

**RATIONAL UNIQUENESS AND RELIGIOUS
DISAGREEMENT**

By

CHRISTOPHER WILLARD-KYLE

For completion of the

MPhil in Philosophical Theology

University of Oxford

June 18, 2014

TABLE OF CONTENTS

ACKNOWLEDGMENTS.....	2
INTRODUCTION TO RELIGIOUS DISAGREEMENT: A PERSONAL QUESTION.....	3
1. A FRAMEWORK FOR RATIONAL DISAGREEMENT	5
1.1 DEFINING THE QUESTION.....	5
1.2 DEFINING THE ANSWERS.....	8
1.3 A MODAL FRAMEWORK.....	10
1.4 RATIONAL PERMISSIVISMS.....	15
1.5 CONCLUSION	17
2. THE PROBLEM OF THE BEST	18
2.1 BESTISM.....	18
2.2 MAXIMIZING OVER MULTIPLE VARIABLES.....	20
2.3 TWO ARGUMENTS: A BRIEF INTRODUCTION.....	23
2.4 THE JAMESIAN ARGUMENT: MAXIMIZING TRUTH AND MINIMIZING ERROR.....	24
2.5 THE ARGUMENT FROM ABDUCTION.....	30
2.6 IS EXTREME PERMISSIVISM TOO PERMISSIVE?.....	34
2.7 CONCLUSION	38
3. RESISTING UNIQUENESS	40
3.1 QUICK AND DIRTY PERMISSIVISM.....	40
3.2 FELDMAN’S PRINCIPLE AND ESCALATING SKEPTICISM.....	41
3.3 CHANCY PILL-POPPING AND EPISTEMIC LUCK.....	44
3.4 ANTI-LUMINOUS RATIONALITY	51
3.5 NEITHER CONFUSION NOR ISOLATION.....	55
3.6 CONCLUSION	57
4. FIRST-ORDER STEADFASTNESS.....	58
4.1 CONCILIATIONISM AND STEADFASTNESS	58
4.2 DOES UNIQUENESS MATTER?.....	59
4.3 KELLY ON HIGHER-ORDER EVIDENCE: THE PERILS OF DOUBLE COUNTING	61
4.4 EVIDENCE OF EVIDENCE AND FUN WITH MARBLES.....	65
4.5 HIGHER-ORDER CREDENCES: UNDER PRESSURE.....	69
4.6 EASY AND HARD QUESTIONS: WHY DIFFICULTY MATTERS.....	76
4.7 SOME CONCERNS ABOUT ASSESSING DIFFICULTY	79
4.8 CONCLUSION	82
EPILOGUE: HUMBLE CONVICTION.....	85

Acknowledgments

One of the motivating intuitions guiding research on disagreement is the belief that how and what other people think matters for how and what I should think. Though in the following pages I shall often talk about disagreement as producing a skeptical problem to be overcome, I've frequently benefited from honest disagreement and criticism from friends and colleagues throughout this project. It seems only right that in recognition of this debt, I should thank those who, whether in agreement or disagreement, have especially impacted this project.

I first owe thanks to Brian Leftow, who, in addition to providing thoughtful feedback, continues to be a fine exemplar of intellectual humility and interpretive charity. Max Baker-Hytch and Joshua Layton-Wood never tired of requests for yet another conversation about my research; or if they did, they generously persevered for friendship's sake. Most importantly, my fiancée, Danielle Willard, whose patient love endured many ramblings on this topic; moreover, no section of the paper is untouched by her eye for detail and clarity.

It would take too long to recognize every contribution properly, but I would be remiss not to name the following: Charity Anderson, Matthew Benton, Paul Billingham, Laura Callahan, Geoff Dargan, Jonathan Erhardt, Hannah Kim, Matt Kirkpatrick, Joanna Lawson, Tim Mawson, Kian Mintz-Woo, Mark Nelson, Josh Roe, Wes Skolits, Gregory Stacey, Richard Swinburne, and Daniel Young

I am thankful for their thoughtful contributions. This project is better for their involvement, and I am better for their friendship.

NB: An updated version of §1-3 is now published as "Do Great Minds Really Think Alike?" *Synthese*, 194(3): 989-1026.

Introduction to Religious Disagreement: A Personal Question

Should I believe in string theory? I *hope* that I am a generally well-educated and intelligent person, but I know that I am a certified ignoramus when it comes to physics. Despite this local incompetence, I accept a number of complicated and abstract theses about physics. I believe that all physical objects attract each other at a rate proportionate to the product of their masses and inversely proportionate to the square of the distance between them. I believe that there are such things as protons and electrons, even though I have never seen any. But I do not—and think I *should* not—believe in string theory.

Why is this? What is the relevant difference between the string theory case and the gravity case? One plausible answer is that all (or most) of the experts agree about the equation for gravity, but the experts disagree about string theory. Some physicists think it is true; others do not. Not only *don't* I believe in string theory, I *should* not. To do so would be rationally deficient. Where the experts of physics disagree about physics, I should be agnostic.

What if, instead of the ignoramus that I am, I were an expert in physics, and if after reflecting on the evidence, string theory seemed true to me? Should this change my answer? Not obviously. It would still be true that the experts disagreed, even if I were one of them. Other expert physicists who had access to the exact same data, whom I would consider to be just as intelligent and scientific as my fictitiously savant-like self, would still interpret the evidence in the opposite way. Why believe that *I* have got the matter right, and *they* have got the matter wrong? If they are just as rational as I am and looking at the same evidence, isn't the rational thing to do to suspend judgment? Wouldn't it exhibit an appropriate amount of intellectual humility to show restraint concerning such a difficult and contentious issue?

Before we give our assent, let us consider another case. We live in a religiously plural society. People disagree over a variety of religious issues, even though they seem to have roughly the same evidence available to them. I am a theist, but I have had conversations with atheists who are more intelligent than I am, who know all of the arguments for theism that I do, and yet remain atheists. They know all about fine-tuning and the cosmological argument; they have read Aquinas, Anselm, and Kierkegaard. They may have even had at some point in their lives a kind of religious experience. But they are atheists and I am a theist. Is not the mere fact that so many rational people disagree with me itself evidence that I should give up my belief in theism and become agnostic? Wouldn't this be the intellectually humble and rationally appropriate course of action? What right have I—rationally speaking—to believe that *I* am right and *they* are wrong?

Recently, a number of epistemologists have argued that in cases of entrenched disagreement between epistemic equals, the right course of action is to suspend belief. For example, in an introspective and thought-provoking article, Richard Feldman argues that it would be rational for him to give up his own atheism and embrace agnosticism in the light of religious disagreement and encourages others to do likewise. This comes as bad news for committed theists and atheists. The epistemic effects of disagreement should be felt across philosophy, politics, ethics, and even science, but perhaps nowhere more poignantly than in religion. More than merely academic, such disagreements are deeply personal.

And so we come to our question: Is continued, rational disagreement about religious beliefs possible given the widespread religious disagreement we inevitably encounter? What should *we* do when we confront epistemic equals who disagree with us about our most deeply held philosophical and religious beliefs?

1. A Framework for Rational Disagreement

I examine the rational uniqueness thesis and place it within the context of the disagreement literature. I show that, when combined with evidentialism, rational uniqueness is a very strong thesis and, as such, bears a substantial burden of proof.

1.1 Defining the Question

Sometimes we use the term, “rational,” to describe a person. This use of “rational” in this sense indicates intellectual competence and mental ability, as in the sentence “Sherlock Holmes is a remarkably rational investigator.” Other times, we use the term adverbially to describe epistemic behavior, as in “Sherlock Holmes rationally inferred that Moriarty was the killer,” or “Holmes rationally believes that Lestrade arrested the wrong suspect.” Both rational and irrational agents can sometimes behave rationally, just as both virtuous and vicious people can sometimes act ethically. Having formed a belief rationally does not entail that one's belief is true, only that one has formed the belief in an epistemically responsible way according to the rational norms. I will not here give an analysis of what the norms of rationality are, though I will make some substantive claims about what rationality is throughout the paper. Most people know fairly well what is rational even if they do not have a theory of rationality, just as most people know fairly well what is ethical even if they do not have a theory of ethics.

Additionally, we might use “rational” to describe beliefs or attitudes. Usually, we say a belief is rational when it is produced by someone who forms that belief *rationally*. At other times, we talk about rational beliefs without having any particular agent in mind, but rather with respect to whether a belief is consistent (logically or probabilistically) with a set of evidence.¹ I suspect that even in these cases, we mean to say that an epistemic agent who had that evidence could rationally form the belief in question. Doubtless, the

¹ This corresponds to propositional rather than doxastic justification.

linguistically observant will notice ordinary uses of “rational,” that are not categorized so neatly; but at any rate, this is what I shall mean when I talk about rational beliefs or attitudes. A rational belief is one that is rationally formed by the agent (when one is specified), or one that *could be* rationally formed by an agent (when one is not).

Is rational disagreement possible? Our answer depends crucially on what sort of rationality we are talking about. Can rational people disagree? The answer seems obvious. Rational people disagree all the time about a variety of issues. Even very rational people make mistakes. The fact that one makes an occasional mistake in reasoning does not mean that person is irrational all things considered.

But we might want to know not just whether rational people can disagree, but whether people can disagree without having made any mistakes in reasoning.² We might wonder, that is, whether two people can disagree rationally. The answer still seems to be “yes.” Sometimes, two people rationally draw different conclusions because they are presented with different evidence. I look at a clock in your house that reads 11:00. I rationally form the belief that it is within a few minutes of 11:00. You, however, though you have all the evidence I have (seeing the clock, knowing that clocks generally tell time within a few minutes of accuracy, etc.) also know that this *particular* clock runs an hour ahead of time, and so rationally form the correct belief that it is within a few minutes of 10:00. I have got the matter wrong, but the error doesn’t seem to have anything to do with my rationality. Our disagreement is explained not by a difference in our rationality, but a difference in our evidence.

² It is crucial that *no* mistakes in reasoning have occurred. Rationality is a graded concept. A person can be more or less rational than another. I suppose we might also say that one event of reasoning can be more or less rational than another. With this in mind, we may think it is fairly obvious that two people (even with the same evidence) can disagree rationally, at least if the problem they are considering is lengthy or complicated enough. Each of them could have made very few mistakes, but still have performed at a high level of rationality.

The more interesting question, however, is whether two people who have made no mistakes in reasoning—that is who have been *perfect* with respect to rationality—can disagree on the basis of the same evidence. In this paper, “rationality” always means “perfect and complete rationality” unless otherwise specified.

But we might want to know not just whether people can disagree rationally, but whether people *with the same evidence* can disagree rationally. The answer to this question is much less obvious, and has generated much disagreement in its own right over the last twenty years. It is common to frame this question by using the concept of an epistemic peer. An epistemic peer is usually defined along the following lines:

Epistemic Peerhood: S is an epistemic peer of T regarding p iff S and T have the same evidence regarding p and S is as competent as T regarding p .

Evidence regarding p includes any and all types of information relevant to determining the truth-value of p . S 's competence includes her rationality, that is, her ability to reason and make inferences from her evidence; however, it also includes non-inferential powers of discrimination, such as the vision of the sighted or the mysterious talents of the chicken-sexer.

One might want to add to the traditional definition that both agents are equally devoted to the investigation of p . If S and T have the same evidence and competency, but T is genuinely seeking the truth-value of p whereas S is only half-heartedly committed to the enquiry, then there is good reason to think that S enjoys an epistemic advantage over T . Let us propose the following amendment:

Epistemic Peerhood*: S is an epistemic peer of T regarding p iff (i) S and T have the same evidence regarding p , (ii) S is as competent as T regarding p , and (iii) S and T are just as earnestly seeking the truth about p .

An epistemic peer, like the physicist's frictionless plane, is an idealization. Rarely if ever do we have perfect epistemic peers. Usually, our fields of evidence do not overlap completely, and degrees of competence are multi-faceted and vary widely. However, we are often confident that our real peers would continue to disagree with us even if they were provided with all the extra bits of relevant information we have, or that we would continue to disagree even if our competence were slightly boosted to match theirs.

The literature on disagreement is really dominated by two questions that are not often fully distinguished. We can call them (1) the problem of unrevealed disagreements, and (2) the problem of revealed disagreements:

- 1) Can two people with the same evidence disagree rationally?
- 2) How, if at all, should we adjust our beliefs when we discover that someone disagrees with us?

The first question can be asked without reference to epistemic peers. All that matters is that both parties have the same evidence. Let us call parties who have the same evidence “evidential peers.” All epistemic peers are evidential peers, but not all evidential peers are epistemic peers. Whether or not the beliefs in question have been acquired rationally from the evidence is distinct from whether or not those beliefs are held by rational agents. Even fools believe rationally sometimes.

The concept of an epistemic peer is essential for the second question, however, for it is not wise to change one’s mind whenever some fool disagrees with you. Whether or not I adjust my belief in the face of disagreement should depend on whether I acknowledge the person with whom I am disagreeing to be a competent master of the evidence. If a child disagrees with me on the answer to a complicated math equation, I will not be tempted to adjust my belief. Of course, if he gives me a *reason* why my answer is wrong, I should evaluate this reason with full seriousness, but the mere *fact* of his disagreement should not incline me to adjust my belief, at least not substantively. On the other hand, if a professor of mathematics disagrees with me, I should (at least in most cases) significantly adjust my belief even if she can’t be bothered to explain the reason she thinks my answer is wrong.

Though we will eventually consider answers to both questions, it is sensible to begin with question (1): can two people with the same evidence disagree rationally? That is, can they both, starting with the same evidence, arrive at different conclusions without making any errors of rationality?

1.2 Defining the Answers

There are two prominent answers to this question: “No, never!” and “Yes, sometimes!”

Those who answer, “No, never!” accept the rational uniqueness thesis.

Rational Uniqueness [plus Evidentialism] (RUE): For all complete sets of evidence E , for all epistemic agents S , and for all propositions p , if S has E , then there is just one doxastic attitude A such that it is rationally permissible for S to bear A toward p .

What is a doxastic attitude? A doxastic attitude (DA) is an evaluative attitude toward a proposition. There is disagreement about what sorts of things doxastic attitudes are fundamentally. The most commonly discussed doxastic attitudes are beliefs. One can believe, disbelieve, or refrain from belief in a proposition. Others insist that our doxastic attitudes are more fine-grained and prefer to speak of credences. An agent’s credence in a proposition can be expressed as a number (or range of numbers) between 0 and 1, and is intimately related to how confident the agent is that the proposition is true.³ Still others wish to consider as separate and irreducible such concepts as knowing, believing, opining, being certain that, being skeptical that, and the like. At most points in my argument, it will not matter how doxastic attitudes are parsed; a doxastic attitude is just whatever it is that is fundamental to our evaluations of propositions.

Those who answer, “Yes, sometimes!” may be called rational permissivists. Permissivists believe that sometimes several doxastic attitudes can be permissible relative to the same evidence. It would be tempting to define rational permissivism as the negation of (RUE); however, as Matthew Lee has shown, permissivism is only a contrary of uniqueness and not a contradictory. For if it is ever the case possible that *no* doxastic attitude is permitted, then neither (RUE) nor permissivism is true.⁴ Nevertheless, it is dubious whether the evidence could ever justify no doxastic attitudes. Plausibly, the DA of

³ There is disagreement among Bayesians about what exactly credences are. Some talk as though they are levels of confidence. Others, however, talk about them more as probability ascriptions (i.e. If I believe there is a 75% likelihood that p is true, my credence in p is .75). I am more inclined to think of them as levels of confidence; however, what I say shall apply equally to either conception.

⁴ Lee, Matthew, “Conciliationism without Uniqueness,” *Grazer Philosophische Studien* (2013), vol. 88.

agnosticism, or refraining from belief, is always permissible when no other DA is. At any rate, all of the arguments I give against permissivism could also be given as arguments against the thesis that *at most* one DA is rationally permissible for an agent given her evidence. Thus we may, for all practical purposes, argue that (RP) is true iff (RUE) is false.

What are the prominent responses to question (2): How, if at all, should we adjust our beliefs when we discover that someone (especially an epistemic peer) disagrees with us? Most answers fall into one of the following camps: conciliationism or steadfastness. Various types of conciliationism exist, but the central claim is that when faced with known peer disagreement, one is rationally obligated to revise her belief significantly. What does it mean to revise one's belief *significantly*? Conciliationists themselves are divided on this. But for now, we can think of conciliating over p as becoming agnostic about the truth of p .

Nearly everyone agrees that you should conciliate in some cases. However, those who believe in steadfastness claim that conciliation is not always rationally required when you face peer disagreement. Sometimes you can continue to believe as you did before you became aware that a peer disagreed with you.

The question concerning conciliationism and steadfastness will be easier to evaluate once we have already formulated an answer about uniqueness. Let us return, therefore, to question (1).

1.3 A Modal Framework

Modally speaking, propositions fall into one of three exhaustive categories. Some propositions are necessarily false. Others are possibly true. Of the propositions that are possibly true, some are necessary, and others are contingent. We define these categories in terms of a single modal operator where p is a proposition:

Necessary: $\Box p$

Impossible: $\Box \sim p$

Contingent: $\sim \Box p \wedge \sim \Box \sim p$

These categories are exhaustive and exclusive. All propositions fall into one and only one of these three categories. The contingent and necessary together are the possible propositions:

Possible: $\sim\Box\sim p$

Deontic logic contains an analogous tripartite division. All actions are permissible or impermissible. The permissible actions can be further divided into the obligatory (OB) and the optional where x is an action:

Obligatory: OBx

Impermissible: $OB\sim x$

Optional: $\sim OBx \wedge \sim OB\sim x$

These categories are also exhaustive and exclusive. All actions fall into one and only one of these three categories. The optional and obligatory together are the permissible actions:

Permissible: $\sim OB\sim x$

It is a notable feature of epistemology that we often apply normative language to doxastic attitudes in a similar way as we do to actions in ethics. We speak of *good* or *bad* inferences. We say things like, “You had no *right* to assume that!” when blaming someone for having a belief, or “You *should have* known!” when blaming someone for not having a belief. As epistemologists, we care greatly whether we are *justified* in our beliefs. When we believe something true, we are *right*, and when we believe something false, we are *wrong*. Let the sentential operator “ \Box_R ” mean “rationally obligatory.” Taking deontic and modal logic as a guide, the following schema is plausible where Bp is belief that p is true:⁵

Rationally Obligatory: $\Box_R Bp$

Rationally Impermissible: $\Box_R \sim Bp$

⁵ Notably, this model does not depend on doxastic voluntarism. We can describe beliefs (or other DA’s) as obligatory, permissible, etc., even if they are actions we cannot voluntarily perform.

Rationally Optional: $\sim\Box_R Bp \wedge \sim\Box_R \sim Bp$

These categories are also exhaustive and exclusive. All beliefs fall into one and only one of these three categories. The rationally optional and rationally obligatory together are the rationally permissible beliefs:

Rationally Permissible: $\sim\Box_R \sim Bp$

(RUE) is a very strong thesis. According to (RUE), there is only one rationally permissible DA given a complete evidence set. But if there is only one rationally permissible DA, then that DA must be rationally obligatory. If only one action is morally permissible, it is not optional but obligatory; all the other actions are impermissible. If only one world is possible, it is not contingent but necessary; all the other worlds are impossible. Similarly, if only one DA is permissible, it must be obligatory; all other DA's are impermissible. For every set of evidence, (RUE) entails that the set of rationally permissible beliefs (or DA's more generally) and the set of rationally obligatory beliefs are coextensive.

But this is a hard teaching. Who can accept it? Few believe that all morally permissible actions are obligatory unless they have already been persuaded of a very strict sort of morality. Even fewer believe that all possible worlds are necessary unless they have been persuaded by strong arguments toward a strict necessitarianism. This suggests that we need arguments of similar strength to convince us that all rationally permissible doxastic attitudes are also rationally obligatory. If the modal and moral cases are even roughly analogous, (RUE) shoulders a significant burden of proof.

