions to the Parity Problem, one inspired by David Hume, the other by Immanuel Kant. I will not be defending either a Humean or a Kantian doctrine *per se*, but rather taking inspiration from (1) the argumentative moves they make, (2) their respective insights into the nature of rationality, and (3) the nature of certain sorts of judgments, particularly judgments of taste.

I have chosen to focus on Hume and Kant for several reasons:

First, with respect to nearly *any* significant epistemological issue, consulting these figures is almost an *a priori* necessity. Kant is the most important contributor to questions about human knowledge at least since Descartes, and his contributions cannot be adequately appreciated apart from their relationship to Hume, who famously aroused Kant from his “dogmatic slumber.” More specifically, there is a prevalent sense in the disagreement literature, noted by both Feldman and Lackey in the last chapter, that peer disagreement is a puzzle with a demonstrable solution. Hume and Kant help us to see that this is unlikely. Note, however, that that does not mean that reasonable disagreement is never possible.

Second, the interaction between the two figures provides a relevant historical backdrop for, and in some ways anticipates, the current divisions in the disagreement debate. Put another way: the disagreement between Hume and Kant about judgments of taste prefigures current disagreement about disagreement. Hume, who seems to think that one should adjust one’s beliefs in accordance with the judgment of recognized experts, can be seen as anticipating the “Conformist,” “Conciliationist,” “Equal Weight,” or

“suspend judgment” views, which hold that the realization of disagreement with a peer gives one a strong reason to alter one’s attitude toward the disputed proposition.¹ On the other hand, Kant, who thinks that it is sometimes appropriate to form judgments which are truth-directed, even when one recognizes that one lacks the ability to demonstrate their truth, may be seen as representing the “Nonconformist,” “Steadfast,” or “Anti-conciliationist” views, which hold that it is rational to maintain one’s prior attitude toward a disputed proposition in the face of peer disagreement.

Third, the current debate, though still fairly young, has already produced something of a stalemate of positions. The Conformist and Nonconformist views all have able defenders, and, with a flourish of irony lost on no one, have generated seemingly intractable disagreement among philosophers. The debate is in need of penetrating insight. I think that Hume and Kant provide options for such a resolution.

8.1 – A Humean Approach to Disagreement

While Hume did not discuss peer disagreement in the contemporary sense, he was self-avowedly motivated by a desire to settle the “endless disputes” of philosophy. He set out to base our understanding of human nature and to determine the extent of human knowledge on strictly empirical foundations.² In addition, his concern with weighing testimonial evidence against one’s experience of the world can be seen in his well-known

¹ See, e.g., the conclusion of Hume’s *Natural History of Religion*, where he says of religious disagreement: “The whole is a riddle, an enigma, an inexplicable mystery. Doubt, uncertainty, suspence of judgment appear the only result of our most accurate scrutiny, concerning this subject.” David Hume, *A Dissertation on the Passions & The Natural History of Religion: A Critical Edition*, ed. Tom Beauchamp (New York: Oxford University Press, 2007), 87.

² See David Hume, *The Letters of David Hume*, ed. J.Y.T Greig, vol. 1 (Oxford: Clarendon Press, 1932), 13.

argument against the rationality of belief in miracles. In that argument he takes competing claims from different religions that a miracle has occurred to count against the rationality of belief in miracles.³ If such competition occurs, “there is,” he says, “a mutual destruction of arguments.”⁴ In addition, Hume famously noted the inability of human reason to non-circularly demonstrate the reliability of its own sources of belief about the world. As he says, once we have “arriv’d at the utmost extent of human reason, we sit down contented,” because we “perceive that we can give no reason for our most general and most refin’d principles, beside our experience of their reality.”⁵

My discussion here is “Humean” in several ways:

(1) It is inspired by the idea that a theory of disagreement is empirically constrained. As Hume says, respecting our knowledge of principles in general:

...tho’ we must endeavour to render all our principles as universal as possible, by tracing up our experiments to the utmost’ and explaining all effects from the simplest and fewest causes, ‘tis still certain we cannot go beyond experience; and any hypothesis, that pretends to discover the ultimate original qualities of human nature, ought at first to be rejected as presumptuous and chimerical.⁶

As Lackey and others have noted, disagreement as an epistemological issue arose because of prolonged philosophical and religious disagreement, and it should remain focused on real people in real situations of peer disagreement. This is what motivates the Parity Problem, and what tells against Nonconformist views which take the objective accuracy of one’s evidential judgments to justify the decision to maintain one’s beliefs.

³ See David Hume, *An Enquiry Concerning Human Understanding*, ed. Peter Millican (Oxford World’s Classics, 2007), Section X, esp. Part 2, paragraph 24, pp. 87-88.

⁴ Ibid., Section X, Part 1, paragraph 6, p. 81.

⁵ David Hume, *A Treatise of Human Nature*, ed. David Fate Norton and Mary J Norton (Oxford: Clarendon Press, 2007), 5.

⁶ Ibid.

Such views are not sufficiently sensitive to the impoverished evidential situation people are typically in when encountering disagreement, nor to their inability to verify the objective veracity of their assessments.

(2) My discussion in this section is “naturalistic” in that knowledge of the limitations of disagreement is based on empirical observation of actual disagreement (including the experience of having a disagreement).⁷

(3) As it was for Hume, this naturalism is motivated by a desire to settle the “endless disputes” of philosophy by distinguishing those disputes that are legitimate from those that stem from mere idle speculation about things beyond the limits of human reason.

(4) My discussion here is methodologically similar to Hume’s discussion of definition as far as its purpose is concerned: to obtain clarity about an obscure idea. In Hume’s case, the obscure idea is “necessary connection”; in mine, it is “peerhood.” One should begin, Hume says, by tracing the idea back to its initial “impressions” (seemings) in order “to fix... the precise meaning of [the] term... and thereby remove some part of that obscurity, which is so much complained of in this species of philosophy.”⁸

Consider, for example, Hume’s account of the source of moral sentiment: sympathy. For Hume, morality is based on (1) feelings of approval and disapproval, and (2) the ability to sympathize with others.⁹ Through sympathy, we move from an idea of

⁷ Cf. Hume’s claim that “as the science of man is the only solid foundation for the other sciences, so the only solid foundation we can give to this science [of human nature] itself must be laid on experience and observation.” *Ibid.*, 4.

⁸ Hume, *Enquiry*, 45.

⁹ See Hume, *Treatise*, Part 3, Section 1, 367-378.

what someone else is feeling to feeling it ourselves. William Morris and Charlotte Brown explain:

There are four steps to this process. I first arrive at the idea of what someone is feeling in any of the usual ways. I next become aware of the *resemblances* between us, so we are linked by that principle of association. While we resemble every human being to some extent, we also resemble some individuals more than others--for instance, those who share our language or culture or are the same age and sex as we are. The associative principles of *contiguity* and *causality* also relate individuals who are located closely to us in time or space or who are family members or teachers. According to Hume, we are able to sympathize more easily and strongly with individuals with whom we have strong associative ties. The stronger the associative relations, the stronger our sympathetic responses. Hume then claims--controversially--that we always have a vivid awareness of ourselves. Finally, he reminds us that the principles of association not only relate two perceptions, but they also transmit force and vivacity from one perception to another.¹⁰

This is analogous to peerhood in important respects. When we judge that someone is a peer, there is a similar stepwise, “sympathetic” process: I first find out how someone sees things in the usual way (conversation, body language, etc.). I then notice that that person resembles me in how I see things. For Hume, this step can invite bias--I am inclined to privilege the perspective of those similar to myself, and so to consider them my equals more readily than those not as similar to me.¹¹ Only through noting this resemblance do I attribute peerhood to another, for it is only then that I have any reason to take his/her opinions seriously. Prior to recognizing the resemblance, I have only evidence about p (first-order), and evidence about the other person’s view of p, which is also first-order, but about the person, not the proposition. But peerhood is not directly about the proposition or the person. It is about the proposition-as-held-by-the-person (Elga phrases

¹⁰ William Edward Morris and Charlotte R. Brown, “David Hume,” in *The Stanford Encyclopedia of Philosophy*, Spring 2017, <https://plato.stanford.edu/archives/spr2017/entries/hume/>, Section 7.2.

¹¹ Cf. Elga’s view of forming peers. See footnote 20 in the previous chapter.

this as having an “equal likelihood to be mistaken” about p). Since I am most familiar with my own view of p and my own reasoning process about p (we need not follow Hume here in the stronger claim that I always have a “vivid awareness” of my own thoughts), I will always be more confident of my own conclusion about p than I am about another’s conclusion about p, at least initially. But through observation of another’s reasoning about p, I may come to see a greater resemblance to me, and then my confidence in my own reasoning may be transferred to him/her, so that I begin to see him/her as an epistemic peer. Crucially, my judgment of peerhood is based on resemblance to me (this is true even if I come to view the other as an epistemic superior or inferior).

For Hume, a belief is an idea that borrows its vivacity from impressions themselves. By doing so, the idea becomes like an impression, for example, by sharing its behavioral effects. I hear thunder, for instance, and react as though there is lightning: I leave the playing field, say, even though I did not see lightning. Similarly, judging another as my epistemic peer confers some of the vivacity of my own evidential seemings onto his/her putative evidential seemings. On this view, to believe that someone is my peer is to view his/her reported evidential experiences about p similarly to how I would if they were *my* experiences. I did not see the lightning, so to speak, but I’m inclined to act as if I had.

