

UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

DISAGREEMENT, AGREEMENT, AND THE CONSENSUS GENTIUM ARGUMENT

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

ZACHARY G. MILSTEAD

Norman, Oklahoma

2021

DISAGREEMENT, AGREEMENT, AND THE CONSENSUS GENTIUM ARGUMENT

A DISSERTATION APPROVED FOR THE
DEPARTMENT OF PHILOSOPHY

BY THE COMMITTEE CONSISTING OF

Dr. Linda Zagzebski, Chair

Dr. Wayne Riggs

Dr. Stephen Ellis

Dr. Cynthia Rogers

*To my wife, Rebekah,
Your love, kindness, understanding, and reassurance
have allowed me to persevere on this journey.*

*To my boys, James and Jonathan,
any philosophical accomplishments I achieve
pale in comparison to the pride I have in both of you.*

Acknowledgements

I owe many thanks to many people, and it is a difficult thing to attempt the inclusion of all the appropriate names and thanks. I made many friends over the years since I began my first semester of grad school, and I have benefited not only from their friendship but from their time and encouragement as well. My degree from OU is the culmination of the completion of three different graduate degrees at three different institutions in three very different parts of this country. While this garnered a significant amount of relocation fatigue, it also led to friendships I would never have expected and for which I will be eternally grateful. Let me begin, though, by thanking those who have shaped my philosophical career.

First and foremost, I must thank Linda Zagzebski for her patience and guidance on this project. It was Linda's suggestions and promptings that led me to the idea for this project and focused and refined it into what it became. Linda is one of those rare individuals who naturally, and with little apparent effort, makes those around her want to be and do better. It was a pleasure to study under her as her student, to serve as her TA and RA, and to have her supervise this project. I count it a great blessing and honor to say that she directed this dissertation and am thankful for the opportunity.

I must also express my gratitude to the rest of my committee. A special thanks to both Wayne Riggs and Steve Ellis, who provided comments on the first chapter. Steve's comments were especially helpful in pointing out some areas that lacked clarity or required more context. Also, many thanks to both of them for the numerous conversations in the philosophy department lounge. Even those short moments of their time contributed to making me a better philosopher. I am thankful for their investment in students like me. Finally, I am grateful to Cynthia Rogers for

serving as the outside member of my committee. Her valuable time and willingness to serve are greatly appreciated.

I found the philosophy department at OU to be a uniquely encouraging and supportive environment. I am thankful to Gabe and Shelley, who were always helpful and kind to the many requests for assistance from all the grad students, including myself. Their patient work and diligence are significant reasons why the department runs as smoothly as it does. I am also thankful to our faculty, who would take time to engage with students and encouraged many wonderings down philosophical rabbit-holes. Many thanks to Monte Cook for the numerous chats and suggestions of resources for historical writings that were relevant to my work. Thanks also to the rest of my professors Heather Demarest, Neal Judisch, Martin Montminy, and Matt Priselac. It was a pleasure to be your student. I must also thank my friends, colleagues, and fellow grad students for the years of fun and great conversations. Specifically, I would like to thank Raymond Stewart and Kyle Burnard-Fullington for the many discussions that led to much of what ended up in this dissertation.

Additionally, I must thank Robert Stewart. Bob was my first philosophy professor, who introduced me to philosophy and mentored me as a budding philosophy student. His encouragement is what initially led me to consider further philosophical studies. I will forever be indebted to him. I must also thank Timothy McGrew, whom I met through Bob Stewart. Tim introduced me to the wonders of the intersection of the traditional and the formal approaches to epistemology and philosophy of science. He also nurtured and encouraged my love of the history of philosophy. Through Tim, I discovered Lydia McGrew's work in formal epistemology. Both Tim and Lydia have influenced my philosophy more than they will ever know.

I must also say a word of thanks to my family and friends. Thank you for the encouragement and support throughout the years to my parents, siblings, extended family, and in-laws. To the many friends who have supported us on this journey, a big thanks to you as well. To my friends Casey Hofberger and Jonathan L'Hoste, who have encouraged me on a monthly, weekly, and sometimes daily basis for the last few years, thank you for being the brothers and friends you have been to me.

Above so many things, I am thankful for my two boys, James and Jonathan. With patience and understanding, they have endured my academic busyness and the general distractedness that comes along with being a graduate student. Being a philosopher is not something you turn on and turn off. My children have graciously endured being the sons of a philosopher and all the lectures, lessons, deep discussions, and constant correction about language and logic that goes along with that. They both make me proud every day, and I am thankful for every moment I have with them.

Finally, I'd like to thank my dear wife, Rebekah. She has sacrificed much, endured much, and loved much as I have completed paper after paper and degree after degree. She has modeled the type of wife that every man should hope to have on his side. She has been my best friend and #1 encourager for the last twenty years. I would not have been able to complete this journey without her. She has proof-read many dry philosophy papers out of her loving-kindness and the knowledge that her husband is a perpetually bad proof-reader. She has endured many nights of me being off in the other room studying, reading, or writing. Through the best parts and the worst, she has been there, a constant source of confidence, determination, and strength. She is the true hero of this story.

Table of Contents

INTRODUCTION	1
CHAPTER 1: RATIONAL DISAGREEMENT VS. RATIONAL AGREEMENT	3
1.1 RATIONAL DISAGREEMENT: SURVEYING THE FIELD & INVERTING THE PROBLEM	8
1.1.1 <i>The Equal Weight View (Conciliationist View)</i>	9
1.1.1.1 Inverting the Problem: The Equal Weight View	12
1.1.2 <i>Steadfast Views - Multiple Variations</i>	17
1.1.2.1 Inverting the Problem: The Steadfast View	21
1.1.3 <i>The Total Evidence View</i>	24
1.1.3.1 Inverting the Problem: The Total Evidence View	26
1.1.4 <i>The Justificationist View</i>	27
1.1.4.1 Inverting the Problem: The Justificationist View	28
1.1.5 <i>Formal Issues in Disagreement</i>	30
1.2 DEALING WITH AGREEMENT	36
1.2.1 <i>A Possible Solution</i>	36
1.2.2 <i>Limitations, Special Considerations, Interesting Developments</i>	41
1.3 CONCLUSION	43
CHAPTER 2: HOW SHOULD WE RESPOND TO AGREEMENT?	45
2.1 PEERS, INFERIORS, AND SUPERIORS	46
2.1.1 <i>Disagreement Between Superiors and Inferiors</i>	48
2.2 AGREEMENT BETWEEN SUPERIORS AND INFERIORS	56
2.2.1 <i>Peer Agreement: Assorted Cases</i>	58
2.2.2 <i>Peer Agreement for Different Reasons</i>	59
2.2.3 <i>Agreement in Testimony</i>	61
2.2.4 <i>Strict Superior Agreement</i>	64
2.3 UPDATING FROM PEERS, SUPERIORS, AND INFERIORS	66
2.3.1 <i>Concerns About Varying Instances of Agreement and Disagreement</i>	71
2.4 CONCLUSION	74
CHAPTER 3: CONSENSUS GENTIUM AND THE ASSURANCE OF AGREEMENT	75
3.1 CONSENSUS GENTIUM ARGUMENTS	76
3.1.1 <i>Common Consent as Inference to the Best Explanation</i>	78
3.1.2 <i>Common Consent and Self-Trust</i>	84
3.2 DEVELOPING A NEW ARGUMENT: EXPLANATION AS A HEURISTIC	89
3.2.1 <i>Conditions for Knowledge as the Best Explanation of Consensus</i>	91
3.2.1.1 Uncoerced Believers	91
3.2.1.2 Diverse Believers	92
3.2.1.3 Non-trivial Consilience of Evidence	93
3.2.1.4 Sufficiently Large Group	94
3.2.2 <i>Common Consent as Explanation</i>	95
3.2.2.1 Common Consent as Explanation and the Non-Trivial Consilience of Evidence	96
3.3 A NEW ARGUMENT: COMMON CONSENT FROM OPINION UPDATING	98
3.3.1 <i>Basic Assumptions for the Argument</i>	98
3.3.2 <i>The Argument from Common Consent based on Opinion Updating</i>	102

3.3.3	<i>A Cumulative CCO Argument</i>	105
3.3.3.1	Focusing on a Specific Account of God	105
3.3.3.2	Updating from Superiors	109
3.4	CONCLUSION	112
WRAPPING THINGS UP		114
APPENDIX A		116
APPENDIX B		118
APPENDIX C		119
BIBLIOGRAPHY		122

Abstract

Since its initial proposal, epistemologists have suggested various theories for how the problem of peer disagreement might be resolved, followed by numerous critical and supporting assessments of these proposals. Little attention, however, has been paid to the other side of this issue. In this dissertation, I propose there is an apparent inverse symmetry of sorts between the problem of disagreement and the flip-side of that issue, which I call the assurance of agreement. I argue that a unified theory is needed to deal with both the disagreement and agreement sides of the overarching picture. This theory should not merely address the question of disagreement and agreement between peers but should also extend to inferiors and superiors as well. Such a theory should also include a means for updating one's confidence in such cases. In contribution to this effort, I propose a model for belief updating that is agnostic to the underlying view and is suitable for cases of agreement (or disagreement). I show how the proposed model can be expanded for use in cases of agreement involving peers, inferiors, and superiors. I then expand this discussion to show how the model can also be used in cases of consensus, which leads to the development of a consensus gentium argument for the reasonable belief in the existence of God. In chapter 1, I make a case for the need for a unified theory, and I propose a theory-agnostic model for updating in cases of agreement and disagreement. In chapter 2, I make a case for updating based on the agreement of superiors and inferiors and show how the model can be adjusted for these cases. Finally, in chapter 3, I set the stage for a consensus gentium argument and develop the argument using a modified version of the previously developed model.

Introduction

In contemporary analytic philosophy, there is an interesting pattern that repeats over and over again. This pattern usually begins with an interesting new problem, the formulation of a novel approach to resolving some existing problem, or maybe recasting some generally assumed term or idea into a new light. The reaction is quite naturally one of great excitement and interest by almost everyone working in the relevant field (and sometimes by those outside of the relevant field) and a desire to say something about this new thing. Every area of discipline has probably seen this sort of thing innumerable times. In epistemology, among others, we saw this with Gettier's paper about justification, with C. A. J. Coady's work in testimony, and more recently with the discussion on the problem of peer disagreement that Richard Feldman initially illuminated. Like many in epistemology, I found the disagreement literature to be quite interesting (and at times quite frustrating). While I do have some ideas about how one should deal with the problem of rational disagreement, the aim of this project focuses on drawing attention to the other side of that issue which revolves around rational agreement.

In chapter one, I set out to make a case for the apparent inverse symmetry between the *problem of disagreement* and the flip-side of that issue, which I call the *assurance of agreement*. Afterward, I attempt to lay out some of the major positions proposed for dealing with the problem of peer disagreement. My goal here is not to defend or give a critique of these positions. Rather, I try to describe them as accurately and charitably as possible. Once I have laid out the basics of the positions, I attempt to show that there may be some additional considerations or difficulties were the same approach taken for addressing the question of how we should deal with rational peer *agreement*. The goal is to show that, no matter the view, a unified theory should be able to deal with both the disagreement and agreement sides of the overarching

picture. Finally, I propose a model for updating confidence in a belief in cases of agreement (or disagreement) that is agnostic to the underlying view and should adapt to any of the views discussed in the chapter (and variations that are not discussed). The ultimate goal of this model is to provide the groundwork for a formal tool that might be useful to traditional epistemologists working in this area.

Chapter two expands on the discussion of the initial chapter. After all, disagreement is not merely limited to instances involving peers. Many times, our cases of disagreement involve a superior and an inferior. This is quite common when we consider instances involving a layperson and someone who is an expert in a specific field. In this chapter, I discuss the specific considerations that must be taken into account for deciding if and when an inferior, a superior, or both should update their confidence in their beliefs. Having established these parameters, I then move on to address how these parameters might be applied in different instances of agreement between superiors and inferiors. Finally, I show how the model developed in chapter one can be expanded for use in cases of agreement involving peers, inferiors, and superiors.

In chapter three, the focus changes to address how the model might be used in instances of consensus. More specifically, I argue that it can be used in a rather unique argument for the existence of God, which is generally referred to as a *consensus gentium* argument. While this type of argument is quite commonly used throughout history, it has fallen somewhat out of favor in contemporary philosophy of religion. I give a brief overview of why this is the case and then discuss in some detail two recent versions of this argument. I propose that, even if one rejects these two recently proposed arguments, there is still reason to adopt my argument. Finally, I layout the groundwork for my argument and make a case for belief in the existence of God based on the opinions of others.

Chapter 1: Rational Disagreement vs. Rational Agreement

The “problem of peer disagreement” continues to lead to a significant amount of discussion among epistemologists. The fundamental conception of the problem is that it is not clear that two people who have the same evidence, cognitive abilities, background knowledge, etc., can rationally come to different conclusions regarding some proposition or hypothesis. Humorously, there has been quite a bit of disagreement about the details of the problem itself, *e.g.*, the reason the problem arises in the first place, and how the problem should be resolved (or at least dealt with). In general, however, two things have been well established: 1) apparent cases of rational peer disagreement occur (or at least have the potential of occurring) and, according to some, 2) this leads to an epistemic problem, *i.e.*, something “bad”, something we should avoid, or something we must resolve. Why is it a problem if my peer and I disagree? Well, the disagreement seems to be indicative of a mistake. Something went wrong in either my reasoning process or my peer's reasoning process. Otherwise, we would have arrived at the same conclusion.

Suppose my peer and I were to come to different conclusions regarding some proposition or some hypothesis. Both of us have access to the same evidence, and both generally end up at the same conclusions in previous instances of such evaluations. In that case, it is not unreasonable to wonder whether some error in reasoning might have occurred. Indeed, it seems to many, the fact that we came to different conclusions is direct evidence that such an error exists. This disagreement does not actually grant us any evidence of where the error lies, merely that there is some error. That some error exists, then, is understood to mean that there is some aspect of irrationality that has arisen in my reasoning processes or that of my peer. This

irrationality, and the lack of clarity regarding which of us is being irrational, is epistemically bad. It is not to be desired and requires a resolution.

If disagreement entails something bad, this might lead one to wonder what would happen were the situation inverted. It seems reasonable to suppose that the inverse of Disagreement → ‘something bad’ would then be Agreement → ‘something good’. We might say there is an inverse symmetry of sorts between the two. This idea does have quite a natural intuitive appeal. *Disagreement* seems to give us some clue that something has gone wrong, but *agreement* appears to provide us with some indication that something has gone right. Thus, we might say the inverse of the *problem* of peer disagreement is the *assurance* of peer *agreement*.

Consider the various instances where agreement appears to be a good thing. We generally think it a good thing when witnesses agree regarding the facts of some crime, at least instances of agreement in which there is no collusion or coercion. We tend to think it is a good thing when experts have independently worked through some difficult topic regarding some complex hypothesis, and we take it as more likely that the hypothesis is true when the experts agree. If historical texts independently agree that some event occurred on a specific date, we take that as a good sign that the event actually happened on that date. In school, when we learn how to work through complex mathematical problems, we generally think of it as a good thing and confirmation that we likely did things correctly if other students end up with the same solution to the problem. All of the above appears to indicate a possible sort of inverse symmetry between agreement and disagreement given the problem of disagreement described above. If that is correct, then we should be able to take different views on the *problem of disagreement* and apply the various claims, requirements, and proposed resolutions to the problem entailed by those views to instances of *peer agreement*. Doing so should illuminate the sort of symmetry between

agreement and disagreement, and some of the results might be somewhat surprising. The purpose of this chapter is to do just that.

To begin, I want to be clear about what I mean by *agreement*. Specifically, I want to make a distinction between mere *agreement* and *consensus*. Consensus is a type of agreement, but it meets certain requirements beyond what I consider in this chapter. In social epistemology, the term consensus may refer to an agreement amongst members of a group. More specifically, the term may be used to describe complete agreement, *i.e.*, there is no dissent whatsoever, in communities of experts, *e.g.*, a scientific or philosophical community (c.f., Miller, 2019). Some might even use more specific terms like *knowledge-based consensus* to refer to instances in which consensus has met some conditions that show it is a mark of knowledge (Miller, 2013, 1294). The term may also refer to an agreement amongst members of a group in which a sufficient number of fundamental points regarding some hypothesis or view grants a more general sense of consensus. In such cases, there may be a few dissenters that disagree on minor issues that are relatively unproblematic to the hypothesis as a whole.¹

While consensus is an interesting subject on its own and may be relevant in an expanded discussion of our current topic, this is not the type of agreement I will consider in this project. For our purposes, we are interested in agreement between two individuals at one point in time (at least) regarding some hypothesis or proposition, in which the two individuals are aware of the agreement. More specifically, we are interested in what could be referred to as *belief-agreement*,

¹ Such instances may be similar to the Rawlsian idea of “overlapping consensus” (Rawls, John. *A Theory of Justice*. (Revised Ed.) Harvard University Press, 1971, 1999, p. 340.).

where the two individuals agree regarding some specific truth claim.² I will elaborate on this further as the chapter progresses.

These two individuals in the agreement must also share some characteristics that make them peers. Generally, this is *at least* an informal recognition regarding access to the same evidence, the same background knowledge, and having the same cognitive abilities. A more formal definition might be something the following:

Epistemic Peers (Factor Version)

R and S are epistemic peers regarding p if and only if:

- I. R and S are cognitive equals, *i.e.*, they are equals with respect to cognitive virtues such as intelligence, thoughtfulness, competence, and freedom from bias, and
- II. R and S are evidential equals, *i.e.*, they are equally familiar with the evidence and arguments that concern p and have considered them with equal care (Ferrari and Pedersen, 2019, 175).³

This definition, of course, limits us somewhat to the various epistemic factors included in the definition. What happens if we missed some critical factors? What if we disagree on the relevant factors for peerhood? Let me offer the following alternative definition that allows us to avoid such concerns and difficulties. If the definition above for peerhood is not deemed appropriate, then we can at least assume the following:

Epistemic Peers (Position Version)

R and S are epistemic peers regarding p if and only if they are equals in epistemic position for judging the confidence one should have in the truth-value of p .⁴

² *Belief-agreement* would be different from *action-agreement* where there is agreement over what should be *done*. For instance, there may be agreement over whether Bill and Sue should or should not pursue a doctorate in philosophy. For further discussion on this distinction, see Matheson and Frances (2018) and also Frances (2014, I.9).

³ Some philosophers have proposed that more factors must be considered, such as relevant biases, intellectual virtues, time devoted to consideration, freedom from distractions, etc. Gutting originally used the term to describe individuals who were alike in “intelligence, perspicacity, honesty, thoroughness, and other relevant epistemic virtues” (Gutting, 1982, 83).

⁴ This is inspired by Frances’s distinctions for superiors and inferiors (Frances, 2014, 54).

What determines an agent's epistemic position is the relevant factors, *e.g.*, cognitive abilities, epistemic virtues, access to evidence, awareness of biases, etc. Generally, epistemic agents who are equals in the relevant factors equals in epistemic position. In some cases, agents may demonstrate some inequalities in the relevant factors, but these inequalities ultimately balance out in such a way where the agents still end up positioned equally. A very strict definition of epistemic peers might not allow this, but for our purposes we will aim for a more permissive definition.

Finally, most cases usually include the stipulation that the two individuals involved in the disagreement must recognize that they are peers. Why? Without the recognition of peerhood, what reason would they have to acknowledge that someone made a mistake? If I disagree with another person and don't recognize them as a peer, I can simply assume that the error is on their side because the other person wasn't in as good of a position as I was to arrive at the correct conclusion. Oddly enough, it is possible to have an instance where two individuals are in disagreement, and they *think* they are peers, though they do not meet the requirements outlined in either of the definitions above. This situation may lead to different (though still rather interesting) problems. For our purposes, however, we are interested in cases where the two individuals at least meet the *subjective* requirements of peerhood. While *objective* peerhood is ultimately desired in some cases, as we will see it is not always clear that is what we have in some of the paradigmatic cases.

Given the above, consider a claim made by two grad students, Bill and Sue. They are familiar with the difficulties of both getting a doctorate in philosophy and of getting a doctorate in art history. Both Bill and Sue hold to be true, "It is more difficult to attain a doctorate in philosophy than it is to attain a doctorate in art history." Bill and Sue both recognize they are in

the same position to judge such a conclusion. This type of agreement, in which two epistemic peers are involved, is the specific type of agreement with which I will be concerned in this project. Going forward, I will simply refer to such instances as *agreement*, *peer agreement*, or *rational peer agreement*.

If there is a sort of inverse symmetry between agreement and disagreement, as I have proposed, then it makes sense to examine relevant views related to the problem of peer disagreement. In what follows, I will discuss some recognizable views on peer disagreement and then evaluate the possible effects of inverting those views. We will see that some of the views may need some additional work to provide a unified theory for addressing both instances of peer disagreement and peer agreement. Finally, I will suggest minimal requirements for any proposed theory that offers a solution for updating in cases of peer agreement. Additionally, I will sketch what a possible solution might look like. In order to proceed, however, let us first consider four archetypal responses to the problem of peer disagreement and the possible inversions of those views.

1.1 Rational Disagreement: Surveying the Field & Inverting the Problem

One specific feature that all philosophers working on disagreement require is that the individuals involved in paradigmatic cases of disagreement must be epistemic peers. More specifically, they must be peers who are in an epistemically good position with regard to the evidence at hand. (No one is particularly interested in a disagreement between two obviously foolish individuals.)

Beyond this, while none of these views are monolithic in nature, we can generally divide positions on peer disagreement into four broad views, each of which may then be divided into further categories. We will begin with the Equal Weight View, one of the early views developed

in the literature, which has garnered a significant amount of approval, but also entails at least one fairly controversial thesis. The remaining views are usually advanced as alternatives to something like the Equal Weight View and are generally held to reject one or more of the Equal Weight View's major theses. Nonetheless, these alternatives have not escaped objections. In what follows, I will provide a summary of the various views as generally presented in the literature on peer disagreement. My aim is not a critique of these views nor to address the variety of objections that could be raised against these positions. Rather, my goal is to present them as charitably as possible while noting the potential results when we invert the views to address the possibility of peer agreement.

1.1.1 The Equal Weight View (Conciliationist View)

The Equal Weight View is probably one of the more prominent views expressed in the Disagreement literature. Initially, Christensen (2007), Elga (2007), and Feldman (2006) were among the first to propose the primary tenets of what is generally recognized as the Equal Weight View. As one might guess from the name, the primary motivation behind this view is the supposition that in instances where epistemic peers disagree, each of the peers has just as much reason, given the position of their peer, to think that their own position is as correct as that it is mistaken. In other words, each of the peers is just as likely to be right given their individual evaluation of the evidence, and because they disagree, each has reason to think he or she may be mistaken. Additionally, the reasons to discount one's peer's position about the proposition at question must be independent of the disagreement itself.

We might say then that this view generally entails something like the following three claims⁵:

Defeat: Upon learning that an epistemic peer disagrees with me about some proposition that p , this gives me a reason to believe I am mistaken about p .

Independence: Any reasons I have to discount another's opinion about p must be independent of the reasoning behind my own initial belief about p that led to the disagreement.⁶

Equal-Weight: The reason to think I am mistaken about the proposition that p in light of my peer's opinion about p is just as strong as the reason to think I am correct about p in light of my opinion about p .

Some slight variations of these three claims may arise (and they may be referred to by different terms), but these are generally accepted by those who adhere to the Equal Weight View.

The classic case provided to demonstrate an instance of peer disagreement is Christensen's Restaurant Check Case.

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

In this case, Christensen clearly stipulates that the friends have a long track record of eating out together and splitting the check in this fashion. The friends have been equally successful in their calculations previously, with them agreeing about the arithmetic the overwhelming majority of the time. On the occasions when they do disagree, each individual is right an equal amount of the

⁵ The three claims framework comes from the work of Matheson and Frances (2018), but it can be found both explicitly and implicitly in Feldman (2006a, 2006b, 2009), Elga (2007), Christiansen (2007, 2011), Bogardus (2009), and Matheson (2015).

⁶ See "Other Definitions" in Appendix A for two variations of Independence.

time. Those who hold to the Equal Weight View generally agree that upon learning of the disagreement regarding the shares, the peers should each have weakened their position to some degree.

At this point, we can further divide the Equal Weight View into two sub-categories by how one chooses to approach the doxastic aspects of the position. Philosophers like Feldman (2006), for example, have proposed that there are three doxastic attitudes one might have towards a proposition in cases of disagreement: a) belief that p , b) belief that $\sim p$, or c) the suspending of judgment regarding that p . I will refer to this as the *all-or-nothing* approach. Christensen and others have proposed that we need to allow for a broader range of doxastic attitudes than what is permitted in the *all-or-nothing* route since each individual may actually have different degrees of confidence in the held belief. For instance, one of the friends from the Restaurant Check Case may have a 0.6 degree of confidence that p , and the other may have a 0.4 degree of confidence that p . I will refer to this as the *divergent credences*⁷ approach.

On the Equal Weight View, many argue that since the peers' positions are weighted equally, the rational response is for the peers to "split the difference" regarding their doxastic attitudes. In the *all-or-nothing* case where one peer believes that p and the other peer believes that $\sim p$, the rational response would be for both peers to suspend judgment regarding that p . In *divergent credences* cases like the one above, the procedure would require some adjustment of confidence. Given the equal weight of the peers' attitudes, where one has 0.6 degree of confidence that p and the other has a 0.4 degree of confidence that p , "splitting the difference" would, by calculating a mean between the two degrees, call for each peer to adopt an attitude of

⁷ For my purposes, I do not use the term "credence" to refer to a degree of belief. Rather, I use it interchangeably with "confidence" to describe the probabilistic confidence one has in a belief.

0.5 confidence that p . More specifically, in the Restaurant Check case, “splitting the difference” doesn’t require a change in the calculations of \$45 and \$43 by the peers. Instead, they must both split-the-difference with regard to their confidence about their belief of what the split check amount should be. This splitting-the-difference will lead them to each end up with a confidence of 0.5. In this case, neither ends up with a strong confidence about how much the individual amounts for the check should be.

1.1.1.1 Inverting the Problem: The Equal Weight View

Recall that the Equal Weight View is generally said to entail something like *Defeat*, *Independence*, and *Equal-Weight*. We are, however, no longer focusing on that side of the picture. Because of this, we need to adjust the claims to create analogous claims regarding agreement. For both *Defeat* and *Equal-Weight*, the challenge of developing an analogous claim is relatively easy. The inversion of *Independence* requires a quick explanation of the motivation behind the claim.