Let us take a step back. Nothing that I have said entails that there are, in fact, rational disagreements, or even that such disagreement is possible. That we can identify three different categories—the rationally obligatory, the rationally optional, and the rationally impermissible—does not entail that any DA's satisfy the concept of being rationally optional. What we have done is show that there is conceptual space for rationally optional

belief. But perhaps what is required to establish a burden of proof is not only conceptual space, but also anti-collapse intuitions. Anti-collapse intuitions are clearly present in the modal and moral cases. We are naturally inclined to think that there are contingent facts. There are ways things are that do not exhaust the way things might have been. Similarly, we are naturally inclined to think that there are morally optional actions. There are permissible ways we can act that do not exhaust the morally permissible actions available to us. In the same way, many report they are inclined to think that there are rationally optional beliefs. In an oft-cited passage, Gideon Rosen surmises, “It should be obvious that reasonable people can disagree, even when confronted with the same body of evidence. When a jury or a court is divided in a difficult case, the mere fact of disagreement does not mean that someone is being unreasonable.”⁶

Nevertheless, even those who share these intuitions might be persuaded to abandon their pre-theoretic views in the light of good arguments to the contrary. Perhaps one thinks, for instance, that there is a God who must of necessity create the best world. Such a person may well end up believing that there is ultimately only one possible world, that the modally possible and modally necessary overlap completely. Again, though we are inclined to think that there are morally optional actions, we might give up this intuition if confronted with arguments against it. One ethical theory that seems to have this consequence is strict act utilitarianism. On strict act utilitarianism, an action is right (permissible) iff it maximizes utility. Since *only* the action that maximizes utility is permissible, that action is obligatory. There may be a few cases of merely permissible actions when multiple actions tie for maximizing utility, but these will be rare.

The defender of (RUE) will insist that it would be wrong to finish our argument here. The burden of proof is not so heavy that no good argument could overturn it. In this,

⁶ Rosen, Gideon (2001). “Nominalism, Naturalism, Epistemic Relativism,” *Philosophical Perspectives* 15, pp. 71.

she would be right. Our argument for permissivism should not (and does not) end here. But perhaps my objector will go further and complain that I have not given much in the way of *any* argument for the possibility of rational disagreement. I have only presented a certain picture that is consistent with rational disagreement. Simply laying out a framework that is consistent with rational disagreement is, by itself, hardly an impressive feat. For whether or not there can be rational disagreement depends not only on a framework for evaluating rationality, but a method for filling that framework in.

In partial answer to this latter accusation, I will briefly indulge in biography. Growing up, I used to believe in a sort of maximizing morality. If an action was not best—the most wise, most prudent, or most loving—then it was not right. There was something morally inappropriate about doing second-best. After all, what moral justification could I give for doing second-best when I could just as well have done the best? These were torturously guilt-inducing thoughts. Here is not the place to identify all that was right and wrong about this attitude. What is relevant is what changed my mind. Importantly, it was not strictly an argument. Rather, it was a certain picture of morality. Someone distinguished the concept of the morally forbidden from the morally permitted and morally obligatory. Just by seeing these categories—or at least so it seemed to me—I also saw that I had been confusing a certain class of permissible actions (especially supererogatory actions) with obligatory ones.⁷ I later learned of arguments that I believe support this conclusion, but these arguments only played the role of confirming an insight that this picture had already bestowed upon me.

If the model I have presented does not constitute an argument, it may nonetheless communicate an important insight about the possibility of rational disagreement. Certainly, it is not a conclusive consideration, just as the modal and moral analogues are not

⁷ Mark Nelson has helpfully identified this sort of persuasion by presentation as “descriptive manifestation,” in an excellent meta-philosophical piece, “Telling It like It Is: Philosophy as Descriptive Manifestation,” *American Philosophical Quarterly*, Vol. 42(3), 2005, pp. 141-156.

conclusive. But even in the modal and moral cases, these considerations are far from impotent. It is a surprising and unwelcome feature of Leibniz's belief that God must create the best that it seems to deny any contingency. It is an equally surprising and unwelcome feature of strict act utilitarianism that there are (almost) no morally optional actions. At the very least, these considerations place a burden on those who would deny the possibility of rational disagreement to show that the rationally permissible and rationally obligatory are exactly coextensive. And if the cases of moral and modal reasoning are at all analogous, this is no small task.

1.4 Rational Permissivisms

(RUE) entails that, relative to some evidence, there are no rationally optional beliefs; there is no set of evidence such that for any proposition p , $\sim\Box_R Bp \wedge \sim\Box_R \sim Bp$. However, many permissivists are not content with simply claiming that some beliefs are rationally optional. Many permissivists also want to say that it is sometimes permissible for one person to believe that p while another person with the same evidence permissibly believes that not- p . There are some sets of evidence such that there are some p such that $(\sim\Box_R Bp \wedge \sim\Box_R \sim Bp) \wedge (\sim\Box_R B\sim p \wedge \sim\Box_R \sim B\sim p)$.

Let us distinguish between what we may term "weak" and "strong" disagreement. Weak disagreement occurs when one person believes a proposition and another person withholds belief from that same proposition. Strong disagreement occurs when one person believes a proposition and another person believes the negation of that proposition. The theist and the atheist strongly disagree, whereas agnostics and theists (or atheists) disagree weakly. Whether a disagreement is weak or strong in this sense has nothing to do with the intensity of the disagreement.

Among weak disagreements, we must be aware of another distinction. There are disagreements of commitment and disagreements of indifference. A committed

disagreement about p is one in which both parties have considered whether or not p . In a disagreement of indifference, one party has not reflected about whether or not p . If neither party has considered whether or not p , then presumably no disagreement exists. Consider again the agnostic. There are really (at least) two types of agnostics. First, there is the committed agnostic who, reflecting on the evidence, neither comes to believe that God exists nor that God does not. In contrast, the indifferent agnostic simply hasn't fully considered the question of God's existence.⁸ Perhaps she has grown up in a context where the subject is rarely if ever brought up, or perhaps she simply doesn't care. Whether a disagreement is committed or indifferent in this sense has nothing to do with how earnestly one is living out her convictions. Accordingly, there cannot be an indifferent theist or an indifferent atheist in this sense. If someone believes p or *not- p* , she must have considered p , even if briefly.

These definitions have been presented in terms of beliefs. If someone prefers to think of DA's as credences, however, we may make similar distinctions. Weak and strong disagreements would be defined by the relative disparity between credences. If one person assigns a credence of .98 and another of .97, they have a relatively weak disagreement. If the second credence were .41, however, it would constitute a relatively strong disagreement. Just how disparate credences must be to be considered a "strong" disagreement may be contextually determined. Or if there is no firm boundary between "strong" and "weak" disagreement, we can still speak of disagreements as "stronger" or "weaker." A disagreement of indifference occurs when one party fails to assign a credence concerning p .

As defined, (RP) says nothing about what kind of rational disagreement is possible. But most interesting cases of disagreement involve not only weak but also strong disagreement, not only disagreements of indifference but also disagreements of

⁸ Mark Nelson helped me fully see this distinction.

commitment. Religious, political, and moral disagreements are often both strong and committed. Some believe that physician-assisted suicide is permissible while others, who seem to have access to the same evidence, believe it is forbidden; some believe one candidate should be elected while others believe her opponent should be elected instead; some believe that God is triune while others believe that no God exists at all.

Earlier we defined (RP) as the claim that (RUE) is false.⁹ Let us now distinguish between moderate and extreme rational permissivism. Moderate permissivists simply believe that (RUE) is false. I will argue for what I shall call extreme rational permissivism. Extreme (RP) holds not only that (RUE) is false, but also that there are, or at least can be, some cases of rational, strong, and committed disagreements among evidential peers.

1.5 Conclusion

We now have the conceptual tools and framework we need to approach the debate over rational uniqueness. In what follows, I argue against (RUE) and for extreme (RP). We have already seen that the defender of (RUE) bears a burden of proof, for she must maintain that the rationally obligatory and permissible are coextensive. It is noteworthy that both the modal and moral analogues of (RUE) rely on some maximizing principle. Leibniz's God chooses the *best* world, and strict utilitarians believe that the only permissible action is the one that *maximizes* utility. Our first task, then, is to examine the possibility of an epistemic principle that operates on a similar maximizing principle. What we find will suggest not only that (RUE) is false, but also give us reason to suppose that there are strong, rational, and committed disagreements among evidential peers.

⁹ More precisely, that (RUE) is false and not false solely because zero DA's are sometimes permissible.

2. The Problem of the Best

I argue that (RUE) is false and extreme (RP) is true. I show that (RUE) is committed to a principle of epistemic bestism and then argue that either bestism is either false or bestism is inconsistent with (RUE). I advance two arguments, the Jamesian argument, and the argument from abduction, before considering and responding to objections.

2.1 Bestism

What doxastic attitude ought we to have toward a proposition? One plausible answer might be, “the *best* doxastic attitude.” We can call this “epistemic bestism.”

Epistemic Bestism: For all S and for all p , S ought to have the best doxastic attitude toward p .

The “ought” should be understood in the epistemic rather than moral sense.¹⁰ Epistemic bestism might not seem very informative. If someone were to quote this principle to me in response to the question, “what ought I believe?” I should probably be annoyed: “If I *knew* what the best doxastic attitude was, I would not have asked!” Unfortunately, however, obnoxious advice is no less true for being obnoxious. Even if this principle does not count as informative advice, it can still be good encouragement when we *do* know what the best doxastic attitude is. If I am tempted to be cowardly, I benefit from the admonishment, “Be brave!” even if I know that is what I should do. When I am tempted as an epistemic agent to succumb to sentimentalism, I could benefit from the admonishment, “Remember, aim for the best doxastic attitude!”

Someone might suggest that the best doxastic attitude toward any proposition just is to believe it if it is true and disbelieve it if it is false. Indeed, there might be an epistemic duty to believe truths and disbelieve falsehoods. However, this is not the sort of duty I have in mind, nor can it be the sort of duty people generally have in mind when considering

¹⁰ Clifford famously argues that such rules are not only epistemically but also morally normative. For the purposes of this essay, we will focus on the epistemic requirements. If an ethical interpretation of the “ought” in epistemic is to be plausible, there must be some bridge principle or argument such as Clifford attempts to give in his *Ethics of Belief*.

rationality, as it is possible that a belief be both rational and false. When we are confronted with misleading evidence, we are also confronted with conflicting duties. On the one hand, we have a duty to believe the truth (and, in addition, to disbelieve falsehoods). On the other hand, we have a duty to respond to the evidence rationally.

This is a puzzle, but not a puzzle unique to rationality. Hawthorne and Srinivasan give the helpful comparison to a soldier who is given a grenade with the pin already pulled along with misleading evidence that there are people to the right, when in fact there are only people to the left.¹¹ On the one hand, the soldier has a moral duty not to throw the grenade toward people. The soldier should throw the grenade to the right, where there are no people. On the other hand, the soldier has a moral duty not to throw the grenade where she believes people are. The soldier should throw the grenade to the left, where she believes no one is present. Both duties are real, and neither is reducible to the other. The same holds for the analogous epistemic duties. It would be an interesting further project to analyze how these duties relate in the case of disagreement, and indeed, this is part of Hawthorne and Srinivasan's project. But for the purposes of this paper, unless otherwise specified, all the epistemic duties in this section concern the duty we have to respond rationally to the evidence we have, not our epistemic duty to believe the truth whatever the evidence may be.

Finally, and most importantly for our purposes, if (RUE) is true, epistemic bestism certainly is. (RUE) does not entail epistemic bestism, although epistemic bestism entails (RUE), at least in cases where there are no ties for "best." But (RUE) without epistemic bestism has severely unattractive consequences. Imagine that a close contender for the best DA, the second or third best, was epistemically permissible. But if (RUE) is true and the

¹¹ Hawthorne, John and Amia Srinivasan (2013). "Disagreement Without Transparency: Some Bleak Thoughts," in David Christensen and Jenifer Lackey (eds.) *The Epistemology of Disagreement: New Essays*, Oxford University Press, pp. 20.

second-best DA is ever permissible, then sometimes having the second-best DA is rationally obligatory. (Recall that if (RUE) is true, the rationally obligatory and the rationally permissible are completely coextensive.) And, if this were not bad enough, it would imply that having the best DA is sometimes rationally forbidden! But this is surely absurd. Therefore, if (RUE) is true, so is epistemic bestism.

Despite the position's initial plausibility, I will argue that epistemic bestism is ultimately untenable. Surprisingly, there is often no uniquely best DA toward a given proposition that is uniquely fixed by the evidence. If there is no uniquely best DA toward a proposition, it is impossible to have the best DA toward that proposition. And if it is impossible, it is not rationally obligatory. The impossibility stems from the problem of maximizing over multiple variables.

2.2 Maximizing over Multiple Variables

Who is the biggest person in the world?

This is a tricky question. "Being big" may refer to mass, height, weight, width, volume, or any other size-related category. Which particular "bigness" is being assessed might usually be clear in some contexts, but the word itself invokes no distinction. Suppose Achilles is taller than Hercules, but Hercules is heavier. Who is bigger? It isn't clear what the right answer is, or even that there is a right answer. The difficulty is that we do not know how much relative value to give height vs. weight. Neither height nor weight can be reduced to the other; neither can they both be reduced to some third quality. Height and weight are *incommensurable*. We cannot answer questions about superlative or comparative ascriptions about "bigness" without contextual constraints because one cannot maximize over multiple, incommensurable variables without an equivalence relation.

Is an equivalence relation anywhere to be found? Perhaps we could stipulate one. Supposing we had stipulated some equivalence relation, we could say whether Achilles or Hercules was bigger. For instance, if we posited that 1 inch equaled 5 pounds, we could

convert Achilles' and Hercules' height and weight into a single unit. The difficulty is that any stipulated equivalence relation seems arbitrary. Why shouldn't 1 inch be worth 6 pounds, or 7 pounds, or 6/7 pounds? A dilemma emerges: (1) if we are to make sense of the question, we must stipulate an equivalence relation, but (2) if we stipulate an equivalence relation, the relation is arbitrary. No one can be the biggest person without a stipulated, arbitrary equivalence relation. More generally, without such an equivalence relation, nothing is or even could satisfy the superlative of an incommensurable, multi-variable expression.

At any rate, *almost* nothing could. Suppose there were someone—let's call him Goliath—who was bigger than every other person with respect to every single size-related category. It might truly be said of Goliath that he is the absolutely biggest person. Or at least, it can truly be said of Goliath that he would qualify as the biggest person on any stipulated equivalence relation, and this seems like a good enough reason to *call* him the biggest. Either nothing can satisfy the superlative of an incommensurable, multi-variable expression or nothing can satisfy the superlative of an incommensurable multi-variable expression unless it satisfies the superlative of every variable within the multi-variable expression.

There is not very much at stake over whether Achilles or Hercules is bigger. But the problem of maximizing over multiple variables is neither new nor without importance to philosophy. Its most famous application is perhaps as a reason to reject some of the early versions of utilitarianism according to which an action is right iff it maximizes pleasure and minimizes pain. Critics have pointed out that it is impossible in some situations both to maximize pleasure and to minimize pain since neither is reducible to the other. (Minimization is no less difficult than maximization in this regard. Minimization is merely maximization of an inverse.) Moreover, utilitarians have for the most part tried to

accommodate this critique, finding ways to evaluate utility by a single variable (preference, for example).

So far, we have been considering the word, “biggest.” Identifying the biggest is difficult because “being big” can indicate height, weight, mass, or a host of other size-related qualities. Another word that often operates similarly is “best.”

Let’s modify an example from Plato’s *Meno*. In Plato’s famous dialogue, Socrates asks Meno what value there is in knowing the road to Larisa over and above having the right opinion about the road to Larisa. This is a difficult question, requiring us to distinguish knowledge from merely true belief. But Socrates’ question would have been even more difficult if he had asked about the *best* road to Larisa, for in order to answer this question, we need not just understand what is knowledge, but how to determine which road is best.

In order to understand what is the *best* road, one must know what the good-making properties for roads are. The difficulty of determining one, unique, *best* road, is that there seem to be several (potentially) conflicting and incommensurable good-making properties for roads. Let us consider just two of these qualities: being quick and being easy. Let us suppose that road A is quick but not easy, road B is easy but not quick, and a third road, road C, is quicker than B but slower than A and easier than A but more difficult than B. Which road is best? There seems to be no answer to this question.

Questions about maximums can be difficult even when multiple variables are not involved. Suppose Opie is 4’6”, and both Andy and Barney are 6’2”. Who is tallest? There is no good answer to this question either. Or rather, there is a good answer to this question, but it is bad English. No one person is uniquely the *tallest*, although Andy and Barney are both tied for being the most tall. However, the difficulty of this question is fundamentally different than the difficulty about the best road. In the question about height, we know that Andy and Barney are tied because they have the same height. The only difficulty is that the use of a superlative in English prompts us to pick out one unique object, and in this case

there is a tie between two. In the question about which road is best, we don't even know whether or not there is a tie. Answering that the three roads are equally good would be no better than answering that one is best indisputably. *Any* answer we give would imply the commensurability of good-making qualities that are, in fact, incommensurable.

Note that this is a different question than "Which road is best for so-and-so?" Perhaps Socrates needs to get to Larisa very quickly, and Meno's leg is cramped. Even if there is no answer to the question of which road is best *simpliciter*, we are still prepared to say that the best road for Socrates is road A and for Meno B. Why is this so? One plausible answer is that it is because Socrates and Meno have different telic and pragmatic values assigned to the various good-making properties of roads. In this case, Socrates values the goodness of speed over ease whereas Meno is presently disposed to value ease over speed. The values that they have actually assigned (or their preferences, or their dispositions) give us a way to weigh the different values, but they only do so relative to an individual's values and goals, not absolutely.

2.3 Two Arguments: A Brief Introduction

These observations play an important role in two arguments against rational uniqueness. But before we develop each argument independently, it will be helpful to note how they operate corporately in the case for extreme rational permissivism.

Recall that we distinguished between weak and strong disagreements. A disagreement is strong if one person believes that p , whereas another believes that $not-p$; conversely, a disagreement is weak if one person believes that p , whereas another refrains from belief concerning p . In what follows, I present two different arguments for the possibility of rational disagreement between evidential peers. The Jamesian argument targets weak disagreements, and the argument from abduction targets strong disagreements.

The Jamesian argument begins with the observation that two, distinct good-making properties of DA's are their orientation towards truth (or some other epistemic success) and their orientation away from falsity (or some other epistemic failure). Drawing on observations we have already made about the impossibility of maximizing over incommensurable variables, I argue that whether or not a belief is rational depends on the cognitive goals of the agent. The abduction argument considers the various good-making properties that play a role in inferences to the best explanation. There can be no "best" explanation absent an individual's valuation of distinct theoretic virtues. If rationality depends on a person's cognitive goals and epistemic values, rational disagreement is possible. (RUE) is false.

2.4 The Jamesian Argument: Maximizing Truth and Minimizing Error

Recall both that (RUE) is committed to epistemic bestism and that it is impossible to maximize over multiple, incommensurable variables without stipulating an equivalence relation (with a possible exception clause for "Goliath" cases). The Jamesian argument shows that there are at least two incommensurable desiderata for DA's with respect to our cognitive goals. Accordingly, though there can be (and are) best DA's relative to some standard or other, there is no best DA absolutely.

One of William James's most important insights in epistemology is that we have epistemic obligations both toward truth and against falsehood.

There are two ways of looking at our duty in the matter of opinion,—ways entirely different, and yet ways about whose difference the theory of knowledge seems hitherto to have shown very little concern. *We must know the truth*; and *we must avoid error*,—these are our first and great commandments as would-be knowers; but they are not two ways of stating an identical commandment, they are two separable laws.¹²

¹² James, William. *The Will to Believe : and other essays in popular philosophy*, New York : Dover, 1956, section 7.

Though James is known for his pragmatism, we must understand these duties, as Chisholm does, to be “purely intellectual requirements.”¹³ The distinction is not merely that it is sometimes more *pragmatically* helpful either to search for truth than avoid error or vice versa, though this may be true. Rather, we have irreducible and distinct epistemic obligations toward correct and against incorrect belief.

It is unclear what precisely these two opposing requirements are. James defines them this way:

James’s Norms

- 1) Know the truth.
- 2) Avoid error.

But perhaps we are unsatisfied with James’s formulation. We might think, for instance, that truth and falsity most fundamentally explain our epistemic duties, or that that our duties should be knowledge-centric, or virtue-centric. Perhaps so, but the same Jamesian duality pops up whichever norms we choose.¹⁴

Truth Norms

- 3) Do believe truths.
- 4) Do not believe falsehoods.