(5) My discussion is also methodologically similar to Hume’s discussion of analogical argument. In criticizing the teleological argument in the *Dialogues Concerning Natural Religion*, Hume argues that disputes regarding “the degrees of any

quality or circumstance” are incurably indefinite.¹² This is because, as Morris and Brown put it,

Analogies are always matters of degree, and the degrees of the qualities involved in the design argument aren’t capable of exact measurement. The controversy thus “admits not of any precise meaning, nor consequently of any determination.”¹³

The result is that with such analogies (i.e., those that resist precise measurement), everything is analogous to everything else in some respect, and if the similarity cannot be specified, then the analogy itself is meaningless. For Hume, this renders all natural theology suspect by virtue of being predicated on such an analogy between the universe and artifacts. Such analogies, Hume thinks, are unintelligible. Similarly, in the case of peer disagreement, if the degree of similarity between one’s own reasoning ability and another’s cannot be reasonably compared, then it too would be subject to the unintelligibility objection, since it would be no more meaningful than comparing any two phenomena.¹⁴

Part of our task, then, is to determine if the “peer-analogy”--“My interlocutor’s reasons respecting p are similar to mine”--can be made precise enough to support a Conformist or Nonconformist stance.¹⁵ Given the argument in the preceding chapter for the priority of the “peer factor” approach (FP) to defining peerhood, assigning precise

¹² David Hume, *Dialogues Concerning Natural Religion*, ed. Dorothy Coleman (Cambridge University Press, 2007), 92.

¹³ Morris and Brown, “David Hume,” Section 8.4, citing Hume, *Dialogues Concerning Natural Religion*, 93.

¹⁴ We need not follow Hume here in insisting that the analogues must be susceptible to “exact measurement.” After all, we can and do reasonably compare things which cannot be quantified, such as the beauty of various works of art. For our purposes, it is enough to note that the imprecise nature of such analogies means that they will not yield a single most rational decision principle which can justify either a Conformist or Nonconformist stance on the disagreement problem.

¹⁵ Note that the Conformists and Equal Weight theorists usually assume that the peer comparison can be exactly quantified; otherwise, I could not assign the proper credences to our views of p and thus calculate the reasonable credence.

probabilities to one's peer status will be difficult but not impossible. I may not be able to say, as the Equal Weight theorist would like, that my peer and I are each 85% successful at evaluating evidential support for propositions like *p*. But I can say that we are *roughly* equal in intelligence (measured either objectively or by our ability to marshal and respond to powerful objections), conscientiousness (measured by our track record of sticking to our own stated norms for evidence assessment), carefulness (measured by how much time and attention we commit to each piece of evidence and each objection), thoroughness (measured by how many objections we consider, and how many sources we consult to find them), and so forth. Collectively, such rough measures will plausibly yield a tolerably clear comparison between interlocutors, and thus a reasonable means of claiming peerhood. But the inherently inexact nature of these measures, especially in real-world contexts, means that one's peerhood-judgments will always be somewhat indefinite. Or put differently, they will be *underdetermined*, an idea which will become important below. Because I cannot precisely measure the extent to which my peer and I compare with respect to the various peerhood factors, in any particular disagreement it could be the case *either* that in this instance, I am slightly more conscientious, careful, thorough, etc., *or* that my peer is. And of course the same will be the case from my peer's perspective. It is thus up to each of us to decide whether to weigh our own evidence assessment more heavily, and so maintain our first-order view of *p*, or to weigh our assessment and our peer's assessment as equal, and so, perhaps, suspend judgment about *p*.

The crucial Humean insight here is that there is no numerical means of comparison that can ensure that we make the unquestionably most rational decision,

because the claim that there is an unquestionably most rational decision about this is itself suspect. In order for there to be a quantifiably “correct” way to proceed, we would need to be able to measure ourselves and our peer with respect to the various peer factors much more precisely than we in fact can. But since all we have to go on in any actual disagreement is an imprecise sense of similarity between ourselves, which decision we make (conciliate or not) must be made on some basis other than what is demonstrably most rational. This levels at one stroke both the Conformist and Nonconformist (including the Equal Weight and Total Evidence) responses to peer disagreement, insofar as they are interpreted as providing a demonstrable, quantifiable assessment of the most rational response to peer disagreement. Hume helps us to see that such a response is not to be had.

One Conformist, David Christensen, has recognized this issue and hinted at a partial reply. In the course of considering a case in which two doctors disagree about the proper treatment for a patient, Christensen says:

... do I think that [the other doctor’s evidence] weighting leads in general to equally accurate beliefs? If so, then why think my belief is likely to be more accurate now? ... On the other hand, if I think her weighting does not lead in general to equally accurate beliefs, why should I grant that it’s just as rational to form beliefs using that weighting?¹⁶

The implication is of course that if I think her evidence-weighting policy is as accurate as mine, then I have no reason to prefer mine, and I should conciliate. While if I think that hers is less accurate than mine, then I should not grant that hers is as rational as mine.

Christensen adds in a footnote:

¹⁶ David Christensen, “Epistemology of Disagreement: The Good News,” *Philosophical Review* 116, no. 2 (2007): 191.

... to say that I believe my weighting policy is more rational than hers need not commit me to being able to give an ultimately non-question-begging defense of my policy. Hume, I think, showed us that this cannot be the standard for rational belief-forming policies. But to the extent that I regard it as my duty to prescribe... treatment, I must regard my colleague in somewhat the way I regard the counterinductivist: she's not completely rational even if I can't demonstrate that. We should not take the impossibility of non-question-beggingly demonstrating that one method of forming beliefs is uniquely rational to show that more than one method of forming beliefs is rationally acceptable.¹⁷

Christensen's point here seems to be that even if Hume is right that we can't non-question beggingly demonstrate the superiority of one method of rational belief formation, it does not follow that just any method is equal to any other, or even that there is more than one rational method. Now, this is true so far as it goes, but it misses the force of the problem in Christensen's own case. Take the second horn of Christensen's dilemma: I think my colleague's evidence-weighting policy "does not lead in general to equally accurate beliefs" as mine (after sharing evidence, discussing it, and so on), but I cannot demonstrate this. Should I still think of my peer as being as rational as me? This will depend, not on what merits I think that policy has independently of who is implementing it, but on whether I take my colleague to be my peer with respect to assessing this kind of evidence. I must assess her, not merely her evidence-weighting method. Now, the evidence-weighting policies she chooses will surely factor into my judgment here, but it will not be the whole of it--I will also need to assess her with respect to the various peer factors discussed above. It could be that her similarity to myself along those factors is such that I come to believe that she is indeed rational in implementing an evidence-weighting policy that I would reject.

¹⁷ Ibid., 191-192, fn. 7.

So while Christensen is correct that a permissive view of rational evidence-weighting policies does not follow from the failure to demonstrate the superiority of a single one, it may yet follow from the features of epistemic excellence which factor into peer-judgments. Now consider the first horn of the dilemma: I think my colleague's evidence-weighting policy *does* lead to equally accurate beliefs as mine. Is it obvious that I should conciliate? I think not, and I think Hume is again helpful here. In particular, Hume's own constructive project regarding causality hints at a means of deciding whether or not to conciliate in a given case of disagreement.

Hume famously argued that causal inference is not demonstratively certain, because a priori reasoning does not in fact yield a notion of necessary connection between cause and effect.¹⁸ In place of reason, Hume suggests that "custom," or habit, is the source of our belief that the future will be like the past.¹⁹ While I do not intend to suggest that custom is also the source of our beliefs that other people are our epistemic peers, I do think that there is a way of understanding peer judgments that is similar in spirit to Hume's constructive account of causal inference. That is, there is a way of justifying the choice to maintain one's own view in the grip of the Parity Problem. This Humean way does not require one to think (1) that there is a "necessary connection" between the evidence and one's own view, or (2) that one is any more rational than one's

¹⁸ See *Treatise* 1.3.6, and *Enquiry*, Section 4.

¹⁹ See *Enquiry*, Section 5: "Sceptical solution of these doubts," e.g.:

Custom, then, is the great guide of human life. It is that principle alone, which renders our experience useful to us, and makes us

expect, for the future, a similar train of events with those which have appeared in the past. Without the influence of custom, we should be entirely ignorant of every matter of fact, beyond what is immediately present to the memory and senses. We should never know how to adjust means to ends, or to employ our natural powers in the production of any effect. (*Ibid.*, 32-3.)

See also *Treatise* 1.3.7 - 1.3.10.

interlocutor. Moreover, this Humean path allows us to say that there is a “reasonable” solution to the Parity Problem, that is, an epistemically appropriate response. It is epistemically appropriate in much the same way that following the “uniformity principle” is an epistemically appropriate thing to do in everyday life. It is appropriate, not in the sense of being demonstrable from *a priori* principles, but in the sense that following our habits of inference has in general led to beliefs that are satisfactory for us. Perhaps our beliefs about the rationality of our peers, and of ourselves, are like this as well: not strictly speaking demonstrable, but nonetheless useful for our overall epistemic goals.

8.2 – *Moffett’s Millian Solution to Peer Disagreement*

One commentator in the recent disagreement literature, Marc Moffett, has come close to this conclusion, but by a different route. His guide is John Stuart Mill, who argues that one ought to submit one’s views to the assessment of one’s disagreeing peers.²⁰ It will be helpful to consider Moffett’s view in detail here, in order to elaborate my own Humean solution to the Parity Problem.

Moffett notes that a Mill-inspired theory of group inquiry has two principles:

- (1) The Millian Platitude (MP): to be fully justified, theoretical beliefs (i.e., those that are non-deductively inferred from other beliefs) must hold up against the strongest counterarguments practically available at the time.
- (2) The Collective Criticism Condition (CCC): one can be adequately justified in thinking that the Millian Platitude is satisfied for some theoretical belief only if

²⁰ Marc Moffett, “Reasonable Disagreement and Rational Group Inquiry,” *Episteme: A Journal of Social Epistemology* 4, no. 3 (2007): 352–367.

there is “free and open critical discussion” of the belief with one’s disagreeing epistemic peers.²¹

In other words, to be fully justified in our theoretical beliefs, we must engage in critical discussion with a community of disagreeing peers--indeed, with the best among those peers that are practically available.

Moffett argues that MP entails CCC, at least for creatures like ourselves with limited epistemic abilities. MP says that the justification of our theoretical beliefs depends partially (i.e., in addition to whatever positive evidence we have for them) on our own critical evaluation of those beliefs. Moffett interprets this to mean that we are only fully justified in these beliefs (and maybe even adequately justified) if we are able to defend them against the strongest practically available objections.²² The strongest practically available objections are those which are *externally* available to one, meaning either that one is aware of them, or that one could be reasonably expected to be aware of them given one’s current social situation (e.g., given the philosophical/scientific theories of one’s community, and given one’s freedom within the community to pursue objections). They must be externally available, because if they need only be internally available, then one could simply avoid learning about objections, or forget them, and thereby improve one’s epistemic situation. Because the emphasis here is on what one is *able* to do (namely, respond to the best counterarguments), one may be justified in one’s

²¹ Ibid., 352-353.