Generally speaking, the motivation behind *Independence* is to “prevent blatantly question-begging dismissals of the evidence provided by the disagreement of others” (Christiansen, 2011, 2). The point is that there seems to be something wrong with reasoning in the following manner, “My peer disagrees with me about p . Since I have reasoned correctly that p is true, then my peer must be wrong about p . Though my peer is generally reliable about reasoning well and evaluating evidence, I do not need to take this instance of disagreement as a reason to doubt my belief that p is true.” The goal then of an analogous claim regarding agreement should be to prevent similar dismissals of the evidence based merely upon the agreement of others. We also want to avoid similarly question-begging acceptance of evidence provided by the agreement of others, such as, “My peer agrees with me about p . Since I have

reasoned correctly that p is true and my peer also reasoned that p is true, then I can take this instance of agreement as a confirmation that my belief that p is true is correct.” Given this, the inversion of *Independence* becomes, “Any reasons I have to accept another’s opinion about p must be independent of the reasoning behind my own initial belief about p that led to the agreement.”

Given this, consider the following as possible candidates for the role of inverted Equal Weight View claims.

Table I. Three Claims of the Equal Weight View and their Inversions⁸

Disagreement Claims	Inverted Claims / Agreement Claims
<i>Defeat</i> : Upon learning that an epistemic peer disagrees about some proposition p , this gives me a reason to believe I am mistaken about p .	<i>Validation</i> : Upon learning that an epistemic peer agrees about some proposition p , this gives me a reason to believe I am correct about p .
<i>Independence</i> : Any reasons I have to discount another’s opinion about p must be independent of the reasoning behind my own initial belief about p that led to the disagreement.	<i>Inverted Independence</i> : Any reasons I have to accept another’s opinion about p must be independent of the reasoning behind my own initial belief about p that led to the agreement.
<i>Equal-Weight</i>: I have just as much reason to trust my peer’s reasons for p as my own reasons for p. His reasons lead him to p, and mine lead me to not p. When we <i>disagree</i>, his conclusion from his reasons weaken my conclusion from my reasons.	<i>Inverted Equal-Weight</i>: I have just as much reason to trust my peer’s reasons for p as my own reasons for p. His reasons lead him to p, and mine also lead me to p. When we <i>agree</i>, his conclusion from his reasons strengthen my conclusion from my reasons.

The resulting inverted claims provided above seem to be a fair interpretation of the reasoning behind the Equal Weight View and what the analogous inverted reasoning should look like in

⁸ For reference purposes while reading, a separate copy of this table is included in Appendix B.

cases of agreement. Because each of the other views discussed below appears to require the denial of at least one of the claims of the Equal Weight View, it seems reasonable that we can evaluate all of the views that follow in light of the inverted claims.

Feldman is generally considered the original *all-or-nothing* adherent among those who advance the Equal Weight View. He makes a relatively strong claim about evidence that applies to agreement just as much as it is to disagreement. Feldman notes, “A point about evidence that plays a role here is this: evidence of evidence is evidence. More carefully, evidence that there is evidence for P is evidence for P. Knowing that the other has an insight provides each of them with evidence” (Feldman, 2006, 208). If Feldman is right about this, then the statement applies in cases of disagreement and cases of agreement. Evidence of evidence is evidence whether the evidence is that you are right or that you are wrong. So, let us consider a case of agreement following Feldman’s lead from the disagreement side of the discussion.

Let us suppose that we have two epistemic peers R and S, who are evaluating the evidence for some proposition p . Upon evaluating the evidence, at t_0 , R comes to believe that p and S comes to believe that p . At t_1 R and S discover that they are in agreement regarding the truth of p . How now should the peers proceed? Well, there is no need to “split the difference” as there would be in cases of disagreement. In fact, *Validation* tells us that the agreement should increase the peers’ confidence in the belief regarding the proposition or hypothesis in question. Generally, splitting the difference is assumed to be something like calculating a mean or linear average. For instance, were one person to have a credence of 0.3 that p and another to have a credence of 0.7 that p , splitting the difference would result in a credence of 0.5. The math is a simple case of $(0.3 + 0.7) \div 2 = 0.5$. One might naturally assume that the inverse of division in cases of disagreement would be multiplication in cases of agreement. On this assumption, the

calculation in a case of agreement would be the following: $(0.6 + 0.6) \times 2 = 2.4$. Of course, this is completely unreasonable. We cannot have a credence that exceeds 1. (Even if we assumed that any calculation that must be limited to 1, it is still incredibly strong to say that agreement with peers could lead us to certainty.) It is reasonable that the number might increase some, but simple multiplication does too much. Let's briefly consider some examples of what this inversion of the "splitting the difference" procedure might look like and some of the absurdities to which it could lead.

In the *all-or-nothing* approach of the Equal Weight View, of course, one really only has the three doxastic attitudes that can be taken towards a proposition, and as such calculations are not readily available. What is clear for someone holding a position like this, is that, in cases of disagreement requiring conciliation, one's belief will change by one degree in the direction towards the belief held by one's peer. Given the Claims of Agreement, in cases of agreement, one's belief should change by one degree in the direction already held. You can see how this could start getting silly rather quickly, especially when you have multiple instances of agreement. You would either have peers with some sort of extreme confidence or they would receive no epistemic benefit at all from agreement.

Obviously, these results are on the absurd side of things and probably not at all what someone holding to an Equal Weight View on disagreement would desire with regards to agreement. Despite this, Feldman's principle that "evidence of evidence is evidence" requires that someone holding to this position must address the argument I have raised. They must provide some reasonable method of dealing with the results of agreement between epistemic peers. That being said, as has been noted by others, this *all-or-nothing* way of thinking about beliefs tends to be quite limited for discussing doxastic attitudes in such scenarios. Hence, we

will evaluate the inversion of the *divergent credences* approach to the Equal Weight View as a more reasonable alternative.

Again, let us suppose that we have two epistemic peers R and S, who are evaluating the evidence for some proposition p . Upon evaluating the evidence, at t_0 , both R and S come to a credence of 0.6 that p . At t_1 , the two peers discover that they are in agreement regarding the truth of p . Given the Agreement Claims above and the natural inversion of “splitting the difference,” the resulting credences for R and S should be 2.4. Once again, the natural inversion of splitting the difference, *i.e.*, multiplication, does too much.

Of course, we might wonder whether R and S can be said to agree if their credences differ. Generally, it has been taken to be the case that if the credences differ (even slightly), then the two peers are in disagreement.⁹ For instance, suppose that R has a starting credence of 0.6 regarding that p , and S has a starting credence of 0.7 for that p . Though both are confident that it is more likely that p than that $\sim p$, they *disagree* about how likely it is the case that p . In such instances thus, the peers would split the difference. Only the cases in which the peers have the same confidence levels can we say the peers agree. This, of course, leads to other rather surprising results. Suppose that the credences of both peers after evaluating the initial evidence is quite low. Perhaps, something like 0.2 for both peers. Upon discovering that they agree about their evaluation of the evidence, given the above, one might think that the peers should increase their confidence. The reasoning that motivates “splitting the difference” does not give us any clear method of dealing with these sort of agreement scenarios. A unified theory should grant us a method of dealing with both sides of the issue, *i.e.*, numerical updating for disagreement and

⁹ For examples, see Christensen (2011, 5) and Kelly (2010, 125-6).

agreement. These are not two separate problems that can be handled independently. Given this, it seems that we might have reason to doubt the numerical model for handling disagreement since there appears to be no proposed reasonable inverse method of handling agreement. This, too, seems to be a rather unexpected result of this inversion process.

If simple multiplication in cases of disagreement does too much and supposing that this symmetry between agreement and disagreement exists, this might be an indication that simple division in cases of disagreement also does too much.¹⁰ The inversion of the Equal Weight View Disagreement Claims leads to some rather odd results with both the *all-or-nothing* and the *divergent credences* approaches. This does not show that the Equal Weight View is wrong. However, it does give those who adhere to the view some reason to consider my argument and provide some response to resolve the problems I have raised.

1.1.2 Steadfast Views - Multiple Variations

One might describe Steadfast Views as the opposing view of the Equal Weight View. Rather than proposing that peers should conciliate in instances of disagreement, those who adhere to the Steadfast Views argue that discovering that one disagrees with an epistemic peer requires no change at all in one's doxastic attitudes. There are various reasons offered for why this is the case, and generally, these reasons involve a denial of some essential claim of the Equal Weight View.

One might hold to the position that there is no reason to believe that one is mistaken in the face of disagreement with an epistemic peer. Such is the position proposed and defended by

¹⁰ Indeed, problems with this method have been raised elsewhere for other reasons. See Fidelson and Jehle (2009) and Easwaran *et al.* (2016) for some discussion of these problems.

Peter van Inwagen (1996). In his article, van Inwagen appeals to private evidence as a reason to reject the Equal Weight View. (More specifically, this is a rejection of the *Defeat* claim because the conditions of peerhood cannot be met.) Recall that *Defeat* and *Validation* are as follows:

<i>Defeat</i> : Upon learning that an epistemic peer disagrees about some proposition that p , this gives me a reason to believe I am mistaken about p .	<i>Validation</i> : Upon learning that an epistemic peer agrees about some proposition that p , this gives me a reason to believe I am correct about p .
--	--

As an explanation of how he could be justified in beliefs that are rejected by people who are aware of all the same evidence and arguments that he is, and who are just as good at evaluating evidence and arguments, he notes,

...I am inclined to think that I must enjoy some sort of incommunicable insight that others, for all their merits, lack. I am inclined to think that “the evidence and arguments I can adduce in support of my beliefs” do not constitute the totality of my justification for these beliefs (van Inwagen, 1996, 34).

After all, it does seem that there are things that cannot be shared with a peer (at least not in the most meaningful sense of *shared*). For instance, *my* experiences and *my* intuitions are precisely that, *mine*. While I may be able to describe them to an extent to a peer, I am also limited in my abilities and cannot duplicate those experiences and intuitions in my peer.

Another variation of the Steadfast View rejects the *Defeat* claim by noting that a peer’s doxastic attitude toward a hypothesis H (whether it is true, false, etc.) is not relevant evidence regarding H . For instance, consider a case where E is the evidence supporting H . We might say that E provides good reasons for believing that H . A peer who is good at evaluating evidence, thus, would believe H on the basis of E . The fact such a peer would believe H on the basis of E is actually evidence that “ E provides good reasons for believing that H .” It is not, however, evidence for the truth of H itself. Given this, gaining information about a peer’s disagreement

regarding H is irrelevant to one's evaluation of the proposition that H . Such is one position discussed by Kelly (2005, 185-8) against the *Defeat* claim of the Equal Weight View.

Consider a slight variation of the view above. For two peers, S and T , we begin by distinguishing two types of evidence. *First-Order* evidence for H is evidence that directly pertains to H . *Higher-Order* evidence for H is evidence about an individual's evidence for H . Given this, we can say the E is first-order evidence for H . S 's judgment that H and the fact that S is a competent evaluator of evidence is higher-order evidence for H . Finally, T 's judgment that $\sim H$ and the fact that T is a competent evaluator of evidence is higher-order evidence for $\sim H$. If S believes H and T believes $\sim H$ at t_0 and they both then become aware of their disagreement and its related higher-order evidence at t_1 , then the evidence at t_1 has changed. Due to the fact, however, that we are supposed to give equal weight to evidence, it seems like the higher-order evidence of the two peers cancels out. We are left only with E . So, at t_1 S 's judgment that H and T 's judgment that $\sim H$ must rationally remain the same (Kelly, 2005, 189-90).

Another option for one holding a Steadfast View and denying *Defeat* is to appeal to the idea that the justified evaluation of evidence is permissive. This view rejects the following thesis:

The Uniqueness Thesis: For a given body of evidence and a given proposition, there is some one level of confidence that it is uniquely rational to have in that proposition given that evidence (Kelly, 2010, 119).

Feldman has argued that the Equal Weight View entails the Uniqueness Thesis, and this thesis motivates peers to conciliate in cases of disagreement (Feldman, 2007, p. 205). If we reject this thesis and hold that evidence is permissive, *i.e.*, multiple justified positions on the same hypothesis or proposition are allowed, then the motivation to conciliate is removed.¹¹ For

¹¹ For reasons one might reject the Uniqueness Thesis, see the following: Dogramici and Horowitz 2016, Greco and Hedden 2016, Kopec 2015, Matheson 2011, Rosa 2012 and 2016, Schoenfield 2019, and White 2005

instance, consider a case where you and I disagree regarding some hypothesis, H . We have both carefully evaluated the evidence and have the necessary skills, cognitive abilities, etc., to do so successfully, and have done so for similar hypotheses in the past. After evaluating the evidence, you form the belief that H is true with a confidence of 0.6, and I form the belief that H is true with a confidence of 0.7. If evidence is permissive, then it seems like we each have no reason to think we have individually made a mistake. Because of this, we can deny *Defeat* and hold to the Steadfast View.

Finally, one may hold to the Steadfast View by denying *Equal-Weight*. Suppose one can provide some good reason that a peer's opinion regarding some hypothesis or proposition does not weigh as much as one's own opinion. In that case, there appear to be no reasons for conciliation. One way that this might be achieved is by appealing to self-trust.¹²

Generally, those who argue for such a position are assuming something like the following when they talk about self-trust:

Self-Trust: Every normal individual is entitled, from the beginning of inquiry, to rely on her own basic belief-forming faculties, at least to some degree (Schafer, 2015, 3).

Along with this, there is some supposition that, at least in some instances (possibly because of some privilege that our own faculties hold or because some specific requirements have been met, there is some doxastic value of self-trust such that "self-trust has sufficient weight to make it the case that it would not be wrong, at least in some cases of peer disagreement, to give significantly less weight to the views of one's epistemic peer" (Pasnau, 2015, 2323). Pasnau claims that we do this sort of thing in philosophy all the time. We regularly form beliefs about some hypothesis, H , even though the odds of being correct about H do not look good from a certain point of view.

¹² For variations on self-trust see Enoch (2010), Foley (2001), Pasnau (2015), Peter (2019), Schafer (2015), Wedgwood (2007; 2010), and Zagzebski (2012).

Despite this, it is not contradictory to believe *H* even while recognizing another point of view from which *H* would seem unlikely.

1.1.2.1 Inverting the Problem: The Steadfast View

Recall that the Steadfast View's various forms will deny either *Defeat*, *Independence*, or *Equal-Weight*. It then seems reasonable to suppose that they should reject the corresponding inverted claims as well. Let us begin by considering Peter van Inwagen's position in light of our new claims.

According to van Inwagen, it could be the case that instances of disagreement should not give me (or any other individual) reason to believe that I am mistaken since I will have access to private evidence to which my apparent peer will have no access. This is a rejection of the *Defeat* claim in such instances because Defeat's requirements cannot be met. Someone employing a position like van Inwagen's appears to be saying that the evidence that an apparent peer disagrees does not give her reason to believe she is mistaken about *p*. If it is the private evidence, along with the initial evidence that both individuals can access, that allows this person to deny the other's belief as a reason that she is (or could be) wrong in her original conclusion, then it appears reasonable to say that she must also deny *Validation*. The requirement of peerhood for *Validation* cannot be met since apparent peers cannot share private evidence. On this account, the apparent peer's agreement should not even count as supporting evidence that one's belief is correct. Only the initial evidence and the private evidence give an individual reason to believe or not to believe that *p*. Given the fact that an epistemic peer (or what least what many people would consider an epistemic peer) agrees with us is a common and useful tool for evaluating our positions in a multitude of formal disciplines, *e.g.*, the sciences and everyday life, some may consider this a rather undesirable position in which to find oneself.

One possible way to avoid this situation while maintaining a similar position regarding disagreement is to argue for an agent-centered account of evidence. Michael Huemer argues for such a position and notes,

Agent-centered theories allow one to assign different (and presumably greater) weight to one's own intuitions and judgments than one assigns to those of other people, even in the absence of any independent evidence for one's own greater reliability. This seems to imply that each individual has justification for taking himself to be *a priori* more likely to make correct judgments than other people (Huemer, 2011, 25).

On this account, private evidence is not used to deny *Defeat* (and, consequently, *Validation*). Instead, the evidence is no longer granted equal weight. Thus, this is a denial of *Equal-Weight* and its inverted claim. The result is that an asymmetry of sorts exists between one's own private evidence and the private evidence of a peer. Given this, in cases of disagreement, a peer's private evidence might still give one reason to conciliate (if only to a smaller degree). In cases of agreement, one's confidence might even increase (if only to smaller degree) due to the peer's private evidence.¹³

In cases where one rejects *Defeat* because one holds that a peer's doxastic attitude toward a hypothesis *H* (whether it is true, false, etc.) is not relevant evidence regarding *H*, there is also no reason to hold to *Validation*. If the first-order evidence is the only factor motivating one's

¹³ One would still have to give an account for this asymmetry. Feldman has argued against something like this position. He gives the following case in support of his objection, "Suppose you and I are standing by the window looking out on the quad. We think we have comparable vision and we know each other to be honest. I seem to see what looks to me like the dean standing out in the middle of the quad. (Assume that this is not something odd. He's out there a fair amount.) I believe that the dean is standing on the quad. Meanwhile, you seem to see nothing of the kind there. You think that no one, and thus not the dean, is standing in the middle of the quad. We disagree. Prior to our saying anything, each of us believes reasonably. Then I say something about the dean's being on the quad, and we find out about our situation" (2007, 207–8). Because neither party can share their private evidence (their personal experiences), Feldman argues that there is still reason to conciliate.

confidence regarding *H*, then the fact that a peer agrees regarding *H* appears to have no impact upon one's confidence regarding *H*. This is, however, quite contrary to our natural intuitions regarding agreement and unlikely to be something proponents of such a view would desire.

Were one to reject a conciliationist position by appealing to *Equal-Weight* and the fact that the conflicting high-order evidence would cancel itself out, the same strategy would have apparently positive results in cases of agreement. In such cases, given *Inverted Equal-Weight*, we could actually conclude that the aggregate of the higher-order evidence is what serves to strengthen the peers' belief regarding the hypothesis or proposition in question. It is not clear to what degree or how precisely the strengthening should be calculated. The result on the inversion, however, does not initially appear undesirable.

On views that appeal to permissive evidence to reject *Defeat*, the inversion would require the rejection of *Validation* for the same reasons. Some might be content with this result since neither peer would change their epistemic status or credence upon finding out that their peer agrees with them. So, there are no complex calculations or updating procedures with which to concern oneself.

Finally, were one to reject the conciliationist position by denying *Equal-Weight* as done in cases where there is an appeal to self-trust because one's own beliefs are weighted more heavily than those of a peer, *Inverted Equal-Weight* must also be rejected. Given that those who advocate for self-trust generally allow that a peer's beliefs should still be considered and not merely dismissed, this perhaps is not as difficult a position as others. There are still, however, some puzzles about how one should weight one's own beliefs in relation to the beliefs of others.

1.1.3 The Total Evidence View

Those who adhere to the Total Evidence View hold that the justification of one's beliefs, especially in cases of peer disagreement, should rely on one's *total evidence*. This view is advanced by Thomas Kelly (2010, 2013). On the Equal Weight View, the peers' higher-order evidence (the fact that they are epistemic peers, have evaluated the same evidence, and they disagree regarding the conclusion of the evaluation) is what motivates the peers to conciliate and adopt a new doxastic attitude toward the hypothesis in question, *H*. Kelly asks, "Why is the original evidence for *H* not a factor in the doxastic attitude that is ultimately adopted?" He proposes that it is the total evidence (the original first-order evidence and the higher-order evidence) that should be evaluated to judge one's final confidence in the belief that *H*.

To motivate why the total evidence must be involved and not simply the higher-order evidence, Kelly provides the following case:

Bootstrapping: At time t_0 , each of us has access to a substantial, fairly complicated body of evidence. On the whole this evidence tells against hypothesis *H*: given our evidence, the uniquely rational credence for us to have in *H* is 0.3. However, as it happens, both of us badly mistake the import of this evidence: you give credence 0.7 to *H* while I give it 0.9. At time t_1 , we meet and compare notes. Because we both accept the Equal Weight View, we converge on credence 0.8 (Kelly 2010, 125–6).

Given a commitment to the Equal Weight View, it seems that the rational thing to do, for the peers in this case, is to conciliate and adopt a confidence level of 0.8 in the belief that *H* is true. This is in spite of the original evidence for *H* that should have led the peers to a confidence level of 0.3. Kelly notes, "It is not only that one's higher-order evidence typically makes a difference to what one is justified in believing about the world; it is also the case that one's first-order evidence makes a difference to what one is justified in believing about higher-level epistemic matter" (Kelly 2010, 159). There is both a phenomenon of *downward epistemic push* from higher-order evidence and of *upward epistemic push* from first-order evidence that must be

considered. Kelly notes, however, that proper evaluation of the first-order evidence should be reflected in the higher-order beliefs. Thus, it seems that we should assume first-order evidence will end up being weighted more heavily than the higher-order evidence. Despite this, in some cases, the Total Evidence View will yield the same results as the Equal Weight View, and because of this, the view appears to some to be a suitable middle ground between the Steadfast View and the Equal Weight View.

Kelly notes, however, that there are clear instances that give us reasons to reject at least one of the claims of the Equal Weight View. To portray what this might look like, Kelly provides the following case for consideration:

Holocaust Denier: R possesses a great deal of evidence that the Holocaust occurred. Having evaluated the evidence, R judges it to strongly support that hypothesis, *H*. Upon adopting a high amount of credence that the Holocaust occurred, R encounters an individual, S (who is quite reliable in all other matters), who denies that the Holocaust ever occurred (perhaps because he is grossly ignorant of some historical facts, historical facts which make it overwhelmingly likely that the Holocaust occurred). (Kelly 2013, 40-1)

In such instances, *Independence* does not allow R's first-order reasons for believing *H* as reasons to discount S's judgment about *H*. Recall that *Independence* and its inversion are as follows:

<i>Independence</i> : Any reasons I have to discount another's opinion about <i>p</i> must be independent of the reasoning behind my own initial belief about <i>p</i> that led to the disagreement.	<i>Inverted Independence</i> : Any reasons I have to accept another's opinion about <i>p</i> must be independent of the reasoning behind my own initial belief about <i>p</i> that led to the agreement.
--	--

Kelly notes, however, that without taking those first-order reasons into account, R has no other recourse for judging that S is indeed grossly ignorant when it comes to the Holocaust. Were R to ignore her assessment that the historical facts strongly confirm the occurrence of the Holocaust, then she would no longer take S's ignorance of them to be a handicap in judging whether the

Holocaust occurred. Such cases, argues Kelly, give us reason to reject *Independence*. Let's consider how this works out on the inversion.

1.1.3.1 Inverting the Problem: The Total Evidence View

Since there is some conciliation allowed on the Total Evidence View, the concerns raised above with the inversion of the Equal Weight View may also be something that will need to be addressed by adherents of the Total Evidence View. Similar to the Equal Weight View, on the Total Evidence View, one will still need to have a unified theory that provides a method for updating in cases of disagreement and agreement. However, a separate problem arises if the Total Evidence View entails the rejection of *Independence*. The reasons that motivate *Independence* also motivate *Inverted Independence*. Consider this slight modification of the *Holocaust Denier* case using the same reasoning. This also appears to motivate the rejection of Inverted Independence:

Holocaust Agreement: T possesses a great deal of evidence that the Holocaust occurred. Having evaluated the evidence, T judges it to strongly support that hypothesis, *H*. Upon adopting a high amount of credence that the Holocaust occurred, T encounters an individual, U (who is quite reliable in all other matters), who concurs that the Holocaust occurred. (Due to bad history teaching, U is grossly ignorant of the actual historical facts. Despite this, the incorrect historical facts he was taught also make it overwhelmingly likely that the Holocaust occurred).

In such instances, without taking her first-order reasons into account, it seems like T also has no other recourse for judging the ignorance of U. Since it appears that Kelly is not going to completely reject *Defeat* (and *Validation* by inversion), then in cases of where they both agree, T's confidence should increase upon the basis of U's agreement. Of course, this seems quite unreasonable in the Holocaust Agreement case. The apparent option for avoiding this unreasonable result is also to reject *Inverted Independence* (as with the rejection of *Independence* in light of the Holocaust Denier case). While I find the view compelling, if this is where the

Total Evidence View leads, it seems like there may still be a puzzle about how one should deal with other instances of agreement on the Total Evidence View.

1.1.4 The Justificationist View

Jennifer Lackey has been the major proponent of the Justificationist View. In many of the cases proposed, Lackey grants that the Equal Weight View proponents judged the situation and the corresponding response correctly. (In other words, there are times when peers should conciliate.)

There are cases, however, where conciliation is not appropriate. Consider the following case:

Elementary Math: Harry and I, who have been colleagues for the past six years, were drinking coffee at Starbucks and trying to determine how many people from our department will be attending the upcoming APA. I, reasoning aloud, say, ‘Well, Mark and Mary are going on Wednesday, and Sam and Stacey are going on Thursday, and since $2+2=4$, there will be four other members of our department at that conference.’ In response, Harry asserts, ‘But $2+2$ does not equal 4.’ (Lackey, 2010b, 283.)

In cases like Elementary Math, Lackey argues that we can be highly justified prior to the disagreement such that there is no real need for conciliation. We can contrast this with Christensen’s Restaurant case. In that instance, Christensen is not as highly justified regarding the restaurant bill’s correct split as Lackey is in her belief that “ $2+2=4$ ”. It is the level of justification one has for believing the proposition or hypothesis in question *before* discovering the disagreement that determines how one responds to the disagreement.

What Lackey’s view does, therefore, is to deny the *Independence* claim of the Equal Weight View. According to *Independence*, anything that is used to discount another’s opinion must be independent of the reasoning behind my own initial belief that led to the disagreement. Still, in these cases, it is the fact that the peer disagrees with a highly justified belief that was held before the disagreement that allows one to discount the peer’s position (Lackey, 2010a, 307). Lackey’s reasoning for this is similar to van Inwagen’s reason for his view. She proposes that individuals will have access to information about themselves, which they lack regarding

their peers. Personal information such as their own open-mindedness, level of alertness, focus on the question at hand, etc., is something that one has easy access to for themselves, but much less so for their epistemic peer. This access to personal information in combination with the highly justified belief prior to disagreement allows one to judge that the other person at the time is not performing as a rational epistemic agent should. To do so, however, one has to reject

Independence. The other person is thus re-evaluated according to the new epistemic standards and, as a result, loses the status of epistemic peerhood. While I think Lackey's position has some strong intuitive appeal, it also appears to lead to some puzzles when considering the inversion of the view.