Knowledge Norms

- 5) Do believe that *p* if by believing that *p* you will come to know that *p*.
- 6) Do not believe that *p* if by believing that *p* you would not come to know that *p*.

Virtue Norms

- 7) Do perform epistemically virtuous behavior.
- 8) Do not perform epistemically vicious behavior.

Note that we could fulfill (4), and (6) by refusing to believe anything, whereas we could fulfill (3), and (5) by believing everything.¹⁵ This is not at all to say that these norms are

¹³ Chisholm, R. M. 1977. *Theory of Knowledge*, 2nd edn. Upper Saddle River, NJ: Prentice-Hall, pp. 14.

¹⁴ Of course, it might be that several of these pairs of norms are fundamental to epistemic normativity.

¹⁵ Brian Leftow helpfully pointed out to me that if someone were to believe every proposition, they would plausibly have a defeater for any knowledge that she otherwise might have had. Knowing that *p* is plausibly impossible if one also believes that *not-p*. Nevertheless, believing every proposition is still a good strategy for satisfying (5). If one believes every proposition, then there will be no proposition such that if one believed it one would come to know it (for every such belief would have a defeater). Thus, the obligation to believe that *p* if by believing that *p* one would come to know it will be satisfied, even if somewhat trivially.

defective, only that they are incomplete. Only both norms together can explain our epistemic obligations. The virtue norms are a bit harder to evaluate, but it is plausible that there will be virtues directed toward acquiring epistemic goods as well as virtues directed against committing epistemic flaws.

Pick any pair or pairs of norms you like. What relevance could this have toward bestism and (RUE)? Let's start with an example from Thomas Kelly.¹⁶ Suppose I tend to emphasize the requirement to believe truths. I will be on the lookout for new information and be careful not to discount any source of evidence that might lead to true belief. By contrast, you emphasize the requirement not to believe falsehoods. You will exhibit caution, constantly wary of errors or misleading evidence. Suppose both of us are presented with the same evidence E concerning and hypotheses H. As per Kelly's scenario, let's further assume that the evidence in question is good but not overwhelming. Naturally, for a number of such hypotheses, I will come to believe H whereas you will refrain from belief. Kelly's judgment on this case is illuminating:

As an H-believer, if I learned that we differed in our cognitive goals in this way, I would be disinclined to conclude that the manner in which you are responding to our shared evidence is unreasonable, even though it differs from my own. In fact, I might even think that if you were responding to the evidence in any other way than you are, then *that* would be unreasonable, given your cognitive goals. Moreover, notice that making such a judgment has no tendency to make me insecure in my conviction that I am also responding to the evidence in a reasonable way, given my cognitive goals. Theupshot: subtly different ways of responding to the same body of evidence seem equally reasonable, given corresponding differences in the weights that we give to our shared cognitive goals.¹⁷

Kelly is suggesting that rationality is not determined just by the evidence, but by how the evidence relates to our epistemic (cognitive) goals. Someone who more highly values the goal of acquiring truth will have a different threshold for evidence than someone who more highly values avoiding error.

¹⁶ Kelly, Thomas. (2013). "Evidence Can Be Permissive," In *Contemporary Debates in Epistemology*, 2nd ed.(ed.) Mathias Steup, John Turri, and Ernest Sosa, Blackwell.

¹⁷ Kelly, 302.

So far, our appraisal of the Jamesian argument has concerned distinct and conflicting good-making properties, but distinct and conflicting duties. These problems are intimately connected. Conflicting duties create a need for trade-offs between irreducible “right-making” qualities in the way that talk of a best road or DA requires trade-offs between irreducible good-making qualities. But rather than developing a parallel problem for duties, we shall translate the Jamesian problem into a dilemma about good-making properties.

We have just learned from James and Kelly not only that there are two different obligations we have toward belief formation, but two different good-making qualities for a DA. A DA is made good by the property “being oriented toward true belief” as well as “being oriented away from false belief.” If these are not good-making properties concerning the DA’s themselves, they are at least good-making properties concerning the formation of DA’s. And it is precisely the appropriate formation of DA’s that is relevant to rationality.

What is the best DA to have? It depends on the relative weight you assign the good-making properties of DA’s. Even if there is a uniquely best DA for an individual given her evidence *in conjunction with* her particular epistemic goals, there is not a universally best DA for all people simply given the evidence. Perhaps it is necessary that the best doxastic attitude should be oriented toward both truth-acquisition and error-avoidance, but what relative weight should these values be given? Recalling what we have said about maximizing over multiple variables, it is doubtful that any single weighting can claim a privileged position. Insisting that there is a universally best DA for each set of evidence looks no better than insisting that there is a universally best road. Note also that no ratio valuing the attainment of epistemic goods relative to the avoidance of epistemic harms can qualify as a “Goliath.” Any increase in valuing the attainment of epistemic goods increases

the risk of accruing epistemic harms, and any increase in valuing the avoidance of harms decreases the chances of attaining epistemic goods.

One possible objection to the Jamesian argument as presented is that it only works when we think of DA's as beliefs but fails if we think of DA's as credences. In borderline cases between belief and unbelief, perhaps the truth-seeker and error-avoider can agree on what credence to assign p even if one is more likely to assert p in ordinary contexts than the other. But then, if credences rather than beliefs are the most fundamental DA's, James's norms make no trouble for (RUE), as both parties can agree about what credence the evidence demands.

This response will only work if it turns out that credences are not subject to the same sort of Jamesian duality of norms as beliefs. But their exemption is dubious. Consider:

- 9) Do assign the correct credence to p .
- 10) Do not assign an incorrect credence to p .

Someone could satisfy (10) but not (9) simply by not assigning credences to propositions. Furthermore, there are Jamesian dualities for credences even assuming that all parties have assigned a credence:

- 11) Do not assign too high a credence to p .
- 12) Do not assign too low a credence to p .

Citing Rachiele, Kelly suggests a duality of his own:

- 13) Minimize the gradational inaccuracies of your credences.
- 14) Lower the variance of gradational inaccuracy in your credences.¹⁸

Kelly rightly points out that these goals, whether formulated for beliefs or credences, are complementary in that satisfying one requirement makes it easier to satisfy the other. Nevertheless, "the fact that they are different goals creates the need for trade-offs; the optimal strategy for the achievement of one is not the optimal strategy for the achievement of the other."¹⁹ Just as in the previous dualities, these epistemic goals can be translated into

¹⁸ Kelly, 10.

¹⁹ Kelly, 10

good-making properties for DA's. DA's have the good-making properties of "being directed away from assigning overly high credences" and "being directed away from assigning overly low credences," for instance. Insisting on credence-language rather than belief-language offers no escape from the Jamesian critique.

Whatever the strength of the Jamesian argument, its effect is somewhat limited. It suggests only that there are weak, but not necessarily strong, rational disagreements. If some evidence is good but not overwhelming for a hypothesis, it may be rational either to believe or refrain from belief. It does not follow from this, however, that it is rationally permissible to *disbelieve* that hypothesis (that is, to believe its negation).

Although uniqueness is usually articulated in a way that rules out both weak and strong disagreement, in the face of the Jamesian argument, the (RUE) advocate might make yet another tactical retreat. "The sort of supposedly rational disagreement that most troubled us," she might say, "was not when disagreement arose simply because one person considered a proposition while her evidential peer did not give it a second thought. Nor was it even borderline cases of belief in which one evidential peer commits to a belief while another abstains. No, the cases that most trouble us are when two people look at the same evidence and come to opposing conclusions. It seems absurd to us that the same evidence could rationally seem to one person evidence for p and, simultaneously, that it could rationally seem to be evidence for *not-p* to her evidential peer. What concerns us is not weak but strong disagreement."

We might doubt whether most defenders of (RUE) would actually be happy to make this concession. Indeed, given that (RUE) is usually formulated as a stronger thesis, we might be suspicious that such a move would be *ad hoc*. The possibility of rational, weak disagreement among evidential peers is important in its own right, for it is incompatible with strong interpretations of (RUE). Nevertheless, we should want to know not just

whether rational, weak disagreement among evidential peers is possible, but whether rational, strong disagreement is possible. After all, religion and philosophy are full of strong disagreements, and wondering whether these kinds of disagreements could be rational was what motivated our investigation in the first place.

The argument from abduction attempts to rule out any such retreat on behalf of (RUE). Rational, strong disagreement is possible, for different agents may rationally infer the truth of incompatible hypotheses by applying different, acceptable weightings of the theoretic virtues.

2.5 The Argument from Abduction

Abduction may be casually defined as “inference to the best explanation.” This definition is too simple. Various restrictions must be put in place. If hypothesis H_1 is the best overall, but H_2 is only slightly worse than H_1 , then an abductive inference to H_1 is invalid. So not only must H_1 be better than all other (considered) explanations, but it must also be significantly better than the second best explanation.²⁰ A related constraint often suggested is that the hypothesis must be all things considered probable in the eyes of the agent. If H_1 is the best explanation and significantly better than its competitors, but improbable overall, then an abductive inference to H_1 is invalid. Although it is important to acknowledge these constraints, abduction nevertheless begins with a search for the best explanation.

How does one determine which theory is best? No canonical list exists, but there is general agreement on some of the core considerations. All else being equal, a theory is better if it...

- 1) Is simple
- 2) Fits the data
- 3) Explains the data
- 4) Prioritizes the important data
- 5) Is elegant or lovely

²⁰ How “significant” the gap must be is also a matter of debate. Plausibly, it could be contextually determined.

Let us call these qualities “theoretic virtues” and examine each further.

Simplicity: The value of simplicity in philosophical thinking is most famously captured by Ockham’s razor. More recently, we can call to mind Quine’s articulated preference for desert landscapes. All things being equal, we prefer a theory that does not posit many things when positing a few would suffice. We call this kind of simplicity “parsimony.” Simplicity itself can be broken into distinct desiderata. Simplicity of types is distinct from simplicity of tokens. David Lewis’s version of modal realism, for instance, has type but not token simplicity.

Fits the Data: Especially when dealing with theories that attempt to account for a wide spectrum of phenomena (as in scientific theories), we often find that no available theory perfectly fits the data. Nevertheless, we value theories that have few outliers and anomalies.

Explains the Data: Explaining the data is distinct from fitting the data. Imagine trying develop a theory for some evidence set $\{E_1, E_2, E_3, \dots E_n\}$, where all the evidence is propositional. One theory that would fit this data very well would be the theory that of all the propositions that constitute the evidence set are true. Perhaps the truth of all these propositions is just a brute fact. However appalling we might find this theory, we must admit that it would fit the evidence perfectly. Nevertheless, such a theory would not *explain* any of the data, but only restate it. We value theories that explain much of the data by relying on relatively few brute facts and primitive concepts.

Prioritizes the Important Data: Often a theory fits or explains part of the data but not all of it. When this is the case, we prefer a theory that answers our questions about the data we consider to be most important or fundamental. We prefer theories that provide answers to the questions that matter to us and that we bring to an inquiry. If one theory of reference provided a good explanation for non-fictional reference, but not for fictional

reference, we would probably prefer it to an incompatible theory that explained fictional but not real reference. Plausibly, this is in part because most of us think a theory of non-fictional reference is more important and fundamental.

Elegance and Loveliness: Simplicity, no doubt, contributes to a theory's elegance, but we also value other aspects of a theory that make it lovely. We value theories that have patterns, symmetry, or parallels, and theories that are united rather than disjointed. It is perhaps not clear why we epistemically favor these theories, but that we favor them cannot be doubted.

Some philosophers would prefer to do away with talk of inference to the best explanation and talk instead about inference to the most probable theory. But it is unclear how probability, subjectively understood, could be determined without relying on appeals to the theoretic virtues listed above, or others like them. If asked why a certain theory seems more probable, it seems natural to respond by saying that it is simpler, or better fits the evidence, or is more lovely, etc. Thus even if inference to the best explanation is ultimately cashed out in terms of probabilities, the same cluster of theoretic virtues are likely to play an important role in our epistemic decisions.

Abduction plays an important role in our ordinary reasoning. When I awoke this morning, I heard voices conversing downstairs. I concluded that some of the people who live in my house were already up. The evidence of my experience was not incompatible with other hypotheses. It could have been that I was having an auditory hallucination, or that trespassers had broken into the kitchen and decided to have a friendly chat over tea. However, though not incompatible with my evidence, these explanations seem significantly worse than the hypothesis I actually endorsed. It is much simpler to suppose that my fellow residents had already woken and were talking downstairs.

Abductive reasoning is also essential to anti-skeptical arguments. Few philosophers believe the existence of the external world can be proven deductively, and Hume's problem

of induction is a serious challenge to attempts to do so inductively. A plausible strategy is to show that the existence of the external world offers the best explanation for our apparently coherent and overlapping experiences. I do not claim that abduction solves the problem of skepticism, but it does seem to give the philosopher more ammunition against certain skeptical problems.

Finally, abductive reasoning is integral to science. Consider the following example from Igor Douven:

At the beginning of the nineteenth century, it was discovered that the orbit of Uranus, one of the seven planets known at the time, departed from the orbit as predicted on the basis of Isaac Newton's theory of universal gravitation and the auxiliary assumption that there were no further planets in the solar system. One possible explanation was, of course, that Newton's theory is false. Given its great empirical successes for (then) more than two centuries, that did not appear to be a very good explanation. Two astronomers, John Couch Adams and Urbain Leverrier, instead suggested (independently of each other but almost simultaneously) that there was an eighth, as yet undiscovered planet in the solar system; that, they thought, provided the best explanation of Uranus' deviating orbit. Not much later, this planet, which is now known as "Neptune," was discovered.²¹

These examples demonstrate what Douven calls the "ubiquity of abduction."²² It is present in ordinary, philosophical, and scientific reasoning. If the nature of abduction creates a problem for rational uniqueness, those problems will be ubiquitous as well.

Indeed, abduction does create a problem for (RUE). Abduction relies on determining a best explanation. But as we have seen, there are several incommensurable criteria for determining which theory is "best." Our previous argument has shown that maximizing over multiple variables is impossible. Accordingly we cannot speak of an absolute best, but only what is best relative to a certain weighting of the theoretic virtues. The distinct theoretic virtues will sometimes be competitive, and as such trade-offs between

²¹ Douven, Igor, "Abduction", *The Stanford Encyclopedia of Philosophy* (Spring 2011 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2011/entries/abduction/>>.

²² Douven, Igor, "Abduction"

virtues will be necessary. Someone who values simplicity especially highly may come to different conclusions than someone who values more highly the strict fitting of the data.

Unlike the Jamesian argument, the argument from abduction results in the rational permissibility of strong disagreements among evidential peers. Suppose Amy endorses H_1 on the basis of evidence E according to her weighting of the theoretic virtues whereas Nico endorses the incompatible H_2 on the basis of the same evidence E and his distinct weighting of the theoretic virtues. Amy is committed to the denial of H_2 just as Nico is committed to the denial of H_1 . There will be propositions that Amy rationally believes such that Nico believes the negation. Yet both may do so rationally if they have appropriately evaluated the hypotheses according to their distinct weighting of theoretic virtues.

The ubiquity of abduction makes it an indispensable part of our reasoning practice. Accordingly, we cannot reject abductive inference because no best explanation exists in the absolute and universal sense that (RUE) requires. Rather, we must interpret the inference to the best explanation as an inference to the best explanation relative to an individual's standards or weighting of the theoretic virtues. Having considered both the Jamesian argument and the argument from abduction, we have good reason to believe that there can be rational disagreement among evidential peers. Whether a disagreement is strong or weak, agents' cognitive goals and theoretic values can lead to diverse and contradictory DA's without either party making an error of reasoning.²³

2.6 Is Extreme Permissivism Too Permissive?

The theory I have sketched is permissive indeed. Some might think that it is a good deal *too* permissive, that in trying to show how some evidence is permissive (as in religious cases), I have opened the floodgates to a multitude of unacceptable doxastic attitudes.

Says the objector:

²³ Though developed independently, this argument bears some resemblance to an argument from Igor Douven in his (2009) "Uniqueness Revisited." *American Philosophical Quarterly* 46: 347–361.

On your view, it seems that almost *any* view could be rationally permissible. At least, any view can be permissible so long as one has the right sort of cognitive goals, and epistemic values to go along with it. It could be rationally permissible for someone to be a solipsist, or to doubt the existence of an external world. But it should be obvious that some views are just plain unacceptable. If your view has the consequence that a view as bizarre as solipsism could be permissible, so much the worse for your theory! Your permissivism is simply too permissive.

My view really does have the consequence that all sorts of beliefs might be permissible given the right epistemic framework. After all, despite its strangeness, solipsism is a very parsimonious theory. It posits only one object, my own mind! And what if someone were to value parsimony far enough beyond the other theoretic virtues that solipsism started to look like a good theory—perhaps even the *best* theory? Such a solipsist would be rational on my view. Why isn't this a crazy thing to say?

The first thing to note is that not every epistemic failure must also be a failure of rationality. As Jonathan Kvanvig rightly warns us, “there is a predilection among epistemologists to identify too many regrettable features of cognitive systems in terms of one favorite term, ‘irrational.’” Solipsism has the epistemic defect of being false. Worse than that, it upholds a way of looking at the world that insures one will come to many more false beliefs (and fail to acquire much knowledge). There is a sense in which solipsism is thus deeply delusory. But we should hesitate to jump immediately from delusion to irrationality. A system of beliefs may be delusory in virtue of dramatically clashing with the way the world is, and not necessarily because the agent has been epistemically irresponsible. Systematically irrational epistemic behavior may be a surefire way to end up with delusory beliefs, but the reverse may not always be true.

If a solipsist's beliefs really are calculated correctly given her cognitive goals and epistemic values, there is no reason to say that they are irrational. Of course they are profoundly wrong, delusory even. But it is not uncommon to form false, rational beliefs. Why should it be surprising that solipsism is among the false beliefs one can rationally believe, given

the right sort of epistemic framework? Nearly no one (perhaps no one at all) actually has the sort of epistemic framework according to which solipsism could be permissible, but it *would* be permissible for someone who *did* have the that framework.

Counters the objector:

Your “solution” only pushes the problem back a step. It seemed crazy that solipsism could be rational. You responded that it could be rational if someone valued a certain kind of parsimony far and above the other theoretic virtues. But this weighting of the theoretic virtues is surely just as crazy! If someone really values parsimony that much, to the extent that he doesn’t believe anything outside his mind exists, he has behaved irrationally. That is an incorrect way to weigh the theoretic virtues.

Once again, I think there is something epistemically deficient about the way the solipsist has weighed the theoretic virtues, but nothing rationally deficient. Any belief-forming method that involves such askew weightings of epistemic goods is unreliable. An agent employing such a method may occasionally strike a true belief, but will not have a good track-record overall. If an agent can have a rational, false belief, it is hard to see why an agent could not also have a rational, unreliable belief. After all, to have a rational, unreliable belief is just to have a rational, false belief enough of the time. Consider: our beliefs about the external world would continue to be rational even if we lived in Descartes’s demon world. Demon world scenarios are fascinating in part because are cases in which what is rational for us remains (nearly) the same even though the world itself changes dramatically. Despite being rational, our beliefs would be unreliable since they would be false most of the time. Thus, there is no forced march from unreliability to irrationality.

More fundamentally, the objector seems not to have fully grasped the difficulty of the various theoretic virtues (or other epistemic goods) being incommensurable. If the values are incommensurable, there just isn’t an objective, universal way to weigh them against one another. It would be arbitrary to identify some of those ways as necessarily irrational.

Counters the objector:

I'm afraid it's *you* who haven't really grasped *my* argument. When I insist that a certain way of weighing epistemic goods is obviously incorrect, I am also insisting that the epistemic goods must really be commensurable. Perhaps we cannot know exactly how they commensurate or exactly the weight given each value ought to be. Nevertheless, it is clearer to me that *some* ways of weighing the epistemic goods are wrong than it is that such goods are truly incommensurable.

This seems the wrong tack to take. How could it be that simplicity and explanatory power, for instance, commensurate? It isn't even clear how simplicity and explanatory power could be broken into units. And even if we could, how many simplicities would a powerandum be worth? These questions seem no better than asking how many inches a pound is worth. We are forced to choose between claiming that the various epistemic goods we have considered are commensurable and claiming that some bizarre ways of weighing such goods, whatever other epistemic defects they may exhibit, are not inherently irrational. I will happily defend the latter—especially since we can still criticize such weightings as epistemically deficient in other ways, e.g. as being delusory or unreliable.