²² If this seems too strong, keep in mind that the Millian principles are restricted to theoretical beliefs. Moffett notes that beliefs may be more or less theoretical “depending on the nature of their evidential support.” Ibid., 364, fn. 1. Still, depending on how “fully” and “adequately” are interpreted here, it may turn out that Mill’s principles restrict justification too much, since they would seem to require one to be able to defeat all objections in order to have one’s beliefs justified. However, nothing in my argument below turns on this interpretation of these principles, nor on the principles themselves.

views, even if one doesn't have reason to believe that one *can* respond to the best counterarguments, because one can be *able* to respond to them without believing that one has considered the best ones or that one is able to respond to them.

This leads to a problem if we are considering only the ability of an individual herself to come up with the best counterarguments to her views: namely, she is unlikely to be able to do so because humans are cognitively limited and prone to bias. Hence, in order to *really* have these objections available to her, she needs to interact with disagreeing peers. Hence, MP entails CCC. As Moffett says, "In sum: from the point of view of individual epistemology, it is an epistemic good and a practical necessity to rely on a network of genuine critics."²³ He continues: "It follows that in the absence of epistemic peers with whom *x* reasonably disagrees, *x*'s epistemic situation will be significantly compromised."²⁴ Indeed, she would not be able to be justified in her theoretical beliefs *at all* with such an absence unless (per improbable) she is unusually gifted at predicting the counterarguments of the strongest possible interlocutors.

Unfortunately, CCC has negative epistemic consequences in the form of the peer disagreement problem. If, for example, Conformists like Elga or Feldman are correct, then the existence of dissenting epistemic peers entails suspension of judgment about our own controversial beliefs. But then, a *requirement* for fully justified belief (i.e., CCC) would entail that our beliefs are *not* even adequately justified (because according to the Conformist view, peer disagreement undermines justification to the extent of requiring suspension of judgment), and this is undesirable, to say the least. But why think that peer

²³ Ibid., 355.

²⁴ Ibid., 356.

disagreement undermines justification? Recall Elga's argument: peerhood, he thinks, is a matter of likelihood to be right. This is because if a person takes herself to be right and an interlocutor to be wrong, then she must think there is some factor F that makes her more *likely* to be correct, which is just the same as saying that she does not think the interlocutor is a peer. Moffett unfortunately accepts Elga's understanding of peerhood, declaring *equal likelihood to be correct* a "necessary condition" for peerhood.²⁵ But, says Moffett, Elga is mistaken to then conclude that peerhood implies that we should accord *equal weight* to our interlocutor's view, because "Elga does not adequately take into account the holistic relation between evidence and theory."²⁶ By this, Moffett seems to mean that Elga does not consider the possibility that revising one's own belief may be much more *costly* (my word) than simply rejecting one's interlocutor's peerhood.

Consider cases of strong evidence, such as the dean-in-the-quad case (i.e., cases of perceptual disagreement). In such cases, says Moffett, it is more reasonable to maintain one's view and reject peerhood because revising one's view would require embracing localized perceptual skepticism. In addition, Elga assumes a false dilemma: either we are not peers, or, if we are, I must give our views equal weight. But there are other options: I might conclude that because we disagree, your evidence must not really be the same as mine (because you have missed some small part of it that I've seen); or I might conclude that you are simply lying about your conclusions. In these cases, we are still *peers* (in the *likelihood to be right* sense), but I do not need to give our views equal weight.

²⁵ *Ibid.*, 357.

²⁶ *Ibid.*

Unfortunately, this strategy will not work in every case. As Moffett says:

... given that the disagreement will persist as we proceed further in our collective inquiry we will, after a very short period of time, have to regard these critics as simply being incapable of attaining the same quality of evidence as us. But this amounts to positing a truth-conducive factor on which we are superior to our critics and so, to not regard them as genuine peers after all!²⁷

The lesson here, according to Moffett, is that the more idealized our cases of disagreement become, the more intractable they are, because our main way of resolving them in particular cases is to fall back on the possibility of some shortcoming in our interlocutor. But in cases like Mill's, "As we idealize away from those limitations, it becomes much more difficult to avoid epistemic stalemate with our peers."²⁸ So then CCC, while providing ideal grounds for the justification of our theoretical beliefs, also simultaneously undermines those beliefs by introducing the problem of peer disagreement.²⁹

Moffett's solution to this problem involves the underdetermination of theory by evidence, which he argues allows us to avoid the dilemma posited by the Conformist: either abandon our own view, or else downgrade the "epistemic competency of our critics."³⁰ Unfortunately, Moffett's explanation of a third option here is not as clear as one would like. He says:

... if one were to accept that one's current evidentiary position (perhaps including whatever theoretical virtues one wishes to posit) does not rationally determine a commitment to a unique theory, then one would be free to maintain the competence of one's critics despite their genuine disagreement. Against this backdrop, an epistemic peer would be construed as someone who is as likely as

²⁷ Ibid., 358.

²⁸ Ibid., 359.

²⁹ Note that while epistemic stalemate in an idealized world may not be much of a threat to us in *this* world, it is a threat to CCC, which Moffett is concerned to defend.

³⁰ Ibid., 360.

we are to judge correctly given the same evidence *and the same background theoretical beliefs*.³¹

He adds in a footnote:

In saying that the individual would judge “correctly” we cannot mean that the individual would judge “truly” unless we are also willing to embrace some form of anti-realism. What is intended, rather, is that *an epistemic peer is as likely as we are to arrive at the most rationally supported conclusion given those features which fix such a conclusion uniquely*.³²

I interpret Moffett as saying that a peer is someone who is as likely as me to be right, *given whatever might make for a uniquely rational conclusion*, including both first-order evidence, and background theoretical beliefs. The way to avoid denying the *peerhood* of an interlocutor (which I take to be synonymous with Moffett’s confusing usage of “competence”) is to recognize that those things which *would* make for a uniquely rational conclusion, namely, sameness of background theoretical beliefs, in this instance *do not*. In other words, our background theoretical beliefs differ, but are equally rational, because they are “theoretical beliefs” and so underdetermined by the evidence for them. On this view, both interlocutors can admit that the evidence is as consistent with the other’s theory as it is with their own, and neither needs to abandon their own view, nor deny peerhood to the other. “In this way,” he says, “underdetermination appears to allow for blameless (and so, reasonable) disagreement.”³³

Moffett here considers an objection from Richard Feldman: underdetermination amounts to the same thing as the Equal Weight View--the result is still suspension of

³¹ Ibid.

³² Ibid., 366, fn. 8. Italics added. The footnote continues: “This way of framing the issue has the evident virtue of making correctness an issue of epistemic rationality rather than truth, which is what we would expect of a discussion concerning reasonable disagreement.” On this idea, cf. my discussion of van Inwagen’s version of the Uniqueness Thesis in Chapter One.

³³ Ibid., 360.

judgment. This is because if there is no uniquely rational theory, then our choice of theory is arbitrary and therefore cannot be preferable to our peer's choice of theory. Moffett responds by arguing that *epistemic conservatism* is true. This is the view that underdetermination need not undermine one's own theory because one has prima facie justification for any belief that one holds.³⁴ This is not to say that one was justified in *forming* any belief that one happens to have formed, but rather that one has prima facie justification for *relying on* one's standing belief. As Moffett says, "For any individual x and proposition p , x 's believing that p is both necessary and sufficient for x 's being prima facie justified in believing that p ."³⁵

The question then is: Does the existence of an equally adequate but incompatible theory provide counterevidence sufficient to outweigh this prima facie justification? Moffett answers no, because believing a theory has *theoretical benefits*. Believing things, Moffett says, is a crucial part of having an overall worldview, and this in turn "is fundamental to our ability to produce and develop new ideas and direct the course of further investigation."³⁶ Note that this benefit is distinct from any concern for the theory's truth. As Moffett says, "what is being claimed is that genuinely believing a theory helps us to elaborate and develop that theory even though we acknowledge that the current evidence in support of it makes it no more likely to be true than some set of

³⁴ A simpler way to respond to Feldman is as follows: if our choice of theory is arbitrary in any meaningful sense in this context, then we can simply embrace a "relativistic" sense of rationality. Instead of "going agnostic," as Feldman would have us do, we say we're both rational, in much the same way that the person who sees a rabbit in Wittgenstein's famous image is equally rational as the one who sees a duck. This is arguably a less controversial claim than the epistemic conservatism Moffett opts for. Thanks to Michael Wreen for this point.

³⁵ *Ibid.*, 361.

³⁶ *Ibid.*, 362.

alternatives.”³⁷ Referencing Kuhn, Moffett notes that elaborating and developing one’s theories “plays a crucial role in overall theoretical progress,” the implication being that genuine belief in one’s view encourages such progress.³⁸ This interpretation of Moffett’s view is confirmed when he quotes Mill to the effect that “it is genuine believers who are best positioned to... ‘do their very utmost’ for a theory.”³⁹

There is another benefit as well. Maintaining one’s standing beliefs is more conducive to one’s own self-understanding as a seeker of knowledge.

From the point of view of one’s intellectual development, our theoretical commitments shape, not just how we see things, but who we are... Abandoning such commitments simply because an attractive competitor, even an otherwise equally attractive one, arises does not do justice to our status as epistemic agents.⁴⁰

I take this to mean that it is reasonable for me to maintain my standing belief in the face of an alternative that I recognize is equally likely to be true, because failing to do so would undermine my own goal of becoming a responsible and able epistemic agent. Or, perhaps more accurately, a *habit* of failing to do so would undermine that goal, because becoming a responsible and able epistemic agent is a long-term project that is hampered by giving up too easily. Moffett here considers the objection that this is not a properly “epistemic” reason to hold onto one’s belief, but is instead a “prudential” or “moral” reason. He replies that, for two reasons, it should be considered epistemic. First, “it is our status as epistemic agents which generate [*sic*] the commitments” to become good epistemic agents.⁴¹ Second, even if a reason is prudential, it is not therefore *not* epistemic.