1.1.4.1 Inverting the Problem: The Justificationist View

Because the Justificationist will agree in most cases with the Equal Weight View, in normal instances of agreement, the Justificationist View will face the same problems upon its inversion as those mentioned above for the Equal Weight View. In cases of highly justified beliefs, the Justificationist denies *Independence*. Given this, there is some reason to think that the Justificationist might also reject *Inverted Independence* in such cases. To see why this could be problematic, let's consider an inverted case similar to the one described above.

Suppose, for instance, S has some belief that is highly justified due to personal information such as her own open-mindedness, level of alertness, focus on the question at hand, etc. In cases of disagreement, the highly justified status allows S to re-evaluate a peer's position as a peer. Given this, can a highly justified belief not be used to evaluate a non-peer, who is in agreement, as a peer? Since the Justificationist rejects *Independence* and the fact that a peer disagrees with a highly justified belief that was held before the disagreement allows one to discount the other as a peer, why shouldn't the fact that a non-peer agrees with a belief that is

highly justified held before awareness of the agreement allow one to accept a non-peer as a peer?

For reference, recall that *Independence* and its inversion are as follows:

<i>Independence</i> : Any reasons I have to discount another's opinion about <i>p</i> must be independent of the reasoning behind my own initial belief about <i>p</i> that led to the disagreement.	<i>Inverted Independence</i> : Any reasons I have to accept another's opinion about <i>p</i> must be independent of the reasoning behind my own initial belief about <i>p</i> that led to the agreement.
--	--

This may be similar to *Holocaust Agreement*, but in this case, S has some strong intuition about the sincerity and reliability of her interlocutor (along with the other reasons listed above) that highly justify her belief.

Suppose S is highly justified and re-evaluates the other person as a peer because they are in agreement regarding the proposition or hypothesis in question. S does this because she rejects *Inverted Independence*. In that case, the fact that this person may not have good reasons for agreeing should not be problematic for the Justificationist. Because they are peers after the re-evaluation and the Justifications will accept the other theses of the Equal Weight View (and by inversion the inverted theses), it seems like S's confidence in her highly justified belief should increase in such an instance. Perhaps this is not the route those holding to a Justificationist view want to take, but it appears there are some questions that they must address to show that is not the case.

It seems that the inverse relationship between disagreement and agreement raises some interesting puzzles and questions for most major views that have been proposed in response to the problem of peer disagreement. In addition to the issues that arise in the more traditional approach to epistemology, there are also some formal concerns involving the views we have discussed so far that should be examined.

1.1.5 Formal Issues in Disagreement

Formal concerns of various types have been raised regarding the Equal Weight View, which resulted in the proposal of various Bayesian approaches for dealing with disagreement.¹⁴ These concerns may also apply to the other views discussed above depending on which of the Equal Weight View's three claims are accepted by the view in question. One concern of particular note, for both disagreement and agreement, involves the commutativity of evidence principle. This principle states the following:

Commutativity of Evidence Principle (CEP): To the extent that what it is reasonable for one to believe depends on one's total evidence, historical facts about the order in which that evidence is acquired make no difference to what it is reasonable for one to believe (Gabbay *et al.*, 2010, 451).

In essence, the principle states that the order in which one gains evidence should not change the degree to which it is reasonable for one to believe upon said evidence once the total evidence is gathered. This is generally considered an uncontroversial principle. Gardiner (2014) has noted quite clearly, however, that this principle could pose a problem for any view that requires conciliation.

Suppose that Ray and Steve are discussing the time at which a favorite local pub opens. Before discussing the issue, Ray has a 0.6 confidence the pub opens at 3:00 pm, and Steve has a 0.2 confidence the pub opens at 3:00 pm. Upon finding out they disagree, the two split-the-difference and adopt a 0.4 confidence that the pub opens at 3:00 pm. Shortly afterward, Ray talks to Jon, who has a 0.8 confidence that the pub opens at 3:00 pm. Given this, the two split-the-difference and adopt a confidence of 0.6 that the pub opens at 3:00 pm. Now, consider what would happen were Ray* to meet the same two peers in a different order.

¹⁴ Cf. Brössel and Eder (2014), Easwaran *et al.* (2016), Fidelson and Jehle (2009), Gardiner (2014), and Wilson (2010).

Let us say that Ray* has a 0.6 confidence the pub opens at 3:00 pm, and Jon, who has a 0.8 confidence the pub opens at 3:00 pm. Upon finding out they disagree, the two split-the-difference and adopt a 0.7 confidence that the pub opens at 3:00 pm. Shortly afterward, Ray* talks to Steve, who has a 0.2 confidence that the pub opens at 3:00 pm. Given this, the two split-the-difference and adopt a confidence of 0.45 that the pub opens at 3:00 pm. Notice the odd results summarized in Table II.

Table II. Summary of non-commutativity of conciliation

Person	Original credence	Credence of 1st Peer	Credence after 1st discussion	Credence of 2nd Peer	Credence after 2nd discussion
Ray	0.6	0.2	0.4	0.8	0.6
Ray*	0.6	0.8	0.7	0.2	0.45

Ray and Ray* both encounter the same evidence. Nonetheless, due to the order in which the evidence is presented and the standard rules of conciliation, each ends up with a different confidence regarding the time at which the pub opens by the time all of the evidence has been evaluated. Obviously, this sensitivity to the ordering of evidence is undesirable and points to some change that needs to be made in how the peers should conciliate. Depending on the method chosen, this same effect could find its way into the inverted process of dealing with agreement. Any method of dealing agreement must take the commutativity of evidence principle into account.

A natural response that some philosophers might want to take is to propose that some sort of Bayesian updating is the only way to deal with agreement, as such proposals have been made regarding dealing with disagreement.¹⁵ There is an equally natural resistance to this, however, by other philosophers who observe that the calculations involved in Bayesian updating can be quite

¹⁵ Cf. Brössel and Eder (2014) and Fidelson and Jehle (2009)

difficult and cumbersome. For this reason, the standard solution for conciliation has been to use a short-cut rule like linear averaging. Unfortunately, such simplistic methods will suffer from various formal problems, such as the non-commutativity of conciliation as described above.

One interesting proposal, made by Easwaran *et al.* (2016), is that some sort of short-cut rule like linear averaging could serve as a *heuristic* in real-world cases where Bayesian updating is not feasible. Such a short-cut rule should be more accurate than simple linear averaging while maintaining similar simplicity. In an attempt to formulate such a rule, Easwaran *et al.* introduced and defended the *Upco* (‘updating on the credences of others’) rule. They argue this rule achieves many of the positive goals of Bayesian updating, and note the rule “is equivalent to conditionalization when the likelihood ratios of the agents take a particular simple form,” the post-update credences respond synergistically where reasonable¹⁶, and the rule “preserves shared agreement about independence in a wide range of cases” (Easwaran *et al.*, 2016, 3). Supposing H is some hypothesis in which P has an initial credence $P(H) = p$ and a peer Q has an initial credence in the same hypothesis $Q(H) = q$, the simplified version of the rule for cases of disagreement with two peers, P and Q , who have credences for H and $\neg H$ states the peers should update to the following:

$$P(H^+) = Q(H^+) = \frac{pq}{pq + (1-p)(1-q)}.$$

We will refer to this simplified¹⁷ instance of the rule as *Upco-S_I*.

¹⁶ This is to account for cases discussed by Christensen (2009, 759) and (2011, 3n).

¹⁷ The more general *Upco* rule, for instances where peers $P, Q, R, \dots$ have credences over the partition $\{H_1, \dots, H_k\}$, with $P(H_i) = p_i$, $Q(H_i) = q_i$, etc., would be $P^+(H_i) = \frac{p_i q_i r_i \dots}{\sum_{j=1}^k p_j q_j r_j \dots}$ (Easwaran *et al.*, 2016, 3.).

$Upco-S_I$ is intended to mimic the Bayesian conditionalization where $P(H)$ is P's initial credence in H , and P learns that proposition E is true (and nothing else). P's new credence then becomes $P(H^+) = P(H | E)$. $P(H)$ is P's prior probability for H , and $P(H^+) = P(H | E)$ is her posterior probability. Upon learning that Q's credence in H is $Q(H) = q$, P treats the proposition $Q(H) = q$ as evidence and updates by conditionalization. P's new credence should be $P(H^+) = P(H | Q(H)=q)$. The Bayesian formulation of this would be the following:

$$P(H^+) = \frac{P(Q(H) = q|H) \cdot P(H)}{P(H) \cdot P(Q(H) = q|H) + P(\neg H) \cdot P(Q(H) = q|\neg H)}$$

$Upco$ is designed to give the same result as Bayesian conditionalization when P's likelihoods¹⁸ take a particular form. Specifically, the form Easwaran *et al.* are interested in is:

$$P(Q(H) = q|H) \propto q;$$

$$P(Q(H) = q|H) \propto 1 - q.^{19}$$

This basically just means that the terms on each side are proportional to each other. In the scenario above, it would be something like P thinks that Q is more likely to have higher credences in H when H is true, and more likely to report lower credences in H when H is false. Consider the simple idea they provide where P's likelihoods for Q's credences are linear to Q's credences (Easwaran *et al.*, 2016, 11-2.). Suppose that P knows that Q will report his credence rounded off to the nearest tenth. P's likelihoods would then be as follows:

$$\begin{aligned} P(Q(H) = 0|H) &= 0; \\ P(Q(H) = .1|H) &= 1/55; \\ P(Q(H) = .2|H) &= 2/55; \\ &\vdots \end{aligned}$$

¹⁸ These are P's beliefs about what credences Q will have, conditional on H and $\neg H$.

¹⁹ See Easwaran *et al.* (2016, 11-55) for further elaboration on this point.

$$P(Q(H) = .9|H) = 9/55;$$

$$P(Q(H) = 10|H) = 10/55;$$

P's likelihoods, conditional on $\neg H$, would follow a decreasing pattern proportional to what is described above.

When comparing the Upco rule to something like linear averaging, it does, at first glance, appear to be a useful heuristic for dealing with disagreement. Consider the results in Table III of the updating calculations with both methods for various instances of peer disagreement. In certain cases, the results are quite close to linear averaging (Case 1), and in others (Cases 3-4) they are much more comparable to Bayesian updating.

Table III. Summary of updating calculations of two peers with respective credences $\{p, q\}$ using Upco- S_1 and Linear Averaging.

	Case 1 {0.6, 0.4}	Case 2 {0.7, 0.4}	Case 3 {0.8, 0.4}	Case 4 {0.97, 0.96}
Upco- S_1	0.5	≈ 0.601	≈ 0.727	≈ 0.998
Linear Avg.	0.5	0.55	0.6	0.965

Case 4 shows that Upco can account for instances such as the one proposed by Christensen where a doctor trying to determine the drug dosage concludes with a confidence of 0.97 that the dosage is correct, while slightly weakening his confidence by taking into account his own fallibility with regards to such calculations. The doctor confers with an equally capable and self-aware colleague who arrives at a 0.96 confidence regarding the dosage. Christensen notes, that in such instances, the rational response is to increase one's confidence rather than splitting the difference (Christensen, 2009, 579). Upco appears, thus, to handle such cases much better than linear averaging. Additionally, Upco also accounts for CEP by allowing a simple modification of Upco- S_1 to include an additional peer for instances like the one summarized in Table II. We could use Upco- S_2 for such instances. Supposing that there are three peers, P, Q, and R and their respective credences for H are p, q, r . The formulation of Upco- S_2 would be:

$$P(H^+) = Q(H^+) = R(H^+) = \frac{pqr}{pqr + (1-p)(1-q)(1-r)}.$$

In cases such as the one with Ray, Steve, and Jon, who disagree about their favorite local pub's opening time, Ray would simply use *Upco-S₁* upon interacting with the first peer and would modify the rule to *Upco-S₂* upon discussing the matter with the second peer.

Unfortunately, Upco does not appear to be a useful solution for updating in some instances of agreement. Consider a case where two peers, P and Q, share a confidence of 0.2 regarding *H*. Upon finding out they agree, if they were to run their credences through *Upco-S₁*, the rule would say their updated credence should be ≈ 0.059 . This is not at all what we would expect. If I believe it is very *unlikely* that *H* is true, the fact that you agree with me will not make me believe that *H* is more *likely* than I previously believed. Quite the opposite, I think. I would actually think there is a good chance that it is slightly even ***more unlikely*** than I had thought previously. To calculate this correctly, we need to understand, instead, that P and Q share a confidence of 0.8 regarding $\neg H$. Were they to update using *Upco-S₁* after finding out they agree on $\neg H$, their updated credence should be ≈ 0.941 . However, this seems like a huge leap for one to make based on the opinion of a peer. Certainly, if one found out that an expert regarding $\neg H$ shared the same opinion about $\neg H$, then one might have good reason to update from 0.8 to 0.94.

This seems an excessive bit of updating merely on the basis of a single peer being in agreement. In the first case, we realize that agreement in instances where confidences are below 0.5 results in unexpected increases in confidence. This can be resolved by calculating on the agreement regarding $\neg H$ instead of *H*, but the increase in confidence is much greater than one would expect. Additionally, *Upco* assumes that peers' opinions should be weighted the same as one's own opinion. It provides no simple means for updating when the opinions of peers should be weighted differently than one's own.

1.2 Dealing with Agreement

All of the views discussed above will likely approach agreement differently, but there are some minimal requirements that any view will need to account for in cases of agreement. Drawing some insights from the discussion above, any proposed solution for dealing with peer agreement should achieve the following minimal requirements. The first requirement is that some sort of strengthening in confidence should occur in *some* instances of rational peer agreement. This is tied to our natural intuition that agreement is a good thing, which strengthens our confidence. The second requirement is that upon finding out that they agree, peers who share a *low* confidence in some hypothesis shouldn't become *more* confident that the hypothesis is true. This idea is tied to the example of two peers who share a low confidence of 0.2 regarding some hypothesis where an increase in confidence upon discovering they agree seems like an odd response to agreement. Finally, any solution should account for the commutativity of evidence principle. That is, the solution should be *diachronic* in nature and adhere to CEP to allow for more than one instance of agreement over time.

1.2.1 A Possible Solution

One point of the discussion above is that, initially, some sort of increase in confidence in cases of agreement seems like the analogous response to *splitting-the-difference* in cases of disagreement. This, however, should be balanced with considerations for cases where the peers agree that the probability of some hypothesis or proposition is quite low. One option worth considering is that, at least in some instances, agreement merely focuses or reassures the peers of the initially established credence. Why might this be the case? Are there some rather natural cases where this seems to be the case? Perhaps. Consider the following.

If I were to remove a coin from my pocket and were to ask myself, “What are the chances that this coin will land on heads were I to flip it?” What is the obvious response? Given my evidence, I should naturally set my confidence in the hypothesis, H = “this coin will land on heads were I to flip it,” to 0.5. There should be very little controversy with my confidence in H given my evidence. Now, suppose I recognize there could be some things that might affect how the coin performs in coin-flipping. The coin could be marred in some way or have some odd property that causes it to land on heads more often than on tails. Because of this, I take the coin out again. I look at it very closely. I weigh it carefully in my hand. I compare it to other coins. I do my due diligence to confirm there is nothing odd about this coin that would cause me to think it would behave differently than expected. After all of this, what confidence should I have in H ? The obvious answer is still 0.5.²⁰ Though I have not gained any new evidence that would change my confidence in H , I have gained a reassurance of sorts regarding my original assessment. This, at least, seems to represent some instances of agreement. More specifically, it lines up nicely with instances of agreement where our confidence is low, *e.g.*, lies between 0 and 0.5, in which we would not expect agreement to lead to increased confidence.

Of course, the analogy above does not capture the instances of agreement where our confidence is high, *e.g.*, lies between 0.5 and 1. Those are the instances where some sort of increase in confidence seems appropriate. Still, we also need some means of limiting the exponential aspects of increasing confidence by large n factors in such cases. By this, I am proposing that our confidence increases at a diminishing rate as the instances of agreement accumulate. Suppose then that we have instances of agreement with peers $P, Q, R, \dots$ who have corresponding confidences in some hypothesis, H , over the partition $\{H_1, \dots, H_k\}$. Were P to

²⁰ This example is inspired from examples often given by Henry Kyburg.

update her confidence in H on the basis of Q 's confidence, followed by R 's confidence, followed by S 's confidence, and so on, the increase in P 's confidence from updating on the basis of S 's confidence should result in less of a change than the increase in P 's confidence from updating on the basis of Q 's confidence.

In an attempt to provide a model that meets the three requirements listed above, consider what I will refer to as the *Agreement Updating* (AU) rule. This model is the mathematical formulation of a simple idea of what updating on the opinions of others generally looks like in the real world. We begin with the confidence of a person, P , about a hypothesis, H . When P encounter someone, Q , who agrees with her about H , we add the product of Q 's confidence in H , P 's uncertainty about H , and Q 's weighting value to P 's original confidence about H . Why? Well, Q 's confidence is required because that is what we think should cause the increase in P 's confidence. However, there should be a relevant limiting factor that prevents how much of a boost Q 's confidence provides. That is the role that P 's uncertainty plays in the model. P 's uncertainty is calculated by subtracting her certainty (or confidence), 0.6, from 1. As P 's confidence increases (and her uncertainty decreases), the boosts from other's opinions will diminish. The weighting factor allows us some control of how much we judge the opinions of others should weigh in comparison to P 's original confidence. If we are assuming that peers' opinions should weigh equally (as with the Equal Weight View), then we simply assign a weight of 1 to the opinion.

So, suppose that P 's confidence is 0.6, Q 's confidence is 0.6, and Q 's weighting value is 1. In this case, P 's uncertainty would be 0.4^{21} , and the calculation would look like this:

$$P's \text{ new confidence} = 0.6 + 1(0.6 \times 0.4) = 0.84$$

²¹ Uncertainty is calculated as $1 - \text{Pr}$. So, in this case $1 - 0.6 = 0.4$.

Suppose that P immediately encounters another peer who agrees with her original confidence. She would follow the same pattern with the confidence and weight values of the new peer. The calculation would follow the same pattern with the only difference that we should assume P's new confidence (and thus her new uncertainty). So, if another peer, R, were encountered who had a confidence of 0.6, the new calculation would be like this:

$$P's \text{ newer confidence} = 0.84 + 1(0.6 \times 0.16) = 0.936$$

Notice that P's confidence boost decrease as her confidence increases. We can therefore have a long series of updates like this but the increase in confidence diminishes as P's confidence draws closer to 1. Thus, we avoid the unrealistic problem of eventually arriving at actual certainty.

Now that we have the basic intuition to build on, we can formalize this intuition. A simplified version of AU, let's call this AU_S is formalized as follows:

$$C_n = C_{n-1} + k \left[\sum_{i=1}^m O_m(1 - C_{m-1}) \right] . \quad (AU_S)^{22}$$

For this rule, let C represent P's confidence regarding H , where C_0 is P's initial confidence before interacting with any peers. The partition of $\{O_1, \dots, O_k\}$ represents the confidence in H of the other peers with whom P interacts. For instance, Q's confidence in H might be $Q(H) = O_1$, R's confidence in H might be $R(H) = O_2$, S's confidence in H might be $S(H) = O_3$, and so on. Also, let l = the number of others being considered with the requirement that $l > 0$. Assume that $n = l$ and $m = l$. Let k = the weight applied to the confidences of other peers. (I will discuss this in more detail shortly.) The variables following the summation will account for the confidences of

²² Many thanks to my friend Kyle Burnard-Fullington for his help in formalizing my ideas with the result being this model. A more complex version of the rule and why it might be needed will be discussed later in the chapter.

the peers and their uncertainty. The AU_S rule can be adapted to instances where only two peers are involved or instances where multiple peers are involved.

Let us say that AU_2 is an instance where the peers P and Q are involved in the updating process. AU_2 is formulated as follows:

$$C_1 = C_0 + k[O_1(1 - C_0)].$$

Were we to adapt the formula to account for peers P, Q, and R. We would simply modify the rule to a rule we will call AU_3 :

$$C_3 = C_0 + k[O_1(1 - C_0) + O_2(1 - C_1) + O_3(1 - C_2)].$$

AU can easily be modified to adapt for a multitude of peers.

As should be quite obvious, AU is additive in nature. Each additional peer's confidence, however, has a diminishing effect upon the increase of P's updated confidence so that, assuming confidences are $0 < \text{and} < 1$, P's updated confidence will never reach 1. Though it may increasingly draw closer, P's updated confidence will never actually reach certainty. In addition, the constant k allows for a determination of how heavily the peers' confidences should be weighted. Some argument could be made that the peers' confidences should be weighted as heavily as P's own confidence. In such a case, $k=1$ might seem to be the appropriate weight. We might also set k to equal something much lower, say 0.25, were we to consider the confidences of the others to be less weighty than P's confidence. I will make no argument at this point as to what value k should receive.

To see how these calculations work out, let us suppose that P's initial confidence regarding H is equal to 0.6 and her peer, Q, who also agrees that H is true, has a confidence of 0.6. Let us also assume that two other peers, R and S, agree regarding H , and they share the same confidence in H as P and Q. In this case, let us also assume that *Inverse Equal-Weight* is true. Given this, we

will let $k = 1$. So, we will end up with the following values: $C_0 = P(H) = 0.6$, $O_1 = Q(H) = 0.6$, $O_2 = R(H) = 0.6$, $O_3 = S(H) = 0.6$, and $k=1$. Assuming these numbers, the results of updating for these peers are as follows:

Table IV. Summary of agreement updating calculations with AU_S

$C_1 = C_0 + k[O_1(1 - C_0)]$	$= 0.84$
$C_2 = C_0 + k[O_1(1 - C_0) + O_2(1 - C_1)]$	$= 0.936$
$C_3 = C_0 + k[O_1(1 - C_0) + O_2(1 - C_1) + O_3(1 - C_2)]$	≈ 0.944

1.2.2 Limitations, Special Considerations, Interesting Developments

As it stands, AU_S is only applicable and works as intended in instances of agreement where the peer agreement should lead to some increase in confidence. As was proposed above, AU_S will be limited to instances of agreement where the peers' confidence lies between 0.5 and 1. It is possible that we could build something into the model, such that upon encountering confidences less than 0.5 the model adapts and responds appropriately. To do this, we might assume the same things above for AU_S to formulate an updated rule, AU , to account for this:

$$C_n = C_{n-1} + k \left[\sum_{i=1}^m O_m(1 - C_{m-1})f(O_m) \right] \quad . \quad (AU)$$

Additionally, we would need to include the following piecewise function

$$f(x) \begin{cases} -1 & \text{if } x < 0.5 \\ 0 & \text{if } x = 0.5 \\ 1 & \text{if } x > 0.5 \end{cases} \quad .$$

With this function's inclusion, the updated rule would automatically adjust for instances where confidences are $0.5 \geq O_x \leq .05$. For the sake of simplicity, in most cases of peer agreement, we can just assume this limitation and continue to use AU_S .

Another limitation of AU_S is that the model does not account for variances in how much individuals are weighted. One might be interested in using AU_S to account for cases of agreement where there is some difference in trustworthiness, knowledge, or some other variable

that applies to the individuals under consideration. For instance, we might like to model agreement between peers, almost-peers (individuals with confidences that are slightly above or below), superiors, or inferiors. The current formulation of the model only allows for the constant k to apply the same weight to all individuals considered in the updating calculation.

There is a basic assumption made for the use of AU_S and what is considered “agreement.” I am supposing instances of agreement are ones in which peers share the same confidence in some hypothesis or proposition. Given this, an instance where two peers have a confidence of 0.6 in H counts as agreement. An instance where one individual has a confidence of 0.6 and another has a confidence of 0.62, one should note, does not count as agreement (though it might be referred to as “close agreement” or something along those lines). Now, an argument could be made that variances in confidence to some degree (say $0.001 \geq$ and ≤ 0.055) should be considered agreement. If such variances are allowed, there may be a concern that problems with CEP might arise in cases of diachronic updating that involve multiple individuals, and this certainly would be the case with AU_S . Fortunately, the full and slightly more complex rule, AU , allows for such considerations and avoids any problems with CEP.

Finally, one interesting aspect of AU is the fact that it does allow for updating in cases of disagreement as well as agreement. Suppose again the case of the four peers, P, Q, R, and S discussed above. In this instance, let us modify the values of their confidence in the truth of H so that the peers disagree. We will assume the following values: $C_0 = P(H) = 0.6$, $O_1 = Q(H) = 0.2$, $O_2 = R(H) = 0.5$, $O_3 = S(H) = 0.8$, and $k=1$ ²³. The resulting updating using AU would be the following:

²³ What value one assigned to k might depend if one holds to something like Equal Weight or some other view. The model is mostly view agnostic in that sense and allows one to vary how heavily the confidences of others should be weighted.

Table V. Summary of disagreement updating calculations with AU

$C_1 = C_0 + k[O_1(1 - C_0)(-1)]$	$= 0.52$
$C_2 = C_0 + k[O_1(1 - C_0)(-1) + O_2(1 - C_1)(0)]$	$= 0.52$
$C_3 = C_0 + k[O_1(1 - C_0)(-1) + O_2(1 - C_1)(0) + O_3(1 - C_2)(1)]$	$= 0.904$

The results appear to be along the lines of what we would desire in cases of disagreement, can account for instances where confidences range anywhere between 0 and 1, and also avoid any problems with CEP.²⁴

1.3 Conclusion

The tremendous focus in the Peer Disagreement literature on how to resolve the problem of disagreement has left somewhat of a void regarding the equally important aspect of peer agreement. I have examined four specific views on the problem of peer disagreement. While this is by no means an exhaustive list, I think the analysis has raised enough questions and brought forth enough puzzles that should motivate some more work on development of a unified theory that accounts for instances of peer agreement and disagreement. All of the positions examined in this paper appear to lead, at least to some degree, to various entailed results that need further consideration and discussion. Though it was not feasible to include all possible views on rational peer disagreement, the ones not included were similar enough to those examined in this paper that I think it very likely that they would also encounter similar questions or puzzles to those I have raised. Much more could (and should) be said by those who hold to the positions mentioned in this paper (or any variations of those positions) in regard to the development of a unified theory and methods for dealing with peer agreement.

²⁴ One concern that might arise in using AU to update in cases of disagreement is that we should not be updating on someone's low confidence that a hypothesis, H , is true. Rather, we should be updating on their confidence that $\sim H$ is true. So, in the case above the value for O_2 should not be $Q(H) = 0.2$. Rather, it should be $O_2 = Q(\sim H) = 0.8$. This could be achieved with a slightly modified piecewise function but is not needed for the current discussion.