Are there no rational constraints at all on the ways one can weigh various epistemic goods? Even the incommensurability of epistemic goods does not entail such a strong thesis. Consider the following principles:

P1) In order to be rational, one must value every good cognitive goal and every theoretic virtue.

If someone doesn't value acquiring the truth at all, or doesn't see that simplicity is a virtue, she isn't being rational.

P2) In order to be rational, one must not value any bad cognitive goals or any theoretic vice.

If someone values acquiring falsehoods as such, or counts it as an epistemic good that a theory justifies drinking exorbitant amounts at the pub every night, he isn't being rational. Furthermore, even given that all and only the epistemic goods are valued, not just any hypothesis could be rationally permissible according to some weighting of values:

P3) If doxastic attitude A concerning p is worse than doxastic attitude B concerning p with respect to one epistemic good, and is worse than or equal to B with respect to all other epistemic goods, then doxastic attitude A concerning p is impermissible.

Consider the theory that I haven't received my coffee yet because the barista has forgotten my order due to the whimsy of an invisible, meddling, memory-erasing wizard. This theory is less simple than the theory that I haven't received my coffee solely because the barista has forgotten my order, and it enjoys no other epistemic advantages over the simpler theory. Accepting the wizard hypothesis is thus rationally impermissible according to any way of weighing the epistemic goods.

(P1)-(P3) is not meant to be exhaustive. They are simply meant to show that accepting the incommensurability of epistemic goods does not entail the absence of any rational constraints on permissible ways of weighing epistemic goods or of rationally permissible hypotheses.

Extreme permissivism is permissive indeed, but not recklessly so. There will be some bizarre hypotheses that, when conjoined with a certain way of weighing epistemic goods, cannot be denounced as irrational; however, this seems less problematic when we consider that they can be epistemically criticized in other ways. Moreover, the extreme permissivists can still place substantive constraints on the acceptable weightings of epistemic goods without violating their intuitive incommensurability.

2.7 Conclusion

In the absence of further argument, the prospects for (RUE) look grim. Over the last several pages, we have considered several reasons to reject the rational uniqueness thesis. We noted that if (RUE) is true, epistemic bestism must be as well. However, because there are several distinct good-making properties of DA's, and because one cannot maximize over multiple variables,²⁴ neither can there be an absolutely best DA. The duality (or dualities) present in

²⁴ Except in "Goliath" cases, which do not appear to be at play in the Jamesian argument or argument from abduction

epistemic goals secures weak, rational disagreement. The abduction argument shows how similar reasoning about the distinct theoretic virtues secures rational, strong disagreements.

Even if there is not an absolutely best road, there can be a best road *for* so-and-so according to her individual needs, context, and weights assigned to the various good-making properties for roads. In the same way, even if there is no absolutely best doxastic attitude, there can be a best doxastic attitude according to the individual's epistemic goals and values. Epistemic bestism, therefore, is either false or must be interpreted in a way that is agent-relative, or at least relative to the epistemic goals and values of the agent. Both options are inconsistent with (RUE), thus (RUE) is false. Moreover, it seems that extreme (RP) is true. Two evidential peers can strongly disagree about the truth of p without having made any mistakes in reasoning.

3. Resisting Uniqueness

I consider arguments for Uniqueness as presented by Roger White and Richard Feldman. Although their arguments force the permissivist to draw some important distinctions, they do not compel the permissivist to abandon her position.

3.1 Quick and Dirty Permissivism

In “Epistemic Permissiveness,” White begins his defense of uniqueness with what he calls the “quick-and-dirty” argument.²⁵ Though we shall also evaluate some of his “long-and-clean” argumentation, this argument is a good place to begin as it captures one of the central intuitions motivating rational uniqueness. White asks us to imagine that we are in the juror’s box weighing the evidence for and against Smith, the defendant:

[T]he evidence cannot support both Smith’s innocence and his guilt. Whatever is evidence for P is evidence for the falsity of not-P and hence is evidence against not-P. Of course, certain elements of or aspects of the total body of evidence might suggest that Smith is guilty, while others suggest the opposite. But it is incoherent to suppose that a whole body of evidence could count both for and against a hypothesis. So then it is impossible that my examination of the evidence makes it rational for me to believe that Smith is guilty but also rational to believe instead that he is innocent. And since neither view was rational apart from any evidence, the proposed radical departure from Uniqueness cannot be right.²⁶

This is a natural and intuitive thought. How could evidence be both for and against a hypothesis? How could evidence make something more and less likely at the same time?

Fortunately, if the arguments for permissivism previously presented are sound, then there is a quick-and-dirty response to White’s objection. Evidence cannot be evaluated apart from standards. Though not necessarily exhaustive, we may suppose that epistemic standards include the epistemic goals and values that we considered earlier. No set of evidence can be evidence for and against *p according to the same standards*. However, a set of evidence can be evidence for *p* according to one set of standards, but evidence for $\sim p$ according standards that are different.

²⁵ White, Roger. “Epistemic Permissiveness,” *Philosophical Perspectives*, 19, 2005, pp. 447.

²⁶ White, 447.

3.2 Feldman's Principle and Escalating Skepticism

In "Reasonable Religious Disagreements," Feldman anticipates this dialectical move from the permissivist. Feldman does not use the language of "standards," but rather "starting points" or "fundamental principles:"

A difficult project, which I will not undertake here, is to identify just what these starting points or fundamental principles might be and to explain how they might affect the sorts of disagreements under discussion. But whatever they are, I do not think that they will help solve the problem. Once people have engaged in a full discussion of the issues, their different starting points will be apparent. And then those claims will themselves be open for discussion and evaluation. ...Once you see that there are these alternative starting points, you need a reason to prefer one over the other.²⁷

In this paragraph, Feldman proposes a principle for determining the rational validity of different starting points. Let's call this "Feldman's Principle":

(FP): If S has starting point α , and if S knows that there is an alternative starting point β , and if S is behaving rationally, then S has a reason to prefer α to β .

Notice that Feldman's Principle only applies when S knows about an alternative starting point. Thus formulated, (FP) is not strong enough to secure rational uniqueness, for in some cases, the agent may not know an alternative starting point exists. To secure Feldman's intended conclusion, we can strengthen the principle thus:

(S-FP): If S has starting point α , and if there exists another starting point β , and if S is behaving rationally, then S has a reason to prefer α to β .

If (S-FP) were true, then the permissivist could not appeal to alternative standards as an explanation of how the same set of evidence can be rationally interpreted as evidence for contradictory hypothesis. Fortunately for the permissivist, both (FP) and (S-FP) are false. Since (S-FP) logically entails (FP), we can restrict our attention to the falsity of (FP).

Rational reasoning must start somewhere. An infinite chain of reasoning might be possible, but it does not reflect our actual intellectual life. Circular reasoning is certainly

²⁷ Feldman, 206

possible, but just as certainly irrational. Even if our reasoning more closely models coherentism, there must be some initial starting points among which coherence can be found. These considerations are normally advanced on behalf of some version of foundationalism concerning propositional justification; however, they also reveal something about epistemic standards. According to what standards does one accept the epistemic standards one uses in weighing evidence? There must either be further standards by which those standards are evaluated or else those standards are in some way basic.

Suppose (FP) is true. If (FP) is true, then a rational agent must have a reason to prefer her standards over every other standard she knows to exist. This is a tall order indeed! For as we have seen, there is reason to think that there are an infinite number of potential standards. There are an infinite number of ways one could weigh the relative cognitive goals of acquiring epistemic goods and avoiding epistemic harms; there are an infinite number of ways one could weigh the relative theoretic virtues of simplicity, explanatory power, and the like. Even if an agent is not cognizant of every alternative standard, or does not understand the content of every alternative standard, the reflective person should quickly recognize how myriad and how disparate various standards might be.

But that something is hard is rarely a decisive objection against it. Perhaps rationality simply is very difficult. A deeper worry is that (FP) and uniqueness seem to commit one to broad skepticism. Suppose Kat and Michelle disagree about p . Kat believes that p and cites standards α in the evaluation of her evidence. Michelle considers the same evidence but believes that $\sim p$ according to her standards β . Applying (FP), each searches for a reason to prefer her own standards. Kat finds a reason through the application of meta-standard γ , whereas Michelle does the same via δ . They continue to disagree, and each becomes aware that yet another alternative starting point exists. They continue for some time. In fact, being stubborn and finding themselves with too much time on their hands, Kat and Michelle continue to exhaust all of their standards until they arrive at standards ψ

and ω respectively. Neither can think of any additional reason or standard, unless it is according to itself (if this is possible), by which they prefer ψ or ω ; they are basic, foundational standards. According to (FP), Kat and Michelle both ought to forfeit their standards and, accordingly, their beliefs about p .

This wouldn't be very problematic if it only meant forfeiting beliefs about p . Indeed, this is just what conciliationists, as well as many defenders of uniqueness, will recommend. What *is* problematic is that once one sees that there are alternatives even to her most basic and foundational standards, the threat of skepticism that began with forfeiting belief that p spreads to all other beliefs that are evaluated by the same standards. Suppose Kat and Michelle discover that they share a belief that q . Each of them evaluates q according to the same standards as she evaluated p . But now neither of them can continue to believe q —even though they agree about q —because it is licensed by the same standards they were already forced to abandon as a result of (FP).

Epistemic standards are relatively consistent across domains. If Kat applied standards α in deliberation about p but some incompatible standards β in deliberation about q , we should think she is exhibiting a rational deficiency. Her standards are inconsistent. Of course, various features of some set of standards may be more salient when considering p than when considering q , but it would be surprising if the standards themselves substantially changed. Accordingly, what began looking like a relatively benign skepticism about p has revealed itself to be an infectious and widespread skepticism. At least, it will spread to all other beliefs that S arrived at by using the same standards as she used to arrive at the belief that p , regardless of whether there is any actual disagreement concerning the relevant belief. We may call this the problem of escalating skepticism.

Uniqueness may be preserved if the story about Kat and Michelle is impossible. If the ultimate, basic standards are such that they admit of no alternatives, then Kat and

Michelle could never rationally arrive at the opposing, basic standards ψ and ω . In this case (FP) need not lead to a radically skeptical conclusion. However, I take it that the onus is on the defender of (RUE) to show what these elusive standards might be. The history of human thought does not seem to have converged upon any such standards and appears to be in no danger of doing so.

3.3 Chancy Pill-popping and Epistemic Luck

White presents his own series of arguments intended to block the permissivist's appeal to alternative standards. White develops several stories involving someone who knows himself to be in an epistemically permissive situation. He is offered a pill that would change his standards so that he would go from rationally believing that p to rationally believing that $\sim p$. The opportunity to take a standards-altering pill invites us to consider the apparent randomness and arbitrariness of having one set of standards as opposed to another—arbitrariness that seems to undermine the reliability of such standards in disputed cases.

White writes:

I think we ought to be suspicious of this position [that different standards make permissivism plausible] for similar reasons as before. First note that if our permissivist takes his own standards to be a reliable guide to the truth, then since the alternative standards deliver very different conclusions, he must judge them to be rather unreliable. So he should judge himself very lucky to have adopted truth-conducive standards, since with full rationality he could have followed ones that would lead him wildly into error. We see here that the same worries about arbitrariness just arise at a different position. How have I come to hold the epistemic standards which lead me from my evidence to conclude that P ? According to this permissivist it was not by virtue of being rational, since it is consistent with my being rational that I adhere to rather different standards that would have me believe not- P instead. But then it seems that my applying the correct standards and hence arriving at the right conclusion is just a matter of dumb luck, much like popping a pill. And hence I ought to doubt that I really have been lucky enough to do so.²⁸

White claims that the permissivist is committed to the following conditional: if one is in a situation that is known to be rationally permissive, one couldn't have arrived at his standards in virtue of being rational. This is, according to White, because "it is consistent

²⁸ White, 451-452.

with [his] being rational that [he] adhere to rather different standards” according to which he would have believed differently.²⁹ But this seems to beg the question against the permissivist. The extreme permissivist accepts that it is possible for a person to rationally believe that p even though, on the same evidence, someone could rationally have believed that $\sim p$. Why shouldn't the permissivist also think that it is possible for a person to rationally arrive at one set of standards even though one might rationally have arrived at a different set? Someone could come to accept a set of standards in virtue of her rationality even if her rationality does not fully determine what standards she accepts. Comparatively, subatomic particles behave in virtue of their microphysical properties even though such properties do not fully determine their behavior.

Nevertheless, White is right that there does seem to be an important sense in which our pill-popping protagonist is extraordinarily lucky (if he gets the right result). Arriving at the correct standards appears to be a matter of luck. And if he is lucky, how can he be credited for a rational achievement? White thus raises an important question for the permissivist. However, simply showing that a belief is lucky need not mean it is irrational. One must show that the luck is malicious with respect to rationality and not benign.

Recent reflections on knowledge shed light on the distinction between malicious and benign luck. An underlying problem in nearly all Gettier cases is that the would-be knower seems to have benefited unduly from luck. Suppose I think I see a sheep in the field and justifiably form the belief, “There is a sheep in the field.” Curiously, the sheep is actually a dog that just happens to look very much like a sheep, but the belief is true in virtue of a real sheep standing in the field on the other side of a hill. My belief is true and justified, but the fact that I am right seems like dumb luck. I was lucky, epistemically

²⁹ White, 452.

speaking, that a real sheep happened to be somewhere else in the field. Our reactions to such cases reveal an anti-luck intuition about knowledge.

But not all luck is incompatible with knowledge. In Dahl's *Charlie and the Chocolate Factory*, Charlie finds one of five golden tickets distributed throughout the world that allows him to visit Mr. Wonka's marvelous chocolate factory. While there, he learns many strange things, among them this proposition: "There exists a chocolate waterfall." Charlie might think to himself, "Wow! I am so *lucky* to be here! If I hadn't happened to pick the fifth golden ticket, I would not be here now. I wouldn't have ever known that there was such a thing as a chocolate waterfall. In fact, I would have even had the false belief that there weren't any chocolate waterfalls." Charlie's belief is indeed lucky in a very real sense. He could easily have had a false belief about the existence of chocolate waterfalls, and he would have failed to know if the world had been only slightly different. But the near-universal appraisal of such situations is that the relevant luck is epistemically benign. Someone can be lucky to find herself in a position to know without contradiction.³⁰

Here is another case that bears more directly on standards. Suppose that in a philosophical apocalypse, skeptics of the worst kind infiltrate the education system of every major government. Children are raised with epistemic standards that assign no evidential weight to information from the senses. You, however, grow up in post-apocalyptic Oxford, one of the few resistance holdouts against the new regime. Accordingly, you do assign evidential weight to the evidence of your senses. One day you think, "Wow! I am so *lucky* to have been born in Oxford! It is just a matter of chance that I was raised with the epistemic standards that allow me to admit the existence of an external world." Indeed, it is a matter of luck that you have the standards that allow you to have knowledge of the external world,

³⁰ I do not mean to imply that Charlie's beliefs are not "safe" in the technical, epistemic sense. The point is rather that certain kinds of luck (e.g. the luck of being in a position to know) are not relevant to the safety condition.

but it is epistemically benign luck—and this despite the fact that you are even luckier in this case than White’s pill-popping case in which the odds are presumably 50/50.

What is the difference between the luck in the Gettier case and Charlie’s luck, or the luck of the denizens in post-apocalyptic Oxford? What explains why luck is malignant in the first case but benign in the others? This question deserves a paper in its own right, but for now let me gesture toward one answer.

Everyone was lucky to arrive at a true belief; however, only in the Gettier case was it lucky *given* the belief-forming method that the belief turned out to be true. Charlie was lucky to be in a position to form the belief, and the Oxonians were lucky to have an appropriate belief-forming method. But holding fixed their beliefs and belief-forming methods, there is nothing lucky about their having arrived at the correct answer. In the Gettier case, one is lucky to have arrived at a *true* belief; the luck comes “in between” the belief and the truth. In the benign cases, one is lucky to have arrived at a true *belief*; the luck does not come between the belief and the truth.

In response to White’s charge, permissivists should maintain that any luck involved in acquiring truth-conducive standards is benign to rational belief. Rational belief involves correct calculation of the evidence given appropriate standards. If White’s protagonist does not take the pill and correctly calculates the evidence upon what are, *ex hypothesi*, appropriate standards, his belief is rational. If those standards are truth-conducive, he may be lucky indeed, but not lucky in such a way that invalidates his rationality. He is lucky to have arrived at standards that produce true and rational beliefs; but, given the rational standards he in fact has, there is nothing lucky about those standards producing true and rational beliefs.

It can hardly escape notice that there are important dissimilarities between White’s case and my philosophical apocalypse. Perhaps most relevant is that those who have grown

up in Oxford may well have a story to tell about *why* their epistemic opponents have come to a different conclusion. The skeptics have grown up in a society that brainwashes children into a certain kind of skepticism. Nothing in White's story suggests that the protagonist would have a similar story to tell about why he would have deviant standards if she took the pill. In fact, because the situation is *known* to be permissive, he knows that if he had taken the pill he would *not* have had deviant standards.

Perhaps in these cases—cases in which one knows that the situation is permissive—one ought to suspend belief. Knowledge of the permissibility may thus invalidate what was previously rationally permissible.³¹ There is a genuine puzzle here. On the one hand, on what basis can someone justify her own belief if she knows that someone else's contrary belief is just as permissible? On the other hand, if we have already stipulated that the agent is in an epistemically permissible position, how can her beliefs become rationally impermissible simply by learning about the range of permissible beliefs the evidence allows—especially since she will find her own belief among them?

An analogy to ethical norms may help alleviate the tension in the first question. I am completely untroubled in my belief that action A is morally permissible upon discovering that action B is also permissible in the same situation. If I can justify my moral actions even with the full knowledge that someone else's contrary action is justified, perhaps the same can be true of beliefs. One salient difference we can consider on behalf of (RUE) is that rational belief, unlike moral behavior often has the phenomenology of being forced. Rational belief is often, if not always, accompanied by the feeling that one cannot help but believe as one does given the evidence.³² Some moral decisions are accompanied by a similar phenomenology (e.g. when I heard about their situation, I just *had* to help), but most are not. I feel that I could justifiably donate some portion of my

³¹ White considers this option on pp. 450.

³² This is not to say that one cannot (or feels that one cannot) do things that indirectly affect one's own belief.

income either to the Red Cross or to Oxfam in a way that I never (or at least very rarely) feel that I could choose to believe *p* or its negation. Nevertheless, it may be that the phenomenology of being forced has more to do with the involuntary nature of beliefs than a difference in the underlying, normative structure. If so, we may have at least a hint of how someone might be justified in her belief even while knowing that another belief is equally permissible.³³

Fortunately, however, few cases are actually like this. Rarely do we take ourselves to know that those who disagree with us are being rational in a particular case. More often we are unsure whether or not someone who disagrees with us may be behaving rationally. And in these cases, all the permissivist must maintain is that the relevant sort of luck is rationally benign.

Though I believe White and Feldman's objections can be rebuffed, they place a considerable explanatory burden on the permissivist, especially in the cases of known permissibility. There is pressure to modify one's belief in some way in these circumstances. One natural response for the permissivist is to search for moderate concessions whereby the permissivist shows sensitivity to the tension without surrendering the full rationality of the disputed belief. Does the permissivist have any way to relieve the tension of rational disagreement for cases in which we suspect we may be in permissive circumstances?

One option available to the permissivist is to suggest that evidence of a rationally permissive situation operates not as first, but second-order evidence for the continued rationality of our beliefs. Higher-order evidence, according to Feldman, is "evidence about the existence, merits, or significance of a body of evidence."³⁴ If we acquire evidence that our situation might be a permissive one, we plausibly acquire evidence that is more directly relevant to the question of whether or not the evidence forcefully supports *p*, or about what

³³ I owe thanks to Brian Leftow for drawing my attention to this potential parallel.

³⁴ Feldman, Richard. "Evidentialism, Higher Order Evidence, and Disagreement," pp. 304.

standards should can be used to evaluate p , than about p itself. Perhaps the permissivist should suggest that in at least some cases of recognized, permissive situations, one can rationally maintain a first-order belief about p while simultaneously suspending reflective judgment about whether or in what way one's belief that p is rational.