³⁷ Ibid.

³⁸ Ibid.

³⁹ Ibid., 363.

⁴⁰ Ibid., 362-363.

⁴¹ Ibid., 363.

This is because “what is morally/prudentially required ought to be epistemically permissible and, conversely, what is epistemically prohibited ought not to be morally/prudentially required.”⁴² So moral/prudential reasons can have epistemic entailments.

So Feldman, Moffett says, is guilty of “construing the epistemic in an inappropriately narrow manner, as reasons which make the belief itself more likely to be true.”⁴³ In other words, practical reasons are not precluded from playing an epistemic role by virtue of their epistemic entailments. In this sense, something can be an epistemic reason for belief even if it doesn’t make the belief more likely to be true.⁴⁴ Because there is this practical benefit to privileging one’s own standing beliefs, one is permitted in holding onto them in the face of a competing, equally well justified theory. In this way, reasonable disagreement with an epistemic peer is possible.

8.3 - Hume Again

We can now elaborate our own Humean response along similar lines, but without Moffett’s problematic view of peerhood.

Recall the Parity Problem:

Parity: From the perspective of either party to a peer disagreement, the epistemic situation is indistinguishable from one in which the accuracy of the assessments are inverted. Things would appear just as they do to the correct party if she were incorrect, and vice versa.

⁴² Ibid.

⁴³ Ibid., 362.

⁴⁴ This is not to say that epistemic and prudential reasons are *equivalent*. Presumably Moffett would grant that there can be purely epistemic reasons, i.e., epistemic oughts which are not prudential or moral.

As we saw above, Parity implies a conflict in one's own evidential seemings: it appears to one simultaneously that the first-order evidence supports *p*, and that one's recognition of that support may be erroneous. Moffett's solution suggests that a person is reasonable to maintain her view in a situation of parity owing to the practical benefits of resisting habits of regular belief revision, benefits which accrue to one's overall theoretical goals as an epistemic agent. So far, however, Moffett is still thinking of peerhood judgments in terms of comparing the reported seemings of another with one's own first-person evidential seemings. His adoption of Elga's "equal likelihood to be correct" account of peerhood ensures that when I judge that another person is my epistemic peer, the content of that judgment is about our comparative capacity to be correct about *p*.

This, however, is not what motivates the Parity Problem. If this were all there is to peerhood judgments, then Wedgwood's egocentric epistemic bias discussed above would be an adequate solution: we could simply say that my own intuitions about the disputed proposition are presumed basically reliable, while yours are not, and justify this by appeal to the larger epistemic goal of minimizing sources of error.⁴⁵ However, the Parity Problem arises only when the reliability of an interlocutor's intuitions become indistinguishable from the reliability of one's own, after long and inconclusive dispute. After long dispute with a potential peer, I come to see that his/her methods of evidence assessment are very similar to mine, and this then prompts the realization that we are epistemic peers. Moffett's solution, intelligent and well-argued as it is, cannot help here, burdened as it is with seeing the peer's perspective as fundamentally different from one's own.

⁴⁵ See my discussion of Wedgwood's view in Chapter Four, pp. 90-92.

This is where Hume is most helpful. Viewed through a Humean lens, peerhood judgments are not just about reports of another's point of view and its likelihood to be correct. They are about the resemblance of another's point of view to my own. To judge that another is my epistemic peer is to judge that when he/she has an evidential seeming, it is as if I had had it, and hence the reasonable response to that seeming is my own response. Put differently, the experience of recognizing an epistemic peer entails that I have competing, equally compelling responses to my total evidence: the response that I would have had without the disagreement, and the response of my peer, which has borrowed its "vivacity" from my trust in my own methods of evidence assessment.

From this Humean perspective, the Parity Problem makes sense: it is the expected outcome of sustained disagreement with an equally qualified interlocutor. And as we saw, a response to the problem based on demonstrative reason is not to be had, owing to the rough nature of peerhood. But Moffett's insight is still helpful: there may be practical benefits to maintaining one's standing belief--the belief one would have had without the disagreement--which make holding that belief the overall reasonable thing to do, even in a situation of parity.

To see as much, consider again the case of two philosophers who disagree about free will (à la van Inwagen and Lewis). Call them A and B. A and B share all of the evidence and arguments related to free will that they can think of over a long period of time, and each observes the other acting as carefully, conscientiously, creatively, fairly, thoroughly, etc., as oneself. Each consequently comes to believe the judgment of the other resembles one's own. A has a clear and vivid evidential seeming that the first-order evidence E (i.e., all the evidence and arguments pertaining to free will that A can think

of, plus the evidence and arguments that B has shared with A) supports p . A also knows that B has a clear and vivid evidential seeming that the same body of first-order evidence supports $\sim p$. After long exposure to B's habits of thought, A comes to consider B an epistemic peer, meaning that A now takes B's reported evidential seeming to be as evidentially weighty as A's own evidential seeming. That is to say, from A's perspective, B's report that E supports $\sim p$ carries the same weight as A's seeming that E supports p . And A knows that if in fact $\sim p$ is correct, things would look just the same to A as they do, because as things presently stand, A has the conflicting seemings that E supports p (from A's own first-order evidential assessment), and that E supports $\sim p$ (from A's higher-order evidential assessment derived from the recognition of B's peer status, a recognition which carries with it the weight of long observation of B's epistemic virtue). Unlike most cases of disagreement, then, A has more to go on than *this* piece of disagreeing testimony alone--A also has the entire past history of observing B's reasoning, which has led A to believe that B reasons as well as A does about this kind of thing. A is thus in a situation of parity with B.⁴⁶

The Hume/Moffett solution to this problem is to see that there is practical, *epistemic* benefit in maintaining one's standing view, even though one may *not* think that one is more likely to be correct about p . Why not suspend judgment in this case, as Conformists such as Feldman would recommend? Because, Moffett says, given our limited understanding of the world, a practice of suspending judgment in that way would

⁴⁶ Note that this situation also satisfies our modified version of Feldman's Meta-Evidence Principle from Chapter Six: "MEP*: For peers S1 and S2, evidence of S2's trustworthy judgment that p is evidence for S1 that p ."

be devastating for our project of developing a coherent worldview. The benefit of standing one's epistemic ground, Moffett says, accrues at two levels:

1. The group level: *Ceteris paribus*, the overall human project of developing a coherent worldview works better if people stick with the theories that they hold (even while admitting that the evidence does not tell solely in their favor).
2. The individual level: *My* goal of becoming a knower involves a long-term project that requires resilience in the face of uncertainty and even epistemic parity.

The underlying insight of both points seems to be that if one wants to know, one must give it one's best shot, which implies sticking with a view until there is good positive reason to abandon it, and not merely disagreement with a person like oneself. Similarly, our Humean approach implies that the habits of considering myself reliable (and my seemings *prima facie* evidential) and of considering those who reason like me to be similarly reliable (and so their seemings also *prima facie* evidential) are ultimately productive of knowledge. In general, I must assume that things are the way they seem to me to be, even if I can't demonstrate that they are, and I am reasonable in trusting those who see things the way that I do.

The danger here is that this will result in a situation in which the justification of *any* belief is unlikely to be challenged because the people that one is likely to consider worthy dialogue partners are those with whom one already agrees.⁴⁷ But this can be mitigated by building in the Millian expectation that one is seeking out the most informed

⁴⁷ See Chapter Seven, fn. 20 for an argument that Elga's Equal Weight View really does have this undesirable consequence.

disagreeing peers for dialogue, as well as doing one's best to maximize one's standing in such dialogues by developing all the relevant cognitive virtues--care, thoroughness, conscientiousness, etc. One will, as a matter of course, encounter evidence, both first-order and higher-order, that counts against one's standing view, especially if one is engaged in this Millian project. In most cases, a new standing view will naturally develop. It is only when one forms an *unjustified* attachment to a view, for reasons *other than* any of the peer factors, that the danger mentioned above becomes salient. But our Hume/Moffett solution neither implies nor recommends such an approach to belief.

So, there is at least one way out of the Parity Problem. Briefly, it is that there are practical reasons to maintain one's view that are properly epistemic reasons without being truth-conducive reasons. Essentially, this is to interpret the notion of "reasonable" as "epistemically blameless" with respect to what can be expected of knowers, given their overall epistemic goals. As Moffett says, "underdetermination appears to allow for blameless (and so, reasonable) disagreement."⁴⁸ While *blameless* does not entail *reasonable* in the usual sense of "believing appropriately"--one might, after all, believe wrongly through no fault of one's own--nonetheless, blamelessness might be the best one can hope for in a situation of parity with underdetermined beliefs. To be reasonable in this sense does not entail being right, or even thinking that one's position is more likely to be right than any other position that one is aware of; it is merely doing the best one can in the circumstances. One is "reasonable" in this weaker sense insofar as one is blameless with respect to the dispute one is involved in, given one's overall epistemic goals and the

⁴⁸ Moffett, "Reasonable Disagreement and Rational Group Inquiry," 360.

norms that govern those goals. To have a reasonable disagreement is to engage blamelessly with an equally blameless peer.