In addition to raising questions that I hope will help advance the literature, I have outlined what should be the minimal requirements for any proposed solution for dealing with instances of peer agreement, and I have proposed a possible method that could be used in such instances. More importantly, I hoped to have provided a model that could be used on any of the views discussed. While there may be some limitations to this rule, it appears to be a friendly solution for the non-ideal epistemic agent. Additionally, the AU rule can deal quite nicely with instances of agreement, instances of close agreement, and instances of disagreement. The rule may be limited to such instances as described above. There are many other considerations that might need to be taken into account regarding other types of agreement outside of the specific focus of rational peer agreement, which might require some modifications to the rule proposed in this chapter (or may require a completely separate solution). Nonetheless, I hope what I have proposed here will serve as a beginning to further discussion and analysis.

Chapter 2: How Should We Respond To Agreement?

It is practically impossible to ignore the extent to which agreement plays a positive role in so much of our belief development. Just as disagreement seems to give us reason to believe that something has gone wrong in the belief-forming process or at least in the estimation of how confident one should be in a belief, agreement seems to indicate that something has gone right in the process. In §1 of Chapter 1, I proposed that there is an inverse symmetry to some degree between the *problem of peer disagreement* and what I call the *assurance of peer agreement*, which has of yet been mostly overlooked in the literature. I have discussed this inverse symmetry in the context of peers and how agreement might affect our confidence in our beliefs, but the discussion of agreement and disagreement, unfortunately, can (and perhaps should) be taken beyond the contexts of epistemic peers. Indeed, some have hinted and argued that it might be impossible to establish whether two individuals are actually peers.²⁵ If this is the case, then we must definitely expand our discussion. In the real world outside of the epistemology classroom, we're actually encountering cases of epistemic superiors, inferiors, and peers (or maybe it's something like "approximate peers" that get as close to peerhood as we could really hope). In this chapter, I will discuss various types of agreement that raise different questions and concerns beyond what we see in standard cases of peer agreement. I will then discuss the model proposed in chapter one, which can be used for updating confidences in beliefs given these instances of agreement. Finally, I will address some concerns regarding this model and respond to those concerns.

²⁵ See Frances (2014), King (2012), Matheson and Frances (2018), and Mulligan (2015) (among others) for some discussion on this.

2.1 Peers, Inferiors, and Superiors

Suppose that a friend and I are in disagreement about some proposition. If my friend and I are both generally correct about judging things of this sort, and although neither of us gets it wrong very often, we get it wrong about the same number of times, then we would generally be considered epistemic peers. This fact makes it a problem when we disagree, and if we recognize this fact, we recognize the problem. Our disagreement indicates that one of us must have made an error in judgment. If one or the other is a better judge of the sort of proposition in question (and we recognize this), then there appears to be reason for us to weigh one position more heavily than the other. But defining the exact requirements for epistemic peerhood has proven to be a significant challenge for epistemologists. The initial conception that philosophers generally have is that peers are something like epistemic equals. Equals in what ways? Generally, the claim is that they must at least be equals in cognitive abilities, in experience or background knowledge, and in access to the relevant body of evidence. While that is a good place to start, there are many other factors that one might easily argue should not be left out of this discussion, *i.e.*, the agent's open-mindedness, ability to focus on the question at hand, etc. Nonetheless, it is the fact that they share some sort of equality and that they recognize this equality, which motivated the key aspects of the discussion surrounding the problem of rational peer disagreement. For the purposes of this project, we will assume the following:

Epistemic Peers: R and S are epistemic peers regarding p if and only if they are equals in epistemic position for judging the confidence one should have regarding the truth-value of p .

I will not argue whether or not some specific set of characteristics are all there is to peerhood, but the definition above should serve as a good enough starting point for what I will argue going forward.

How and whether epistemic peers should adjust the confidence they have in their beliefs in cases of rational disagreement is the central “sticking point” of the problem of peer disagreement. I have argued that whatever the solution for the problem may be, the method of dealing with disagreement should lead us to something similar to an inverse method for adjusting the confidence in beliefs in instances of peer agreement, given the apparent inverse natures of the *problem of rational disagreement* and the *assurance of rational agreement*. If disagreement and agreement do indeed share some sort of inverse symmetry, then philosophers should be able to apply strategies for dealing with agreement and disagreement that retain consistency in regard to the theses that motivate these strategies. For instance, those who hold to the Equal Weight View (EWV) regarding peer disagreement, *e.g.*, Christensen (2007), Elga (2007), and Feldman (2006), maintain specific theses that motivate the requirement of conciliation between peers in cases of disagreement. In §1.1 of Chapter 1, I proposed that in inverse cases of agreement, those holding to EWV should require an increase in confidence for peers who are in agreement. Likewise, those who hold to a strong Steadfast View (SV) regarding disagreement, *e.g.*, van Inwagen (1996), reason that there is no motivation for peers to conciliate in cases of disagreement. Given this, those who hold to SV have no motivation for an increase in confidence for peers who are in agreement.²⁶ The topic of disagreement between superiors and inferiors has generated somewhat less discussion. We might hope for a similar inverse relationship between disagreement and agreement between superiors and inferiors. Before attempting to establish this, we need a working definition of epistemic superiors and inferiors.

²⁶ Given this, what I say for the rest of this chapter (unless explicitly stated otherwise) assumes that some degree of conciliationism is the view being assumed.

The following definitions of *superior* and *inferior* should suffice for the discussion that follows:

Epistemic Superior: R is a superior to S in regard to p if and only if R is in a superior epistemic position to S for judging the confidence one should have regarding the truth-value of p .

Epistemic Inferior: S is an inferior to R in regard to p if and only if S is an inferior epistemic position to R for judging the confidence one should have regarding the truth-value of p .

I concede that there may be various noteworthy epistemic factors in which one might be superior or inferior. Hopefully, the definitions above will allow us to side-step discussions about the exact mix needed of those factors, and we can instead consider how superiors and inferiors should each respond to cases of disagreement and (more importantly for this chapter) agreement.

2.1.1 Disagreement Between Superiors and Inferiors

Before investigating the proper response to agreement between superiors and inferiors, we should first discuss disagreement in such instances. For many, there is a quite natural assumption that an epistemic inferior will defer to an epistemic superior. This is most commonly described in cases involving novices (or lay-persons) and experts. In cases of disagreement, the inferior's confidence in p is expected to update in the direction of the superior's confidence. Depending on the specifics of the situation, this may mean that the inferior merely updates his confidence, but it could also mean the inferior updates his confidence to such a degree that he should abstain from belief or change his belief. Consider a case in which Bill is considering the truth of p . Bill does his due diligence. He first attempts to recall all that he's heard regarding the truth of p and attempts to come up with any good reasons why he shouldn't believe that p . Bill then asks around amongst his peers to see what general information can be found about the truth of p from that source. Finally, he then plugs a search query, "Is it true that p ?", into Google – the fount of

all wisdom – and gathers more evidence there. Bill then concludes that there is a very good chance that p . The following day, Bill unexpectedly encounters his friend Rufus who has just returned from a trip. Bill is quite aware that Rufus is much more intelligent than he is, has access to much more evidence than Bill does, and in general, Rufus tends to end up arriving at the correct conclusions regarding topics in the domain of p much more often than Bill does. When Bill asks Rufus about it, he learns that Rufus believes that $\sim p$. For some of those who hold to some sort of *conciliationist* view, it might seem quite clear that Bill should at least decrease his confidence in the belief that p . Perhaps, he should even withhold belief or move closer to Rufus's view on p .

Should we suppose that it is always the case that inferiors should defer to superiors? Alvin Goldman (2001) has noted, specifically in the literature on expert testimony, that some philosophers have argued for a sort of blind epistemic trust in cases in which novices (inferiors) defer to experts (superiors).²⁷ Goldman proposes that experts should meet certain requirements before such authority is granted to them. He notes that novices have five sources for gathering evidence as to whether or not an expert should be trusted when the novices themselves do not have enough knowledge in the domain to evaluate the credibility of the expert on their own:

- A) Arguments presented by the contending experts to support their own views and critique their rivals' views.
- (B) Agreement from additional putative experts on one side or other of the subject in question.
- (C) Appraisals by “meta-experts” of the experts' expertise (including appraisals reflected in formal credentials earned by the experts).
- (D) Evidence of the experts' interests and biases vis-a-vis the question at issue.

²⁷ Cf Burge (1993), Foley (1994), and Hardwig (1985, 1991).

(E) Evidence of the experts' past "track-records" (Goldman, 2001, 93).²⁸

This still leads us, however, to inferiors updating in the direction of superiors once evidence of the experts' credentials has been confirmed. There are instances of disagreement between inferiors and superiors, however, that deserve a closer look.

For a superior to be in a better epistemic position regarding p , one of the factors generally required is *competence* in the domain in of p . While competence is a somewhat tricky concept to define, we could say something like the following for a working definition of competence:

Competence: R is competent in subject matter M when R is sufficiently qualified to make judgements in which the central question being asked mostly involves M.

Thus, competence can serve as an epistemic indicator that an agent's beliefs should carry some epistemic weight. This is one reason why experts' opinions, specifically regarding issues in the domain over which they hold a high level of competency, are generally held to carry a certain amount of authority. Competency, however, does not prove that one will judge correctly. Rather, the competent person will generally judge correctly more often than an incompetent person. A competent mechanic will be able to judge whether a car she has never seen before is running as it should. Whereas I might be able to judge that *my* car is running as it should (since I am fairly familiar with it), it would be a complete guess if I were to judge the health of some random car with which I had no experience. There are just too many subtle issues that I don't have the training and experience to spot. Nonetheless, I could on occasion get it right in such a situation.

That being said, it is perfectly natural for some inferiors to be as competent in the same domain, to varying degrees, as someone who might generally be considered a superior. Consider the somewhat familiar academic scenario of a veteran professor interacting with graduate

²⁸ Goldman reiterates this in a more recent publication (2018) though he does not go into the same level of detail.

students. While the professor may be said to have a higher level of competence *about* her overarching field of study, *i.e.*, she may be able to identify specific concepts, theories, positions, etc. and she may be qualified judge questions *about* the domain, that does not mean she will be more competent than her students in judging specific questions involving that domain. Nevertheless, the professor may be equally as competent or incompetent as her students regarding some questions involving her field of study. Perhaps there are too many unsettled questions in the domain. Perhaps the domain is rather new and unexplored. Perhaps some surprising and unexpected changes develop in the field that throw previous assumptions into doubt. Perhaps something new and never seen before arises in the domain and there is insufficient evidence upon which to anticipate accurate judgements. In such cases, the professor could quite simply overlook or ignore something that is grasped by the students. In instances such as these, it may be more accurate to describe the professor (the superior) as merely less incompetent than it is to say she is competent. Since no one involved is competent and no one's opinion greatly outweighs others due to a competence factor, were a disagreement to arise, the students' opinions should be considered high-order evidence that should be factored into the professor's confidence in her belief regarding the proposition over which they disagree.

Another way of motivating such updating is by noting that there are plausible cases in which a superior may only be in a better epistemic position due to evidence, cognitive abilities, or some other positive epistemic factor, and she is not able to "account for the disagreement via her disagreeing party's evidential [or cognitive] inferiority"²⁹ (Priest, 2019, 271). Objectively speaking, the superior may have no reason to take an inferior's opinion seriously. However, if

29 Or inferiority due to some other epistemic factor.

she cannot explain the dispute from her own superiority, this might not be the case.³⁰ Thomas Kelly makes a similar point regarding scenarios involving peers by noting:

One should give some weight to one's peer's opinion even when from the God's eye point of view one has evaluated the evidence correctly and he has not. But why? Exactly because one does not occupy the God's eye point of view with respect to the question of who has evaluated the evidence correctly and who has not (Kelly, 2010, 57).

Even though Kelly is specifically referencing instances involving peers, this also applies to instances involving superiors and inferiors. According to this line of thinking, there may be cases where her superiority is clearly evident to the superior, but when that is not the case and the superior does not recognize their own superiority, she should update in the direction of her inferior.

Priest (2016) makes a useful distinction between *weak superiority* and *strict superiority*. In case of weak superiority, the superior is only superior in one factor. For instance, Bill and Ted may be equals in all epistemic factors except evidence. Bill has access to certain evidence that Ted lacks. In this case, Bill is the weak evidential superior. Alternatively, Bill and Ted could be equals in all epistemic factors except cognitive ability, and Ted just happens to have greater cognitive abilities than Bill. In this case, Ted is the weak cognitive superior. Let's say, however, that Rufus has greater cognitive abilities and access to more evidence (evidence that is relevant to some proposition in question). In this case, we would say that Rufus is the strict superior of Bill and Ted.

Consider a rather plausible instance where two friends, Ray and Jon, are planning to meet up at a local restaurant for dinner. Ray and Jon are cognitive equals. They are both familiar with the local restaurant's hours of operation, since they have been there many times, both

30 Maura Priest offers one of the few analyses made of this position and makes a pretty compelling case for the position in her 2019 article.

independently and together. In trying to schedule what time they are going to meet at the restaurant, they cannot agree as to what time the restaurant closes. Ray believes that the restaurant closes at 6:00 pm, and Jon believes that it closes later than 6:00 pm. Ray believes that the restaurant closes earlier because he heard an announcement on the news that local restaurants were running at limited hours due to a global pandemic. It was a pretty widely circulated announcement of which most people in town are aware. It just so happens, however, that Jon has been unusually distracted this week and totally missed the announcements. In this case, Ray is Jon's weak evidential superior. Despite this, Ray is not aware of Jon's ignorance of this matter and naturally expects his friend to have heard about the change in hours. In addition, it's quite possible that the original local mandate had changed. (After all, such inconsistencies and changes are common during a global pandemic.) Given this, it seems like Ray has some reason to lower his confidence in the belief that the restaurant closes at 6:00 pm.

It does not take a lot of imagination to come up with other variations in which the two agents are equals in all aspects except one (cases of weak superiority) that might lead to the superior updating. More interesting, perhaps, are cases where the superior is strictly superior that leads to similar updating. Suppose for instance that a veteran professional software developer at a large software company is giving a talk on software development to a group of high school students who are all aspiring developers. Our veteran developer is incredibly bright, has degrees in software development, and has been in the development industry for quite a while. He has seen the rise and fall in popularity of a number of important coding languages, and he has worked with all of those languages to a reasonable degree. He understands the best coding methodologies and the different approaches to solving problems one would take in each of the languages, and he understands which language works best for different types of applications. He

even helped in the conception of a couple of these coding languages. We would say the veteran developer has *strict superiority* in the domain of software development in comparison to the students.

Now suppose that as part of his talk to the high school students, our veteran developer is describing a scenario involving a particular problem that software developers often face. His goal is to get his audience to think through the scenario and offer possible solutions. He has faced this problem himself many times and knows the best way to deal with it. He has also given this talk before on multiple occasions, and he has helped coax his audience through thinking about the problem and arriving at the most effective way of addressing the problem. Over the years, high school students have proposed various solutions that “get the job done”, but they eventually come to realize that none of those solutions work quite as well as the solution our veteran developer uses. Each time the veteran eventually guides them to proposing his better solution on their own. On this occasion, however, when our veteran gets to the point in the talk where he has raised the problem and asked for suggestions for the best way to solve the problem, after some thought, a young student, who overall is pretty average for the group, raises her hand and offers her suggestion for the best solution. She then proceeds to propose a completely unorthodox solution to the problem that our veteran developer had never actually considered. Upon initial reflection on the proposed solution, our veteran developer cannot think of any reason that the proposed solution is any worse than his “best” solution. Indeed, there is a chance that it could be better. If this is the case, then the veteran may have reason to lower his confidence in his belief regarding the best solution and to reevaluate his evidence.

Of course, all of this does not mean all instances of inferior-superior disagreement should lead to updating by the superior. There are various scenarios in which an inferior is incompetent

with regard to the domain of p . An inferior may just lack competence due to having limited access to the relevant evidence. In such a case, we would say the inferior is incompetent due to ignorance. An inferior may have access to the relevant evidence, though, and still remain incompetent. There are those who have access to evidence and are not lacking in cognitive abilities, yet they willfully or carelessly believe in the most unreliable way. Such people are just as likely to have false beliefs as they are to have true beliefs, and there seems little rhyme or reason for holding the beliefs they hold. The terms *madmen* and *fools* come to mind as apt descriptors for epistemic agents of this sort. Yet, as odd as it might sound, incompetence might also be intentional in nature. In the 21st century world of online interactions and social media, a common fixture of many online platforms is the Internet troll. Though they might often have access to relevant evidence and are often not lacking in cognitive abilities, trolls can purposely disagree out of spite, boredom, or for other malicious purposes, *i.e.*, without good reason.³¹ Purposeful incompetence of this sort, just as much as incompetence due to ignorance or other factors, grants superiors grounds for ignoring the opinions of inferiors. While the disagreement between inferiors and superiors is interesting on its own, I want to focus on its inverse, agreement between inferiors and superiors, going forward.

Given what I have said so far in this section, we can sum up the appropriate response to disagreement between superiors and inferiors with the following:

- 1) In cases of *disagreement* between superiors and incompetent inferiors, the inferiors have reason to lower their confidence in their currently held belief.

³¹ While there are other variations of trollism that may not entail such incompetence, what I have described is quite common. For more on Internet Trolls in this context, see the following: Fuller, Glen, Jason Wilson, and Christian McCrea. 2013. "Troll Theory?: Issue 22: Trolls and the Negative Space of the Internet." *Fibreculture Journal* 22: 1–15., and March, Evita. 2019. "Psychopathy, Sadism, Empathy, and the Motivation to Cause Harm: New Evidence Confirms Malevolent Nature of the Internet Troll." *Personality and Individual Differences* 141: 133–37.

2) In cases of *disagreement* between superiors and competent inferiors, the superiors and inferiors have reason to lower their confidence in their currently held belief (though the degree to which they each do this may not be commensurate).

We must now consider if the following is also the case:

1a) In cases of *agreement* between superiors and incompetent inferiors, the inferiors have reason to increase their confidence in their currently held belief.

2a) In cases of *agreement* between superiors and competent inferiors, the superiors and inferiors have reason to increase their confidence in their currently held belief (though the degree to which they each do this may not be commensurate).

2.2 Agreement Between Superiors and Inferiors

I argued in Chapter 1 that, if disagreement (at least in cases of peer disagreement) should lead to lowering one's confidence in a belief, then, *ceteris paribus*, agreement should lead to increasing one's confidence in a belief. In the case of disagreement between superiors and incompetent inferiors, it is natural to consider the superior's better epistemic position as higher-order evidence³² that should be considered in light of the disagreement and should motivate an incompetent inferior to lower her confidence. Likewise, in cases of agreement, we might expect a superior's better epistemic position to serve as a confirmation of the inferior's belief that *p*, which should lead the inferior to increase her confidence in the belief that *p*. So, let us suppose that Mary is a novice car enthusiast who is looking into buying a new car. She wants something that is easy to modify for increased performance. Although she is intelligent and picks up new skills quickly, Mary has little experience in modifying cars. She's considering two different cars,

³² By this I mean something similar to Whiting's definition in which *higher-order evidence* is "evidence which bears on whether one is able to assess or respond to one's evidence concerning a proposition" (Whiting, 2019, 246). This stands in contrast with *first-order* evidence that is considered "evidence which bears on whether a proposition is true, that is, which indicates or makes it likely that a proposition is (or is not) true" (*ibid*). Similar definitions and usages are found in Christensen (2010), DiPaolo (2018), Feldman (2005), Kelly (2010), Lasonen-Aarnio (2014), Schechter (2013), Sliwa and Horowitz (2015), and many others.

a Subaru and a Nissan. After looking at the two cars and based on some vague background information she has heard about Subarus and Nissans, Mary believes that the Subaru will be easier to modify. Mary calls up a friend, Lisa, who is a professional mechanic who works on imported Japanese cars and who spends a significant amount of time modifying those types of cars for customers. Lisa tells Mary that the Subaru is definitely an easier car to modify. While it could be the case that Lisa is wrong, Mary has good reason to increase her confidence that the Subaru will be easier to modify. Due to Mary's incompetence regarding the question at hand, Lisa has no reason to change her position regarding the question of which car is easier to modify. These results will likely seem uncontroversial to most people and match what was proposed in 1a above. Let us consider, however, cases where the inferior is competent.

Recall that in disagreement between a superior and a competent inferior, when the superior has no reason to think her own superiority explains the inferior's disagreement, then confidence reduction by both parties is the appropriate response, although the superior would reduce confidence to a lesser degree than the inferior. Now let's consider a case parallel to the one above with Ray and Jon. In this scenario, Rebekah and Jan are planning to meet up at a local restaurant for dinner. Rebekah and Jan are cognitive equals. They are both familiar with the local restaurant's hours of operation. In trying to schedule what time they are going to meet at the restaurant, they agree that the restaurant closes at 10:00 pm. Rebekah believes that the restaurant closes at 10:00 pm because she heard an announcement on the news that local restaurants were no longer running at limited hours due to global pandemic restrictions being lifted. The announcement was pretty widely circulated, and most people in town were aware of it. It just so happens, however, that Jan is unaware of the announcement, and she also did not realize that the restaurant had previously been closing at an earlier time. Rebekah is not aware of Jan's

ignorance of this matter and naturally expects her friend to have heard about the change in hours. Rebekah appears to have no reason to think her own superiority explains Jan's agreement. Given this, it seems that both Rebekah and Jan have reason to increase their confidence in the belief that the restaurant will close at 10:00 pm. That being said, it does not seem that they have reason to both update by the same amount. In such cases, the superior would increase her confidence to a lesser degree than her inferior would. This corresponds with what Priest says regarding disagreement. She notes that, in cases of disagreement with competent inferiors, superiors should downgrade *less* than they would in cases of peer disagreement (Priest, 2016, 274). This, however, only shows what we might expect in cases of *weak superiority* agreement, where the superior is only in a better cognitive position, only in a better evidential position, etc. What happens in cases of strict superior agreement? Before considering this, there are relevant cases of peer agreement that might give us some insight into the appropriate response to strict superior agreement.

2.2.1 Peer Agreement: Assorted Cases

For the most part, cases of peer disagreement have generally been proposed with mostly ordinary scenarios involving relatively straightforward propositions or hypotheses in dispute. Consider, for instance, Christensen's "Restaurant Check Case" (2007, 193), Kelly's "Holocaust Denier" case (2013, 40-1), or Lackey's "Elementary Math" case (2010, 283). These are all relatively mundane cases, and it is easy to propose similar mundane scenarios for cases of peer agreement. For instance, we could invert the "Restaurant Check Case":

Inverted Restaurant Check Case: Suppose that I go out to dinner with several friends. It's time to pay for dinner and the question we're interested in is how much we each owe. We can all see the bill total clearly, we all agree to give a 20 percent tip, and we further agree to split the whole cost evenly without worrying over who asked for what specific item. I do the math in my head and become highly confident that our shares are \$44 each.

Meanwhile, my friend does the math in her head and also becomes highly confident that our shares are \$44 each.

In this case, suppose also that my friend and I have a long track record of eating out together and splitting the check in this fashion. We have also been equally successful in their calculations previously, with them agreeing about the arithmetic the overwhelming majority of the time. Given this, it seems like I should be even more confident about the conclusion that the appropriate share is \$44. There are, however, interesting cases that are less mundane but are still quite common in the real world.

2.2.2 Peer Agreement for Different Reasons

For this scenario, let us suppose that there are two colleagues, M and T, who work in the same academic discipline. They both have similar educational backgrounds. They have both taught and published extensively in the same academic subfield. They also both have the same number of years serving as faculty at the same institution, and both have attained the same tenured positions. It is not far-fetched to say that these two individuals are epistemic peers. Suppose also that both of these peers hold that a specific thesis, Z, is true. Z is somewhat of a hot-topic item in their particular subfield. Many others in the same field draw a proverbial “line in the sand” regarding Z and end up taking positions on either side of the line. What is interesting about our two colleagues is how they arrive at believing that Z is true. Z relies upon other foundational theses. One can either arrive at Z by holding to V_1 , W_1 , X_1 , and Y_1 , or by holding to V_2 , W_2 , X_2 , and Y_2 . Regarding these theses, the matched pairs are incompatible with each other. So, V_1 is incompatible with V_2 , W_1 is incompatible with W_2 , and so on.³³ Interestingly, our two colleagues both believe that Z is true, they have access to all the same evidence, but they each arrived at this

³³ You might think of these as being analogous to incompatible epistemological theories like Classical Foundationalism vs Coheretism or Internalism vs Externalism.

conclusion via a different path. While M holds to V_1, W_1, X_1 , and Y_1 , T holds to V_2, W_2, X_2 , and Y_2 .

We can look at this probabilistically to gain some insights as to how the peers should react upon realizing that they agree about Z. We know that it has to be the case that $\frac{P(Z)}{P(\sim Z)} = 1$.

Suppose that V_1, W_1, X_1 , and $Y_1 = E_1$ and V_2, W_2, X_2 , and $Y_2 = E_2$. Let us also suppose the following regarding the evidence given the truth of Z:

$$\frac{P(E_1|Z)}{P(E_1|\sim Z)} = \frac{3}{2} \quad \text{and} \quad \frac{P(E_2|Z)}{P(E_2|\sim Z)} = \frac{3}{2}$$

Using the odds form of Bayes' theorem, we know the following:

$$\frac{P(Z|E_1)}{P(\sim Z|E_1)} = \frac{P(Z)}{P(\sim Z)} * \frac{P(E_1|Z)}{P(E_1|\sim Z)} = \frac{3}{2} \quad \text{thus} \quad P(Z|E_1) = \frac{3}{5} = 0.6$$

We similarly can say:

$$\frac{P(Z|E_2)}{P(\sim Z|E_2)} = \frac{P(Z)}{P(\sim Z)} * \frac{P(E_2|Z)}{P(E_2|\sim Z)} = \frac{3}{2} \quad \text{thus} \quad P(Z|E_2) = \frac{3}{5} = 0.6$$

So, given the evidence M and T used to arrive at the conclusion, Z, they should each have a confidence of 0.6. We can, however, calculate the probability of Z given both E_1 and E_2 . By once again using the odds form of Bayes' theorem, we can calculate the following:

$$\frac{P(Z|E_1 \& E_2)}{P(\sim Z|E_1 \& E_2)} = \frac{P(Z)}{P(\sim Z)} * \frac{P(E_1|Z)}{P(E_1|\sim Z)} * \frac{P(E_2|Z)}{P(E_2|\sim Z)} = \frac{9}{4} \quad \text{thus} \quad P(Z|E_1 \& E_2) = \frac{9}{13} \approx .6923$$

Given different evidence used for arriving at the conclusion, it seems then that the agreement should be reason for both M and T to increase their confidence in Z.³⁴

³⁴ Thanks to Timothy McGrew for suggesting this example during discussion of the topic of agreement between peers.