Is such suspension of judgment at the second-order level consistent with maintaining rational belief? Can one believe that p on the basis of the evidence while simultaneously suspending judgment about whether that evidence supports p ? More than that, could someone believe that p on the basis of the evidence while simultaneously disbelieving that the evidence supports p ? Feldman argues that one cannot. Against the claim that, possibly,

- ii) One is justified in maintaining one's belief that P and also justified in suspending judgment or in disbelieving that the first-order evidence supports P .³⁵

Feldman objects:

While there is no logical inconsistency in the attitudes one has in (ii), it is odd to maintain attitudes that have the kind of levels incongruity this implies. This view seems to have the implication that reflection on epistemic support relations has no impact on the justification of one's first order beliefs in these cases. Imagine this attitude applied to another kind of case, one in which a person does not encounter disagreement but instead is simply reflecting alone. The person reasonably comes to the conclusion that his evidence does not support P . And he then stops believing P . Approach (ii), applied to this case, would have the proper analysis be that this person was reasonable in drawing his conclusion about his evidence, but not reasonable in forming a first order belief in light of that evidence. This isolation of levels seems a clear mistake to me. There would be something extremely odd about criticizing a person who justifiably believes that his evidence does not support a proposition for failing to believe that proposition.³⁶

In his critique, Feldman focuses on cases in which one believes that p but disbelieves that the first-order evidence supports p . *That* does seem very odd; however, it seems much less odd that one could rationally believe p while suspending judgment about whether or not her belief in p is rational.

The relationship between first and second-order levels of rationality merits closer inspection. In the next section, I contend that rationality is anti-luminous, allowing for

³⁵ Feldman, 307.

³⁶ Feldman, 307-308.

sizable differences between our credences at the first and second-order levels. One can reflectively suspend judgment about whether one's beliefs is rational while maintaining rational belief at the first-order level.

3.4 Anti-Luminous Rationality

Knowledge is anti-luminous. One can know without knowing that one knows. More than that, one can know without *being able* to know that one knows. But is rationality anti-luminous? Can one rationally believe that p without rationally believing that one rationally believes that p ? More than that, are there some p such that one can rationally believe that p without being able to rationally believe that one rationally believes that p ?

I answer that rationality, like knowledge, is anti-luminous. (I refer those unconvinced that knowledge is anti-luminous to Alston [1980] and Williamson [2000]).³⁷ Those unpersuaded by arguments for the anti-luminosity of knowledge are unlikely to be swayed by my arguments for the anti-luminosity of rationality. However, I will argue that *if* some of the arguments for anti-luminosity concerning knowledge are any good (and they *are*), then there is also good reason to believe that rationality is anti-luminous.

What is luminosity? Tim Williamson defines the luminosity of a condition C in the following way:

(L): For every case α , if in α C obtains, then in α one is in a position to know that C obtains.³⁸

When a child whines from the back seat, "Are we there yet?" and his Mom retorts, "When we get there, you'll *know*," she isn't only hushing him; she is also saying that the condition of "being there" is luminous for him. Do I know that p ? "When you know, you'll *know* you know," responds the luminist. "Or at least you'll be able to." Knowledge is luminous if the following conditional obtains in all cases:

³⁷ That is, Alston's "Level Confusions in Epistemology" (1980) and Williamson's *Knowledge and its Limits* (2000) (especially chapters 4 and 5) respectively.

³⁸ Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press, pp. 95.

(KK): Necessarily, if *S* knows that *p*, then *S* is in a position to know that she knows that *p*.

What does it mean to be in a position to know? There is some flexibility for the luminosity theorist to say what precisely this means. As a first pass, however, let us say that being in a position to know means that if one were to reflect upon whether she knows that *p* in normal circumstances, she would be capable of knowing that *p* without any additional evidence or special effort.³⁹

Here we are primarily interested not in the luminosity of knowledge but the luminosity of rationality. More specifically, we are interested in whether the following conditional is true:

(RR): Necessarily, if *S* rationally believes that *p*, then *S* is in a position to rationally believe that she rationally believes that *p*.

The reader may have already noticed that, as it stands, (RR) is not technically a claim about luminosity. Luminosity concerns being in a position to know, whereas (RR) concerns being in a position to rationally believe. Strictly speaking, we are evaluating whether rational belief meets the following closely related condition:

(Rational-L): For every case α , if in α *C* obtains, then in α one is in a position to rationally believe that *C* obtains.

However the parallels to questions about the luminosity of knowledge are evident. I will continue to refer to (RR) as a luminosity condition, trusting the reader to note the distinction. Incidentally, on the plausible assumption that knowledge entails rational belief, rejecting the (RR) principle is tantamount to rejecting the luminosity of rationality anyway.

The anti-luminosity of rationality is motivated by the “different evidence” argument. If knowledge and rational belief are luminous, one ought to be just as confident (or at least to be able to be just as confident upon reflection) that one rationally believes

³⁹ There are plenty of abnormal circumstances in which KK would fail to be satisfied for unimportant reasons. Suppose an evil demon stops Joanna from forming beliefs any time she reflects on her knowledge. There is a sense in which she is not in a position to know that she knows, but these are not the sorts of cases luminosity theorists are targeting.

that p or that one knows that p as one is confident that p . Additionally, we ought to proportion the strength of our beliefs to the weight of the evidence by the lights of our epistemic values and goals.⁴⁰ (Since it is cumbersome always to include “according to the lights of our epistemic values and goals,” “evidence” will sometimes be used where “evidence according to the lights of our epistemic values and goals” is intended.) Accordingly, if rational belief is luminous, so is the “Same Weight” principle:

Same Weight: The evidence for the hypothesis that p and the evidence for the hypothesis that one rationally believes that p have the same weight.

How shall we evaluate whether the same weight principle is true? A good start would be to see whether we can think of any evidence that is either (i) evidence for p or for RBp , but not both, or (ii) weightier evidence for p or for RBp than the other. As it turns out, there is evidence that meets both of these criteria, and thus, we have reason to doubt the same weight principle.

Does p itself count as evidence for p ? There are two distinct ways of thinking about evidence that suggest opposite answers to this deceptively simple question; both are problematic for the same weight principle. The first way of thinking about evidence takes the metaphor of evidential support seriously. It seems bizarre to say that a proposition could evidentially support itself, just as it seems bizarre to say that a table-top could stand by mechanically supporting itself. If a table-top is to stand, it must be supported by legs. Else it is resting on the ground. If a justified proposition is to be evidentially supported, it must be supported by other propositions. Else it is foundational. True, we sometimes speak of certain propositions as self-evident. However, it seems more likely that, in these cases, we *really* mean that the proposition requires no evidential support rather than that the proposition is self-supporting. (Compare: a person who supports herself is really someone

⁴⁰ The “ought” here should be understood as epistemic, and not necessarily as *ultima facie*.

who manages just fine while being unsupported. Or if she is supporting something, it isn't strictly herself, but a certain lifestyle or standard of living.) The "supports" relation eschews reflexivity.

The second way of thinking about evidence takes evidence to be fully explicable in terms of conditional probability. Evidence is just that which makes a hypothesis more probable when conditionalized upon. Is $P(h|e \wedge k) > P(h|k)$? Then e is evidence for h . Is p evidence for p on this model? Yes, and especially awesome evidence at that. For $P(p|p) = 1$. A proposition is always perfect evidence for itself (except, perhaps, when the prior probability was 0 or was already 1).

However, on either model of evidence, we have an apparent problem for the same weight principle. If nothing can be evidence of itself, then p is not evidence for p . But this is no reason to think p cannot be evidence for the claim that one knows that p (KP) or that one rationally believes that p (RB p).⁴¹ On this model, p and RB p are supported by different evidence sets. On the other hand, if p is really awesome evidence for itself, then p can be evidence both for p and for RB p . But, the evidence will have different weight. $P(p|p) = 1$, whereas $P(\text{RB}p|p) < 1$.⁴² Either p is evidence for RB p but not p , or p is weightier evidence for p than for RB p . Either way creates trouble for the same weight principle.

Let us next consider the evidence that one believes that p (B p). How much evidential support does B p provide for p ? The answer to this question is controversial, but arguably it is none at all (see sections 4.3-4.4), at least when taken by itself. Even if B p provides some evidential support for p , it does not provide evidential support in the same way that B p does for RB p . For B p is a necessary condition for RB p , but not for p . It may

⁴¹ If this seems counter-intuitive, first think about it in third-person situations. If I know that p is true and that Danielle is rational, that provides some weak support for the hypothesis that Danielle believes that p (at least if p is within Danielle's expertise and domain of interest). Plausibly, we can usually do this in first-person cases as well.

⁴² Perhaps for some agents and for some p , $P(\text{RB}p|p) = 1$, as when the agent is omniscient (God) or when a situation is perfectly luminous, if this is possible (perhaps when I am in sharp pain). But for most p , $P(\text{RB}p|p) < 1$.

be simpler to evaluate the evidential weight of the negation. $P(RBp | \sim Bp) = 0$, but $P(p | \sim Bp) > 0$.⁴³ Yet again, we see that not all evidence has the same weight concerning both p and RBp .

Finally, consider evidence about what knowledge or rationality is. When considering Kp or RBp , such evidence is extremely relevant. It matters whether knowledge requires safety, whether rationality is internal or external, and the like. But these considerations don't seem to matter much (if at all) when evaluating p itself, at least when p does not contain content concerning the nature of knowledge or rational belief. Thus there is a whole class of evidence, evidence about knowledge and rational belief, that is relevant for Kp and RBp and yet not for p . This might be one reason we intuitively feel that reflective knowledge and rational belief are more difficult than mere knowledge and mere rational belief. Reflective knowledge and rational belief require that we take a stance not only about some proposition p , but the nature of knowledge and rational belief itself.

It is possible that, even though p and RBp have individual pieces of evidence that are different or that bear different weights, it always works out that the whole evidence for p has the same weight as the whole evidence for RBp . But although possible, it stretches credulity to think it is actually so. What a great coincidence it would be if neither having the same evidence nor the same support-relations between evidence and hypothesis, p and RBp had the same amount of evidential support for all p ! Accordingly, the argument from different evidence gives us good reason to suppose that rationality is anti-luminous.

3.5 Neither Confusion nor Isolation

The regress argument and argument from different evidence provide reasons to think that rationality is anti-luminous. However, they don't give us a reason to suppose that one can rationally believe a hypothesis while disbelieving that one rationally believes that

⁴³ Once again, there may be some cases in which the $P(p | \sim Bp) = 0$, but these will be rare.

hypothesis. Some have defended that stronger claim,⁴⁴ but here we can be content with the following, more modest conclusion:

- I) Rational belief that p does not entail that, even upon reflection, one also believes that one's belief that p is rational.

Principle (I) keeps us safe from falling into what Alston calls a "level confusion"⁴⁵ and is entailed by the arguments presented for the anti-luminosity of rationality. However, we might be concerned not only about level confusion, but what Feldman calls "level isolation." One's second-order beliefs are relevant to one's first-order beliefs. In this spirit, let's grant Feldman principle (II):

- II) Rational belief that p is incompatible with belief that one's belief that p is irrational.

If (I) and (II) are assumed, it should be fairly clear that one can rationally suspend certain second-order beliefs about one's beliefs without giving up their rationality. An agent reflecting on the rationality of her belief that p will

- a) believe that her belief that p is rational,
- b) believe that her belief that p is irrational, or
- c) suspend judgment about whether her belief that p is irrational.

Assumption (I) entails that, in addition to (a), either (b) or (c) is compatible with rational belief that p . (II) requires that (b) is incompatible with rational belief that p . Accordingly, (c) is compatible with rational belief that p . Given (I) and (II), one can, even upon reflection, suspend judgment about whether or not one's belief that p is rational while still maintaining rational belief that p . Presumably, something similar is true in the case of credences. One can rationally have a high credence that p while maintaining a middling credence about whether having a high credence that p is rational.

⁴⁴ See especially Maria Lasonen-Aarnio's 2014 article, "Higher-Order Evidence and the Limits of defeat," *Philosophy and Phenomenological Review*, where her argument entails that (II) is false. I don't intend to take a stand on (II) here. My immediate claim is that even if (II) is true, extreme (RP) and first-order steadfastness are still viable views.

⁴⁵ Alston, William (1980). "Level Confusions in Epistemology," *Midwest Studies in Philosophy*, 5: 135–50.

It is consistent with what we have said so far to insist that being agnostic about whether one's belief is rational makes a difference to the overall first-order credence one ought to have toward a proposition. For instance, consider the following plausible principle:

PP) All else being equal, if *S* has a middling credence about whether or not her credence that *p* is rational, her credence that *p* ought to be lower than it would have been if *S* had a high credence that her belief that *p* was rational.

(PP) may well be true, so long as the credence need not be so much lower as to make a high credence that *p* impossible. By endorsing (II) and perhaps also (PP), the permissivist avoids the complete "isolation of levels"⁴⁶ against which Feldman rightly cautions us; whereas by accepting (I), she avoids the equally dangerous pitfall of committing a level confusion.

3.6 Conclusion

The arguments for uniqueness from its best proponents have forced us to draw careful distinctions. Evidence cannot be both for and against some hypothesis *simpliciter*, but it can be for a hypothesis according to one set of standards and against it according to another. A certain amount of luck is present in permissive circumstances, but the permissivist can consistently maintain that such luck is benign. Known permissive situations are the most troubling of all, but even these can be resolved if we make use of level distinctions between first and second-order beliefs. Though sometimes subtle, these distinctions enable the permissivist to avoid the escalating skepticism that faces defenders of uniqueness.

In evaluating both Feldman's principle and cases of known permissive circumstances, we have blurred the lines between revealed and unrevealed disagreements, cementing the intuition that although distinct, these problems are intimately connected. The next section explores the consequences of our discussion of uniqueness to debates about conciliationism and steadfastness.

⁴⁶ Feldman, 308.

4. First-Order Steadfastness

I develop and defend a variant of steadfastness that I call “First-Order Steadfastness.” Disagreement is entirely mediated through higher-order evidence, which is swamped by the first-order evidence in peer disagreements. Although some cases of peer disagreement make it irrational to preserve one’s initial beliefs, it is often the case that all parties of a dispute may maintain their first-order beliefs without significant revision, especially when considering hard questions.

4.1 Conciliationism and Steadfastness

The debate about what to do when confronting disagreement exhibits a number of diverse views, but most can be situated within two broad camps: conciliationism and steadfastness. “Conciliationism” and “steadfastness” do not name particular philosophical positions so much as identify two poles of a spectrum. Those on the conciliationist side of the spectrum believe that when faced with peer disagreement, one always ought to revise one’s initial beliefs (or credences). On theories of steadfastness, it is at least sometimes permissible to maintain one’s belief in the face of peer disagreement.

Within these broad camps, diversity abounds. One prominent version of conciliationism is the equal weights view. According to the equal weights view, one must give equal evidential weight both to one’s own belief and that of a known, epistemic peer. If a theist and an agnostic were to compare their initial credences in theism at, say 97% and 49% respectively, both should update to 73%. According to other versions of conciliationism, one should become agnostic in the face of peer disagreement. On this view, the theist should become agnostic, but the agnostic need not change her view.⁴⁷ Yet other “moderate” conciliationisms insist that one’s belief should always be revised in the face of peer disagreement, but lay down no rule about how much revision is required or whether

⁴⁷ This still qualifies as conciliationism, despite the fact that the agnostic does not have to revise her position, because agnosticism (about the existence of God or any other issue) is not strictly a belief, but abstention from belief. All the beliefs involved are still revised.

such peers ought to be in full agreement at the end of their inquiry. Such moderate conciliationisms can be difficult to distinguish from theories of steadfastness. In the context of this project, “conciliationism” will always refer to something more demanding unless otherwise stated.

Proponents of steadfastness may be divided into two further subgroups. Moderate versions of steadfastness only claim that in a given dispute, one party is rationally permitted to maintain their initial belief. According to uniqueness-friendly versions of the “right reasons” view, for instance, the party that interpreted the evidence correctly need not conciliate, whereas all others should revise their beliefs to match the group with the right reasons.⁴⁸ In contrast, extreme steadfastness is committed to the claim that sometimes several parties in a dispute may rationally stand their ground. Kelly’s total evidence view, for example, is usually presented as a version of extreme steadfastness.⁴⁹

I defend a version of extreme steadfastness, whereby multiple parties in a dispute may preserve their first-order beliefs without significant revision in the face of peer disagreement. However, I also try to account for conciliationist intuitions by showing how disagreement often makes an important difference at the second-order level. Sometimes, the change in second-order beliefs constitutes a doxastic defeater for first-order beliefs, but they often leave the rationality of the first-order beliefs (almost) completely intact.

4.2 Does Uniqueness Matter?

Debates about uniqueness are interesting in their own right. Discovering an answer to uniqueness means discovering whether unrevealed, rational disagreements are possible.

⁴⁸ If there are multiple sets of “right reasons,” the view becomes more complicated. On moderate versions of the right reasons view, when multiple parties have distinct sets of right reasons, they must conciliate. Steadfastness is only permissible when one’s party alone has arrived at a rational conclusion. On extreme versions of the view, when multiple parties have distinct sets of right reasons, only parties without right reasons must revise their beliefs.

⁴⁹ Kelly explicitly endorses extreme steadfastness, although he claims that the total evidence view is consistent with both moderate and extreme steadfastness. See “Peer Disagreement and Higher Order Evidence,”

However, apart from its intrinsic interest, we might wonder whether and to what extent an answer about unrevealed disagreements is useful for determining an answer about revealed ones.

Kelly argues that the conciliationist is committed to uniqueness. Imagine that we find ourselves disputing in a recognized, epistemically permissive situation in which I have a higher credence that hypothesis H is true than you do:

According to The Equal Weight View [or most conciliatory views for that matter], you are rationally required to increase your credence while I am rationally required to decrease mine. But that seems wrong. After all, *ex hypothesi*, the opinion that I hold about H is within the range of perfectly reasonable opinion, as is the opinion that you hold. Moreover, both of us have recognized this all along. Why then would we be rationally required to change?

One sympathetic to The Equal Weight View might attempt to heroically defend the idea that you and I are rationally required to revise our original credences in these circumstances. However, a more promising line of resistance, I think, is to deny that [such a case] is possible at all. That is, an adherent of The Equal Weight View should endorse the uniqueness thesis.⁵⁰

Kelly overestimates the force of this argument. The permissive conciliationist may insist that, although the circumstances were permissive *before* the disagreement was revealed, awareness of the dispute requires that the disputants conciliate to a common position. Awareness of a dispute is plausibly a principled reason to conciliate.

Nevertheless, permissivism does place an additional burden on the conciliationist. If uniqueness is false, the conciliationist's explanation for why peers must conciliate in revealed disagreements cannot merely be that they have the same evidence. There must be something special about the evidence that a peer disagrees with you—something not true of evidence in general—that mandates convergence toward a common position. Accordingly, though permissivism is consistent with conciliationism, it does undercut one natural motivation for the view.

The debate about uniqueness also matters when comparing moderate and extreme versions of steadfastness. Suppose uniqueness is true. If two agents steadfastly disagree

⁵⁰ Kelly, "Peer Disagreement and Higher Order Evidence," 119.

about p , at least *one* of them is being irrational; thus, at least *one* of them ought to revise her belief significantly. After all, if she does not revise her belief, she will continue to be irrational. Extreme steadfastness, then, is committed to extreme permissivism.

An answer to the question of uniqueness does not entail an answer to questions about conciliationism and steadfastness. It does, however, have a strong impact on the dialectical alternatives available to us. The truth of (RUE) would have explained why strong versions of conciliationism demand that disputants arrive at one, common view; moreover, it would have ruled out extreme steadfastness as a live option. Having established extreme permissivism in the preceding chapters, we can approach steadfastness more favorably.

4.3 Kelly on Higher-Order Evidence: The Perils of Double Counting

A central issue in the debate concerning disagreement is whether and in what way discovering that an epistemic peer disagrees with you constitutes higher-order evidence.