Chapter Nine – Proposed Solutions II: Kant’s Contribution

9.1 – Kant’s Categories of Judgments and Disagreement

Just as the Humean solution to the Parity Problem argued for above was inspired by a contemporary figure (Moffett), so too the Kantian solution is inspired by another recent philosopher: William P. Alston. Specifically, it is inspired by a remark about Alston that Thomas Kelly makes in a footnote in “Peer Disagreement and Higher-Order Evidence.”⁴⁹ At the place in question, Kelly is considering a potential objection from the Equal Weight theorist regarding a situation of peer disagreement in which one party has in fact evaluated the evidence more accurately than another. Representing the Equal Weight View, he says:

Even if things are asymmetrical at the bottom level (one’s belief reflects the evidence better than one’s peer’s belief...), things are symmetrical one level up: one has no justification for thinking that one’s belief better reflects the evidence... One is thus no more justified in thinking that one’s own belief accurately reflects the evidence than one’s peer is in thinking that *his* belief accurately reflects the evidence. Therefore, given the higher level normative symmetry, it would be unreasonable to favor one’s own belief over the belief of one’s peer.⁵⁰

Kelly then adds this suggestive footnote:

In response to this objection, a proponent of the Total Evidence View might contend that it rests on a ‘level confusion,’ in the sense of Alston (1980): in particular, that it falsely assumes that, in order to be justified in believing *p*, one must be justified in believing that one is justified in believing *p*. In effect, such a response concedes, at least for the sake of argument, that there is a higher-level normative symmetry between the peers but denies that anything directly follows from this about the epistemic statuses of their first-order beliefs. Here I simply

⁴⁹ Thomas Kelly, “Peer Disagreement and Higher-Order Evidence,” in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford University Press, 2010), 111-74.

⁵⁰ *Ibid.*, 155.

want to note the possibility of such a response without exploring its prospects; the response that I offer in the main text proceeds along quite different lines.⁵¹

I think this is a potentially fruitful avenue of response. Just as, according to Alston, in order to be justified in believing *p*, I do not need to also be justified in believing that I am justified in believing *p*, I can judge that *p*, based on reasons, without needing to be able to demonstrate that I have judged correctly, even if I think it's possible that I haven't judged correctly. Applying this to a situation of parity, I can believe simultaneously that proposition *p* is true because of evidence *E*, and that I cannot demonstrate this, nor provide any more reason to think *p* is true than my peer can provide to think that *p* is false, because we both appeal to the same shared evidence *E*. So, to return once more to one of the key questions of this study: Is it rational for me to hold onto my first-order view of *p* when I recognize that I am in a situation of parity, and that things would therefore seem just the same to me if I were mistaken as they would if I were correct? An Alston-inspired response would suggest that the answer is yes, owing to the level of confusion involved in assuming that higher-order symmetry must alter apparent first-order asymmetry. But we should like to know more about why this is a "confusion," since, especially in Parity cases, it seems that the higher-order symmetry *should* make a difference to the lower-order asymmetrical justification for *p*.

Kant is helpful here, as he provides a theoretical framework for possible analyses of peerhood judgments--judgments that so-and-so is my peer--which provide two distinct ways of explaining how it may be appropriate (and in what sense) to maintain one's view in the face of higher-order epistemic parity. Kant thinks that judgments, while always

⁵¹ Ibid., fn. 33. Kelly is referring to Alston's "Level Confusions in Epistemology," *Midwest Studies in Philosophy* 5 (1980): 135–50.

essentially propositional, can play various roles that are not *reality-focused*. In other words, some judgments may not be directed at “mapping the world,” that is, recording facts about the world. Robert Hanna describes such non-reality-focused roles as “uses” of judgments, which include pragmatic, moral, aesthetic, and teleological uses.⁵² According to Hanna, a “non-truth-theoretic,” or simply *non-theoretical*, use of a judgment has an overall purpose other than knowing the truth about the world.

...non-theoretical judgments... are used *non-truth-theoretically*. For example, the overall rational purpose of a non-theoretical judgment might be to make instrumental or non-instrumental free choices or decisions... Or it might be to evaluate the beauty or sublimity of phenomenal objects (aesthetic judgments of taste)... Or it might be to treat natural things as if they had goal-directed or purposive structure, as a heuristic guide for the construction of better mechanistic explanations of physical phenomena (teleological judgments)...⁵³

The important thing to note about a non-theoretical use of judgment is that its “rational purpose or function... can be realized even if its propositional content is false. What matters for a non-theoretical judgment is how things *seem* to the judger, not how they actually *are*.”⁵⁴ In support of this claim, Hanna cites the passage in the third *Critique* in which Kant explains the disinterestedness of judgments of taste.⁵⁵ There, Kant argues that the existence of an object is irrelevant to the judgment of its beauty, and that all that matters is that the representation of the object gives one satisfaction. “If the question is whether something is beautiful,” Kant says,

One only wants to know whether the mere representation of the object is accompanied with satisfaction in me, however indifferent I might be with regard

⁵² Robert Hanna, “Kinds of Use” Supplement to “Kant’s Theory of Judgment,” in *The Stanford Encyclopedia of Philosophy*, Winter 2018, <https://plato.stanford.edu/archives/win2018/entries/kant-judgment/>, Introduction.

⁵³ *Ibid.*, Section A.

⁵⁴ *Ibid.*

⁵⁵ Immanuel Kant, *Critique of the Power of Judgment*, ed. Paul Guyer (Cambridge University Press, 2000), 90-1 (5: 205).

to the existence of the object of this representation. It is readily seen that to say that it is **beautiful** and to prove that I have taste what matters is what I make of this representation in myself, not how I depend on the existence of the object.... One must not be in the least biased in favor of the existence of the thing, but must be entirely indifferent in this respect in order to play the judge in matters of taste.⁵⁶

This would make judgments of taste analogous to peerhood judgments, as these too are rooted in how things *seem* to the judge. How things *are*, at least in Parity cases, is of no justificatory importance, since that information is unavailable to both parties.

On the other hand, Kant also distinguishes between various uses of theoretical judgments, judgments that *are* directed at recording facts about the world. The primary distinction, for our purposes, is that between *constitutive* and *regulative* uses of theoretical judgments. According to Hanna, a constitutive use of judgment is one for which “its meaning, truth, or legitimacy as a ‘principle’ (*Grundsatz, Prinzip*) is not based on any further assumptions, hypothetical conditions, or suppositions.” Alternatively, a regulative use of judgment is one for which “its meaning, truth, or legitimacy as a principle is based on some further assumptions, hypothetical conditions, or suppositions.”⁵⁷ Michael Friedman helps to clarify this distinction within the context of Kant’s larger critical project:

In the *Critique of Pure Reason* this distinction marks the division between the faculty of reason and the faculty of understanding. The understanding--together with its pure concepts or categories--is constitutive of the possibility of experience. All experience must conform to the concepts and principles of the understanding, which, accordingly, are necessarily realized or instantiated in experience: experience necessarily contains substances, causal connections, and so on. The faculty of reason, by contrast, is merely regulative in relation to experience. Although reason too plays an indispensable role in experience, the concepts proper to it--the so-called ideas of reason, such as the idea of God, or the idea of the world as a complete totality--cannot be realized or instantiated in

⁵⁶ Ibid.

⁵⁷ Hanna, “Kinds of Use,” Section B.

experience at all. Nevertheless, ideas of reason--the idea of a highest intelligence or wise Author of the world, for example--still function legitimately to guide empirical enquiry into the objects that can be given in experience.⁵⁸

Friedman here quotes Kant:

... it is said, e.g., that the things in the world must be considered **as if** they had gotten their existence from a highest intelligence. In such a way the idea is only a heuristic and not an ostensive concept; and it shows not how an object is constituted but how, under the guidance of that concept, we ought to **seek after** the constitution and connection of objects of experience in general.⁵⁹

Friedman continues:

Regulative concepts and principles therefore present us, not with objects corresponding to them, but rather with a task: the never ending progress of empirical enquiry whose ideal *terminus*--the *complete* understanding of "the constitution and connection of the objects of experience"--can only be approached asymptotically.⁶⁰

So constitutive judgments pertain to the necessary features of all possible experience.

Regulative judgments, on the other hand, "fill in the general form of pure natural science with actual empirical content, which can only be done in a progressive and asymptotic fashion."⁶¹

A further feature of regulative judgments is that they may be *empty*, in the sense of "lacking objective validity," even though they have the same theoretic *form* as judgments which are not empty (like objectively valid synthetic a priori judgments). For example, the proposition "God exists," while having no proper constitutive use--since it

⁵⁸ Michael Friedman, "Regulative and Constitutive," *The Southern Journal of Philosophy* XXX (1991): 73.

⁵⁹ Immanuel Kant, *Critique of Pure Reason*, ed. Paul Guyer and Allen Wood (Cambridge University Press, 1998), 606 (A671/B699).

⁶⁰ Friedman, "Regulative and Constitutive," 73. Friedman goes on to complicate this distinction, contextualizing it within the state of the current understanding of physics and chemistry at the time of Kant's writing. None of these complications need concern us here.

⁶¹ *Ibid.*, 75.

refers to a concept of a noumenal object--nonetheless has a legitimate regulative use in morality. As Hanna says, the proposition “God exists” is

... a “postulate of practical reason,” which in turn is a basic propositional target of moral belief or *Glauben*, according to which our moral lives are to be conducted as if we could epistemically believe or scientifically know that God benevolently guarantees that all and only the happy people are morally virtuous, in order to strengthen our otherwise naturally flawed and shifting moral resolve, and make our moral lives fully meaningful (A633-634/B661-662, A812-819/B840-847).⁶²

Cases like this, Hanna says, show that

... otherwise empty or non-objectively-valid judgments are used regulatively when their legitimacy as principles depends on their being adopted solely for the purpose of making scientific inquiry or moral life into a coherent, meaningful whole.⁶³

With these distinctions in mind, two Kantian solutions to the Parity Problem are possible. The first assumes that peerhood judgments, judgments of the form “so-and-so is my epistemic peer,” are theoretical; the second assumes that they are not.

The first possible Kantian solution is this. Assume that peerhood judgments are theoretical. They thus may be construed as possibly *empty regulative* theoretical judgments. Due to underdetermination and the vague nature of peerhood, such judgments lack objective validity, but we treat them *as if* they were objectively valid for the purpose of making our epistemic lives into “coherent, meaningful wholes.” How is this accomplished? It is important for me to be able to judge that another is my epistemic peer on the basis of our similar past reasoning (keep in mind that we are considering situations of Parity, which are rooted in such observations), since this helps to justify my beliefs about the world by giving me reason to believe that others who implement habits of

⁶² Hanna, “Kinds of Use,” Section B.