2.2.3 Agreement in Testimony

Let us consider another rather common real-world scenario. Suppose there is a hit-and-run involving a taxicab. Afterward, two witnesses, Rose and Mickey, are interviewed by members of the press representing various local and regional papers. For this scenario, we will assume that Rose and Mickey are reliable testifiers, *i.e.*, they are not habitual liars, they are not colluding about the events that occurred, they are not under duress, their senses are not impaired, etc.³⁵ Below are the reports of the event that are later published by four media sources, as written by reporters who are all present while Rose and Mickey responded to questions by the press.

The Regional Mail

“We knew something was wrong as soon as we heard the screeching tires,” said 19-year-old student, Mickey, after the tragic event. “It’s not something that happens very often around here. It’s a quiet area.”

Were you scared? Did you get a good look at the car involved in the hit-and-run?

“Of course, we were,” said 27-year-old math teacher, Rose. “Yes, I saw the car,” she replied confidently.

“It was a Blue cab that was a Ford Taurus.”

The Post

The two witnesses of the hit-and-run were shocked by what they saw. A young middle-school teacher, Rose, was asked if she was scared and what she saw.

“Course,” she shakily replied.

“The lighting wasn’t too bad,” she noted. Rose responded that she could tell the car was a Blue cab and that it was a sedan.

The City Sun

Rose could barely contain the obvious shock from what she witnessed.

³⁵ Basically, let’s assume they are each as ideal of a testifier as we could get in the real world.

A resident of the neighborhood, she was on her way home from the local middle school.

“It was a scary experience,” said Rose after giving a statement to the police about the hit-and-run.

The Transcript

Rose, one of the two eyewitnesses of the surprising hit-and-run in a quiet suburban neighborhood, commented on the experience.

“I agree. This is a quiet neighborhood,” noted the 27-year-old local resident. “This isn’t the type of thing we see very often at all.”

Rose identified the car involved in the incident as a blue Ford sedan with taxicab decals, tags, and topper.

While there are some obvious similarities in these reports, they are just as obviously each distinct from the others.

Now suppose that another person, Rory, is also present while Rose and Mickey are being interviewed. Based on his recollections of the interview, Rory believes the following:

- a) There are two witnesses, Rose and Mickey, who witnessed the hit-and-run incident.
- b) The car involved in the hit-and-run is a blue Ford Taurus.
- c) Rose and Mickey’s neighborhood is quiet, and hit-and-runs are uncommon there.
- d) Rose was shocked and scared by the experience.
- e) Rose is a middle-school teacher.

The morning after the interview, Rory reads each of the reports published by the four media sources. In this case, since Rory and the reporters are not eyewitnesses of the event, let’s assume they are peers (or at least very close to being peers). Rory has pretty high confidence in the beliefs listed above. If the articles agree with Rory’s beliefs, then Rory has reason to increase his confidence. The question is, do they agree with what Rory believes?

If we assume a strict view of the beliefs as propositions, then only some of the media sources appear to agree with some of Rory's beliefs. The Regional Mail is the only source report to agree with **a** because it is the only source to note the names of the two witnesses. The Post reports that there were two witnesses and The Transcript reports that Rose was one of the two, but they do not strictly agree with **a**. The City Sun only names Rose as a witness. Similarly, The Regional Mail is the only source to report that the car involved in the hit-and-run was a blue Ford Taurus in agreement with **b**. The City Sun says nothing at all about the car. The Transcript reports that it was "a blue Ford sedan with taxicab decals, tags, and topper," and The Post merely notes a cab that was a blue sedan. Regarding **c**, The Regional once again is the only source to agree with the belief in its entirety. The Transcript gets rather close, but there is no mention of Mickey. So, we cannot say there is agreement there. The City Sun once again says nothing at all about this proposition, and the same goes for The Post. The City Sun is the only source that proposition **d**, though others do mention something about shock and being scared. Finally, The Post is the only source that says that Rose is a middle-school teacher in accordance with **e**. The Regional Mail does note that Rose is a math teacher, and The City Sun does report that she was on her way home from the middle school, but that does not grant us agreement under this strict view. Fortunately, such strict limitations aren't always appropriate.

In cases of multiple witnesses, a certain amount of variation is to be expected when the witnesses are recounting the events they witnessed. Indeed, varied evidence can be a very good thing.³⁶ In the instance above, we have a useful case of coherence that lends itself towards what we might consider *coherence* for the truth of H_R . For our purposes here, let's assume, regarding

³⁶ For helpful discussion on this matter, see Fitelson 1996 and 2001 (pp. S130-S134). Also, see McGrew 2016 and 2020.

the case above, H_R is the conjunction of Rory's beliefs, **a-e**. Borrowing a definition from Lydia McGrew, we will say that "a set of reports is coherent for H only if their conjunction constitutes a cumulative case that confirms H , even if each individual report does not separately confirm H " (McGrew, 2020, 801). McGrew provides a probabilistic analysis showing that varied evidence can be more valuable than unvaried evidence for the confirmation of hypotheses and showing that increased consilience given H or decreased consilience given $\sim H$ are both helpful for confirmation.³⁷

Assuming that the case of Rory and the reporters is an instance of peer agreement (or something quite close to peer agreement), then we need to recognize that each of the reports would give Rory reason to increase his confidence in H_R due to the partial agreement found there-in. The agreement found in H_R and the report from The Regional Mail, though, is the only instance where we find almost complete agreement. Given this, we wouldn't expect Rory to have reason to increase his confidence in H_R as much after reading the report from The City Sun as he would after reading the report from The Regional Mail. We wouldn't expect that, at least, if he *only* read one report or the other. The increased probability of the truth of H_R given the varied cumulative evidence provided by the four reports, however, could give us reason to treat this as we would other simpler instances of agreement. For now, for the sake of simplicity, I will assume that is the case.

2.2.4 Strict Superior Agreement

Suppose now that we add a third witness, Amy, to the hit-and-run scenario described above. Unlike the other two witnesses, Amy is an off-duty police detective. She is a trained observer

³⁷ We wouldn't expect a high level of consilience give $\sim H$ unless there were collusion or copying in sources of the reports.

who is familiar with the challenges of witnessing and accurately recalling shocking events. She also knows the details investigators will need for tracking down the perpetrator of the hit-and-run. Amy was not present for Rose and Mickey's interview with the reporters since she was giving her statement to the police investigators at that time. She also didn't get a chance to ask Rose and Mickey about what they had witnessed. In all of the excitement, the reporters missed that Amy was a third witness to the event. Amy formed the following beliefs,

- f)** There were two other witnesses, Rose and Mickey, to the hit-and-run incident.
- g)** The car involved in the hit-and-run was a Ford Taurus.
- h)** The car appeared to be a taxicab.
- i)** Rose and Mickey was shocked by the experience.
- j)** The license plate of the car ended in "3EA".

We can refer to the conjunction of **f-j** as H_A .

The morning after the event, Amy reads the reports from the four media sources giving Rose and Mickey's account of the hit-and-run. Now, it should be quite obvious that Amy is in a better epistemic position than the reporters regarding the events of the hit-and-run. (We might even say that she is in a better epistemic position than Rose and Mickey.) Due to her training and conditioning as detective, she is a cognitive superior. Due to her close access to the evidence, she is an evidential superior as well. For our purposes, this serves to meet the requirements of strict superiority. In this case, the reports from the four media sources are coherent for H_A . If the reports are coherent for H_A , then they confirm H_A and Amy should increase her confidence in H_A . While she may have no reason to increase her confidence in H_A much after reading only the report by The City Sun, it seems that she does have reason to increase her confidence after reading all four reports.

Recall the distinctions between Amy and Rory in this case. Rory is in agreement with the media reports. Since the reporters experienced the same interview that Rory experienced, this is more akin to a case of peer agreement, though it is somewhat different since each report is not specifically in agreement with H_R . Since the reports are coherent for H_R , Rory has reason to increase his confidence in H_R . Amy also agrees with the reports. She, however, is an epistemic superior because she is in a better epistemic position with regards to the evidence and in a better epistemic position with regards to cognitive capabilities due to her training and experience as a detective. Since the reports are coherent for H_A , Amy has reason to increase her confidence in H_A . Because Amy is already in a better position, both cognitively and evidentially, than Rory, she does not have reason for the same boost in confidence that Rory has after reading all four reports. A lesser boost in confidence for Amy, nonetheless, is appropriate for the same reasons. This a noteworthy distinction between how peers and superiors should react upon the realization of agreement.

2.3 Updating from Peers, Superiors, and Inferiors

In §1.2 of Chapter 1, I proposed that we can develop a model for updating confidences based on the opinions of peers who are in agreement. This model can serve as a heuristic for appropriate updating. The original model I proposed assumed the agreement under consideration was between peers. The following is the initial simplified formulation of this model from chapter 1:

Agreement Updating (AU):

l = Number of other agreeing opinions considered

$l > 0$

C_0 = The initial confidence of epistemic agent S

O_i = Confidence level of other agent's opinion i where $i \in \mathbb{N}^*$ and $i \leq l$.³⁸

k = Weight on opinion of peers

$$f(x) \begin{cases} -1 & \text{if } x < 0.5 \\ 0 & \text{if } x = 0.5 \\ 1 & \text{if } x > 0.5 \end{cases}$$

$$C_n = C_{n-1} + k \left[\sum_{i=1}^m O_m (1 - C_{m-1}) f(O_m) \right]$$

Recall that, while this model might look somewhat complicated, it works in a fairly simple fashion. We begin with original confidence of S. Upon encountering a peer with whom S agrees, S's confidence should increase by a degree in accordance with her peer's confidence. We add the original confidence of S to the product of the other's confidence and S's uncertainty (or certainty if the peer doesn't believe the hypothesis in question to be true). This product is weighted by k so as to account for the view agreement/disagreement that one is assuming. In the case of the Equal Weight View, k would equal 1. One might worry that the additive aspect of the model would allow S's confidence to reach certainty with enough updating, but that is not the case. Upon each instance of agreement S encounters, the degree by which she increases (or decreases) her confidence will grow incrementally smaller. So, suppose that S's confidence in H is 0.6, and she encounters three peers with a confidence in H of 0.6. Upon encountering the first peer, she will update her confidence by 0.24. After encountering the second peer, she will update her confidence by 0.096. After encountering the third peer, S would update her confidence by approximately 0.008. While this model captures intuitions about updating based on the opinions

³⁸ $\mathbb{N}^* = \{1, 2, 3, \dots\}$

of peers, it is unfortunately too simplistic to account for the differences between peers, superiors, and inferiors. Fortunately, some slight modifications will allow us to account for these changes.

In the original formulation of AU, we were accounting for the assumed view that would guide how peers' opinions should be weighted, *e.g.*, the Equal Weight View, the Justificationist View, etc., by using k . The weight of k was applied to all of the opinions being considered. If we are to account, not only for peers, but also for inferiors and superiors, we need for each opinion to be weighted individually. Building on the original proposed model, let's assume the following:

Opinion Updating[†] (OU[†]):

l = Number of other agreeing opinions considered

$l > 0$

C_0 = The initial confidence of epistemic agent S

O_i = Confidence level of other agent's opinion i where $i \in \mathbb{N}^*$ and $i \leq l$

k_i = Weight i on opinion of others (peers, superiors, or inferiors) where $i \in \mathbb{N}^*$ and $i \leq l$

$0 < k \leq 1$

$$f(x) = \begin{cases} -1 & \text{if } x < 0.5 \\ 0 & \text{if } x = 0.5 \\ 1 & \text{if } x > 0.5 \end{cases}$$

$$C_n = C_{n-1} + \left[\sum_{i=1}^m O_m(1 - C_{m-1})f(O_m) k_m \right]$$

Notice that now the weighting variable, k , will need to be assigned for each peer, superior, or inferior for which S must update.

Now, let's see how these calculations for $OU^\dagger$ work out. Suppose that P's initial confidence regarding H is equal to 0.6. P finds out that her peer, Q, also agrees that H is true and P has a confidence of 0.6. We have to limit the weighting by k to $0 < k \leq 1$. Value's greater than 1 will obviously skew the number excessively. Let's suppose the peers' opinions should be weighted equally, but less than superiors. Superiors can have maximal weighting of at most 1. Given this, Q's weight for k will be .5. Let us also assume P encounters a superior, R, and an inferior, S, who agree regarding H , and they share the same confidence in H as R and S. Because R is P's epistemic superior, we will assume that her opinion should be weighted more heavily. R's value for k will be 1. Since S is P's epistemic inferior (and we're assuming S is competent regarding H), S's value for k will be 0.25. So, we will end up with the following values: $C_0 = P(H) = 0.6$, $O_1 = Q(H) = 0.6$, $k_Q = .5$, $O_2 = R(H) = 0.6$, $k_R = 1$, $O_3 = S(H) = 0.6$, and $k_S = 0.25$. Assuming these numbers, the results of updating on the basis of these other opinions is as follows:

Table I. Summary of agreement updating calculations with $OU^\dagger$ ³⁹

$C_1 = C_0 + [O_1(1 - C_0)1]$	= 0.84
$C_2 = C_0 + [(O_1(1 - C_0)1) + (O_2(1 - C_1)0.5)]$	= 0.888
$C_3 = C_0 + [O_1(1 - C_0)1 + O_2(1 - C_1)1.5 + O_3(1 - C_2)0.25]$	= 0.9048

This rule appears to model quite nicely how agreement tends to work given what has been described above.⁴⁰ Let's now consider how $OU^\dagger$ would look with the cases of Rory and Amy that are described above.

³⁹ For the sake of simplicity, I have not included the value of f in the calculations since it would be equal to 1 for each of the calculations. The full calculation for C_3 , for example, would be

$$C_3 = C_0 + [(O_1(1 - C_0)1 \cdot 1) + (O_2(1 - C_1)1 \cdot 1.5) + (O_3(1 - C_2)1 \cdot 0.5)].$$

⁴⁰ I am assuming that the cases considered for the model are ones in which the inferiors are competent and the superiors are not aware of their own superiority as a reason for not updating.

Let's assume that Rory's confidence in H_R is 0.85 and that the reporters share the same confidence. Because the reports are coherent for H_R , we could simplify and assume agreement of H_R in the formalization, and since the reporters are Rory's epistemic peers, we could use 0.5 as the k value. However, since Rory is updating on the basis of the reports and the reports aren't individually in complete agreement with H_R , let's assume a lower k value that is between what we are assuming for peers and inferiors. We will end up with the following values: $C_0 = \text{Rory}(H_R) = 0.85$, $O_1 = \text{Mail}(H_R) = 0.85$, $k_{\text{Mail}} = 0.375$, $O_2 = \text{Post}(H_R) = 0.85$, $k_{\text{Post}} = 0.375$, $O_3 = \text{Sun}(H_R) = 0.85$, and $k_{\text{Sun}} = 0.375$, and $O_4 = \text{Transcript}(H_R) = 0.85$, $k_{\text{Transcript}} = 0.375$. The results of Rory's updating based on $OU^\dagger$ are the following:

Table II. Summary of H_R agreement updating calculations with $OU^\dagger$

$C_1 = C_0 + [O_1(1 - C_0)0.375]$	≈ 0.8978
$C_2 = C_0 + [(O_1(1 - C_0)0.375) + (O_2(1 - C_1)0.375)]$	≈ 0.9304
$C_3 = C_0 + [O_1(1 - C_0)0.375 + O_2(1 - C_1)0.375 + O_3(1 - C_2)0.375]$	≈ 0.9526
$C_4 = C_0 + [O_1(1 - C_0)0.375 + O_2(1 - C_1)0.375 + O_3(1 - C_2)0.375 + O_4(1 - C_3)0.375]$	≈ 0.9677

Amy's updating will look quite similar to Table II, but there are some differences worth mentioning.

Recall that in Amy's case, Amy is a superior to the reporters. Thus, any reports by the reporters should be weighted at most with a k value equal to what we are using. An argument could be made that, since the reports aren't in complete agreement with H_A , we should assign a lower k value. For our purposes, we will use a k value of 0.125 for each of the reports. We will also assume that Amy and the reporters share a confidence of 0.85. Given the above, we end up with the following values: $C_0 = \text{Amy}(H_A) = 0.85$, $O_1 = \text{Mail}(H_A) = 0.85$, $k_{\text{Mail}} = 0.125$, $O_2 = \text{Post}(H_A) = 0.85$, $k_{\text{Post}} = 0.125$, $O_3 = \text{Sun}(H_A) = 0.85$, and $k_{\text{Sun}} = 0.125$, and $O_4 = \text{Transcript}(H_A) = 0.85$, $k_{\text{Transcript}} = 0.125$. The results of Amy's updating on the basis of $OU^\dagger$ are as follows:

Table III. Summary of H_A agreement updating calculations with $OU^\dagger$

$C_1 = C_0 + [O_1(1 - C_0)1]$	≈ 0.8659
$C_2 = C_0 + [(O_1(1 - C_0)0.175) + (O_2(1 - C_1)0.175)]$	≈ 0.8802
$C_3 = C_0 + [O_1(1 - C_0)0.175 + O_2(1 - C_1)0.125 + O_3(1 - C_2)0.125]$	≈ 0.8929
$C_4 = C_0 + [O_1(1 - C_0)0.125 + O_2(1 - C_1)0.125 + O_3(1 - C_2)0.125]$ $+ O_4(1 - C_3)0.125]$	≈ 0.9043

Despite the fact that Amy is updating on the basis of the opinions of inferiors, notice that those opinions can certainly make an impact given enough of them.

2.3.1 Concerns About Varying Instances of Agreement and Disagreement

The purpose of the piecewise function in $OU^\dagger$ is to model the intuitive decrease in confidence when agreement is below 0.5 regarding H and an increase in confidence when the agreement regarding H is above 0.5. The multiplier for instances of agreement which should decrease confidence is -1, and the multiplier for instances agreement which should increase confidence is 1. While this is quite intuitive and easy to read, it could become burdensome as it adds complexity to the required calculations. This isn't as significant of an issue when all agents involved are in strict agreement about H with either confidences greater than 0.5 or less than 0.5, but it becomes more problematic in instances of disagreement (even if the disagreement is very slight) when confidences range at different intervals between 0 and 1. (While the original purpose of this model was to resolve the challenge of appropriate updating given agreement, if we can successfully use the same method to update given disagreement, then we can count that as a win.)

To resolve this issue, we can use a slightly less intuitive version of the model to achieve less cumbersome results. Consider the following updated model:

Opinion Updating (OU):

Retain the same assumptions as those for $OU^\dagger$ with the exceptions below.

$$f(x) \begin{cases} (-1 + C_{m-1}) & \text{if } x < 0.5 \\ 0 & \text{if } x = 0.5 \\ (1 - C_{m-1}) & \text{if } x > 0.5 \end{cases}$$

$$C_n = C_{n-1} + \left[\sum_{i=1}^m O_m f(O_m) k_m \right]$$

Given the modifications above, let's consider the results of updating using the confidences of the sort mentioned above. Let's suppose that R is the agent doing the updating, S is a superior, and T is a competent inferior. We will use the following values: $C_0 = R(H) = 0.5$, $O_1 = S(H) = 0.501$, $k_S = 1$, $O_2 = T(H) = 0.499$, and $k_T = 0.5$. The results of updating based on $OU^\dagger$ and OU are the following:

Table IV. Summary of near-agreement updating calculations with $OU^\dagger$

$C_1 = C_0 + [O_1(1 - C_0)1 \cdot 1]$	$= 0.7505$
$C_2 = C_0 + [(O_1(1 - C_0)1 \cdot 1) + (O_2(1 - C_1) - 1 \cdot 0.5)]$	≈ 0.6883

The results of updating based on OU are the following:

Table V. Summary of near-agreement updating calculations with OU

$C_1 = C_0 + [O_1(1 - C_0)1]$	$= 0.7505$
$C_2 = C_0 + [(O_1(1 - C_0)1) + (O_2(-1 + C_1)0.5)]$	≈ 0.6883

While this is a minor change, it does decrease the complexity when updating based on the opinions of multiple other people. Finally, there is one major concern that might be raised against OU that must be discussed.

Recall that the initially proposed model held the k variable constant for all other agents for which an update was required. Thus, the variable would be used to weight all opinions under consideration. This allowed the model to adhere to the *Commutativity of Evidence Principle* which says:

Commutativity of Evidence Principle (CEP): To the extent that what it is reasonable for one to believe depends on one's total evidence, historical facts about the order in which

that evidence is acquired make no difference to what it is reasonable for one to believe (Gabbay *et al.*, 2010, 451).

By adherence to CEP, if a person were to update her confidence regarding H on the basis of three different peers, the order in which she encountered those peers and updated her confidence would not matter. Additionally, it wouldn't matter if those other persons agree or disagreed with the originally held confidence. In the table below is an example of the adherence to CEP with AU (the original model). Each of the sets updates begins with a person who has a confidence of 0.6 in H . This person then encounters three peers with differing confidences regarding H .

Table VI. Summary of updating calculations of two peers with respective confidence $\{p, q\}$ using AU.

Order of Updating 1	Update 1 $\{0.6, 0.7\}$	Update 2 $\{0.88, 0.4\}$	Update 3 $\{0.832, 0.3\}$
Results	0.88	0.832	0.7816
Order of Updating 2	Update 1 $\{0.6, 0.3\}$	Update 2 $\{0.48, 0.7\}$	Update 3 $\{0.844, 0.4\}$
Results	0.48	0.844	0.7816

Ideally any model would demonstrate results of this sort.

With the transition to OU, the k value was moved so that it could account for different weights for each other person in the updating process. A major concern is that this modification might skew the numbers so that the model no longer adheres to CEP. Consider the same example as proposed above with the same confidences and the same order. However, in this instance we will need to include a value for k for each of the individuals in the updating process.

Table VII. Summary of updating calculations of two peers with respective confidence $\{p, q, k\}$ using OU.

Order of Updating 1	Update 1 $\{0.6, 0.7, 0.5\}$	Update 2 $\{0.74, 0.4, 0.25\}$	Update 3 $\{0.714, 0.3, 1\}$
Results	0.74	0.714	0.6282
Order of Updating 2	Update 1 $\{0.6, 0.3, 1\}$	Update 2 $\{0.48, 0.7, 0.5\}$	Update 3 $\{0.662, 0.4, 0.25\}$

Results	0.48	0.662	0.6282
---------	------	-------	--------

Fortunately, it appears the model maintains its adherence to CEP. This should alleviate concerns regarding the effectiveness of the model as heuristic for updating with a multitude of persons over time.

2.4 Conclusion

While the problem of peer disagreement has garnered a significant amount of attention in the more highly theoretical discussions in the literature, not enough attention has been paid to some very practical philosophical questions that lie in the same domain. Indeed, there are many challenges surrounding the issue of updating confidences given disagreement between superiors and inferiors that deserve more attention. While much has been discussed in the literature regarding the relationship between experts and novices, more could be said on this topic in light of the work being done on peer disagreement. Additionally, there is the more positive flip-side of the problem of disagreement for peers, superiors, and inferiors, which I call the assurance of agreement. This too deserves more attention and analysis.

What I have outlined in this chapter is the minimal considerations for the discussion of the assurance of agreement with peers, superiors, and inferiors. While the OU rule does not perfectly model updating between peers, superiors, and inferiors, it does appear to be a good place to start. Additionally, it does overcome some problems that other solutions might run in to. My hope, in this chapter, is to have provided an initial contribution to that discussion that can be expanded upon for further analysis of these sort of questions.

Chapter 3: *Consensus Gentium* and the Assurance of Agreement

It is undoubtedly the case that large groups of people have agreed to some rather preposterous beliefs throughout history. Likewise, a minority of individuals have dissented and have been correct about the truth of unpopular ideas. Despite this, the value of consensus throughout history as a recognized indicator of truth is hard to overlook. In our contemporary society, we are all quite familiar with people referring to something like “the scientific consensus on p ” to justify a belief about p . Such occurrences are, however, a somewhat modern thing. Throughout history, the appeal to consensus as a source of justification outside of more formal settings could probably be more accurately referred to as *tradition*. There are often two aspects associated with tradition; one has to do with practices, and the other with beliefs. In this case, I am referring to the latter. I assume something along the lines of the definition used by Edward Shils when he says,

Traditions are beliefs with a particular social structure; they are a consensus through time. In their content they might well be atemporal (i.e. they might make no reference to past or future) even have a temporal (traditional) legitimation. But even then, they have a temporal structure. They are beliefs with a sequential social structure. They are beliefs which are believed by a succession of persons who might have been in interaction with each other in succession or at least in a unilateral (even if not intergenerationally continuous) chain of communication (Shils, 1971, 126).

It is traditions of this sort that are often the foundations for a particular sort of theistic argument which, until recently, has somewhat fallen out of favor (or at least use) among philosophers.

Consensus gentium, or common consent, arguments for the existence of God generally rely upon the idea that the popularity of some shared opinion provides noteworthy reasons for thinking the opinion is true. For the most part, the normative value of the opinions of those with

little philosophical sophistication or competence has not garnered a significant amount of support in philosophical circles. There have been recent developments in *consensus gentium* arguments, nonetheless, that make a case for the value of such opinions. I will examine and evaluate two of these recent arguments and will examine some recent philosophical discussion on the value of consensus that is separate from such arguments. Finally, I will also provide the foundation and formulation for a somewhat novel (I hope) take on a *consensus gentium* argument.

3.1 Consensus Gentium Arguments

Common consent arguments have taken on various forms throughout the years. The most basic form generally follows something like the argument given by Clinias in Plato's *Laws* when he observes, "Well sir, don't you think that the gods' existence is an easy truth to explain? ...you know that all Greeks and all foreigners are unanimous in recognizing the existence of gods" (*Laws*, Book X, 886). Cicero provides a similar argument in *On the Nature of the Gods* where he states,

We must therefore come to the realization that gods must exist because we have an implanted, or rather an innate, awareness of them. Now when all people naturally agree on something, that belief must be true; so we are to acknowledge that gods exist. Since it is agreed by virtually everyone—not just philosophers, but also the unlearned... (*The Nature of the Gods*, Book I, 44).

Similar arguments have been formulated where it is the universal (or near-universal) belief in the existence of God (or gods), which should give one reason to believe in the existence of God (or gods). Why it is that this gives us reason to believe varies somewhat depending on the argument? In its crudest form, the argument could be formulated as a basic inference in the following form:

1. Everyone believes in the existence of God.

Therefore,

2. God exists.

Of course, even in antiquity, it is doubtful that Premise 1 could be true. Thus, proponents of this argument probably meant something more like “a great majority of people believe in the existence of God,” “an overwhelming number of people believe in the existence of God,” or at least “most people believe in the existence of God.”