The following excerpt from Kelly is worth quoting:

Consider the epistemic proposition that
(1) E is good evidence that H is true
...[I]f I discover that you believe that H on the basis of E, I should treat this discovery as confirming evidence for (1). Should I also treat it as confirming evidence for H itself? If I discover instead that you believe that not-H on the basis of E, this discovery would constitute disconfirming evidence for (1). Would it also constitute evidence against H?⁵¹

Though he ultimately refrains from committing to an answer, Kelly presents one argument that might incline us toward answering “no.” First, imagine you are weighing your evidence regarding some hypothesis H, and you come, on the basis of that evidence, to assign a high credence to H. Your friend, seeing that you have done this says, “Ah, well now you have more evidence for H. For *you* believe that it is likely that H, and you are a generally rational person who has seriously considered the evidence. Given that you know that you believe

⁵¹ Kelly, “The Epistemic Significance of Disagreement,” 186-87

H is likely, you should raise your credence even more.” This seems irrational. As Kelly points out, no one would list her own belief that p as a reason why she should believe that p . If rational, the belief that p should be based on the evidence. But then counting the resultant belief as evidence, in addition to the evidence that motivated the belief in the first place, seems to count the original evidence twice over. It is evidential double counting.

This is a problem for conciliationism. One of the guiding intuitions for conciliationism is that you should give equal weight to the belief of an epistemic peer as you do your own. But then if you should give your own belief *no* evidential weight, this intuition won’t yield conciliationism but steadfastness!

Jonathan Matheson tries to defend conciliationism against Kelly’s objection. Matheson notes that we often give abridged evidential stories when asked to share our reasons for believing a proposition.⁵² This is certainly true. If I were asked why I thought the defendant murdered the victim of a crime, I might answer simply “Because I saw the smoking gun in his hand!” Jurors would no doubt find it dull and patronizing if I were also to explain that my reason for belief was partially based on the evidence that guns are weapons, that a smoking gun usually indicates that the gun has been recently fired, or that solipsism is false. It is (usually) pointless to restate evidence that everyone already knows. And if you are asking me *why* I hold a certain belief, you already know *that* I hold it.

Despite these plausible observations, Matheson’s position is difficult to maintain. The difficulty is brought into clear focus when we remove the other evidence from consideration. Imagine that your friend Chloe tells you that she believes some hypothesis H. Upon being asked what reasons she has for this view, she simply says, “because I believe that H.” You press her a little. Perhaps she means that H appears true to her upon reflection, or that she has an intuition that H is the case. Does she have any sort of *seeming* about p

⁵² Matheson, Jonathon (2009) “Conciliatory Views of Disagreement and Higher-Order Evidence,” *Episteme* 6 (3): 272

that would satisfy even the minimal constraints of phenomenological conservatism? No. She is insistent. The only reason she believes H is her belief that H itself.⁵³ We can assume for this story that her belief that H is not self-evident, analytic, or in any other way properly basic. This seems like an especially appalling kind of circularity. Not only do we think Chloe's belief that H is not *good* evidence for H, we don't think it is evidence *at all*.

Perhaps Matheson has a response to this objection as well. Matheson concedes that an agent's belief that *p* is not evidence for *p* in and of itself. However, he contends that taken together with the fact that the agent is reasonable and tends to have true beliefs, it does constitute evidence.⁵⁴ This is especially plausible when we do not have access to the other person's first-order evidence (or, as in the case above, her lack thereof). After all, her belief is evidence that there is evidence for *p*. Of course, in the odd case where she does not actually have good evidence for *p*, such evidence is prone to be misleading. But misleading evidence is evidence nonetheless; and plausibly, evidence that there is evidence for *p* is itself evidence for *p*.

However, this is *not* the case when evaluating the evidential support of our own beliefs. After all, when evaluating our own beliefs, we *do* have direct access to the evidence upon which our belief is actually based. The evidential force of the underlying evidence upon which the belief is based swamps any evidential force one's consequent belief might otherwise contribute. Something similar seems to be happening in our story about Chloe. Even if you think she is *generally* reasonable, since she has told you her full set of evidence (or, rather, her lack thereof) for her belief, you can make a *particular* judgment about the

⁵³ Several people have commented in response to this story that it must be impossible. Could someone really believe a hypothesis for no other reason than that she believes it? I think, however, that if there is any impossibility of the kind some people find in this story, it must be psychological and not metaphysical. Hopefully, at least on our better days, none of us is even capable of making the sort of blunder that Chloe does. But it is not so difficult to conceive of someone less psychologically restrained committing the same mistake.

⁵⁴ Matheson, 273.

individual case. Because you judge correctly that the evidence does not support her conclusion, her belief gives you no reason to believe that H.

David Christensen offers a different response to the double counting objection than does Matheson. What Kelly's question has shown, Christiansen insists, is that higher-order evidence has the "somewhat unusual feature ... [that] its evidential bearing is often relative to the thinker."⁵⁵ Although my belief that p bears no evidential weight concerning p for me, it may nonetheless bear significant evidential weight for you.

Christensen must posit the agent-relativity of evidence to explain why one's own belief in a hypothesis is usually of no evidential value to oneself for that hypothesis, but often of great evidential value for others when evaluating that same hypothesis. However, if what we have said so far is correct, there is a much more natural explanation than stipulating this kind of evidential relativity: an agent has access to her own first-order evidence for a hypothesis.⁵⁶ Evidence of evidence is not useful if one already has the evidence that the evidence of evidence makes evident. However, one does not usually have access to all of another person's first-order evidence for belief.

Another advantage of this explanation over Christensen's is that it explains why we need not afford any evidential weight to Chloe's belief once she has disclosed her reasons for so believing. We have access to all of her first-order reasons. We may evaluate those directly rather than relying on the second-order evidence. With such a feasible, alternative explanation available to us, we need not stipulate that evidence has the "unusual feature" of thinker-relativity.

In cases of known peer disagreement, both agents have the same evidence. Accordingly, both agents have access to the underlying evidence that forms the basis of her own and her disputant's respective beliefs. Since knowledge of the underlying evidence

⁵⁵ Christensen, David. (2010) "Higher-Order Evidence," *Philosophy and Phenomenological Research* 81 (1): pp. 189.

⁵⁶ By saying that an agent has access to her first-order evidence, I do not mean that she knows that she has it, simply that she has it.

swamps the evidential force that the second-order evidence of disagreement might otherwise have produced, peer disagreement does not directly bear evidential weight against one's initial position. Just as one's own belief concerning p provides no immediate evidential support for or against p , neither does a peer's belief concerning p provide any immediate evidential support for or against p .⁵⁷

Kelly stops short of fully defending the thesis that peer disagreement carries no direct, evidential weight for our first-order beliefs. Instead he hastens to defend the claim that even if evidence of disagreement bears directly on first-order beliefs, one need not always conciliate. The evidential weight of disagreement need not swamp the weight of the total evidence.

In what follows, I construct and defend first-order steadfastness, which claims that peer disagreement has no immediate bearing on first-order beliefs. The view is complementary to, rather than in competition with, Kelly's total evidence view. By thinking more specifically about the interaction of epistemic levels, we can come to a clearer understanding of why disagreement sometimes acts as a defeater for our beliefs and why it sometimes leaves them (nearly) unscathed.

4.4 Evidence of Evidence and Fun with Marbles

Here is a very natural thought: evidence of evidence is evidence. In fact, if evidence of evidence for p were not often evidence for p , it would be difficult to see why disagreement about p should influence our beliefs about p at all. For disagreement constitutes second-order evidence that someone else considers some body of evidence to support a conclusion contrary to one's own.

⁵⁷ The language of "direct" or "immediate" evidential support is essential. As we shall see, even if peer disagreement does not directly bear on our first-order beliefs, it sometimes indirectly constitutes a doxastic defeater via a change in our second-order beliefs.

The following examples may initially seem to support conciliationism, but they actually reveal a truth about second-order evidence that is hard to square with most conciliationist theories. Suppose you and I each reach our hands into an urn filled with black and red marbles, and each of us pulls out 10. I draw 6 black marbles and 4 red marbles. I may not have any full-fledged belief about what percentage of marbles are black. I may not have a large enough sample size. However, if I were forced to guess, the rational thing would be to estimate that 60% of the marbles are black.⁵⁸ You also draw 10 marbles, 4 of which are black and 6 of which are red. You estimate that 40% of the marbles are black. I find out that you have drawn 10 marbles, and you tell me that you estimate 40% of the marbles are black. If I know that you are generally rational, I now have a good reason to recalibrate my estimation to 50%. And if you know that I am generally rational, you have good reason to recalibrate your estimation to 50% as well. If we are really convinced of each other's honesty and rationality, we split the difference and give the other person's opinion an equal weight in our calculations.

In a slightly different situation, you draw half as many marbles (5) and estimate that 40% are black. Because you have half the evidence, I assign half the weight to your estimate, resulting in an updated estimate that 53.33% are black. I give half of the weight to your estimate because it is based on half the amount of the evidence. If instead of drawing 5 you had drawn 20 marbles, I would have given twice as much weight to your estimate as compared to mine.

On the basis of these ordinary observations, we might be tempted to some principle like the following:

(EE): All else being equal, If E2 is evidence of E1 bearing on hypothesis H, then the evidential weight that an epistemic agent *S* rationally ought to assign E2 should be proportional to the projected evidential weight of E1.

⁵⁸ It would be more accurate to say that one should endorse a spread of credences for the expected percentage that center on 60%. If one were confident that the correct percentage was exactly 60%, one would be committed to the claim that the total number of marbles in the urn had a factor of 5. And it does not seem that anyone can rationally be confident of this.

Something like (EE) seems to be doing some of the motivational work behind conciliatory theories. Why should you conciliate with a peer? In no small part, it is because you have the same amount of evidence. However, (EE) is false. (EE) gets (at least) two things wrong. First, it gets the wrong results in cases of partially shared evidence.⁵⁹ Second, it has the false consequence that disagreement can never be confirming evidence, nor agreement disconfirming evidence.

Note that in the cases above, no first-order evidence is actually shared. Neither of us has revealed our marbles, and all I have from you is the second-order evidence of how many marbles you estimate are black. Let's change the example a bit so that some, but not all, evidence is shared. An onlooker draws five marbles and shows them to both of us: 4 black and 1 red. I then draw five more that you cannot see (2 black and 3 red), and you do likewise (4 black and 1 red). Before comparing estimates, I guess that 60% are black, whereas you say 80%. Then we reveal our estimates. Your estimate is based on having seen 10 marbles, as is mine. On this basis, (EE) mandates that I give equal weight to both of our estimates, so that I adjust my estimate to 70% being black.

But this is not a rational thing to do. Though we each base our initial estimates on having seen 10 marbles, each has only seen five marbles that the other has not. Both of us accounted for the 5 marbles that were drawn in our estimates, and we must be cautious not to double count the overlapping evidence. Believing you to be rational, I will infer that since you estimated 80% of the marbles were black after already having seen 4 black and 1 red marbles, it is very likely that, of the five you drew, 4 more were black and 1 was red. When I add both of these to the 5 that I drew (2 black and 3 red), I will infer that 10 of the 15 were black. You, through similar calculations, will arrive at the same conclusion, and each of us will modify our estimates to 66.67% of the marbles being black.

⁵⁹ By a part of one's evidence, I mean an improper part. If two people have all the same evidence, they have (improperly) partially shared evidence.

One curious feature in cases of partially shared evidence is that agreement can constitute disconfirming evidence and disagreement confirming evidence. Suppose that the first 5 marbles drawn—the ones we both can see—are all black. I then draw 1 black and 4 red marbles, rationally estimating that 60% of the marbles are black. (I tell you nothing about my estimate.) You also draw 5 marbles and reveal your estimate that the urn is 80% filled with black marbles. If I were to split the difference, I would estimate that 70% of the marbles are black. However, this isn't what I ought to do. To do so would be to double count the 5 black marbles that were shown to both of us. If you think that 80% of the marbles are black, I can infer that, in addition to the 5 black marbles we both see, you probably drew 3 blacks and 2 reds: 60% black, which matches the ratio of marbles that I have already counted. In fact, this is an odd—though by no means anomalous— case in which your disagreement is confirming evidence for my original position! I should continue to estimate that 60% of the marbles are black.

Imagine that the situation was exactly as before except that, of the five marbles you drew from the urn, 1 was black and 4 were red. You would rationally estimate that 60% of the marbles were black. I had originally thought that 60% of the marbles were black as well, but in this odd case, your agreement counts as disconfirming evidence for my original estimate. Since I will not double count the five black marbles drawn first, I should estimate the percentage of black marbles at 46.67%.

These examples of partially shared evidence should force us to reevaluate (EE). In every case, each of us has seen 10 marbles; but, I know that you have seen some of the same marbles that I have, and I know which ones those are. First I figure out what it is likely your estimate would have been had we not shared any marbles. Then I weigh *that* estimate the same amount that I would have if you had only drawn 5. In other words, we ought to replace (EE) with this:

(EE*): All else being equal, If E2 is evidence of E1 bearing on hypothesis H, then the evidential weight that an epistemic agent *S* rationally ought to assign E2 should

be proportional to the difference between the projected evidential weight of E1 and the weight of the subset of E1 to which *S* has immediate access.

(EE*) is bad news for conciliationism. In the case of known peer disagreement, both agents *ex hypothesi* have access to exactly the same evidence. Since peers have the same evidence, they have the same first-order evidence. One peer's disagreement constitutes second-order evidence that the first-order evidence (E1) supports a contrary hypothesis ($\sim H$). According to (EE*), the amount of weight the disputant should assign this second-order evidence is proportional to the difference between the evidential weight of E1 and the weight of the subset of E1 to which she has access. But the disputants both have access to all of E1, for they are epistemic peers. Thus the difference between the evidential weight of E1 and the weight of the subset of E1 to which she has access is zero. In cases of disagreement between exact epistemic peers, one ought to assign *no* evidential weight to the disagreement of a peer. In fact, contrary to much that is presupposed in the literature, (EE*) suggests that we ought to be more inclined to significantly adjust our first-order beliefs when a disputant's evidence differs dramatically from our own.

Admittedly, these marble cases are just toy models. They are cases in which it is clear what counts as evidence, how much evidence my disputant and I have, how much of our evidence is overlapping, which parts of our evidence are overlapping, and whether my disputant is generally rational. Real world disagreements are almost never so simple. Nevertheless, these simple cases are a good place to begin our theorizing about disagreement. If one need not always conciliate with a peer even in the simple cases, it is unlikely that a consistent requirement to do so will emerge in the complicated ones.

4.5 Higher-Order Credences: Under Pressure

One might think that in arguing that second-order evidence is of no evidential weight in peer disagreements, I have surely bitten off more than I can chew. For isn't it obvious that in some cases of peer disagreement, second-order evidence *can* rationally affect our

beliefs? Conciliationists have proposed cases in which allowing oneself to be moved by the second-order evidence seems not only permissible but mandatory. Perhaps the most famous of these is Christensen's check-splitting case. Christensen asks us to imagine that you and I have each tried to sum a check and to divide it equally between ourselves after a meal. I initially conclude that each of us owes \$43, whereas you conclude that we each owe \$45. How, if at all, should we adjust our beliefs? Christensen makes the following judgment:

I think that if we set the case up right, the answer is obvious. Let us suppose that my friend and I have a long history of eating out together and dividing the check in our heads, and that we've been equally successful in our arithmetic efforts: the vast majority of times, we agree: but when we disagree, she's right as often as I am. So for the sort of epistemic endeavor under consideration, we are clearly peers. Suppose further that there is no special reason to think one of us particularly dull or sharp this evening—neither is especially tired or energetic, and neither has had significantly more wine or coffee. And suppose that I didn't feel more or less confident than usual in this particular calculation, and my friend reports that she didn't either. 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.⁶⁰

Any steadfast account of disagreement worth its salt must make sense of Christensen's check-splitting cases and the real pressure that disagreement bears on our epistemic decisions. This may be done in chiefly one of two ways: (1) by arguing that Christensen's judgment of the particular case is mistaken, or (2) by arguing that examples like the check-splitting case do not sufficiently generalize to motivate a robust conciliationism. In this section, I attempt a response via strategy (1). In many cases of peer disagreement, the appropriate response does not involve substantially reducing one's first-order credence, but only one's second-order credences. Though I will ultimately concede that this response might not be satisfactory for Christensen's check-splitting case, the framework sufficiently models a rational approach for many peer disagreements. In the next section, I adopt strategy (2) and argue that even if one ought to become agnostic about one's first-order

⁶⁰ Christensen, David. (2007) "Epistemology of Disagreement: The Good News," *Philosophical Review* 116:2): pp. 193.

beliefs in check-splitting scenarios, we need not conciliate over most of our religious and philosophical beliefs.

Earlier, we concluded that peer disagreement does not provide direct evidential weight for first-order beliefs. However, disagreement can provide significant evidential weight for higher-order claims about the evidence. What is important when evaluating p is the evidence itself. Evidence about the evidence is significant, but only insofar as it gives indirect access to the evidence. The indirect nature of this access is the very quality that makes it possible for the weight of higher-order evidence to be swamped when all the first-order evidence is shared. But when evaluating claims about the evidence for p , what was indirect evidence about the evidence for p becomes direct evidence for higher-order propositions about p . Accordingly, even when I have direct access to E , someone else's belief that p on the basis of E gives me an additional reason to believe second-order claims such as, "E is good evidence for p ," even if it gives me no additional evidential weight with respect to p .

Suppose someone who believes that p on the basis of E encounters peer disagreement. She continues to believe that p , but significantly reduces her credence in the proposition "E is good evidence for p ." In other words, suppose someone takes the advice of this project seriously. Would such a person end up with a consistent set of beliefs? Or has she acted irrationally by maintaining a high credence that p while significantly lowering her credence that E supports p ?

Once again, when evaluating the consistency of doxastic attitudes on different epistemic levels, we must avoid the twin errors of level confusion and level isolation.

Recall the twin theses we postulated in an earlier section:

- I) Rational belief that p does not entail that, even upon reflection, one also believes that one's belief that p is rational.
- II) Rational belief that p is incompatible with belief that one's belief that p is irrational.

But what is rationally believing that p other than believing that p on the basis of the evidence? And what is believing that one's belief that p is rational other than believing that one has evidence that adequately supports p ? Accordingly, if (I) and (II) are true, we may also assert the following:

- III) Even if p requires evidential support, then rational belief that p does not entail that, even upon reflection, one also believes that one's evidence supports p .
- IV) Even if p requires evidential support, then rational belief that p is incompatible with the belief that one's evidence does not support p .

The antecedent "if p requires evidential support" accounts for the possibility of beliefs that are, in Chisholm's terminology, directly evident. In such a case, a belief might have no evidential support and yet be perfectly rational. Any belief that is not directly evident requires evidential support. Thus qualified, we may happily endorse (III) and (IV) for the same reasons we endorsed (I) and (II).⁶¹

(I) and (II) led us to the conclusion that one can, even upon reflection, suspend belief about whether or not one's belief that p is rational while still maintaining rational belief that p . (III) and (IV) likewise lead us to the conclusion that, even upon reflection, one can rationally believe that p without believing that one's evidence supports p . Presumably, what holds for beliefs and abstention from belief holds for high and middling credences respectively. One can rationally maintain a high credence that p while having a middling credence about whether one has good evidence for p ; thus, we avoid level confusion. But a high credence that p is incompatible with a low credence that the evidence supports p ; thus, we avoid level isolation.

⁶¹ I have deliberately skirted an issue here about whether "having evidence" is sufficient and not merely necessary for the rationality of inferential beliefs. Plausibly, rational belief requires not only having the right evidence, but also believing on the basis of that evidence. e.g. If I have good evidence that Susan has cheated on her spouse, but the reason I believe this is because I am sexist and think all women cheat, we do not want to say that my belief is rational—despite the fact that I do in fact have good evidence that she has cheated. Careful readers will no doubt want to modify (III) and (IV) to include this feature. Nevertheless, I have systematically ignored this aspect of rational belief, not because I believe it to be unimportant, but because the criterion of believing on the basis of the evidence does not seem to play a decisive role in debates about disagreement.

Here are two different ways to think about defeat. On one picture, defeat is epistemically discrete. This picture is most natural when we think in terms of coarse-grained beliefs. Being agnostic about whether or not one's belief that p is rational is compatible with rational belief that p . Disbelieving the same is not. There is nothing "in-between" being rationally permitted to believe p and being rationally forbidden. On another picture, defeat is epistemically continuous. This picture is most natural when we think in terms of fine-grained credences. A relatively high credence that p may be permissible when the agent has a middling credence about whether assigning a high credence is rational, even if it isn't permissible when the second-order credence is low. On this conception, we might sensibly ask how high of a credence is permissible if the agent's second-order evaluation is between low and middling levels. A plausible, if not very informative, answer might be, "higher than the highest credence permissible with a low second-order credence, and lower than the highest credence permissible with a middling second-order credence." This answer would be consistent with a principle we considered earlier in the paper:

PP) All else being equal, if S has a middling credence about whether or not her credence that p is rational, her credence that p ought to be lower than it would have been if S had a high credence that her belief that p was rational.