⁶³ Ibid.

thought similar to mine come to similar conclusions.⁶⁴ Via this route, peerhood judgments are one way in which my understanding of “the constitution and connection of the objects of experience” are approached *asymptotically*, as Friedman says. Nonetheless, in the case of prolonged disagreement with a recognized epistemic peer I have recourse to the fact that my judgment of his/her peerhood is merely regulative and empty, and so may fail to guide my beliefs about how the world is in any given *instance*, even as it helps me to justify such beliefs *overall*. By way of analogy, while for Kant the judgment that God will reward moral virtue plays an important regulative role, nonetheless it may be epistemically appropriate that one should doubt more specific claims about how God metes out reward and punishment. Similarly, while it is important for the reasons just mentioned that I judge that some others are my epistemic peers, it is epistemically appropriate that I should doubt in a specific instance of disagreement that someone’s view of *p* in this case is just as good as mine. Moreover, the justification for the doubt is similar in both cases: the claims conflict with how things seem to me given my own processes of reasoning about the propositions in question--processes I must trust to deliver justified beliefs about the world.

But there is also the second Kantian possibility. Assume that peerhood judgments are non-theoretical. It may then be reasonable to maintain one’s first-order view in the context of Parity because the judgment that someone is a peer does not entail that his/her view should make a difference with respect to the truth of a disputed proposition. By way of analogy, consider judgments of taste, which Kant also takes to be non-theoretical. As

⁶⁴ In other words, peerhood judgments lend my beliefs a measure of intersubjective validity that they would not otherwise have by making me aware that believing various things on the basis of certain habits of thought is not just a fluke in my case.

we saw above, for Kant what matters in a judgment of taste is “what I make of this representation in myself, not how I depend on the existence of the object.”⁶⁵ As Hanna says, “A visual experience of that thing over there, which you take to be a beautiful rose, can be a genuine experience of the beautiful, even if in fact it is not a rose or does not even exist.”⁶⁶ The analogy with peerhood judgments is that they, like aesthetic judgments, may be appropriate (rational) even if false, because their function is not to produce true beliefs about the world.

This is not, however, to say that such judgments make no demands on others. For Kant, aesthetic judgments are both subjective (based on feelings of pleasure) and *universal*, in the sense that they call for universal agreement. As Kant says:

It would be ridiculous if... someone who prided himself on his taste thought to justify himself thus: “This object (the building we are looking at, the clothing someone is wearing, the poem that is presented for judging) is beautiful **for me**.” For he must not call it **beautiful** if it pleases merely him. Many things may have charm and agreeableness for him, no one will be bothered about that; but if he pronounces that something is beautiful, then he expects the very same satisfaction of others: he judges not merely for himself, but for everyone, and speaks of beauty as if it were a property of things. Hence he says that the **thing** is beautiful, and does not count on the agreement of others with his judgment of satisfaction because he has frequently found them to be agreeable with his own, but rather **demands** it from them. He rebukes them if they judge otherwise, and denies that they have taste, though he nevertheless requires that they ought to have it; and to this extent one cannot say, “Everyone has his special taste.” This would be as much as to say that there is no taste at all, i.e., no aesthetic judgment that could make a rightful claim to the assent of everyone.⁶⁷

⁶⁵ Kant, *Critique of the Power of Judgment*, 90-91 (5: 205).

⁶⁶ Hanna, “Kinds of Use,” Section A. Note that this is not to say that one cannot be *wrong* about having experienced a real beautiful object. Rather, the basis of the judgment that something is beautiful--whether that judgment is true or not--is the pleasure or “satisfaction” that is caused in one by the representation of the object, not the existence of the object, nor any attachment one may have to the existence of the object.

⁶⁷ Kant, *Critique of the Power of Judgment*, 98 (5: 212 - 5: 213).

But because the judgment is also subjective, it is not *really* about a feature of the world, but rather about a relation between ourselves and something in the world. But the fact that I find it pleasurable is also an indicator that the beautiful object has a certain *fit* with my cognitive faculties, and so I can demand that others with similar faculties also find it pleasurable. This is the difference between a beautiful object and a merely pleasing object for Kant. With a beautiful object, I can reasonably demand a similar response from others with similar faculties.

This is also the place at which the analogy with peerhood judgments arises: the judgment that another is my peer is the judgment that he/she has similar faculties as myself with respect to assessing evidence. I can thus reasonably demand agreement about p on the basis of our shared evidence E. However, as with judgments of taste, if a dispute cannot be settled, this does not entail that either party is judging irrationally. This is because of the vague nature of the concepts involved in such judgments. As David Stern notes,

This claim to necessity [of an aesthetic judgment] is a sure sign that concepts play a role in the judgment. But this is immediately counterbalanced by the claim that the concepts in question cannot be determined--and so they do not permit a proof that the object is beautiful or not. And that, in outline, is the approach Kant adopts: aesthetic judgments are conceptual, and so can claim universal validity, but the concepts are intrinsically indeterminate, and so we cannot settle aesthetic disputes.⁶⁸

Something similar can be said for peerhood judgments based on “disagreement factors.”

These concepts (conscientiousness, thoughtfulness, carefulness, etc.) are not

indeterminate, but they are vague, and so even though the judgments based on them are

⁶⁸ David Stern, “Are Disagreements about Taste Possible? A Discussion of Kant’s Antinomy of Taste,” *The Iowa Review* 21, no. 2 (1994): 69.

conceptual (i.e., “S is my peer” is a claim that aspires to universal validity), they cannot be demonstrated in the way that a theoretical judgment can be.⁶⁹ Owing to this inability to demonstrate the conceptual bases of one’s judgments, the parties in a dispute about taste may both be reasonable to persist in their disagreement, both because--contra Hume--the truth about whose taste judgments are best cannot be established, and because the awareness of another’s differing aesthetic experience does not outweigh my own.⁷⁰ Similarly, in Parity situations, the parties in a dispute about what their shared evidence supports may both be reasonable to persist in their disagreement, and their peerhood judgments do not entail that either should revise his/her view in the direction of the peer.

9.2 – Conclusion

Thus, reasonable disagreement in Parity situations is possible in three ways:

(1) The Humean Solution. Justified belief is maintained in order to pursue our larger epistemic goals, even though we don’t have sufficient reason to believe that we are right, and indeed do have sufficient reason to believe that things would look just the same to us if we were *not* right.

(2) The First Kantian Solution. We classify peerhood judgments as empty regulative theoretical judgments, with the result that they do not entail that a peer’s view must be as weighty as our own view.

⁶⁹ That is to say, I cannot demonstrate for my interlocutor: “You are my peer based on factors x, y, z; I reasoned this way; therefore, you ought also to have reasoned this way.”

⁷⁰ For discussion of the disagreement between Hume and Kant regarding the possibility of an objective standard of taste, see Jens Kulenkampff, “The Objectivity of Taste : Hume and Kant,” *Nous* 24, no. 1 (1990): 93-110, esp. 99-101.

(3) The Second Kantian Solution. We take peer disagreement to be analogous to disagreement about taste, with the result that we believe appropriately, and are justified in expecting others to believe as we do, even though we cannot demonstrate that, and even though we know that our peers are similarly justified.

There are objections that could be brought against all three of these solutions to the Parity Problem. With respect to (1), one may press the objection that even though one may be “blameless” in one’s disagreement with an epistemic peer (given one’s overall epistemic goals and so on), it does not follow that it is *reasonable* to maintain one’s belief. Perhaps a better view of Parity cases is this: while no one should be expected to act any differently than they do, nonetheless one party must be less reasonable than the other, even if we cannot tell which.

With respect to (2), one may object that considering peerhood judgments to be regulative leads to negative epistemic consequences. Recall that the thought here is that the recognition that others who implement habits of thought similar to mine come to similar conclusions would help to confirm my beliefs, thus strengthening my resolve about them. This is similar to believing that God rewards dutiful action, for that can “strengthen [my] otherwise naturally flawed and shifting moral resolve.”⁷¹ The objector might say, however, that this view suffers from a similar problem to the one earlier leveled against Elga’s Equal Weight View: peer status will be granted to those who reason like myself, which is likely to be determined by who agrees most often with me. That in turn will lead to an insular, parochial system of beliefs that is unlikely to be seriously challenged by new evidence.

⁷¹ Hanna, “Kinds of Use,” Section B.

With respect to (3), one may deny the analogy between peerhood judgments and judgments of taste, arguing that the former are clearly theoretical, i.e., clearly claims about the world. While many would grant that there may not be an objectively valid answer to the question whether a particular painting is beautiful, can we say the same for the question whether two people are epistemic peers with respect to a proposition? As with the Humean solution, perhaps there *is* an answer to this question, even if no one is in a position to know what it is.

Where does this leave us with respect to peer disagreement? Some response can be made to these objections. With respect to (1), the objector should be reminded that what motivates interest in the peer disagreement problem is the desire to *respond well* to capable interlocutors who disagree with us. Whether or not responding well amounts to the same thing as being reasonable is largely a semantic matter, and at any rate would not tell us *what to do* in a real-life case of disagreement under conditions of Parity. The Humean solution at least provides some help here.

With respect to (2), we can reiterate a point made above: avoiding a problematically self-referential belief system depends largely on a good faith effort to practice the intellectual virtues which one uses to determine peerhood in the first place.⁷² While it is probably true that taking peerhood judgments to be regulative is consistent with a narrow, self-referential system of belief, I do not think that this view makes such a system epistemically superior.⁷³

⁷² See pp. 184-185 above.

⁷³ Contrast this with Elga's Equal Weight View, which apparently *recommends* choosing peers based only on similarity of past conclusions.

With respect to (3), I confess that I am not sure whether the question “Is so-and-so my epistemic peer?” has an objectively valid answer, or whether there is a *fact of the matter* about peerhood in a given case (i.e., whether peerhood judgments are theoretical). As such, I am not sure how far the analogy between peerhood judgments and judgments of taste should be carried. But I am also not sure that this question *does* have an objectively valid answer, because the factors used to determine peerhood are often quite indefinite indeed, even in the best of real-life prolonged disagreement cases. Perhaps there are other pertinent differences between taste judgments and peerhood judgments that would ultimately undermine this analogy, but recall that the point of the analogy was to illustrate one way that non-theoretical judgments could have implications for the relationship between higher-order and lower-order beliefs. This point is not defeated even if the analogy with taste cannot be carried very far.