Despite this, it was not merely the fact that a great many people shared this belief that provided the reason for adopting the belief in the existence of God. Up until the rise of British Empiricism, the concept of innate ideas was widely accepted. The reality that so many people share this apparently innate idea of some higher supernatural power (i.e., God or gods) lends an aspect of authority to the argument that it lacks otherwise.⁴¹ (Or at least, that is what those who advocate innate ideas seem to have in mind.) For example, consider how John Calvin quickly moves from the innateness of the idea that God exists to the reasonableness of believing in God based on common consent.

That there exists in the human minds and indeed by natural instinct, some sense of Deity, we hold to be beyond dispute, since God himself, to prevent any man from pretending ignorance, has endued all men with some idea of his Godhead, the memory of which he constantly renews and occasionally enlarges, that all to a man being aware that there is a God, and that he is their Maker, may be condemned by their own conscience when they neither worship him nor consecrate their lives to his service. Certainly, if there is any quarter where it may be supposed that God is unknown, the most likely for such an instance to exist is among the dullest tribes farthest removed from civilisation. But, as a heathen tells us, there is no nation so barbarous, no race so brutish, as not to be imbued with the conviction that there is a God. Even those who, in other respects, seem to differ least from the lower animals, constantly retain some sense of religion; so thoroughly has this common conviction possessed the mind, so firmly is it stamped on the breasts of all men. Since, then, there never has been, from the very first, any quarter of the globe, any city, any household even, without religion, this amounts to a tacit confession, that a sense of Deity is inscribed on every heart...All men of sound Judgment will therefore hold, that a sense of Deity is indelibly engraven on the human heart. And that this belief is naturally engendered in all, and thoroughly fixed as it were in our very bones, is strikingly attested by the contumacy of the wicked, who, though they struggle furiously, are unable to extricate themselves from the fear of God (Calvin, 1536 [1845], 43-5).

⁴¹ See Kelly (2011, 146-7) and O'Briant (1985, 74-5) for more on this.

However, arguments of this sort lost significant traction after criticism from the likes of Locke, Hume, and Mill, and because the disfavor of innate ideas advanced by the empiricists.⁴² Despite this, contemporary philosophers have made some attempts to resurrect the argument in slightly different formulations than in the past.

3.1.1 Common Consent as Inference to the Best Explanation

Thomas Kelly (2011) proposed a more modest formulation of the common consent argument, which is abductive in nature, that could be used as the foundations for a new argument from common consent. Thus, the argument could be stated as the following inference to the best explanation (IBE):

Common Consent argument from IBE (CCI)⁴³

1. The dominant opinion among the world's population is that God exists.
 2. The best explanation for this phenomenon is that God exists.
- Therefore,
3. God exists.

The argument does not rely upon the innateness of the idea of God nor that “widespread belief in God renders such belief reasonable all things considered” (Kelly 2011, 144). Instead, the fact

⁴² Locke, *An Essay Concerning Human Understanding*, Book I, section IV; Hume, *The Natural History of Religion* section XII; Mill, ‘Theism,’ in *Three Essays on Religion* (Peterborough, Ontario, Broadview Editions, 2009), pp. 155–8.

⁴³ While Kelly does not provide the argument in this exact formulation, it does not seem that he would not reject it since this formulation follows the spirit of his argument. The only formal argument he provides is the starting place for his discussion, which is as follows:

(Premise) Everyone believes that God exists.

(Conclusion) God exists (Kelly, 2011, 135).

that there is widespread belief in God serves as strong *evidence* supporting the truth of that belief, and it is the evidence that renders the belief reasonable.⁴⁴

Kelly spends a significant amount of time developing support for why we should believe something based on others' opinions. After all, it does not take much work to think up scenarios where this might be problematic. Consider, for instance, a scenario involving a Principled Serial Killer (PSK). PSK only kills other killers. He lives by a specific code and doesn't hurt innocents. Now suppose that the police are on the trail of PSK and have evidence that he has killed people in the past. Rather than turning himself in, PSK plants evidence so that it looks like another killer, K, was the one who killed all the people that PSK actually killed. Suppose that PSK's evidence planting skills are beyond exceptional. Because of the evidence, the police investigators and all of their forensic specialists come to believe that K was the one who did all the killing. It wouldn't be unreasonable for someone to believe that K was responsible for the killings based on the evidence of the shared belief of K's guilt held by the police investigators and the forensic specialists. Reasonable, but still wrong. Despite this, Kelly argues that we can make a case for an IBE argument from common consent by showing that "it scores higher than any other, rival potential explanation when evaluated by the usual criteria of plausibility, fit with background knowledge, explanatory power, simplicity, and so forth" (Kelly, 2011, 143).⁴⁵

Kelly notes that there are two main strategies that an opponent of the CCI argument might consider for resisting the argument. First, the opponent might offer a competing/debunking

⁴⁴ This is what makes Kelly's argument "modest". Even if one agrees that the CCI should rationally move one to belief in God, there might be other evidence that could weigh against the CCI in the other direction.

⁴⁵ Of course, simply scoring higher than rival explanations does not guarantee the associated inference will lead to a true conclusion.

explanation, *e.g.*, a naturalistic explanation. Second, an opponent might attempt to minimize the datum upon which the argument relies⁴⁶. This second strategy could take three different routes:

1. Attempt to show that the datum is somewhat insignificant due to the disproportionate number of uneducated, uninformed, unenlightened, etc., among theists.
2. Attempt to show that the datum is somewhat insignificant since the strong supermajority wasn't produced by independent convergence.
3. Attempt to show that the datum is false.

Obviously, it seems like the most ideal strategy would lead to the argument never getting off the ground at all. So, let's begin by considering the merits of the first strategy.

As soon as an opponent offers a plausible debunking explanation, Kelly notes that we are forced into judging the comparative plausibility of the explanations. As Kelly says, Durkheim, Freud, Marx, Weber, and many others have attempted to provide these sorts of explanations (Kelly, 2011, 147-8). How plausible one thinks they are in comparison to the theistic explanation really depends on who is doing the evaluating. It is quite apparent there is no *prima facie* winner. Kelly grants this but goes on to say,

If it is conceded that the *only* reason why the datum does not provide evidence for theistic belief is because there is some formidable competing explanation in the field, then to the extent that doubt is cast on that explanation on independent grounds, one's credence in the 'truth hypothesis'—and therefore, in the proposition that God exists—should increase accordingly (Kelly, 2011, 149).

This statement leads to a reasonable question. On what aspect of the explanation would the opponent need to cast doubt? Its precision, simplicity, likelihood, consilience, or some other theoretical virtue? After all, these are generally accepted measures by which we judge explanations (and they are only a few of the theoretical virtues by which we might judge an

⁴⁶ This datum is that a "strong supermajority of the world's population believes that God exists", where a "strong supermajority" is used to refer to >60% of the population (Kelly, 2011, 146-7).

explanation).⁴⁷ I expect that Kelly knows this is an issue that would eventually have to be addressed to push CCI forward and is merely side-stepping the issue to continue developing his modest case. Any serious proponent of an IBE-based common consent argument would need to address this issue for the explanation to be genuinely granted evidential status. Nonetheless, we can move on to evaluate the other areas an opponent of CCI might criticize.

An initially tempting approach for criticizing the datum CCI relies upon is to claim that the people in question (those who are in consensus) are not epistemically competent or are less epistemically competent than those who disagree. Basically, the claim would be that more competent believers hold that God does not exist than those who do. One may be more or less competent due to differences in access to evidence, cognitive abilities, intellectual virtues, etc. Kelly points out that this argument would be somewhat presumptive on the opponent's part, and it would require a significant amount of data to support the claim. Secondly, even if the claim was demonstrated to be true (or even likely), higher levels of epistemic competence do not guarantee correct beliefs.⁴⁸ Actually, we have good reasons to believe that, in greater numbers, epistemic inferiors (those with lower epistemic competence) will end up with the correct answer more often than a small number of epistemic superiors.⁴⁹ The classic example of this is Jack Traynor's jellybean experiment, where the average of the group guesses as to how many jellybeans were in a large jar was much more accurate than all but one or two individual

⁴⁷ For more on this, see Harman (1965), Forster, Malcolm, and Sober (1994), Keas (2018), Lipton (2003, 2004), McGrew (2003), and Thagard (1978) for further discussion. Obviously, there is a much larger literature regarding this topic, but this is a good place to start.

⁴⁸ I have discussed this in Chapter 2 with regards to one-to-one instances between epistemic inferiors and superiors.

⁴⁹ See James Surowiecki's *The Wisdom of Crowds* (2004) for a multitude of examples.

student's guesses. Kelly recounts a lesson learned from the popular show *Who Wants to be a Millionaire?* On the show, contestants were allowed to get help on tough questions by using 'lifelines', two of which were 'poll the audience' and 'phone a friend'. The 'friend' was an epistemic superior chosen by the contestant prior to the show. The 'audience' was filled with random ordinary people. It turns out that the audience was correct 91 percent of the time. While the expert 'friend' was only correct 65 percent of the time (Surowiecki, 2004, 4). Thus, Kelly concludes, "The cumulative opinion of a large number of people often provides significant evidence even when it conflicts with the opinion held by a comparatively smaller group of people who are in a better position to judge" (Kelly, 2011, 152).

In Traynor's jellybean experiment mentioned above, the students were not working in groups or sharing information. Thus, the resulting group guess was an average of the individuals' guesses. In other words, the opinions were arrived at *independently*. For the most part, this was also the pattern for the 'poll the audience' answers. In the case of the group belief in the existence of God, it seems pretty obvious that we do not have that level of independence. This is probably the strongest objection to the CCI argument. If we compared this to Traynor's experiment with his students, one could argue that, in this case, the analogous 'students' the CCI relies upon would not have provided guesses independently. Rather, all the 'students' would have copied off of another 'student' (or perhaps a handful of other 'students'). Some serious empirical work would be needed, however, to support that claim since it seems like we can make a good case that belief in God was arrived at independently by different groups.

It is the *persistence* in religious belief, once it has been attained, that Kelly thinks provides some distinction under the truth hypothesis rather than a naturalistic debunking explanation. Consider how it is that we learn basic arithmetic as children. We learn that $2+3=5$

from parents, teachers, peers, etc. There is a relevant lack of independence in this process. Yet, as we grow in understanding and experience, we retain the belief that $2+3=5$ because it is clearly true that otherwise, we would abandon the belief. While we might allow that some religious believers might retain belief in the existence of God for non-epistemic reasons, it seems uncharitable to say that all (maybe even most?) do so. Given this, Kelly proposes the truth hypothesis addresses things that the naturalistic explanations do not.

The final criticism that an opponent of the CCI might raise is that the datum is false (or that we don't know that it is true, or that we do not have good evidence that it is true). At first glance, this might seem somewhat easy to dismiss. After all, recent research (conducted by an atheist scholar) reports an estimate that 88-92% of the population believe in the existence of God (Zuckerman, 2006, 61).⁵⁰ As Kelly notes, though, the objection could rely on some philosophically interesting concerns. For instance, when Christians, Muslims, or Jews say that they believe in God, are they referring to the same entity?⁵¹ There is no undisputed positive or negative response to this question. Going down the road of answering that question will also likely lead to numerous other concerns that lie outside of epistemology. The question gets even more complex when we consider polytheistic traditions. One approach to resolving this is to

⁵⁰ Percentages based on the numbers provided by Zuckerman and the 2006 world population of approximately 6.594 billion.

⁵¹ For recent discussion on this topic, see the following: Bogardus, Tomas, and Mallorie Urban. 2017. "How to Tell Whether Christians and Muslims Worship the Same God." *Faith and Philosophy* 34 (2): 176–200, Schwartz, Beckwith, Francis J. 2021. "A Christian Account of Why David Novak Is Right About the Same God Question." in *The Achievement of David Novak: A Catholic-Jewish Dialogue*. Eugene, OR: Pickwick Publications, W.A., J.B. Cobb, F.J. Beckwith, G.R. McDermott, J. Walls, J. Cumming, D.W. Shenk, et al. 2019. *Do Christians, Muslims, and Jews Worship the Same God?: Four Views. Counterpoints: Bible and Theology*. Zondervan Academic., and Sullivan, Meghan. 2012. "Semantics for Blasphemy." In *Oxford Studies in Philosophy of Religion*, Vol. IV, edited by Jonathan L. Kvanvig. Oxford University Press.

allow a certain amount of permissivism regarding belief in God. By this, I mean something along the lines of what John Hick described:

...the idea that we are considering is not that any and every conception of God or of the transcendent is valid, still less all equally valid; but that every conception of the divine which has come out of a great revelatory religious experience and has been tested through a long tradition of worship, and has sustained human faith over centuries of time and in millions of lives, is likely to represent a genuine encounter with the divine reality (Hick, 1988, 141).

If this sort of permissivism is allowed, then perhaps this objection is not as much of a challenge to overcome as it could have been. Nonetheless, there are other ways of formulating a common consent argument that avoid the majority of the issues described above entirely. In what follows, we will consider another noteworthy alternative.

3.1.2 Common Consent and Self-Trust

Linda Zagzebski has proposed a very different take on a *consensus gentium* argument from the CCI. To begin building her case, she notes two things about human beings. The first is that we have a natural desire for truth. This desire, like many others, is pre-reflective. With some notable exceptions (*e.g.*, the Pyrrhonian skeptics), this is a desire we all seem to share. The second thing Zagzebski notes that we also share is a natural belief that this desire for truth can be (and often is) satisfied. She observes:

I trust my faculties and my environment sufficiently to think that I believe many truths, and I rely upon those faculties and those truths in the normal conduct of my life. The faculties I use in forming beliefs may operate on the environment directly, or they may operate on the testimony of others. It is not only natural to rely upon my own faculties, but it is also natural to believe what other people tell me. So the natural trust with which I start includes trust in the relation between my faculties and my environment, and trust in the faculties of many other people (Zagzebski, 2011, 23).

This pre-reflective trust (or self-trust) is the starting point and the key to overcoming epistemic circularity.

Following a similar path as Richard Foley (2001) and William Alston (1986, 2005), Zagzebski agrees that the problem of epistemic circularity is inescapable. She notes, “We can do everything epistemically that we are supposed to do, including following the evidence scrupulously, and we have no assurances that the results will give us the truth or even make it probable that we will get the truth” (Zagzebski, 2011, 24). Like Foley and Alston, Zagzebski argues that self-trust is the only way to avoid circularity. However, she takes a somewhat different approach.

Foley is led to self-trust by reflection upon skeptical questions and the failure to answer them. Similarly, Alston is forced into accepting self-trust due to reflection on the unsuitability of the human epistemic condition for full reflective justification. For both, self-trust is a fallback option that would not have been chosen were some other option possible. For Zagzebski, it is the starting point upon which we can build reflective justification. Thus, she argues that not only is it reasonable to trust one’s faculties, but it is also reasonable to believe they are trustworthy (2011, 26). More specifically, we can trust ourselves when we are epistemically conscientious; that is, when we are self-consciously trying our best to get to the truth that we desire. Hence, Zagzebski states, “This means that a self-reflective person trusts herself in advance of reasons to trust herself, but she trusts herself more when she is conscientious and for the reason that she is conscientious. She trusts reasons and evidence because of her trust in herself when she is conscientious” (Zagzebski, 2011, 28).

The question then comes up quite naturally, if we have this natural epistemic trust in ourselves, should we have a similar trust in others? Or should we strive to be intellectually self-reliant? Well, we should expect that an epistemically conscientious person will gain beliefs regarding other people. Some of those beliefs will be that some people are being just as

conscientious in their pursuit of truth, some the same, and some less. When a person discovers that others are more or equally conscientious as she is in the pursuit of truth, it seems reasonable that similar trust should be extended to those people, given the above. Zagzebski notes,

My commitment to trust my conscientiously formed beliefs about other people means that I cannot consistently expect evidence of the reliability of another person before thinking of them as epistemically trustworthy. In so far as I believe myself to be trustworthy when I am conscientious, and I conscientiously believe that many other people are at least as conscientious as I am in certain respects, I must think of them as epistemically trustworthy in those respects. I am committed to that by the principle that I should think of like cases alike (Zagzebski, 2011, 29).

Thus, self-trust leads one also to trust in others. Of course, this trust in others is defeasible, just like self-trust. It is, nonetheless, perfectly reasonable on Zagzebski's account.

Having established self-trust as a basis upon which we can show the reasonableness of our own beliefs and the beliefs of others, Zagzebski gives the following take on a *consensus gentium* argument that is quite different from what we have seen before (Zagzebski, 2012, 186):

Common Consent argument from Self-trust (CCS):

1. I must have a general attitude of self-trust in my epistemic faculties as a whole. This trust is both natural and shown to be rational by philosophical reflection
2. The general attitude of epistemic self-trust commits me to a general attitude of epistemic trust in the faculties of other human beings.
3. So the fact that someone else has a belief gives me *prima facie* reason to believe it myself.
4. Other things equal, the fact that many people have a certain belief increases my *prima facie* reasons to believe it, and the reason is stronger when the beliefs are acquired independently.
5. The fact that other people believe in God is a *prima facie* reason to believe that God exists, and the fact that many millions of people constituting a very high majority believe or have believed in prior ages that God exists increases my *prima facie* reason to believe in God. Even though many of those beliefs were acquired from others, I have reason to trust the beliefs of trusted others who acquired their beliefs from persons they trust. That is, I have reason to trust their trust. Furthermore, even if we completely discount for dependence, there are still many millions of people who independently believe or have believed in past ages in the existence of God.

While this argument obviously takes a different route for arriving at the same conclusion, one might question whether it avoids the concerns associated the CCI.

Obviously, since the CCS does not rely upon an IBE, we do not have to worry about competing explanations. However, one possible option is to object to Zagzebski's account of self-trust or at least to object to the extension of trust to others. One could agree that we cannot escape epistemic circularity (which is a bad thing) and also argue that self-trust isn't a realistic option. That path probably leads to the unfortunate conclusion of accepting skepticism, but perhaps some people are willing to accept that option. One could also argue that we cannot escape epistemic circularity, but in the end that actually isn't a bad thing.⁵² As an alternative, it is possible that one might agree to self-trust for similar reasons as Foley or Alston. If that is the case, it is not clear to me that holding to self-trust for these reasons motivates the same move to trust in others. Given that someone following in Alston or Foley's steps would accept self-trust mainly as means of escaping epistemic circularity (and not for the reasons proposed by Zagzebski), it seems that they could do so and remain committed to some form of *ethical egoism*, *i.e.*, the sort of epistemic self-reliance that Elizabeth Fricker describes as *epistemic autonomy* (Fricker, 2006, 243).

For instance, one might argue that there is actually some significant difference between trusting myself when I am being epistemically conscientious and trusting others when they appear to be epistemically conscientious. Say that I judge the probability that my senses are reliable, $\text{Pr}(M)$, to be 0.95. I also judge the probability of another person's senses as reliable, $\text{Pr}(O)$, to be the same. However, I must first trust my senses to trust theirs. So, I am not simply

⁵² See the following for more on this: Bergmann, Michael. 2004. "Epistemic Circularity: Malignant and Benign." *Philosophy and Phenomenological Research* 69 (3): 709–27. Steup, Matthias. 2013. "Is Epistemic Circularity Bad?" *Res Philosophica* 90 (2): 215–35.

trusting the probability that their senses are reliable. I am trusting the conjunction of the probability that my senses are reliable and the probability that their senses are like mine. We calculate the probability of that conjunction by multiplying the two original values. Thus, $\text{Pr}(M) \times \text{Pr}(O) = 0.9025$. While this is still quite a high number, it is noticeably less than the probability at which I judge my own senses to be reliable. Given that, one could propose an argument along these lines to make the case that these things are not enough alike that I am motivated to “treat like cases alike” as proposed (Zagzebski, 2011, 29).⁵³

Even if one were to accept Zagzebski’s account of self-trust, it looks like there still might be room to raise some of the same objections that one could raise with the CCI. Zagzebski’s argument avoids any concerns about theists needing to be a super-majority by simply noting the large numbers of people in the past and the present who have believed in the existence of God. There is no real question about this. As mentioned before, one could of course, raise the question about whether or not all of these people actually believe in the existence of the same thing. This is the same objection that could be raised against the CCI. Zagzebski grants a certain amount of permissivism about this question to avoid the objection. She notes:

What we get from widespread agreement is closer to the “half glimpse” mentioned by Aristotle than to a clear view of the object. The people of the world give different descriptions of the object glimpsed, but they agree that there is a genuine glimpse of something important. Perhaps the vagueness of the idea of God makes the *consensus gentium* argument uninteresting to some people, but I agree with Aristotle that even a half a glimpse of the profound is more valuable than a complete view of lesser things (Zagzebski, 2012, 188).

Finally, one could raise an objection regarding the lack of independence of the believers who are in consensus. Note, however, the last line of her argument, “...even if we completely discount

⁵³ For other objections, see Dobrzeniecki, Marek. 2018. “Is the Fact That Other People Believe in God a Reason to Believe? Remarks on the Consensus Gentium Argument.” *European Journal for Philosophy of Religion* 10 (3): 143-8.

for dependence, there are still many millions of people who independently believe or have believed in past ages in the existence of God” (Zagzebski, 2012, 186). Zagzebski’s claim certainly seems plausible. If one were going to push harder against this point, the only option seems to be to pull the ‘I’m skeptical about this claim’ card and demand some sort of empirical support.

As far as I am concerned with this project, my goal is not to provide a critique of either the CCI or the CCS. (I actually find aspects of both of these arguments quite persuasive.) My goal is to point out that there may be some reasons why one *might not* accept either of these arguments. Even if it is the case that one rejects both the CCI and the CCS, we still have grounds for belief from common consent in the existence of God, and we can develop those reasons based on a somewhat different approach than those used for both the CCI and CCS. Let us then consider a third option.

3.2 Developing a New Argument: Explanation as a Heuristic

The challenge of developing a new *consensus gentium* argument requires something that avoids the objections that could be raised against both something like the CCI and the CCS. In addition to the problems of dependence and possible reference concerns about the descriptor “God”, one would need to provide an argument that doesn’t face problems of competing explanations for an IBE and one that could be applied on most contemporary epistemological theories, *i.e.*, one that doesn’t rely on the concept of self-trust. The question, as always, is “Where do we begin to build our argument?” Oddly enough, my suggestion is that we go back to the basis of the first argument we examined.

Over the years, there have been many objections to the value of inferences to the best explanation, and of course, there have been many noteworthy replies to those objections. I would

like to avoid that discussion and instead consider the use of explanatory reasoning for heuristic purposes rather for inferential purposes.⁵⁴ It seems perfectly reasonable to use explanation to point us in the right direction even if it doesn't lead us to a precise answer to our question. After all, the idea that consensus can be a sign of reasonable belief (if not a sign of knowledge) is not found solely in the domain of philosophy of religion.⁵⁵ Regarding the value of consensus, Nicholas Rescher noted,

Agreement is indicative of an *invariance* of sorts: it reflects the fact that, notwithstanding the variation of the particular epistemic stance of different individuals, people's efforts at epistemic problem-solving favour one particular outcome. And this sort of invariance is clearly a positive evidential factor. A 'consensus theory of reasonable belief' along such lines accordingly has something to be said in its favour (Rescher, 1993, 61).

What we need then is a set of conditions by which we can determine that knowledge (or reasonable belief) is the best explanation for a consensus. These conditions can serve as a filter useful in discarding explanations we can reasonably overlook. If these conditions are met, then we have good reason to proceed with developing an argument that supports this explanation but does not rely upon the explanation for inferential purposes.

⁵⁴ For more on this see McGrew, Timothy. 2003. "Confirmation, Heuristics, and Explanatory Reasoning." *British Journal for the Philosophy of Science* 54 (4): 553–67.

⁵⁵ For recent works on this see the following: Miller, Boaz. 2013. "When Is Consensus Knowledge Based? Distinguishing Shared Knowledge from Mere Agreement." *Synthese* 190 (7): 1293–1316. Miller. 2015. "Why Knowledge Is the Property of a Community and Possibly None of Its Members." *Philosophical Quarterly* 65 (260): 417–41. Miller, 2019. "The Social Epistemology of Consensus and Dissent." In *The Routledge Handbook of Social Epistemology*, edited by David Henderson, Peter Graham, Miranda Fricker, and Nikolaj Jang Lee Linding Pedersen, 228–37. Routledge., Tucker, Aviezer. 2003. "The Epistemic Significance of Consensus." *Inquiry: An Interdisciplinary Journal of Philosophy* 46 (4): 501–21., and Tucker, 2020. "The Inferences of Common Causes Reduced to Common Origins." *Studies in History and Philosophy of Science Part A* 81: 105–15.,

3.2.1 Conditions for Knowledge as the Best Explanation of Consensus

Boaz Miller (2013, 1294) and Aviezer Tucker (2003, 504) have both independently argued for such a set of conditions as those described above. Their conditions overlap somewhat, but each includes something not required by the other. I will draw inspiration from both of their sets of conditions to hopefully provide something that will grant us suitable results. Let us say then that knowledge (or reasonable belief) is the best explanation of a consensus when the following conditions are met:

1. *Uncoerced Believers* – The individuals involved in the consensus must hold the belief in question freely and not by unwilling acquiescence.
2. *Diverse Believers* – The group of individuals involved in the consensus must demonstrate a high level of heterogeneity.
3. *Non-Trivial Consilience of Evidence* – The consensus is based on varied evidence that all seems to agree.
4. *Sufficiently Large Group* – The group of individuals involved in the consensus must be large enough to avoid issues of unintended results, *e.g.*, consensus due to hidden biases, social connections, etc.

Let's now give each of these some brief consideration to support their use in our filtering process.

3.2.1.1 Uncoerced Believers

Coercion can take many forms. Coerced belief, however, cannot provide the foundations for epistemically reasonable belief. Many of us are familiar with the power of “peer pressure” and how easily some give way to such pressures. Other social pressures related to shame and honor are common in some cultures. Such pressures might achieve similar acquiescence. Economic dependence can also be used for coercion. Tucker notes, “An average contemporary academic environment in a democratic country contains many positive and negative sanctions for publicly espousing this or that belief. These sanctions probably skew the aggregate public reporting of

views of individual academics, in comparison with what honest reporting of their beliefs would have looked like” (Tucker, 2010, 505). In contemporary Western culture, we see this in the form of *cancel culture*. One might lose customers, followers, or even be fired for saying or doing what is perceived as the wrong thing. The most common historical method of coercion is intimidation through threat of violence or through violence itself. Despite the great variety of coercion methods, achieving even a local coerced consensus requires much more than social and economic intimidation. In addition to those things and institutional dominance, it usually requires an extreme degree of violence and use of force. Tucker observed that:

Both the Inquisition and the Soviet NKVD secret police required more than intimidation and institutional hegemony, they required killing and imprisoning those who would not be coerced by other means...Even an all powerful totalitarian state and its secret police can only coerce a local involuntary acquiescence by the extensive use of extreme violence (Tucker, 2003, 505).