If defeat is continuous, (PP) is just one instance of a more general feature of defeat we may call "Continuity":

(Continuity): All else being equal, for all n such that $0 < n < 1$, if S has a credence of n that her credence that p is rational, then S 's credence that p ought to be higher than it would have been if her credence that p is rational had been lower than n , and S 's credence that p ought to be lower than it would have been if her credence that p is rational had been higher than n .

If epistemic defeat is continuous, how might that affect the epistemology of disagreement?

If epistemic defeat is continuous, our picture becomes more complicated. A middling second-order credence will still be compatible with a relatively high first-order credence, but perhaps not with the very highest first-order credences. Why think this? If defeat is

continuous, it seems arbitrary to stipulate that there is some exact second-order credence less than 1 at which the second-order credence starts to become a defeater for one's first-order beliefs. (By "defeater" in this context, I merely mean something that rationally requires one to lower one's credence in a proposition, even if slightly. A belief being defeated in this sense does not require agnosticism.) If defeat is continuous, then disagreement may be a minor defeater for first-order beliefs in virtue of lowering the second-order credence to a middling level. When a first-order belief had previously been on the border between being a high or middling credence, this may be enough defeat to make a high first-order credence irrational. Nevertheless, if the initial first-order credence is high enough, there is no reason to think that a middling second-order credence will rationally obligate the agent to lower her confidence outside of the high range.

These cases of borderline belief aside, can peer disagreement never be a defeater for one's first-order beliefs? In one-on-one peer disagreements, it is natural to think that disagreement should at most reduce one's credence in the rationality of one's belief to .5, and so will not be a significant defeater for one's first-order belief. However, sometimes disagreement will make it rational to decrease one's second-order belief much lower than .5. Take a variant of Christensen's check-splitting case, for instance, in which I estimate that the sum of the tab is \$43, but three of my peers each independently conclude that the sum is \$45. The probability that three of my peers independently arrived at the same, wrong answer is so low that I should downgrade my credence in the rationality of my original answer substantially below .5. In that case, I would have a significant defeater for my first-order belief.

However, let us return to the one-on-one version of Christensen's case. Christensen first claims, "that I should lower my confidence that my share is \$43 and raise my confidence that it's \$45."⁶² So far, our accounts predict the same result. On the continuous

⁶² Christensen, 193.

conception of defeat, at least, if lowering one's second-order credence that one's evidence supports a given proposition, one also ought to somewhat reduce one's first-order credence that the proposition is true. But Christensen goes further, claiming, "that I should now accord these two hypotheses roughly equal credence." In this, our accounts diverge. Even if peer disagreement obligates one to suspend judgment at the second-order level, one may steadfastly retain one's first-order beliefs, so long as the first-order belief was permissible in the first place.⁶³

But does this response match our actual intuitions in check-splitting cases? Most who have written in the disagreement literature seem to share Christensen's judgment of the particular case. Here is one reason to agree. Suppose you rationally retain your belief that the sum is \$43. Assuming that the belief is true, this belief will usually amount to knowledge, as peer disagreement does not seem to present a challenge for knowledge except insofar as it presents a challenge for justification or rationality. Plausibly, when one knows that p , it is pragmatically rational to act as if p , all things being equal. If one knows that the sum is \$43, then one can rationally leave \$43 on the table and leave the restaurant. But this response seems inappropriate in check-splitting cases. Isn't the obviously rational choice to double-check the sum? And why would one need to do this if one already knows what the sum is?

The steadfast theorist is not without a response. Perhaps knowledge that the sum is \$43 does give you a pragmatic reason to pay \$43 and leave, but that reason is defeasible. It would be rude after all to leave without paying what your friend believes is your fair share. Rechecking the sum even though you already know the answer is a small price to pay to avoid insulting your friend. Perhaps the intuition that we ought to suspend belief is tracking

⁶³ Of course, in check-splitting cases, at least one person has arrived at the answer irrationally. Whoever has tabulated the sum incorrectly ought to revise her belief; but the reason she ought to revise her belief has nothing to do with the fact that she encountered peer disagreement. She ought to revise because she did her math incorrectly.

social rather than epistemic norms. Alternatively, perhaps this is just a case in which our intuitions have got the matter wrong. Our theories should account for our intuitive judgments about as many cases as possible, but we should also be willing to revise our judgments about cases in the light of our best theories.

However, these responses on behalf of steadfastness are unnecessary. Even if Christensen's judgment of check-splitting cases is correct, such examples do not generalize to most cases of religious or philosophical disagreements.

4.6 Easy and Hard Questions: Why Difficulty Matters

All that we have said so far is that nothing about the nature of first and second-order doxastic attitudes requires that one cannot have a high first-order credence that p while simultaneously having a middling second-order credence that one's first-order credence that p is rational. However, this does not mean that for every proposition, one's first and second-order beliefs can be quite so disparate. Sometimes, in particular cases, we have additional *level-connecting* information available to us. That is, sometimes we have additional beliefs that further constrain the permissible combinations of lower and higher-order credences. I argue that the ease or difficulty of determining the truth of a proposition often functions as level-constraining data. The gap in difficulty between Christensen's check-splitting cases and our typical religious and philosophical beliefs creates an important disanalogy that prevents our intuitions about check-splitting cases from motivating a robust conciliationism.

Let us distinguish between easy and hard questions. The more probable it is that I rationally believe the correct answer (upon due consideration), the easier the question is. The less probable it is that I rationally believe what is correct, the more difficult the question.

Easy: $P(\text{RB}p|p) = \text{high}$ ⁶⁴

⁶⁴ "RB" is an operator for "rationally believes that."

Hard: $P(\text{RB}p|p) = \text{low}$ ⁶⁵

Ease and difficulty are agent and kind-relative. An equation that is somewhat hard for me to solve may be easy for a mathematician.^{66 67}

Suppose that Emma is considering a question that she believes is easy for her. She has a high credence that p , but only a middling credence about whether she rationally believes that p . Let's fill in some numbers to make her individual case easier to evaluate.

Emma assigns the following subjective probabilities:

$$P(\text{RB}p|p) = 0.9$$

$$P(p) = 0.9$$

$$P(\text{RB}p) = 0.45$$

It should be apparent that Emma's credences are inconsistent. For if $P(p) = 0.9$ and $P(\text{RB}p) = 0.45$, then the highest value Emma can assign $P(\text{RB}p|p) = 0.5$, which contradicts her assessment that the question is easy for her, as expressed by her subjective probability, $P(\text{RB}p|p) = 0.9$.⁶⁸

Emma is now faced with a choice. She may lower her credence that p . She may raise her credence that she rationally believes that p . Or she may conclude that, in fact, the question was more difficult than she had previously thought. Or she may do some combination of these. In any case, Emma must substantially revise her beliefs. If she

⁶⁵ This may require some modification. Some easy questions might be incorrectly classified as hard on this schema if the probability that an agent rationally believes that p is low primarily because it is unlikely that the agent will consider whether p . The question of how much change I have in my pockets (£2, let us say) is an easy question. But the probability that I will rationally believe this is low, because it is unlikely I will care to check the change in my pocket at any given time. We could amend this by slightly altering the formula. (e.g. Easy: $P(\text{RB}p|p \wedge Cp) = \text{low}$, where " Cp " means that the agent has considered whether p). However, in the strong disagreements of commitment that we are most interested in, we can take it for granted that both agents have considered whether p .

⁶⁶ Importantly, one thing that might determine how difficult a question is for some agent may depend on the amount or quality of evidence available to her.

⁶⁷ Really, there is a whole family of nearby concepts relating to ease and difficulty. For instance, we could evaluate $P(\text{K}p|p)$ or $P(\text{B}p|p)$.

⁶⁸ $P(\text{RB}p) = P(\text{RB}p|p)P(p) + P(\text{RB}p|\sim p)P(\sim p)$. The highest value for $P(\text{RB}p|p)$ is obtained when $P(\text{RB}p|\sim p) = 0$. Thus, when $P(\text{RB}p) = 0.45$ and $P(p) = 0.9$, $0.45 = P(\text{RB}p|p)(0.9) + (0)(0.1)$. Accordingly, $0.45 = P(\text{RB}p|p)(0.9)$. Therefore, $P(\text{RB}p|p) = (0.45) \div (0.9) = 0.5$.

maintains her credence that p and her estimation of how easy the question is, her second-order credence that she rationally believes that p must also be fairly high. Anything less constitutes a defeater for her belief that p .

Contrast Emma's situation with that of Matt, whose situation is identical to Emma's except that he believes that the question of whether p is moderately difficult for him:

$$P(\text{RB}p|p) = 0.5$$

$$P(p) = 0.9$$

$$P(\text{RB}p) = 0.45$$

Matt's credences require no corresponding revision. Because he does not expect the question to be answered as easily, there is a lower threshold for second-order credences to be consistent with his first-order credence that p . Of course, Matt's first-order credence may be defeated in the ordinary way. A high credence that p is still incompatible with a low credence that his belief that p is rational, no matter how hard or easy he estimates the problem is. Nevertheless, assuming Emma and Matt hold fixed their credences that p and their estimation of the difficulty of rationally believing the truth that p , many second-order credences about whether one is rationally believing that p that are available to Matt without downward belief revision are unavailable to Emma.

We have found that the subjective difficulty of a question plays an important role in determining at what point a given second-order credence in the rationality of a belief constitutes a defeater for a high, first-order belief. This ought not to surprise us. There is epistemic dissonance when someone is agnostic about whether her belief is rational and simultaneously believes the question is an easy one.

We are now in a position to see why Christensen's check-splitting cases do not readily generalize so as to motivate a robust conciliationism. Check-splitting is easy. It is elementary math. If " p " expresses the correct sum, then $P(\text{RB}p|p)$ = very high. Therefore, a second-order credence in the rationality of one's belief must also be high if it is not to defeat

the first-order belief. Christensen's judgment is correct, at least insofar as one's initial belief is defeated by disagreement.

However, most religious, moral, and philosophical questions about which we disagree are not easy questions at all.⁶⁹ Most of them are so difficult that they have been debated for millennia. When disagreeing about hard issues, there is nothing wrong, epistemically speaking, with maintaining a high first-order credence and a middling second-order credence simultaneously. Even if we ought to conciliate over easy questions, we may remain steadfast about the hard ones.

4.7 Some Concerns about Assessing Difficulty

Concern (1): If someone has already come to rationally believe that p , then rationally believing that p should then be perfectly easy for her. But if the question is easy for someone, then her first and second-order credences ought to be very closely aligned, and so disagreement ought to have a substantial impact on her first-order beliefs. The easy-hard distinction thus cannot stop Christensen's check-splitting cases from generalizing to motivate a robust conciliationism because once someone arrives at a rational belief, rational belief is easy for her.⁷⁰

Imagine that a mountain climber is scaling Mount Everest. This is a hard task. Granted, it might be relatively easier for a good mountain climber than it is for you or me,

⁶⁹ Some theists and atheists think that the question of God's existence is easy, not hard. I don't have much to say in this limited space about why I think this is (in general) an unhelpful way to think about religious disagreement. However, a few brief points are worth making. If some theists or atheists think that the question is an easy for nearly everyone, then on my account, their beliefs ought to be defeated by sustained religious disagreement. Alternately, it might be that one side tries to offer an explanation for why the question is easy for a select in-group, though hard for outsiders (e.g. that Christians have properly functioning *sensi divinitatis* whereas those outside the church do not, or that those who have grown up in a religious community are blinded by psychological or sociological needs cultivated by the community, etc.). To the extent that a group can coherently and non-circularly develop such an account, they need not treat their disputants as peers. I have abstracted away from these error-theoretic accounts in part because they tend to overlook parties who do not already belong to one camp or another. For instance, suppose there is a good argument for the proper basicity of certain theistic beliefs. At least in the absence of defeaters, certain theists will be rational in holding onto their theism. But this account doesn't straightforwardly have anything to say about whether agnostic conversion to theism could be rational. After all, agnostics don't have a belief that God exists, and so can't have a properly basic belief that God exists.

⁷⁰ I owe awareness of this problem to Tim Mawson.

but it is a hard task nonetheless. On one natural way of thinking, the journey becomes easier along the way. How hard is it to get to the top of Mount Everest? If you start from the base, it is very hard indeed. What if, having started from the base, you make your way part way up. How hard is it *then* to get to the top? The task is still very hard, although less difficult than climbing from the base. As you keep going, the task ahead becomes continuously easier. What about when you are one step away from the top? Or when you are on the top itself? How difficult is the task then? *Very easy.*

Indeed, there is a sense according to which we might say that climbing Mount Everest gets easier as one goes, but this can be misleading. In reality, there are a series of overlapping tasks that have different difficulties. Or, if one prefers, there are overlapping parts of the same task that have different difficulties. But that the final task (or the final part of the task) is easy is not a reason to think that climbing Mount Everest as a whole is easy.

This comes into focus when we compare the relevant probability ascriptions. Let “S” represent a success condition, one’s successfully climbing Mount Everest, and let “A” represent an attempt, one’s attempting to climb Mount Everest. Yes, there is a sense according to which climbing Mount Everest is perfectly easy once one has reached the top. But the probability this is tracking is not $P(S|A)$ but $P(S|A \wedge S)$. Of course, $P(S|A \wedge S) = 1$, but $P(S|A)$ remains much lower.

Analogously, in the case of rational belief, there is a sense according to which it is easy to have rational belief once one has already achieved rational belief. But this is not tracking $P(RBp|p)$ but $P(RBp|p \wedge RBp)$. Only the former subjective probability ascription gives us level-connecting information that impacts the size of the permissible gap between our first and second-order credences.

Concern (2): Ease and difficulty aren’t just about the probability of accomplishing a task given the desirability of the task accomplished. It also matters how much effort the task requires.

We should think of difficulty as evaluable in (at least) two ways: according to the chances of success and according to the amount of effort required to succeed. Though both kinds of difficulty often co-exemplify, neither one entails or is entailed by the other. Imagine that Max is a superb mountain climber. In fact, Max has already scaled Mount Everest a dozen times. Max considers how likely it is, if he tries to scale Mount Everest again, that he will successfully do so. Max figures the odds are pretty good. There is a sense according to which climbing Mount Everest is easy for Max. Nevertheless, climbing Mount Everest still takes a lot of effort. Max will have to work very hard, even though he knows that his hard work will probably lead to success. In this sense, climbing Mount Everest is difficult. For Max, the task is easy with respect to the chances of success but hard with respect to the amount of effort required.

Contrast Max's situation with Hannah's. Hannah has found a secret elevator that goes straight to the top of the mountain. But there is a catch: because of the adverse conditions, the elevator breaks down more than half the time. It is a difficult matter according to the chances of success for Hannah to reach the top. Nevertheless, there is also a sense according to which the task is easy. After all, sitting in the elevator takes no great effort. Just as in Max's case, there can be difficulty of effort without difficulty of success, so in Hannah's case, there can be difficulty of success without difficulty of effort.

All the probability ascriptions for ease and difficulty that I have given should be understood only according to the chances of success, though doubtless many of the hard questions we face in philosophy and religion are difficult in both senses.

Concern (3): This view has the bad result that when two peers are considering a question that is hard for one of them but easy for the other, the one for whom the question is easy ought to conciliate whereas the one for whom the question is hard may remain steadfast.

That *would* be a problem if (3) were possible. However if a question really is significantly easier for one disputant than another, they simply aren't peers.

Concern (4): The ease-difficulty distinction gets the wrong results in a revised version of Christensen's check-splitting case: the hard check-splitting case. Imagine that you and your friend are presented with an extremely long check to split at the restaurant. At some length, the easy check-splitting case would become a difficult one. But we still have the same conciliatory intuitions in hard check-splitting cases as we do in easy ones. Therefore, the ease-difficulty distinction cannot do the work required to support first-order steadfastness.

Even very long sums don't reach the level of difficulty that is present in the theoretical questions in philosophy and theology—or of the theoretical questions in mathematics for that matter. Long sums often require more effort than short sums. But, given that one makes the requisite, attentive effort, is it less likely that one will come to a rational answer to a long sum than a short one? It is only minimally less likely, if it is less likely at all. Even if every sum is more difficult than some sum with one fewer addend, it does not follow that a sum can be as difficult as the hard questions in philosophy and religion. The increase in difficulty for each additional addend is plausibly subject to a law of diminishing returns in such a way that there is a limit to how difficult a sum can be. (Note also that certain kinds of infinite sums can be solved quite easily so that there is no straightforward inference from the length of a sum to its difficulty.) In short, the first-order steadfastness can deny the possibility of a hard check-splitting scenario. Even when we speak of a sum's being hard, we mean that it is hard relative to other sums and not hard relative to the deep, theoretical questions of philosophy and religion.

4.8 Conclusion

This view leaves our religious and philosophical beliefs in a strange place. For it seems that though we are sometimes in a position to rationally believe, and indeed, to know religious

and philosophical truths, not often are we in a position to know whether we know.⁷¹ This is odd. As philosophers and theologians, one of the things we seek is to know the truth about important and fundamental questions. The quest for knowledge is near the heart of the academic life.⁷² And yet, though we may sometimes achieve this goal, rarely can we know whether we have succeeded. In fact, since many people disagree with the conclusions of this paper, I cannot even know whether or not I know my own thesis!

I do not say that we can never know whether we have acquired knowledge in philosophy or theology. Peer disagreement is weighty second-order evidence, but that is not to say that it can never be overridden. A good error theory may undercut it. Distinct evidence about the rationality of the belief may countervail it. Independent convergence of majority opinion can outweigh the disagreement of a few. And sometimes, it is just not clear who one's peers are and, accordingly, in what way someone's disagreement should count as evidence at all. This is especially true for debates in which some of the relevant evidence seems to have been acquired through unrepeatable, personal experience (as when evaluating some religious experiences), or through individual introspection (as in certain debates about the nature of consciousness, the self, or the will). If not incommunicable, facts about such experiences or insights are at least difficult to convey, and make it correspondingly difficult to assess the weight such disagreements should have on one's own judgments.

However, the full development of these countervailing considerations is best left to another project. This paper has sought to show that one may rationally retain belief in peer disagreement even when those special circumstances are not realized. Though peer

⁷¹ This assumes that knowledge entails rational belief.

⁷² Of course this is not the only goal of the academic life. Another, perhaps more important, goal of philosophy is seeking wisdom, and theologians working from within a tradition may be seeking the love of God as much as the knowledge of God. Nevertheless, that knowledge is one of the central goals of philosophers and theologians cannot sensibly be denied.

disagreement often defeats rational belief that one's first-order beliefs are rational, rarely is it a substantial defeater for one's first-order beliefs themselves, especially when they concern the difficult questions of philosophy and religion.

Epilogue: Humble Conviction

We began this project with a personal question: what should *we* do when we confront epistemic equals who disagree with us about our most deeply held religious beliefs? One worry we had was that the appropriately humble response to disagreement might be to conciliate. And so it is appropriate to finish this project by reflecting on the virtue (or lack thereof) following its advice might engender.

Conciliationists often cite as an advantage the intellectual virtues that their view (allegedly) fosters. Christensen terms his view “epistemic modesty,”⁷³ and Feldman claims epistemic humility as a virtuous result of applying conciliationism, noting the particular need for humility and tolerance in religious discourse.⁷⁴ Such humility is not *only* important in religious cases; however, since religious histories are too frequently blemished with intolerance, it is especially appropriate to reflect on the virtue of intellectual humility in this context.

Do uniqueness and conciliationism really promote epistemic humility? In “Reasonable Religious Disagreement,” Feldman considers several clearly intolerant attitudes in religious discourse of the garden variety for which, unfortunately, one does not have to look very hard to find. Admitting that his own view may seem a bitter pill to swallow, Feldman suggests that, nevertheless, it is a virtuous prescription, and that the bitterness is sweetened by the view’s rejection of religious intolerance:

This may see [*sic*] to be a distressing conclusion. It implies that many of your deeply held convictions are not justified. Worse, it implies that many of my deeply held, well-considered beliefs are not justified. Still, I think that this is the truth of the matter. And perhaps the conclusion is not so distressing. It calls for a kind of humility in response to the hard questions about which people so often find themselves in disagreement. It requires us to admit that we really do not know what the truth is in these cases. When compared to the intolerant views with which we began, this is a refreshing outcome.⁷⁵

⁷³ Christensen, David (2013). “Epistemic Modesty Defended,” in David Christensen and Jennifer Lackey (eds.) *The Epistemology of Disagreement: New Essays*, Oxford University Press.