In conclusion, I believe that peer disagreement in Parity contexts presents a real epistemic problem that has not yet been adequately addressed in the literature, and that the Humean and Kantian solutions lightly canvassed above go some way towards mitigating this problem. Nonetheless, even if all three of these solutions are ultimately untenable, it is worth keeping in mind the points made above about the skeptical nature of the disagreement problem.⁷⁴ While the possibility of Parity may yet entail that many controversial beliefs are unjustified, it remains true that Parity is likely a rare phenomenon for most people, including experts.

⁷⁴ See pp. 92-93.

Conclusion

In the Introduction, I distinguished between several questions in the literature on peer disagreement, and distilled them into three Main Questions:

(MQ1) What expectations can one reasonably have of a body of evidence for a proposition?

(MQ2) What is the relevance of private evidence to peer disagreement?

(MQ3) What are one's epistemic responsibilities in the context of disagreement?

In this study I have also argued for answers to each of them. In response to (MQ1), I argued in Chapter One that Feldman and van Inwagen's respective versions of the Uniqueness Thesis--the view that "a body of evidence justifies at most one proposition out of a competing set of propositions"--are untenable.⁷⁵ Consequently, we do not have prima facie reason to rule out the possibility of rational disagreement by virtue of the nature of evidence alone.

In response to (MQ2), I argued in Chapters Two and Three that there is unshareable, "private," evidence that is relevant to the justification of one's belief in a proposition, and that it is best understood in terms of evidential seemings--the felt sense that some evidence supports a proposition. Further, these seemings are analogous to perceptual seemings in important ways that have not previously been spelled out in the disagreement literature. I argued that such seemings have evidential significance in a way similar to that of perceptual seemings, and that this significance does not disappear on the discovery of disagreement, though the seemings do of course remain defeasible. I then

⁷⁵ Richard Feldman, "Reasonable Religious Disagreements," in *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, ed. Louise Antony (Oxford University Press, 2007), 205.

considered how evidential seemings might be overcome by the higher-order evidence against one's view provided by peer disagreement.

In Chapters Four through Six, I considered and rejected several ways in which disagreement may defeat one's evidential seemings, such as by providing evidence of the contingency of one's belief (à la Gideon Rosen), by creating epistemic symmetry between peers so that a peer has no way to identify which peer's seemings are correct (à la David Christensen and Adam Elga), or by constituting *meta-evidence*: evidence that there is evidence for p (à la Richard Feldman). I concluded that the fact of disagreement itself does not necessarily provide defeating evidence regarding what one should believe about p , and that meta-evidence may not in fact constitute evidence for a subject that is directly relevant to the disputed proposition. However, I argued that there is something important about higher-order evidence: insofar as one is dealing with a perspective that is constituted by certain features of epistemic excellence, the views of a person with that perspective may count as evidence for a disputed proposition. In the context of Parity, I argued, prolonged disagreement with such a person may provide evidence sufficient to counterbalance one's own evidential seemings.

I also considered the related issue of the relationship between the epistemological problems of testimony and disagreement, as well as the problem of merely possible rational disagreements. With respect to the former, I argued that thinking of testimony in a virtue-theoretic way helps to clarify when disagreement, as known through testimony, carries evidential weight (namely, when the peer is recognized as trustworthy). Regarding the latter, I argued that the problem of possible rational disagreement, while initially

puzzling, does not actually create a unique epistemic problem, nor does it worsen the standard problem of disagreement rooted in real-world disputes between peers.

In response to (MQ3), I argued that once the notion of epistemic peerhood is clarified, both the Conformist and Nonconformist positions on the disagreement problem can be seen to be inadequate to deal with the problem of prolonged disagreement under Parity. To this end, in Chapter Seven I argued that an influential view of peerhood advanced by Adam Elga is mistaken. I then argued that Conformist views (e.g., those defended by David Christensen, Richard Feldman, and Adam Elga) place too much emphasis on the evidential value of disagreement itself, mistakenly taking the fact of disagreement to have some defeating power that it does not in fact have. They also tend to assume (e.g., Christensen) that peer disagreement is epistemically similar to disagreement with a superior, which is unwarranted. On the other hand, Nonconformist views (e.g., those defended by Thomas Kelly, Marc Moffett, and Michael Bergmann) tend to underestimate the significance of epistemic parity produced by long disagreement with someone one takes to be a peer.

In Chapters Eight and Nine, I proposed and discussed three possible solutions to the Parity Problem, one inspired by David Hume, and two inspired by Immanuel Kant. The Humean solution (Chapter Eight) held that practical considerations stemming from the underdetermination of one's views by the available evidence, as well as one's larger epistemic goals, can justify maintaining one's view in a Parity situation. I argued that Hume helps us to see that it can be epistemically appropriate to maintain a view even when one knows that it cannot be demonstrated to be the uniquely rational conclusion to draw on the basis of some evidence set. Alternatively, in Chapter Nine I drew on Kant's

theory of judgment to argue that judgments that someone else is an epistemic peer may be classified as either (1) empty regulative theoretical judgments, or (2) non-theoretic judgments (analogous to judgments of taste). In either case, I argued, rational disagreement under Parity may be possible.

We are left, then, with the view that rational disagreement between recognized epistemic peers who have shared their evidence with one another is indeed possible. In fact, in a typical case of peer disagreement, such as most people may find themselves in with a colleague, a family member, or a stranger, rational disagreement is not even very difficult to achieve. It requires merely that one is justified in *not* weighing the second- or third-person testimonial report of a peer's evidential seeming, in the present case, as heavily as one's own first-person evidential seeming to the contrary. I argued above (§ 2.2, § 4.2) that one is typically so justified. In most cases of peer disagreement, either the disagreement itself or the granting of peer status to one's interlocutor is short-lived. This is because considering someone else an epistemic peer is an inexact approximation, and given the choice between weighing someone's evidential seemings the same as my own, or simply not considering them a peer in *this* case, it is nearly always rational to temporarily revoke peer status.

With respect to typical cases of peer disagreement, then, Nonconformists like Kelly are generally correct: one is under no epistemic obligation to alter one's views in most peer disagreement situations. On the other hand, exceptional cases can always be dreamt up by a clever philosopher. For this reason, Feldman, a Conformist, is also correct to point out that there are no general principles regarding what is rational that hold for all

possible peer disagreement cases. Nonetheless, the Nonconformist approach is likely sufficient for most real-life cases.

However, the peer-revoking policy recommended by Nonconformism is not justified in cases of prolonged disagreement about the same issue, in which one continues to recognize the peer status of the other about that issue. While relatively unusual, such cases do occur, and they are, I argue, the seminal cases in the peer disagreement debate, which has its roots in discussions of religious pluralism within the philosophy of religion. I have argued that these cases present a unique challenge that has been largely overlooked, or paid only scant attention, in the disagreement literature. By creating a situation of Parity, prolonged disagreement with a recognized peer, where both parties are doing their best to share their evidence and to understand the point of view of the other, can lead to a situation in which it is rationally required that one weigh the evidential seemings of a peer the same as one's own (§ 4.2, § 8.3). The Hume and Kant-inspired approaches to peer disagreement are meant to ameliorate this worry, and to explain how we can maintain a Nonconformist stance even in situations of Parity.

I must confess here that I am less than fully confident about my own solutions to the Parity Problem, for the reasons I gave in the last chapter. I have personally experienced at least one disagreement that I would describe as peer disagreement under Parity, and I find that the Humean and Kantian solutions I've offered, while somewhat helpful, do not ultimately remove the discomfort I feel in continuing to believe as I do, when I know that my peer understands the issue as well as I do and disagrees. Here I believe it is worth remembering something that commentators on all sides of the peer disagreement debate agree about: most people *are* too confident about their controversial

beliefs. Insofar, then, as the epistemology of disagreement discussion encourages some epistemic humility, it is a worthwhile endeavor.

BIBLIOGRAPHY

- Adler, Jonathan. "Epistemological Problems of Testimony." In *The Stanford Encyclopedia of Philosophy*, 2017. <https://plato.stanford.edu/archives/win2017/entries/testimony-episprob/>.
- Alston, William P. "Epistemic Circularity." *Philosophy and Phenomenological Research* 47, no. 1 (1986): 1–30.
- . "Level Confusions in Epistemology." *Midwest Studies in Philosophy* 5 (1980): 135–50.
- . *Perceiving God: The Epistemology of Religious Experience*. Ithaca: Cornell University Press, 1991.
- . "Perceiving God." *The Journal of Philosophy* 83, no. 11 (1986): 655–65.
- Audi, Robert. *Epistemology: A Contemporary Introduction to the Theory of Knowledge*. 3rd ed. New York: Routledge, 2011.
- Austin, J.L. "A Plea for Excuses: The Presidential Address." *Proceedings of the Aristotelian Society, New Series* 57 (n.d.): 1–30.
- Bergmann, Michael. "Rational Disagreement After Full Disclosure." *Episteme* 6, no. 3 (2009): 336–53.
- Biro, John, and Fabio Lampert. "'Peer Disagreement' and Evidence of Evidence." *Logos & Episteme* 9, no. 4 (2018): 379–402.
- Bolos, Anthony, and Kyle Scott. "Reformed Epistemology." In *Internet Encyclopedia of Philosophy*, n.d. <https://www.iep.utm.edu/ref-epis/>.
- BonJour, Laurence. *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press, 1985.
- Burge, Tyler. "Content Preservation." *The Philosophical Review* 102 (1993): 457–88.
- Chignell, Andrew. "The Ethics of Belief." In *The Stanford Encyclopedia of Philosophy*, 2018. <https://plato.stanford.edu/archives/spr2018/entries/ethics-belief/%3E>.
- 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.
- . "Higher-Order Evidence." *Philosophy and Phenomenological Research* LXXXI, no. 1 (2010): 185–216.
- Clifford, W.K. "The Ethics of Belief." In *The Ethics of Belief and Other Essays*, edited by Timothy Madigan, 70–96. Amherst, MA: Prometheus, 1999.
- Cohen, G.A. *If You're an Egalitarian, How Come You're So Rich?* Cambridge: Harvard University Press, 2000.
- Comesaña, Juan, and Eyal Tal. "Evidence of Evidence Is Evidence (Trivially)." *Analysis*