Thus, we should take care to be mindful of the influence that coercion can have on group belief, but we must also be aware of the difficulty with which this is obtained and maintained (especially over long periods of time and across large geographic areas).

3.2.1.2 Diverse Believers

Claiming that we are all subject to and struggle with various biases and problematic assumptions should be mostly uncontroversial. It should be just as uncontroversial to note that people from different backgrounds will also often vary in their respective biases and problematic assumptions. Because of this, individuals from diverse backgrounds may produce ideas or solve problems that escape their counterparts who face a different set of biases and assumptions. Given this, it should not be much of a surprise that a certain level of heterogeneity increases the

accuracy of group consensus beliefs.⁵⁶ Additionally, with more diverse populations, alternative explanations become less plausible. Tucker notes,

In the absence of coercion, the more heterogeneous the consensus group, the easier it is to refute alternative hypotheses that link properties that only some members of the consensus group share with the consensus. Alternative hypotheses may attempt to explain a concrete consensus on beliefs with particular power relations, or political interests, or ideological convictions, or gender bias, or cultural contexts. But if the consensus group is composed of some persons who are not connected to each other in any power relationship, others who do not share political interests and ideologies, of both genders and many cultures, such hypotheses are less probable than the knowledge hypothesis (Tucker, 2003, 506).

This diversity also serves to decrease the likelihood of various forms of coercion. For instance, certain public institutions may be subject to certain types of political coercion that are not common in private institutions. On the other hand, private institutions could be subject to economic pressure that is not common in public institutions.

3.2.1.3 Non-trivial Consilience of Evidence

Generally, one might expect that a diverse group will produce diverse evidence. Indeed, Tucker argues for a *unique heterogeneity* condition in attempts to achieve this end on his account. If the uniquely heterogeneous group agrees to the truth of some proposition, it is likely the members arrived at a conclusion in significantly different ways due to a variety of social, economic, and geographic factors (Tucker, 2003, 506-12). While acknowledging that this could serve as a proxy condition of sorts, Miller notes that simply requiring unique heterogeneity as a condition is not enough because it doesn't guarantee what we really need. He notes, "Different social groups and different disciplines do not always use different methods and different types of evidence" (Miller, 2013, 1310). Evidence that is generated by multiple techniques and that relies upon

⁵⁶ For more on this: Intemann, Kristen. 2009. "Why Diversity Matters: Understanding and Applying the Diversity Component of the National Science Foundation's Broader Impacts Criterion." *Social Epistemology* 23 (3): 249–66.

different background assumptions provides better support for proposed hypotheses.⁵⁷ While achieving *actual* consilience of evidence poses significant challenges that would make the condition too demanding⁵⁸, specifying non-trivial consilience of evidence as a condition works more as a common-sense requirement. As long as available evidence isn't being ignored, overlooked, or suppressed, and there is enough consilience for *prima facie* recognition, the condition can be met.

3.2.1.4 Sufficiently Large Group

Of the conditions, this is probably the simplest to understand. Without a sufficiently large group, it is difficult to attain the other conditions. In statistical sampling, it is difficult to avoid accidental results from things like hidden biases with groups of minimal size. It is also difficult to avoid social connections that may introduce unintended dependences or unintentionally coerced (or at least influenced) beliefs. For instance, suppose a group of 100 students was polled in a college building hallway regarding some specific proposition p . Surprisingly, 95 of those students held that p was true. Oddly enough, 90 of those students had all just left a class where the professor had been arguing for the truth of p , and they all thought this might be some sort of weird pop-quiz. So, they all replied to the poll that they believed p was true. A sufficiently large group helps avoid odd events of that sort from a small sampling. Specifying a number that is 'large enough' is rather tricky. However, Tucker notes that, "when the consensus involves hundreds of people who are geographically, institutionally and professionally dispersed, it is safe

⁵⁷ For more on this topic, see Stegenga, Jacob. 2009. "Robustness, Discordance, and Relevance." *Philosophy of Science* 76 (5): 650–61.

⁵⁸ Miller gives some noteworthy reasons for not requiring this (2013, 1310–11). I agree with his position on this and assume his reasons in my conditions.

to assume that it is large enough” (Tucker, 2003, 512). For our purposes, this guideline will suffice.

3.2.2 Common Consent as Explanation

Now we have come to the point where we ask, does knowledge (or at least reasonable belief) serve heuristically as the best explanation for common consent regarding the existence of God? Well, have our conditions been met? Following Kelly's lead in his argument, I would argue that they have. Let's review the available support for our conditions.

Is the group *sufficiently large*? Even if we consider only those living today (while ignoring believers in the past who belong to the group), there are approximately at least somewhere between 5.8 to 6.6 billion people in the world who believe that God exists.⁵⁹ It seems like we can answer that the condition was indeed met. Additionally, for this number to be so large, an incredibly diverse population must be accounted for in this group. Thus, we can also conclude that the *diverse believers* condition has also been met. This number also gives us some idea about *uncoerced believers*. While it is entirely possible that some of those people were coerced into this belief, it is highly unlikely that anything close to a majority (or anything resembling a significant portion) of those in question could have been coerced. The only country with a large enough population and something approaching the right sort of conditions to attempt such a thing might be China, and encouraging belief in the existence of God does not really

⁵⁹ See, e.g., the data compiled in the following: Keysar, Ariela, and Juhem Navarro-Rivera. 2013. “A World of Atheism: Global Demographics.” In *The Oxford Handbook of Atheism*, edited by Stephen Bullivant and Michael Ruse, First, 553–86. Oxford and New York: Oxford University Press UK., and Zuckerman, Phil. 2006. “Atheism: Contemporary Numbers and Patterns.” In *The Cambridge Companion to Atheism*, edited by Michael Martin. New York, NY: Cambridge University Press. Zuckerman estimates that 88% of the population believes in the existence of God. Keysar and Navarro-Rivera estimate the number closer to 93%. My numbers are based on the world population in the years those estimates were made.

appear to be on the agenda of the Chinese government. How well the final condition is met is somewhat more difficult to ascertain.

3.2.2.1 Common Consent as Explanation and the Non-Trivial Consilience of Evidence

For the *Non-Trivial Consilience of Evidence* condition to be met, recall that the consensus should be based on varied evidence that all seems to agree. More specifically, we are looking for evidence that is generated by multiple techniques and that relies upon different background assumptions. While this is somewhat of a different task than determining whether or not people believe in the existence of God, we do have some data on this. Before we begin looking at the evidence, however, we need to classify the varieties of evidence that we can expect to find for the existence of God. Also, we need to specify some assumptions about what counts as evidence. Let's begin with a suitable definition for 'evidence':

Evidence: *E* is evidence for some hypothesis *H*, when *E* makes it more likely that *H* is true (in other words, *E* confirms *H* or *E* gives reason to increase one's confidence in *H*).

This is a fairly broad understanding of evidence, but it should be mostly familiar and unproblematic for those working in epistemology or philosophy of science.⁶⁰

Assuming the definition above of evidence, we can use the following categories, which seem appropriate and mostly uncontroversial, for the evidence provided by a broad population for belief in God. This evidence can generally be divided into the following four categories.

- a. *Philosophical Argument* – Theistic arguments for the existence of God
- b. *Religious Experience* – Personal experiences in which God (or some spiritual entity) personally interacts with humanity
- c. *Religious Testimony* – Tellings regarding religious beliefs by believers/testifiers to some audience

⁶⁰ See, e.g., Thomas Kelly's discussion on this topic under the section "Evidence as a Guide to Truth: Evidence as Sign, Symptom, or Mark" in his article on evidence (Kelly, 2016).

- d. *Special Divine Action* – Events beyond the ordinary which exceed the productive power of nature and are generally referred to as ‘miracles’⁶¹

While it is somewhat challenging to get precise numbers on each of the categories of evidence, it is not difficult to find accounts of all four categories. We find examples of philosophical arguments for the existence of God (or gods) in Plato, Aristotle, Anselm, Aquinas, Avicenna, Averroes, amongst modern and enlightenment thinkers, and into contemporary philosophical literature. We have records of people referencing God interacting with them in some internal or spiritual way.⁶² (Consider, for example, the Quaker movement. This is an essential aspect of their faith.) Religious testimony of the sort I am referring to is (and has been) fairly common amongst evangelists or proselytizers throughout history. In 20th century America, you could even see large-scale examples of this in televised services, *e.g.*, the Billy Graham crusades during the late part of the century. Finally, the Special Divine Action category would include accounts of what people would generally refer to as ‘miracles’. For a fairly thorough documentation of such reported events in contemporary settings, see Craig Keener’s project on miracles (2011). Much more could be said on this, but a brief consideration gives us a good idea that allows us to think we probably have enough non-trivial consilience of evidence. Obviously, for a stronger case a much more thorough evaluation of the evidence would be necessary. However, the above gives us a general idea of what evidence we are after and what evidence is available.

Based on initial considerations, I think we have good reason to think the existence of God is the best explanation from the wide-spread belief in God’s existence. To avoid the problems

⁶¹ I borrowed this term from the Oxford Special Divine Action Project. “SDA.” n.d. Accessed January 11, 2021. <https://sda.bodleian.ox.ac.uk/sda/#!/about>.

⁶² See, *e.g.*, some recent data on these types of experiences: Baker, Joseph O. 2009. “The Variety of Religious Experiences.” *Review of Religious Research* 51 (1): 39–54.

described above regarding IBEs, we need to develop an argument that does not rely upon an inference from that explanation. I will proceed to develop my argument below.

3.3 A New Argument: Common Consent from Opinion Updating

The truth hypothesis gives us (at least) a good enough explanation of the common consent on the existence of God that we should pursue the question further and give it more serious consideration than other explanations. In support of this explanation, I propose that we have reason for belief in the existence of God from common consent on the basis of rules that govern rational disagreement and agreement. To begin, let me discuss the basic assumptions of the argument.

In the literature on peer disagreement, various responses have been given on the appropriate response to rational peer disagreement. In Chapter 1, I proposed there is an inverse to the problem of disagreement, which I call the assurance of agreement, and that there is also an inverse response to rational peer agreement. The appropriate response to peer agreement is connected to the appropriate response to peer disagreement. For my purposes, as long as one adheres to *some* degree of conciliation in cases of disagreement, then we can also assume a model of updating upon agreement that I refer to as Opinion Updating (OU). While not as accurate as something like Bayesian updating, I have proposed that OU models the natural way in which individuals increase their confidence based on the opinions of others in a somewhat less complex manner than the Bayesian approach.

3.3.1 Basic Assumptions for the Argument

The argument relies upon the value of the OU model as a means for updating confidences in a belief given that others share that confidence. The OU model is outlined below:

Opinion Updating (OU):

l = Number of other agreeing opinions considered

$l > 0$

$m = n$

C_0 = The initial confidence of epistemic agent S

O_i = Confidence level of other agent's opinion i where $i \in \mathbb{N}^*$ and $i \leq l$

k_i = Weight i on opinion of others (peers, superiors, or inferiors) where $i \in \mathbb{N}^*$ and $i \leq l$

$0 < k \leq 1$

$$f(x) = \begin{cases} (-1 + C_{m-1}) & \text{if } x < 0.5 \\ 0 & \text{if } x = 0.5 \\ (1 - C_{m-1}) & \text{if } x > 0.5 \end{cases}$$

$$C_n = C_{n-1} + \left[\sum_{i=1}^m O_m f(O_m) k_m \right]$$

The initial intention for this model was for its use as an option for updating in cases of peer agreement (and peer disagreement), which could then be extended to cases of non-peer agreement, *e.g.*, agreement between peers, superiors, and inferiors (Chapter 2). While the model looks somewhat complicated, it is just the mathematical formulation of a simple idea. We begin with the confidence of a person, R, about a hypothesis, H. When R encounter someone, S, who agrees with her about H, we add the product of S's confidence in H, R's uncertainty about H, and S's weighting value to R's original confidence about H. So, suppose that R's confidence, is 0.6, S's confidence is 0.6, and S's weighting value is 0.5. The calculation would look like this:

$$R's \text{ new confidence} = 0.6 + (0.6 \times 0.4 \times 0.5) = 0.72$$

If R encounters a series of people who agree, she would follow the same pattern with the confidence and weight values of the other people. Those would each similarly take into account

R's uncertainty and would be added to the product before it is added to R's confidence. So, if another person, T, were encountered who had a confidence of 0.6 and a weight of 0.25, the calculation would like this:

$$R's\ newer\ confidence = 0.72 + (0.6 \times 0.28 \times 0.25) = 0.762$$

We can therefore have a long series of updates like this but notice that the increase in confidence diminishes as R's confidence draws closer to 1. Thus, we avoid the unrealistic problem of eventually arriving at actual certainty. Below, I will show how this model can be used in cases of consensus.

Before discussing the model's use in cases of consensus, we must make distinctions about consensus itself. I take it that consensus can be thought about in various ways. Suppose we say that a group of people hold to a consensus regarding the truth of H . The first option for what we mean by this is that every member of the group holds that the likelihood of H is some specific number they all agree upon (no less, no more). Thus, they might all agree that $\Pr(H) = 0.8$. Let's call this a *strict consensus*.⁶³ A second option is that they all merely agree that H is likely to be true. So, they might all agree that $\Pr(H) > 0.5$. We can call this a *broad consensus*. Finally, they might agree that H is at least a certain likelihood, but it could be more likely. In this case, they might say that $\Pr(H) \geq 0.8$. We will refer to this as a *general consensus*. Generally, most people mean something like *strict consensus* and *general consensus* when they talk about a group being in consensus regarding some hypothesis. Given this and that OU will work in cases of *strict consensus* and *general consensus*, and since we can assume that each member of the

⁶³ A non-probabilistic view where the group agrees that " H is true", would also fit best under this label.

group has *at least* the same confidence, we will assume one of these senses of the term is being used for this argument.

OU relies upon the weighting of individuals' opinions with the k_i variable. We might assign a higher weight to experts, *i.e.*, superiors, than we would for peers or those whom we deem less competent in some way, *i.e.*, inferiors. This may vary depending on one's underlying assumptions about how disagreement and agreement should be handled. For the sake of simplicity, let's assume that inferiors receive a weight of 0.25, peers receive a weight of 0.5, and superiors receive a weight of 0.75. Here is where we run into a challenge. What if we don't know the competency of the believers in the group? In this case, we can estimate on the conservative side. However, what if we do not know how many competent believers there are in the group? This is somewhat more challenging. We can't just assign some random competency value, even if it is on the conservative side of things. Where does that leave us?

If the group is large enough and competency isn't exorbitantly difficult to attain, we may *at least* be able to estimate the minimum number of somewhat competent believers (inferiors) in the group. If we can estimate this number, we can spread the weight (that should be applied to the one competent believer) across the entire group and still end up with a reasonably accurate number. Remember, we are trying to establish the *minimal level of competence* in the group to remain on the conservative side in our calculations. So, suppose that we have a group where j = 'the total number of believers'. Also, suppose that g = 'the number of competent believers in the group'. We can estimate that there is a least one (1) inferior believer in the group. Assuming the weight values discussed above, the weight for that inferior believer (the k value) should be 0.25.

In this case, we can calculate the weight for the group to be $\frac{(k \times g)}{j}$.⁶⁴ If the total number of believers is 100,000 and we assume that there is at least one competent believer in the group (who should be weighted at 0.25), we end up with a weight of 0.0000025 applied to each member of the group. Fortunately, we can modify OU to account for this with a minor change. We will retain all the other assumptions listed above for this variation (OU^C). Now we can update the formulation for OU^C as follows:

$$C_n = C_{n-1} + \left[\sum_{i=1}^m O_m f(O_m) \frac{(k \times g)}{j} \right]$$

This formulation allows us to calculate an individual's final confidence after updating based on a large consensus in which we cannot easily assign weighting values to specific individuals.

3.3.2 The Argument from Common Consent based on Opinion Updating

Given the above, we can formulate an argument from common consent based on Opinion

Updating for belief in the existence of God. The argument is as follows:

Common Consent argument from Opinion Updated (CCO):

1. Given the right circumstances (see above), I should update my confidence in the existence of God if others agree regarding the existence of God.
2. There is an enormous number of people in the world (both in the past and the present) who believe God exists, and we can establish that a significant number of these are competent in their belief.

⁶⁴ Unfortunately, this somewhat weakens our results. Ideally, our model would ignore incompetent believers and only update on the competent believers. Because we do not know which believer is competent, we cannot do so. For instance, running OU where the two believers share a confidence of 0.6 and the second believer is an inferior weighted at 0.25, the updated confidence is 0.66. If we introduce another 99 incompetent believers and use OU^C, the updated confidence is approximately 0.65576.

3. Updating my confidence based on the OUC model and the number of believers in the world that share this belief, leads me to increase my confidence enough to hold quite securely that God exists.

This argument, of course, should be supported by some actual numbers and calculations.

In 2006, Phil. Zuckerman estimated that 88-93% of the world population believed in God. Let's suppose that Zuckerman's estimate was a bit off, and only 80% of the world population believes in the existence of God. The world population at the beginning of 2021 was approximately 7,733,988,000.⁶⁵ This puts us at approximately 6.19 billion people in the world who believe that God exists. This is a rather large number, and it does not include any of those in the past who have believed in the existence of God. For now, we can work with the number of those living in the world today. The number of people in our consensus group is 6.19 billion, but obviously, not all of those people are competent. How should we estimate the number of competent believers (or at least those who are less incompetent than most)? Well, we can start in academic philosophy departments. More specifically, we can begin with philosophers who are theists *and* who specialize in the arguments for the existence of God. We can assume for argument's sake that philosophers of religion have attempted to evaluate the evidence for God's existence more rigorously than the general population. For this reason, we will assume such individuals are competent or at least less incompetent than the general population. While it is difficult to estimate the numbers of theistic philosophers of religion worldwide, an initial estimation puts us at approximately 2,708.⁶⁶ To make the calculations easier, let's round up to 3,000. With such a large consensus group, that rounding up will be fairly insignificant. Recall that this number, which seems quite small in comparison to the whole group, serves the purpose

⁶⁵ "Population Clock." n.d. Accessed January 2, 2021. <https://www.census.gov/popclock/>.

⁶⁶ See "Estimates on Philosophers of Religion" in Appendix A.

of helping us calculate the minimal competence level across the entire group. The actual competence level could easily be much higher were we to include other competent believers that meet some minimum requirements, *e.g.*, other philosophers who are theists, other PhD's who are theists, etc.

Now, let's suppose that R holds the belief that God exists with a confidence of 0.6. Additionally, R comes to discover a *general consensus* in 80% of the world's population that God exists. R knows that there are at least 3,000 competent theists in the world. Given their training and area of expertise, we could weigh the competency of these theists at 0.75. After updating using OU^C , R's confidence should increase to approximately 0.8963.⁶⁷ This is a noteworthy increase in R's confidence. Even if we decreased the weight applied to the group to 0.25, R's confidence should still increase to approximately 0.775. Changing the scenario slightly, suppose that R starts with a confidence of 0.5. After updating, his confidence should still be 0.8704. So, even with some conservative numbers, we can arrive at a strong confidence that God exists.

This argument, unfortunately, still faces one of the objections that can be raised against the CCI and the CCS. It is not clear that all these people believe in the existence of the same God. One option is to hold a somewhat permissivist view and grant that these people get close enough to the same mark. Another option is to build a variation of the CCO aimed at a particular God. Let me sketch out briefly how we might attempt this with a specific idea of God.

⁶⁷ I developed a small custom application to run these calculations. The code can be found at <https://github.com/DoctorZatchmo/OU-App>.

3.3.3 A Cumulative CCO Argument

First, rather than supposing that 80% of the world population believes that God exists, let's assume that 80% of the world population believes in the existence of the supernatural. This should account for any and all variations of what we mean by "God". We can modify the CCO to conclude that the supernatural exists. We can use the modified version of CCO and the numbers above to conclude that R should have reason to believe in the existence of the supernatural confidently. What we need now is 1) a reason to prefer one account of God over others and 2) an evaluation of how much confidence we should have based on those who agree with that account. These two things should at least get us confident belief in one specific God. It may still be the case that those who believe in the supernatural and claim belief in some other God (or gods) may still believe, either wholly or partly, in the God identified in the argument below though they may refer to or think of this God by some different name or conception. We will leave others to address that question.

3.3.3.1 Focusing on a Specific Account of God

Various traditions accept the existence of the supernatural and claim belief in a God or gods. What we need is some epistemically acceptable reason for narrowing that list down to one account of God or gods. We are moving beyond something like the "God of the philosophers" and onto a more specific account. Much of what supports the reasons for belief in the truth of these different views can be found in the traditions, writings, and testimonial accounts associated with these different views. To evaluate which of these views is more likely to be true, we need some minimum standards for reliable historical and testimonial evidence that will increase the likelihood of one account over the others. I will sketch out some requirements below that should be a good beginning for doing just that.

Since we have already given an argument for reasonable belief in the supernatural on the basis of the CCO described above, if we can establish some reasonable support for a supernatural event tied to a specific tradition, then we have reason to believe the claims of that tradition. I will sketch the basics of how we could go about doing this below. Let's say that any such event must meet the following requirements⁶⁸:

Rule 1: The event must be such a thing that can be experienced by the outward senses, *i.e.*, it must be something that people can at least see and hear.

Rule 2: The event must be such a thing that is done publicly to be experienced by many witnesses.

Rule 3: Some public memorial, writings, or tradition must be established in reaction to the event, and that this memorial involved some sort of outward action that can be witnessed and recorded.

Rule 4: The memorial/tradition and action in rule 3 must be established and initiated from the time when the event took place.

The reasons behind the need for these rules are pretty straightforward. Rules 1 and 2 are in place to eliminate anyone from simply claiming that a supernatural event took place at a certain time, which can't be supported by multiple witnesses. Supposed I was to claim that, last Friday afternoon, I supernaturally froze Lake Thunderbird and transported the frozen lake over Norman, OK and then back to the lakebed. What would happen? Well, there are enough people that should have witnessed this event. If the claim was true, they could attest to it. If the claim was false, they could obviously call me out and say that I am lying. In addition to this, rules 1 and 2 also eliminate cases where there are "multiple witnesses" (supposedly), but the event was held in such a place that only people who would *not* question the event were present. This prevents the

⁶⁸ Requirements of this sort were initially proposed by Charles Leslie (1841) for a similar purpose.

inclusion of events where people attempt manufacturing “multiple eye-witnesses” support for their claim.

Rules 3 and 4 serve a similar role in assuring the genuine origins of the event. While it is relatively difficult to make undisputed claims when there are people still alive who should be able to attest to the truth of the claim (such as the frozen lake example), it is less difficult to conjure up some event that happened in the remote past when there are no people from that time to naysay the claim. Conjuring up some event that also meets the requirements of rules 3 and 4 is much more difficult. Recognition of specific memorials being observed since the event and having records of specific observances or actions that were specifically associated with the event, *i.e.*, corroborating evidence, is something that grants legitimacy to events in the remote past. Now that we have established our requirements, we can apply them to a specific account.

3.3.3.1.1 Applying Our Requirements

I will not endeavor to evaluate all views that hold to a belief in the existence of a God or gods. For my purposes in this project, I merely need to establish that one view has met our reliability requirements. From there, we can move on with our argument. In this case, I will make the case that the Christian account meets our reliability requirement. While the Christian tradition shares some facets with other major religions, one unique and essential piece of this tradition is the death and resurrection of Jesus of Nazareth. The dead ceasing to be dead and reverting back to life is generally accepted as a supernatural event. If this miraculous thing done by Jesus of Nazareth, which is essential to this tradition, meets our requirements, then the other things he is reported to have said and done are much more believable. (One of those things is that the Christian God exists.)

Now, let's ask the question, "Does the resurrection meet rules 1– 4?" Below are the responses to the rules based on common Christian tradition and writings:

Response 1: Jesus died by crucifixion on a Roman cross. This was an event that could be witnessed with the senses. Three days after his death, this same Jesus was witnessed alive. The witnesses saw him, heard him, and some even touched him.

Response 2: Jesus' crucifixion on a cross was a public event with many Jewish and non-Jewish eyewitnesses who attested to his death. Those who witnessed Jesus being alive after three days were numerous, and they witnessed him at numerous times.

Response 3: Jesus' disciples and followers instituted two traditions in response to the resurrection of Jesus. The first is called by various names: Eucharist, Communion, the Lord's Supper. This tradition involves the breaking of bread and drinking of wine is meant for remembering Jesus' sacrifice on the cross. The second tradition is baptism with water. This is a practice that all new followers of Jesus Christ adhered to from the beginning.

Response 4: Both the practice of the Eucharist and baptism by the followers of Jesus can be traced in origin back to the time immediately after his death and resurrection. Not only that, but they can be traced throughout history to this day.

The responses to Rules 3 and 4 are fairly uncontroversial. You will find few historians who reject them. The initiation of the practices of the Eucharist and Baptism were explicitly tied to the belief and the claim that Jesus was resurrected. It is a fairly bold claim to say that someone you know is back from the dead. An even bolder claim is that this person died in a public place where many people witnessed the death and that many people also saw this person alive again three days later. If these reports were fabricated, we would expect to find someone denying such reports as false, but only local Jewish leaders, who were at risk of humiliation and possible negative political fallout, made any objections to the claims. We would expect to find at least one non-interested party from that time rejecting the claims. Instead, it is not until around 400 years later, well after any original witnesses or immediate descendants were no longer around, that we

find the first recorded case of someone questioning the original reports.⁶⁹ Overall, this dramatically increases the likelihood of responses 1 and 2 being true. What if, however, I were to conclude that I should believe that the supernatural exists with a high confidence (based on the CCO) and that I should also believe that what Jesus said and did is true, which entails the existence of God. However, my confidence in the latter is only 0.6. Should the argument end there? No, one more important step should be added.