⁷⁴ Feldman, “Reasonable Religious Disagreements,” 195-97.

⁷⁵ Feldman, 213.

Feldman's call to humility is certainly a welcome alternative to intolerance; however, it is unclear that such virtue belongs especially to proponents of (RUE) and conciliationism.

The steadfast permissivist can readily parody Feldman's speech. Consider:

This, namely the falsity of (RUE) and conciliationism, may seem to be a distressing conclusion. It implies that I am justified in disagreeing with you about many of your deeply held convictions. Worse, it implies that you are justified in disagreeing with me about many of my deeply held convictions. Still, I think that this is the truth of the matter. And perhaps the conclusion is not so distressing. It calls for a kind of humility in response to the hard questions about which people so often find themselves in disagreement. It requires us to admit that we alone are not the sole possessors of rationality in these cases. When compared to the intolerant views with which we began, this is a refreshing outcome.

Moreover, though humility is important, it is not the only virtue we care about. We also admire people who are epistemically courageous: persons of conviction. We admonish those who, in their moral behavior, too easily surrender to peer pressure. Similarly, we might worry that those who surrender their epistemic behavior to peer pressure fail to exhibit a proper epistemic conviction.

I believe my account offers a satisfactory way to affirm both epistemic humility and epistemic courage in religious disagreements. One's beliefs ought to display a real sensitivity to the opinions and disagreements of others. Sustained disagreement, absent special considerations, makes it impossible to know that one knows. Thus, upon finding oneself in religious disagreement, one cannot be dogmatically certain. On the other hand, one need not abstain from belief whenever one encounters sustained peer disagreement. It is possible for one to maintain one's religious beliefs, despite the disagreement of others. To do so exhibits epistemic courage, but courage without arrogance; it exhibits humble conviction.⁷⁶

⁷⁶ I am especially thankful to Brian Leftow, Max Baker-Hytch, Joshua Layton-Wood, and Danielle Willard for never tiring of discussing this project with me. I also owe thanks to Charity Anderson, Matthew Benton, Paul Billingham, Laura Callahan, Hywel Clifford, Geoff Dargan, Jonathan Erhardt, Hannah Kim, Matt Kirkpatrick, Joanna Lawson, Tim Mawson, Kian Mintz-Woo, Mark Nelson, Josh Roe, Wes Skolits, Gregory Stacey, Richard Swinburne, and Daniel Young. Whether in agreement and in disagreement, I have greatly benefited from conversations with each of them.

Bibliography

- Adams, Robert Merrihew. 1987. *The Virtue of Faith and Other Essays*. Oxford University Press.
- Alston, William P. 1980. "Level Confusions in Epistemology." *Midwest Studies in Philosophy* 5: 135–50.
- . 1991. *Perceiving God: The Epistemology of Religious Experience*. Ithaca: Cornell University Press.
- Audi, Robert. 2013. "Cognitive Disparities: Dimensions of Intellectual Diversity and the Resolution of Disagreements." In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford University Press.
- Aumann, Robert J. 1976. "Agreeing to Disagree." *The Annals of Statistics* 4: 12–36.
- Baker-Hytch, Max. 2014. "Religious Diversity and Epistemic Luck." *International Journal for Philosophy of Religion*. doi:10.1007/s11153-014-9452-7.
- Ballantyne, Nathan, and E.J. Coffman. 2011. "Uniqueness, Evidence, and Rationality." *Philosophers' Imprint* 11 (18).
- . 2012. "Conciliationism and Uniqueness." *Australasian Journal of Philosophy* 90: 657–70.
- Barnett, Zach, and Han Li. 2014. "Selection and Merely Possible Disagreement". presented at the Rutgers-Princeton Philosophy Graduate Conference, New Brunswick.
- Bergmann, Michael. 2009. "Rational Disagreement after Full Disclosure." *Episteme* 6: 336–53.
- Bogardus, Tomas. 2009. "A Vindication of the Equal-Weight View." *Episteme* 6 (3): 324–35.
- Bonjour, Laurence. 2002. *Epistemology: Classic Problems and Contemporary Responses*. Rowman & Littlefield.
- Broome, John. 2001. "Are Intentions Reasons? And How Should We Cope with Incommensurable Values?" In *Practical Rationality and Preference: Essays for David Gauthier*, edited by Christopher Morris and Arthur Ripstein, 98–120. Cambridge University Press.
- Carter, Adam. 2014. "(Undercutting) Epistemic Defeat and the 'Conciliatory' Road to Agnosticism." In Oxford.
- Chisholm, R. M. 1977. *Theory of Knowledge*. 2nd ed. Upper Saddle River, NJ: Prentice-Hall.
- Christensen, David. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116 (2): 187–217.
- . 2009. "Disagreement as Evidence: The Epistemology of Controversy." *Philosophy Compass* 4: 756–67.
- . 2010a. "Higher-Order Evidence." *Philosophy and Phenomenological Research* 81 (1): 185–215.
- . 2010b. "Rational Reflection." *Philosophical Perspectives* 24: 121–40.
- . 2011. "Disagreement, Question-Begging and Epistemic Self-Criticism." *Philosophers' Imprint* 11 (6): 1–22.
- . 2013. "Epistemic Modesty Defended." In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford University Press.
- Christensen, David, and Lackey, eds. 2013. *The Epistemology of Disagreement: New Essays*. Oxford University Press.

- Clark, Kelly James. 2004. "Religious Epistemology." *Internet Encyclopedia of Philosophy*. <http://www.iep.utm.edu/relig-ep>.
- Clifford, William K. 1886. *Lectures and Essays*. Edited by Leslie Stephen and Frederick Pollock. London: Macmillan and Co.
- Coady, David. 2006. "When Experts Disagree." *Episteme* 3: 68–79.
- Cohen, Stewart. 2013. "A Tentative Defense of the (Almost) Equal Weight View." In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford University Press.
- Conee, Earl. 2009. "Peerage." *Episteme* 6 (3): 313–23.
- . 2010. "Rational Disagreement Defended." In *Disagreement*, edited by Richard Feldman and Ted Warfield. Oxford University Press.
- Conee, Earl, and Richard Feldman. 2004. *Evidentialism: Essays in Epistemology*. Oxford: Clarendon Press.
- De Cruz, Helen. 2014. "Can People Be Genuine Epistemic Peers." *The Prosblogion*. February 27. <http://prosblogion.ektopos.com/2014/02/27/can-people-be-genuine-epistemic-peers>.
- . "I Just Know I'm Right-What to Do When You Feel This in Religious and Philosophical Disagreements." *The Prosblogion*. <http://prosblogion.ektopos.com/archives/2013/10/i-just-know-im-.html>.
- De Cruz, Helen, and Johan De Smedt. 2013. "The Value of Epistemic Disagreement in Scientific Practice. The Case of Homo Floresienis." *Studies in History and Philosophy of Science* 44: 169–77.
- Dougherty, Trent, ed. 2011. *Evidentialism and Its Discontents*. Oxford University Press.
- Douven, Igor. 2009. "Uniqueness Revisited." *American Philosophical Quarterly* 46: 347–61.
- . 2011. "Abduction." Edited by Edward N. Zalta. *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/spr2011/entries/abduction>.
- Egan, Andy. 2010. "Disputing about Taste." In *Disagreement*, edited by Richard Feldman and Ted A. Warfield. Oxford University Press.
- Egan, Andy, and Adam Elga. 2005. "I Can't Believe I'm Stupid." *Philosophical Perspectives* 19: 77–93.
- Elga, Adam. 2007. "Reflection and Disagreement." *Noûs* 41 (3): 478–502.
- . 2010. "How to Disagree about How to Disagree." In *Disagreement*, edited by Richard Feldman and Ted Warfield. Oxford University Press.
- Elgin, Catherine Z. 2010. "Persistent Disagreement." In *Disagreement*, edited by Richard Feldman and Ted A. Warfield. Oxford University Press.
- Feldman, Richard. 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60: 667–95.
- . 2006. "Epistemological Puzzles About Disagreement." In *Epistemology Futures*, edited by Stephen Heather. Oxford University Press.
- . 2007. "Reasonable Religious Disagreements." In *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, edited by Louise Antony. Oxford University Press.
- . 2009. "Evidentialism, Higher-Order Evidence, and Disagreement." *Episteme* 6 (3): 294–312.
- Feldman, Richard, and Ted A. Warfield. 2010. *Disagreement*. Oxford University Press.
- Forrest, Peter. 2013. "The Epistemology of Religion." Edited by Edward N. Zalta. *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/sum2013/entries/religion-epistemology>.
- Frances, Bryan. 2008. "Spirituality, Expertise, and Philosophers." In *Oxford Studies in Philosophy of Religion*, edited by Jonathan Kvanvig, 1:44–81.
- . 2010. "The Reflective Epistemic Renegade." *Philosophy and Phenomenological*

- Research* 81 (2): 419–63.
- . 2013. “Philosophical Renegades.” In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford University Press.
- Fumerton, Richard. 2010. “You Can’t Trust a Philosopher.” In *Disagreement*, edited by Richard Feldman and Ted A. Warfield. Oxford University Press.
- Gettier, Edmund L. 1963. “Is Justified True Belief Knowledge?” *Analysis* 23: 121–23.
- Goldberg, Sanford. 2013. “Disagreement, Defeat, and Assertion.” In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford University Press.
- Goldman, Alvin I. 2001. “Experts: Which Ones Should You Trust.” *Philosophy and Phenomenological Research* 63: 85–110.
- . 2010. “Epistemic Relativism and Reasonable Disagreement.” In *Disagreement*, edited by Richard Feldman and Ted A. Warfield. Oxford University Press.
- Graham, Peter J. 2012. “Epistemic Entitlement.” *Noûs* 46 (3): 449–82.
- Greco, John. 2012. “Religious Disagreement and Religious Epistemology.” In . <http://www.philosophical-religious-knowledge.com/a-panel-presentation-on-religious-disagreement-by-john-greco>.
- Hawthorne, John, and Amia Srinivasan. 2013. “Disagreement Without Transparency: Some Bleak Thoughts.” In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford University Press.
- James, William. 1956. *The Will to Believe : And Other Essays in Popular Philosophy*. New York: Dover.
- Jehle, David, and Branden Fitelson. 2009. “What Is the ‘Equal Weight View’?” *Episteme* 6 (3): 280–93.
- Kelly, Thomas. 2005. “The Epistemic Significance of Disagreement.” In *Oxford Studies in Epistemology*, edited by John Hawthorne and Tamar Szabó Gendler, 1:167–96. Oxford University Press.
- . 2006. “Evidence.” Edited by Edward N. Zalta. *Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/entries/evidence/>.
- . 2007. “Peer Disagreement and Higher Order Evidence.” In *Disagreement*, edited by Richard Feldman and Ted Warfield, 2010:111–74.
- . 2008. “Disagreement, Dogmatism, and Belief Polarization.” *The Journal of Philosophy* 105 (10): 611–33.
- . 2011a. “Consensus Gentium: Reflections on ‘Common Consent’ Arguments for the Existence of God.” In *Evidence and Religious Belief*, edited by Kelly James Clark and Raymond J. VanArragon. Oxford University Press.
- . 2011b. “Following the Argument Where It Leads.” *Philosophical Studies* 154 (1): 105–24.
- . 2013a. “Disagreement and the Burdens of Judgment.” In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford University Press.
- . 2013b. “Evidence Can Be Permissive.” In *Contemporary Debates in Epistemology*, edited by Matthias Steup, John Turri, and Ernest Sosa, 2nd ed. Blackwell.
- Kornblith, Hilary. 2010. “Belief in the Face of Controversy.” In *Disagreement*, edited by Richard Feldman and A. Warfield, Oxford University Press Warfield. Oxford University Press.
- Kraft, James. 2006. “Religious Tolerance through Religious Diversity and Epistemic Humility.” *Sophia* 45 (2): 101–16.

- . 2007. “Religious Disagreement, Externalism, and the Epistemology of Disagreement: Listening to Our Grandmothers.” *Religious Studies* 43: 417–32.
- . 2012. “Religious Epistemology of Religious Disagreements Based on Religious Testimony and Seemings.” In. <http://www.philosophical-religious-knowledge.com/religious-epistemology-of-religious-disagreements-based-on-religious-testimony-and-seemings>.
- Kratzer, Angelika. 1977. “What ‘Must’ and ‘Can’ Must and Can Mean.” *Linguistics and Philosophy* 1: 337–55.
- Kuhn, Thomas. 2012. *The Structure of Scientific Revolutions*. 3rd ed. Chicago University Press.
- Kvanvig, Jonathan L. 2013. “Perspectivalism and Reflective Ascent.” In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford University Press.
- Kyle, Christopher. 2012. “Epistemicism and the Problem of Arbitrariness for Vagueness.” *Dialogue* 55 (1): 54–64.
- Lackey, Jennifer. 2010a. “A Justificationist’s View of Disagreement’s Epistemic Significance.” In *Social Epistemology*, edited by A. Haddock, A. Millar, and D. Pritchard. Oxford University Press.
- . 2010b. “What Should We Do When We Disagree.” In *Oxford Studies in Epistemology*, edited by Tamar Szabó Gendler and John Hawthorne, 274–93. Oxford University Press.
- . 2012. “Taking Religious Disagreements Seriously.” In . <http://www.philosophical-religious-knowledge.com/video-on-religious-disagreement-and-epistemic-peerhood-by-jennifer-lackey>.
- . 2013. “Disagreement and Belief Dependence: Why Numbers Matter.” In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford University Press.
- Lasonen-Aarnio, Maria. 2013. “Disagreement and Evidence Attenuation.” *Noûs* 47 (4): 767–94.
- . 2014a. “Higher-Order Evidence and the Limits of Defeat.” *Philosophy and Phenomenological Review* 88 (2): 314–45.
- . 2014b. “Learning to Love Mismatch.” In Oxford.
- Lee, Matthew. 2013. “Conciliationism without Uniqueness.” *Grazer Philosophische Studien* 88.
- Leftow, Brian. 1994. “From Jerusalem to Athens.” In *God and the Philosophers*, edited by Thomas Morris, 189–207. Oxford University Press.
- . 1995. “Can Philosophy Argue God’s Existence.” In *The Rationality of Belief and the Plurality of Faith*, edited by Thomas Senor, 40–70.
- Lehrer, Keith. 1976. “When Rational Disagreement Is Impossible.” *Noûs* 10: 327–32.
- Lewis, David. 1996. “Elusive Knowledge.” *Australasian Journal of Philosophy* 74: 549–67.
- . “Academic Appointments: Why Ignore the Advantage of Being Right.” In *Papers in Ethics and Social Philosophy*, edited by David Lewis, 2000:187–200.
- MacFarlane, John. 2007. “Relativism and Disagreement.” *Philosophical Studies* 132: 17–31.
- Matheson, Jonathon. 2009. “Conciliatory Views of Disagreement and Higher-Order Evidence.” *Episteme* 6 (3): 269–79.
- McNamara, Paul. 2010. “Deontic Logic.” Edited by Edward N. Zalta. *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/fall2010/entries/logic-deontic>.
- Moffett, Marc. 2007. “Reasonable Disagreement and Rational Group Inquiry.” *Episteme* 4 (4): 352–67.

- Nelson, Mark T. 2005. "Telling It like It Is: Philosophy as Descriptive Manifestation." *American Philosophical Quarterly* 42: 141–56.
- . 2010. "We Have No Positive Epistemic Duties." *Mind* 119: 83–102.
- Pettit, Philip. 2006. "When to Defer to Majority Testimony-and When Not." *Analysis* 66: 179–87.
- Pittard, John. Forthcoming. "Conciliationism and Religious Disagreement." In *Challenges to Moral and Religious Belief: Disagreement and Evolution*, edited by Michael Bergmann and Patrick Kain. Oxford University Press.
- . 2014. "Fundamental Disagreements and Defeat." In *Oxford*.
- Plantinga, Alvin. 1981. "Is Belief in God Properly Basic." *Noûs* 15: 41–52.
- Plantinga, Alvin, and Nicholas Wolterstorff, eds. 1983. *Faith and Rationality*. University of Notre Dame Press.
- Pritchard, Duncan. 2009. "Safety-Based Epistemology: Whither Now?" *Journal of Philosophical Research* 34: 33–45.
- Rosen, Gideon. 2001. "Nominalism, Naturalism, Epistemic Relativism." *Philosophical Perspectives* 15: 60–91.
- Roush, Sherrilyn. 2009. "Second Guessing: A Self-Help Manual." *Episteme* 6 (3): 251–68.
- Sider, Ted, and David Braun. 2007. "Vague, So Untrue." *Noûs* 41: 133–56.
- Skolits, Wes. 2014. "Religious Experience: Epistemic Safety and the Challenge from Diverse Contents". MPhil Thesis at the University of Oxford.
- Sosa, Ernest. 1980. "The Raft and the Pyramid: Coherence Versus Foundations in the Theory of Knowledge." *Midwest Studies in Philosophy* 5: 3–25.
- . 2007. *A Virtue Epistemology*. Oxford University Press.
- . 2009. *Reflective Knowledge: Apt Belief and Reflective Knowledge*. Oxford University Press.
- . 2010. "The Epistemology of Disagreement." In *Social Epistemology*, edited by A. Haddock, A. Millar, and D. Pritchard, 278–97. Oxford University Press.
- . 2013. "Can There Be a Discipline of Philosophy? And Can It Be Founded on Intuitions." In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford University Press.
- Swinburne, Richard. 2001. "Review of Warranted Christian Belief by Alvin Plantinga." *Religious Studies* 37: 203–14.
- . 2004. *The Existence of God*. Oxford University Press.
- Talbott, William. 2011. "Bayesian Epistemology." Edited by Edward N. Zalta. *The Stanford Encyclopedia of Philosophy*.
<http://plato.stanford.edu/archives/sum2011/entries/epistemology-bayesian>.
- Van Inwagen, Peter. 1993. "Critical Studies of the New Testament and the User of the New Testament." In *Hermes and Athena: Biblical Exegesis and Philosophical Theology*, edited by Eleonore Stump and Thomas Flint, 159–90. University of Notre Dame Press.
- . 1994. "Quam Dilecta." In *God and the Philosophers*, edited by Thomas Morris, 31–60. Oxford University Press.
- . 1996. "It Is Wrong, Everywhere, Always, and for Anyone, to Believe Anything on Insufficient Evidence." In *Faith, Freedom, and Rationality: Philosophy of Religion Today*, edited by Jeff Jordan and Daniel Howard-Snyder, 137–53. Rowman & Littlefield.
- . 2009. "Listening to Clifford's Ghost." *Royal Institute of Philosophy Supplement* 84: 15–35.
- . 2010. "We're Right, They're Wrong." In *Disagreement*, edited by Richard

- Feldman and Ted A. Warfield, 10–28. Oxford University Press.
- Van Wietmarschen, Han. 2013. “Peer Disagreement, Evidence, and Well-Foundedness.” *The Philosophical Review* 122 (3): 395–425.
- Velleman, J. David. 2009. *The Possibility of Practical Reason*. MPublishing, University of Michigan Library.
- Weatherston, Brian. 2013. “Disagreements, Philosophical and Otherwise.” In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford University Press.
- Wedgwood, Ralph. 2007. *The Nature of Normativity*. Oxford University Press.
- . 2010. “The Moral Evil Demons.” In *Disagreement*, edited by Richard Feldman and Ted A. Warfield. Oxford University Press.
- White, Roger. 2005. “Epistemic Permissiveness.” *Philosophical Perspectives* 19: 445–49.
- . 2009. “On Treating Oneself and Others as Thermometers.” *Episteme* 6 (3): 233–50.
- . 2013. “Evidence Cannot Be Permissive.” In *Contemporary Debates in Epistemology*, edited by Matthias Steup, John Turri, and Ernest Sosa, 2nd ed. Blackwell.
- Williamson, Timothy. 1994. *Vagueness*. Routledge.
- . 2000. *Knowledge and Its Limits*. Oxford University Press.