- 75, no. 4 (2015): 557–59.
- . “Is Evidence of Evidence Evidence?” *Nous* 51, no. 4 (2017): 95–112.
- Conee, Earl. “Rational Disagreement Defended.” In *Disagreement*, edited by Richard Feldman and Ted A Warfield. Oxford University Press, 2010.
- Decker, Jason. “Disagreement, Evidence, and Agnosticism.” *Synthese* 187, no. 2 (2012): 753–83.
- Eder, Anna-Maria A, and Peter Brössel. “Evidence of Evidence as Higher Order Evidence.” In *Higher-Order Evidence: New Essays*, edited by Mattias Skipper and Asbjørn Steglich-Petersen. Oxford University Press, n.d.
- Elga, Adam. “Reflection and Disagreement.” *Noûs* 41, no. 3 (2007): 478–502.
- Enoch, David. “Why I Am an Objectivist about Ethics (And Why You Are, Too).” In *The Ethical Life: Fundamental Readings in Ethics and Moral Problems*, edited by Russ Shafer-Landau, 4th ed. New York: Oxford University Press, 2018.
- Faulkner, Paul. “A Virtue Theory of Testimony.” *Proceedings of the Aristotelian Society* CXIV, no. 2 (2014): 189–211.
- Feldman, Richard. “Evidence of Evidence Is Evidence.” In *The Ethics of Belief*, edited by Jonathan Matheson and Rico Vitz, 284–300. Oxford University Press, 2014.
- . “Evidentialism, Higher-Order Evidence, and Disagreement.” *Episteme* 6, no. 3 (2009): 294–312.
- . “Reasonable Religious Disagreements.” In *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, edited by Louise Antony, 194–214. Oxford University Press, 2007.
- . “Reasonable Religious Disagreements.” *Philosophers Without Gods, Meditations on Atheism and the Secular Life*, 2007, 194–214.
- Fitelson, Branden. “Evidence of Evidence Is Not (Necessarily) Evidence.” *Analysis* 72, no. 1 (2012): 85–88.
- Frances, Bryan. *Disagreement*. Cambridge: Polity Press, 2014.
- . “Scepticism and Disagreement.” In *Skepticism: From Antiquity to the Present*, edited by Diego Machuca and Baron Reed, 581–91. New York: Bloomsbury Academic, 2018.
- Frances, Bryan, and Jonathan Matheson. “Disagreement.” In *Stanford Encyclopedia of Philosophy*, 2018. <https://plato.stanford.edu/archives/spr2018/entries/disagreement/>.
- Fricker, Elizabeth. “Against Gullibility.” In *Knowing from Words*, edited by B.K. Matilal and A. Chakrabarti, 125–61. Boston: Kluwer, 1994.
- Fricker, Miranda. “Group Testimony? The Making of A Collective Good Informant.” *Philosophy and Phenomenological Research* LXXXIV, no. 2 (2012): 249–76.
- Friedman, Michael. “Regulative and Constitutive.” *The Southern Journal of Philosophy* XXX (1991): 73–102.

- Goldman, Alvin. "Epistemic Relativism and Reasonable Disagreement." In *Disagreement*, edited by Richard Feldman and Ted Warfield, 187–215. Oxford University Press, 2010.
- Gutting, Gary. *Religious Belief and Religious Skepticism*. Notre Dame: University of Notre Dame Press, 1982.
- Hanna, Robert. "Kant's Theory of Judgment." In *Stanford Encyclopedia of Philosophy*, 2018. <https://plato.stanford.edu/archives/win2018/entries/kant-judgment/>.
- Hume, David. *A Dissertation on the Passions & The Natural History of Religion: A Critical Edition*. Edited by Tom Beauchamp. New York: Oxford University Press, 2007.
- . *A Treatise of Human Nature*. Edited by David Fate Norton and Mary J Norton. Oxford: Clarendon Press, 2007.
- . *An Enquiry Concerning Human Understanding*. Edited by Peter Millican. Oxford World's Classics, 2007.
- . *Dialogues Concerning Natural Religion*. Edited by Dorothy Coleman. Cambridge University Press, 2007.
- . *The Letters of David Hume*. Edited by J.Y.T Greig. Vol. 1. Oxford: Clarendon Press, 1932.
- Inwagen, Peter van. "Listening to Clifford's Ghost." *Royal Institute of Philosophy Supplement* 65 (2009): 15–35.
- Inwagen, Peter van. "We're Right. They're Wrong." In *Disagreement*, edited by Richard Feldman and Ted Warfield, 10–28. Oxford University Press, 2010.
- James, William. *The Will to Believe and Other Essays in Popular Philosophy*. New York: Dover Publications, 1956.
- Johnston, Mark. "The Authority of Affect." *Philosophy and Phenomenological Research* 63, no. 1 (2001): 181–214.
- Kant, Immanuel. *Critique of Pure Reason*. Edited by Paul Guyer and Allen Wood. Cambridge University Press, 1998.
- . *Critique of the Power of Judgment*. Edited by Paul Guyer. Cambridge University Press, 2000.
- Kelly, Thomas. "Peer Disagreement and Higher-Order Evidence." In *Disagreement*, edited by Richard Feldman and Ted Warfield, 111–74. Oxford University Press, 2010.
- . "The Epistemic Significance of Disagreement." In *Oxford Studies in Epistemology, Volume 1*, edited by John Hawthorne and Tamar Gendler, 167–96. New York: Oxford University Press, 2005.
- Kripke, Saul. *Naming and Necessity*. Cambridge: Harvard University Press, 1972.
- Kulenkampff, Jens. "The Objectivity of Taste : Hume and Kant." *Nous* 24, no. 1 (1990): 93–110.

- Lackey, Jennifer. "A Justificationist View of Disagreement's Epistemic Significance." In *Social Epistemology*, edited by Adrian Haddock, Alan Millar, and Duncan Pritchard, 298–325. Oxford University Press, 2010.
- Lampert, Fabio, and John Biro. "What Is Evidence of Evidence Evidence Of?" *Logos & Episteme* 8, no. 2 (2017): 195–206.
- Machuca, Diego, ed. *Disagreement and Skepticism*. New York: Routledge, 2013.
- Moffett, Marc. "Reasonable Disagreement and Rational Group Inquiry." *Episteme: A Journal of Social Epistemology* 4, no. 3 (2007): 352–67.
- Morris, William Edward, and Charlotte R. Brown. "David Hume." In *The Stanford Encyclopedia of Philosophy*, 2017.
<https://plato.stanford.edu/archives/spr2017/entries/hume/>.
- Moser, Paul K. *Knowledge and Evidence*. Cambridge University Press, 1989.
- Nozick, Robert. "Newcomb's Problem and Two Principles of Choice." In *Essays in Honor of Carl G. Hempel*, edited by Nicholas Rescher. Dordrecht: Reidel Publishing Company, 1969.
- . *Philosophical Explanations*. Cambridge, MA: Harvard University Press, 1981.
- Pittard, John Steven. "Disagreement, Rationality, and Religious Belief," 2013.
- Plantinga, Alvin. "On 'Proper Basicity.'" *Philosophy and Phenomenological Research* 75, no. 3 (2007): 612–21.
- . *Warranted Christian Belief*. New York: Oxford University Press, 2000.
- Pollock, John. *Contemporary Theories of Knowledge*. 1st ed. Towota, NJ: Rowman and Littlefield Publishers, 1986.
- Pust, Joel. "Intuition." In *The Stanford Encyclopedia of Philosophy*, 2017.
<https://plato.stanford.edu/archives/sum2017/entries/intuition/>.
- Reisner, Andrew. "Peer Disagreement, Rational Requirements, and Evidence of Evidence as Evidence Against." In *Epistemic Reasons, Norms and Goals*, edited by Pedro Schmechtig and Martin Grajner, 95–114. De Gruyter, 2016.
- Roche, William. "Evidence of Evidence Is Evidence Under Screening-Off." *Episteme* 11, no. 1 (2014): 119–24.
- . "Is Evidence of Evidence Evidence? Screening-Off Vs. No-Defeaters." *Episteme* 15, no. 4 (2018): 451–62.
- Rosen, Gideon. "What Is Constructive Empiricism?" *Philosophical Studies* 74, no. 2 (1994): 143–78.
- Rosenthal-Schneider, Ilse. *Reality and Scientific Truth: Discussions With Einstein, Von Laue, and Planck*. Edited by Thomas Braun. Detroit: Wayne State University Press, 1980.
- Rowling, J.K. *Harry Potter and the Order of the Phoenix*. Scholastic, 2004.
- Sosa, Ernest, and Barry Stroud. "Philosophical Scepticism." *Proceedings of the*

- Aristotelian Society, Supplementary Volumes* 68 (1994): 263–307.
- Stern, David. “Are Disagreements about Taste Possible? A Discussion of Kant’s Antinomy of Taste.” *The Iowa Review* 21, no. 2 (1994): 66–71.
- Wedgwood, Ralph. “Disagreement and the A Priori.” In *The Nature of Normativity*, 248–66. Oxford University Press, 2007.
- . *The Nature of Normativity*. Oxford University Press, 2007.
- Weiss, Bernard G. “Taqlid.” In *The Oxford Encyclopedia of the Islamic World*. Oxford University Press, 2009.
<http://www.oxfordreference.com/view/10.1093/acref/9780195305135.001.0001/acf-9780195305135-e-0785>.
- Yablo, Stephen. “Is Conceivability a Guide to Possibility?” *Philosophy and Phenomenological Research* 53, no. 1 (1993): 1–42.