3.3.3.2 Updating from Superiors

Recall in Chapter 2, I discussed two scenarios involving a hit-and-run. In the first scenario, Rory had reasons to update his confidence in his belief regarding what happened after reading reports the next day from four media outlets that covered the event. His reason for updating was that the reports were all *coherent* for the truth of his belief. Because Rory and the reporters all witnessed the same interview, I argued that this was something like peer agreement (though perhaps slightly weaker since Rory was updating on the writings of the reporters that were coherent for the truth of his belief but were not in exact propositional agreement). Similarly, in the second scenario, I argued that Amy, an off-duty detective who witnessed the hit-and-run but was not interviewed by the reporters, also reason to update her confidence in her belief regarding what happened the following day after she read the reports by the four news outlets. Amy, however, is a *strict-superior* in this scenario since she has more evidence and better cognitive abilities. It should not be unreasonable to say that we can have similar cases where the one doing the updating based on the reporting is an epistemic inferior. After all, that is the position in which we would expect most readers of the reports to find themselves. That is also a similar situation in

⁶⁹ See Augustine's *Against Faustus* [*Contra Faustum Manichaeum*], AD 400.

which I would find myself after updating based on CCO regarding the existence of the supernatural and then concluding that I should have a 0.6 confidence in the truth of the resurrection of Jesus of Nazareth (and what the entails about God). Let me explain further.

Who would be my epistemic superior regarding the information in replies 1-4? Well, anyone who had greater access to evidence or had greater relevant cognitive abilities (or both), and someone who is at least not incompetent with regards to 1-4. Finding someone who has greater relevant cognitive abilities probably would not be too hard. Many brilliant people in the world would support replies 1-4. However, it is difficult to find a living person who would have significantly greater access to the evidence. Fortunately, we do have reports from people who were chronologically closer to the relevant evidence. In this case, we have reports from various sources only 90-175 years removed from the events. (One of these sources would have been close enough in time to have known, and likely did know, at least one of the eyewitnesses.) Let us consider the following five historical figures:

- Tertullian of Carthage (*c.* 155,/160 – *c.* 220+)
- Clement of Alexandria (*c.* 150 – *c.* 211,/215)
- Irenaeus of Lyons (*c.* 120,/140 – *c.* 200,/203)
- Justin Martyr (*c.* 100 – *c.* 165)
- Papias of Hierapolis (*c.* 70 – *c.* 163)

Recall again the hit-and-run scenario described in Chapter 2. The reports described in those scenarios are *coherent for* Amy's and Rory's sets of beliefs related to the hit-and-run. Similarly, the writings of the five historical figures listed above are *coherent for* replies 1-4. Given this, we can update on the basis of their reports. Since we would consider them to be epistemic superiors who should be weighted more heavily than peers, I had previously suggested that 0.75 might be a good number if peers are weighted at 0.5. However, as with the hit-and-run case, we will weaken the weight slightly since each report does not agree exactly with replies 1-4. Let's then assume that each report will be weighted at 0.625.

Suppose I were to start at a confidence of 0.6 regarding replies 1-4 and use $OU^\dagger$ to update based on the reports from the five historical figures listed above. We could assume that each individual's confidence in replies 1-4, H_J , is at least 0.6. (Given the reports, we could probably assume it is much higher, but we can just start with 0.6.) Because the reports are coherent for H_J , we could simplify and assume agreement of H_J in the formalization, and since the individuals are my epistemic superiors, we will use 0.625 as the k value. We will end up with the following values: $C_0 = \text{My}(H_J) = 0.6$, $O_1 = \text{Tertullian}(H_J) = 0.6$, $k_1 = 0.625$, $O_2 = \text{Clement}(H_J) = 0.6$, $k_2 = 0.625$, $O_3 = \text{Irenaeus}(H_J) = 0.6$, and $k_3 = 0.625$, and $O_4 = \text{Justin}(H_J) = 0.6$, $k_4 = 0.625$, and $O_5 = \text{Papias}(H_J) = 0.6$, $k_5 = 0.625$. Recall the formulation of $OU^\dagger$ is as follows:

Opinion Updating[†] ($OU^\dagger$):

l = Number of other agreeing opinions considered

$l > 0$

C_0 = The initial confidence of epistemic agent S

O_i = Confidence level of other agent's opinion i where $i \in \mathbb{N}^*$ and $i \leq l$

k_i = Weight i on opinion of others (peers, superiors, or inferiors) where $i \in \mathbb{N}^*$ and $i \leq l$

$0 < k \leq 1$

$$f(x) = \begin{cases} -1 & \text{if } x < 0.5 \\ 0 & \text{if } x = 0.5 \\ 1 & \text{if } x > 0.5 \end{cases}$$

$$C_n = C_{n-1} + \left[\sum_{i=1}^m O_m(1 - C_{m-1})f(O_m) k_m \right]$$

The results of my updating based on $OU^\dagger$ would be the following:

Table I. Summary of H_J agreement updating calculations with $OU^\dagger$

$C_1 = C_0 + [O_1(1 - C_0)1]$	$= 0.7500$
$C_2 = C_0 + [(O_1(1 - C_0)0.625) + (O_2(1 - C_1)0.625)]$	≈ 0.8438
$C_3 = C_0 + [O_1(1 - C_0)0.625 + O_2(1 - C_1)0.625 + O_3(1 - C_2)0.625]$	≈ 0.9023
$C_4 = C_0 + [O_1(1 - C_0)0.625 + O_2(1 - C_1)0.625 + O_3(1 - C_2)0.625 + O_4(1 - C_3)0.625]$	≈ 0.9390
$C_4 = C_0 + [O_1(1 - C_0)0.625 + O_2(1 - C_1)0.625 + O_3(1 - C_2)0.625 + O_4(1 - C_3)0.625 + O_5(1 - C_4)0.625]$	≈ 0.9619

Given the above, not only should I believe that the supernatural exists, but I should also believe that God exists. More specifically, I should believe that the God taught about by Jesus of Nazareth exists.

3.4 Conclusion

While some beliefs held by large groups of people can be problematic, irrational, and at times even dangerous, there are many such beliefs, which we refer to as traditions, that are indeed good, useful, and ultimately rational. We should be cautious about discounting too quickly long-held traditions of this sort. The fact that I believe something with a high degree of confidence because many other people believe that same thing should not undermine the confidence that I have in that belief. A *consensus gentium* argument aims to show that belief in God is a shared belief or tradition of this sort. While traditional versions of such arguments have fallen somewhat out of favor in contemporary philosophy of religion due to their apparent weakness in comparison to other more popular theistic arguments, the project is not without merit. After all, there seems to be something uniquely valuable about clarifying and showing the reasonableness of a line of thought echoed a multitude of times over the last two millennia (at least).

I have discussed two recent versions of this argument that I expect many will ultimately find quite persuasive. Nonetheless, some may find reasons to reject both of these arguments. For those people, I have offered an alternative argument. My argument relies upon the rather natural

intuition that when others agree with us, that agreement should boost confidence for the belief in question. One unique aspect of this argument is that it includes a model for updating one's confidence based on specific numbers instead of more general references to very large numbers. The upside to this approach is that we can get specific degrees of confidence that should result from updating on the model. The downside is that it requires specific numbers that may be somewhat hard to come by. What I have attempted to show is that, even with some fairly conservative numbers, we can arrive at a significant degree of confidence for belief in the existence of God.

Wrapping Things Up

At the outset of this project, I initially had little interest in interacting with the Peer Disagreement literature. After all, much has already been said with which I agree, and much has been said with which I disagree. Much of what I wanted to say has already been said to some extent, and it seemed like what I wanted to say would simply add to the cacophony of voices discussing the subject. A pair of simple questions led me to undertake this project: “Is there another side to the problem of peer disagreement? Is there an inverse to the problem?” As it turns out, though there have been hints here and there since it is pretty evident that agreement is a related topic, but there has been little focused discussion on the related symmetries between agreement and disagreement. More specifically, there has been minimal discussion on the need for a unified theory that encompasses both disagreement and agreement. My hope is that this project will draw some attention to that need.

Chapter 1 laid the groundwork for drawing attention to the need for a unified theory, and Chapter 2 expanded upon that need beyond the focus on peers. My analysis of the challenges that each of the prominent positions on disagreement may face when attempting to address the agreement side of the issue was by no means exhaustive. At best, I merely scratched the surface. Merely scratching the surface, however, should be enough to draw attention to the need for further discussion. More significantly, my hope is to provide a useful tool for updating credences that can be used by those who do take on the challenge of developing a unified theory of rational agreement and disagreement. While Bayesianism has been embraced by many in formal epistemology, traditional epistemologists are reasonably hesitant to wade into those waters. Nonetheless, the discussion has a natural tendency to wander into the territory of credences/confidences, which are probabilistic in nature. My aim with developing the OU model

was to provide a significantly robust tool that could at least be used for heuristic purposes by those working in this area without requiring them to become Bayesians since the OU model mainly relies upon basic multiplication, addition, and subtraction, and it requires no knowledge of probability theory. Perhaps the next step is to translate the model into a Bayesian framework for those who desire a more idealized model.

Chapter 3 reflects the impact that work in Social Epistemology has on other areas of philosophy. While I did not initially set out to develop a *consensus gentium* argument, the OU model and the discussion of agreement between peers, inferiors, and superiors and the expansion of that into questions regarding consensus seems to lend itself to that argument quite effortlessly. The fruit of this was what I hope is a new and interesting take on the *consensus gentium* argument that might be appealing to those who find other variations of the argument unappealing. That being said, work on the requirements for consensus-based knowledge is ongoing and it seems plain that the discussion of rational agreement spills over into the former. I limited my argument to degrees of confidence based on consensus. Nonetheless, it seems quite natural that one of the next steps in this area will be analysis on the relevant requirements for consensus-based knowledge and how that plays out in this discussion. That, of course, will have some direct relevance on an updated version of the argument from Common Consent based on Opinion Updating.

Appendix A

Other Definitions:

Independence: In evaluating the epistemic credentials of another person's belief about p , in order to determine how (if at all) to modify one's own belief about p , one should do so in a way that is independent of the reasoning behind one's own initial belief about p (Christensen, 2009, 758).

*Independence**: In evaluating the epistemic credentials of another person's belief about p , in order to determine how (if at all) to modify one's own belief about p , one should do so in a way that is independent of one's assessment of those considerations that led one to initially believe as one does about p (Kelly, 2013, 40).

*Inverse-Independence**: In evaluating the epistemic credentials of another person's belief about p , in order to determine how (if at all) to modify one's own belief about p , one should do so in a way that is independent of one's assessment of those considerations that led one to initially believe as one does about p (Kelly, 2013, 40).

Estimates on Philosophers of Religion:

Philosophers who are also theists often have a focus in the philosophy of religion or they have a significant interest in it. We can assume for argument's sake that philosophers of religion or philosophers who have an interest in philosophy of religion have attempted to evaluate the evidence for God's existence more rigorously than the general population. For this reason, we will assume such individuals are competent, or at least less incompetent than the general population. While it is difficult to estimate the numbers of theistic philosophers worldwide. We can find some numbers for philosophers in the US and Canada. So, let's start there.

Based on a survey by PhilPapers⁷⁰, approximately 5.5% of philosophers specialize in philosophy of religion. Of those, approximately 73.3% hold to theism. In the Fall of 2017, there

⁷⁰ "The PhilPapers Surveys." 2009. Preliminary Survey Results. 2009.
<https://philpapers.org/surveys/results.pl>.

were 6,735 faculty members and 24,970 in US philosophy departments.⁷¹ We can estimate approximately 2,040 academic philosophers in Canada, of which approximately 297 would be estimated to be theists.⁷² That gives us approximately 1,354 competent theists in the US and Canada.⁷³ Let's assume we can get double that number by including recently graduated philosophy students who aren't in full-time positions and by including philosophers outside of the US and Canada. We land on a total of 2,708 of philosophers who specialize in philosopher of religion and are theists.

⁷¹ Porter, Anne Marie, Jack Pold, and Susan White. 2018. "The State of Philosophy Departments in Four-Year Colleges and Universities (2017)." American Academy of Arts & Sciences. https://www.amacad.org/sites/default/files/media/document/2020-05/HDS3_Philosophy%20Profile.pdf.

⁷² "Philosophy in Canada - Department of Philosophy - Arts and Science - University of Saskatchewan." Accessed January 11, 2021. <https://artsandscience.usask.ca/philosophy/departments/Resources/philosophy-in-canada.php#WhoStudiesPhilosophy>.

⁷³ Whether they are competent or merely less incompetent could be up for some debate. The important part is that they are generally more competent than others.

Appendix B

Table I. Three Claims of the Equal Weight View and their Inversions

Disagreement Claims	Inverted Claims / Agreement Claims
<i>Defeat</i> : Upon learning that an epistemic peer disagrees about some proposition p , this gives me a reason to believe I am mistaken about p .	<i>Validation</i> : Upon learning that an epistemic peer agrees about some proposition p , this gives me a reason to believe I am correct about p .
<i>Independence</i> : Any reasons I have to discount another's opinion about p must be independent of the reasoning behind my own initial belief about p that led to the disagreement.	<i>Inverted Independence</i> : Any reasons I have to accept another's opinion about p must be independent of the reasoning behind my own initial belief about p that led to the agreement.
<i>Equal-Weight</i>: I have just as much reason to trust my peer's reasons for p as my own reasons for p. His reasons lead him to p, and mine lead me to not p. When we <i>disagree</i>, his conclusion from his reasons weaken my conclusion from my reasons.	<i>Inverted Equal-Weight</i>: I have just as much reason to trust my peer's reasons for p as my own reasons for p. His reasons lead him to p, and mine also lead me to p. When we <i>agree</i>, his conclusion from his reasons strengthen my conclusion from my reasons.

Appendix C

The AU_s Model:

$$C_n = C_{n-1} + k \left[\sum_{i=1}^m O_m(1 - C_{m-1}) \right]$$

The AU Model:

$$C_n = C_{n-1} + k \left[\sum_{i=1}^m O_m(1 - C_{m-1})f(O_m) \right]$$

Let:

n = Number of Other Options Considered

$n > 0$

C_o = The initial confidence of epistemic agent S

O_x = Confidence level of other agent's opinion

$$f(x) \begin{cases} -1 & \text{if } x < 0.5 \\ 0 & \text{if } x = 0.5 \\ 1 & \text{if } x > 0.5 \end{cases}$$

The OU Model:

$$C_n = C_{n-1} + \left[\sum_{i=1}^m O_m f(O_m) k_m \right]$$

Let:

n = Number of Other Options Considered

$n > 0$

C_o = The initial confidence of epistemic agent S

O_x = Confidence level of other agent's opinion

$$f(x) \begin{cases} (-1 + C_{m-1}) & \text{if } x < 0.5 \\ 0 & \text{if } x = 0.5 \\ (1 - C_{m-1}) & \text{if } x > 0.5 \end{cases}$$

The OU[†] Model:

$$C_n = C_{n-1} + \left[\sum_{i=1}^m O_m (1 - C_{m-1}) f(O_m) k_m \right]$$

l = Number of other agreeing opinions considered

$l > 0$

C_0 = The initial confidence of epistemic agent S

O_i = Confidence level of other agent's opinion i where $i \in \mathbb{N}^*$ and $i \leq l$

k_i = Weight i on opinion of others (peers, superiors, or inferiors) where $i \in \mathbb{N}^*$ and $i \leq l$

$0 < k \leq 1$

$$f(x) \begin{cases} -1 & \text{if } x < 0.5 \\ 0 & \text{if } x = 0.5 \\ 1 & \text{if } x > 0.5 \end{cases}$$

The OU^C Model:

$$C_n = C_{n-1} + \left[\sum_{i=1}^m O_m f(O_m) \frac{(k \times g)}{j} \right]$$

Let:

l = Number of other agreeing opinions considered

$l > 0$

C_0 = The initial confidence of epistemic agent S

O_i = Confidence level of other agent's opinion i where $i \in \mathbb{N}^*$ and $i \leq l$

k_i = Weight i on opinion of others (peers, superiors, or inferiors) where $i \in \mathbb{N}^*$ and $i \leq l$

$$0 < k \leq 1$$

j = The total number of believers in the group.

g = The number of competent believers in the group.

$$f(x) \begin{cases} (-1 + C_{m-1}) & \text{if } x < 0.5 \\ 0 & \text{if } x = 0.5 \\ (1 - C_{m-1}) & \text{if } x > 0.5 \end{cases}$$

Bibliography

- Alston, William P. *Beyond "Justification": Dimensions of Epistemic Evaluation*. Vol. 81. Cornell University Press, 2005.
- . "Epistemic Circularity." *Philosophy and Phenomenological Research* 47, no. 1 (1986): 1–30.
- Baker, Joseph O. "The Variety of Religious Experiences." *Review of Religious Research* 51, no. 1 (2009): 39–54.
- Beckwith, Francis J. "A Christian Account of Why David Novak Is Right About the Same God Question." In *The Achievement of David Novak: A Catholic-Jewish Dialogue*. Eugene, OR: Pickwick Publications, 2021.
- Bergmann, Michael. "Epistemic Circularity: Malignant and Benign." *Philosophy and Phenomenological Research* 69, no. 3 (2004): 709–27.
- Bogardus, Tomas. "A Vindication of the Equal Weight View." *Episteme* 6, no. 3 (2009): 324–35.
- Bogardus, Tomas, and Mallorie Urban. "How to Tell Whether Christians and Muslims Worship the Same God." *Faith and Philosophy* 34, no. 2 (2017): 176–200.
- Brössel, Peter, and Anna-Maria A. Eder. "How to Resolve Doxastic Disagreement." *Synthese* 191, no. 11 (2014): 2359–81.
- Burge, Tyler. "Content Preservation." *Philosophical Review* 102, no. 4 (1993): 457–88.
- Calvin, John. *Institutes of the Christian Religion*. Translated by Henry Beveridge. Christian Classics Ethereal Library, 1536. <https://ccel.org/ccel/calvin/institutes/institutes>.
- Christensen, David. "Disagreement as Evidence: The Epistemology of Controversy." *Philosophy Compass* 4, no. 5 (2009): 756–67.
- . "Disagreement, Question-Begging and Epistemic Self-Criticism." *Philosophers' Imprint* 11 (2011).
- . "Epistemology of Disagreement: The Good News." *Philosophical Review* 116, no. 2 (2007): 187–217.
- . "Higher Order Evidence." *Philosophy and Phenomenological Research* 81, no. 1 (2010): 185–215.
- De Cruz, Helen. "Believing to Belong: Addressing the Novice-Expert Problem in Polarized Scientific Communication." *Social Epistemology* 34, no. 5 (2020): 440–52.
- . *Religious Disagreement*. Cambridge University Press, 2019.

- DiPaolo, Joshua. "Higher-Order Defeat Is Object-Independent." *Pacific Philosophical Quarterly*, forthcoming.
- Dobrzeńiecki, Marek. "Is the Fact That Other People Believe in God a Reason to Believe? Remarks on the Consensus Gentium Argument." *European Journal for Philosophy of Religion* 10, no. 3 (2018): 133.
- Dogramaci, Sinan, and Sophie Horowitz. "An Argument for Uniqueness About Evidential Support." *Philosophical Issues* 26, no. 1 (2016): 130–47.
- Easwaran, Kenny, Luke Fenton-Glynn, Christopher Hitchcock, and Joel D. Velasco. "Updating on the Credences of Others: Disagreement, Agreement, and Synergy." *Philosophers' Imprint* 16 (2016): 1–39.
- Elga, Adam. "Reflection and Disagreement." *Noûs* 41, no. 3 (2007): 478–502.
- . "Reflection and Disagreement." *Noûs* 41, no. 3 (2007): 478–502.
- Elkin, Lee, and Gregory Wheeler. "Resolving Peer Disagreements Through Imprecise Probabilities." *Noûs* 52, no. 2 (2018): 260–78.
- Feldman, Richard. "Epistemological Puzzles about Disagreement." In *Epistemology Futures*, edited by Stephen Hetherington, 216–36. Oxford University Press, 2006.
- . "Evidentialism, Higher-Order Evidence, and Disagreement." *Episteme* 6, no. 3 (2009): 294–312.
- . "Reasonable Religious Disagreements." In *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, edited by Louise Antony, 194–214. Oxford University Press, 2006.
- . "Respecting the Evidence." *Philosophical Perspectives* 19, no. 1 (2005): 95–119.
- Fitelson, Branden. "A Bayesian Account of Independent Evidence with Applications." *Philosophy of Science* 68, no. S3 (2001): S123–40.
- . *Studies in Bayesian Confirmation Theory*. Dissertation, 2001.
- . "Wayne, Horwich, and Evidential Diversity." *Philosophy of Science* 63, no. 4 (1996): 652–60.
- Fitelson, Branden, and David Jehle. "What Is the 'Equal Weight View'?" *Episteme* 6, no. 3 (2009): 280–93.
- Foley, Richard. "Egoism in Epistemology." In *Socializing Epistemology: Social Dimensions of Knowledge*. Lanham: Rowman and Littlefield, 1994.
- . *Intellectual Trust in Oneself and Others*. Cambridge University Press, 2001.

- Forster, Malcolm, and Elliott Sober. "How to Tell When Simpler, More Unified, or Less Ad Hoc Theories Will Provide More Accurate Predictions." *British Journal for the Philosophy of Science* 45, no. 1 (1994): 1–35.
- Frances, Bryan. "Disagreement." In *The Routledge Companion to Epistemology*, edited by Duncan Pritchard and Sven Bernecker. Routledge, 2010.
- . *Disagreement*. Polity, 2014.
- . "Discovering Disagreeing Epistemic Peers and Superiors." *International Journal of Philosophical Studies* 20, no. 1 (2012): 1–21.
- Fricker, Elizabeth. "Testimony and Epistemic Autonomy." In *The Epistemology of Testimony*, edited by Jennifer Lackey and Ernest Sosa, 225–253. Oxford University Press, 2006.
- Fricker, Miranda, Peter Graham, David Henderson, and Nikolaj Jang Pedersen. *The Routledge Handbook of Social Epistemology*. New York, USA: Routledge, 2019.
- Fuller, Glen, Jason Wilson, and Christian McCrea. "Troll Theory?: Issue 22: Trolls and the Negative Space of the Internet." *Fibreculture Journal* 22 (2013): 1–15.
- Gabbay, Dov, Stephan Hartmann, and John Woods. *Handbook of the History and Philosophy of Logic, Vol. 10: Inductive Logic*. Elsevier, forthcoming.
- Gardiner, Georgi. "The Commutativity of Evidence: A Problem for Conciliatory Views of Peer Disagreement." *Episteme* 11, no. 1 (2014): 83–95.
- Goldman, Alvin. "Expertise." *Topoi* 37, no. 1 (2018): 3–10. <https://doi.org/10.1007/s11245-016-9410-3>.
- Goldman, Alvin I. "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research* 63, no. 1 (2001): 85–110.
- . *Knowledge in a Social World*. Vol. 36. Oxford University Press, 1999.
- Greco, Daniel, and Brian Hedden. "Uniqueness and Metaepistemology." *Journal of Philosophy* 113, no. 8 (2016): 365–95.
- Gutting, Gary. *Religious Belief and Religious Skepticism*. University of Notre Dame Press, 1982.
- Hardwig, John. "Epistemic Dependence." *Journal of Philosophy* 82, no. 7 (1985): 335–49.
- . "The Role of Trust in Knowledge." *Journal of Philosophy* 88, no. 12 (1991): 693–708.
- Harman, Gilbert H. "The Inference to the Best Explanation." *Philosophical Review* 74, no. 1 (1965): 88–95.
- Hick, John. *God and the Universe of Faiths: Essays in the Philosophy of Religion*. Macmillan Press, 1988.

- Huemer, Michael. "Epistemological Egoism and Agent-Centered Norms." In *Evidentialism and Its Discontents*, edited by Trent Dougherty, 17. Oxford University Press, 2011.
- Intemann, Kristen. "Why Diversity Matters: Understanding and Applying the Diversity Component of the National Science Foundation's Broader Impacts Criterion." *Social Epistemology* 23, no. 3 (2009): 249–66.
- Inwagen, Peter van. "It Is Wrong, Everywhere, Always, for Anyone, to Believe Anything upon Insufficient Evidence." In *Faith, Freedom and Rationality*, edited by Jeff Jordan and Daniel Howard-Snyder, 137–54. Savage, Maryland: Rowman and Littlefield, 1996.
- Keas, Michael. "Systematizing the Theoretical Virtues." *Synthese* 195, no. 6 (2018): 2761–93.
- Keener, C.S. *Miracles : 2 Volumes: The Credibility of the New Testament Accounts*. Baker Publishing Group, 2011. <https://books.google.com/books?id=oCrSpYJvGakC>.
- Kelly, Thomas. "Consensus Gentium: Reflections on the 'Common Consent' Argument for the Existence of God." In *Evidence and Religious Belief*, edited by Kelly James Clark and Raymond J. VanArragon. Oxford University Press, 2011.
- . "Disagreement and the Burdens of Judgment." In *The Epistemology of Disagreement: New Essays*, edited by David Phiroze Christensen and Jennifer Lackey. Oxford University Press, 2013.
- . "Evidence." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2016. Metaphysics Research Lab, Stanford University, 2016. <https://plato.stanford.edu/archives/win2016/entries/evidence/>.
- . "Evidence Can Be Permissive." In *Contemporary Debates in Epistemology*, edited by Matthias Steup and John Turri, 298. Blackwell, 2013.
- . "Peer Disagreement and Higher Order Evidence." In *Social Epistemology: Essential Readings*, edited by Alvin I. Goldman and Dennis Whitcomb, 183--217. Oxford University Press, 2010.
- . "The Epistemic Significance of Disagreement." In *Oxford Studies in Epistemology, Volume 1*, edited by John Hawthorne and Tamar Gendler, 1:167–96. Oxford University Press, 2005.
- Keysar, Ariela, and Juhem Navarro-Rivera. "A World of Atheism: Global Demographics." In *The Oxford Handbook of Atheism*, edited by Stephen Bullivant and Michael Ruse, First., 553–86. Oxford and New York: Oxford University Press UK, 2013.
- King, Nathan L. "Disagreement: What's the Problem? Or A Good Peer Is Hard to Find." *Philosophy and Phenomenological Research* 85, no. 2 (2012): 249–72.
- Kopec, Matthew. "A Counterexample to the Uniqueness Thesis." *Philosophia* 43, no. 2 (2015): 403–9.

- Lackey, Jennifer. "A Justificationist View of Disagreement's Epistemic Significance." In *Social Epistemology*, edited by Adrian Haddock, Duncan Pritchard, and Alan Millar, 298–325. Oxford University Press, 2010a.
- . "What Should We Do When We Disagree?" In *Oxford Studies in Epistemology*, edited by Tamar Szabó Gendler and John Hawthorne, Vol. 3. Oxford University Press, 2010b.
- Lasonen-Aarnio, Maria. "Higher-Order Evidence and the Limits of Defeat." *Philosophy and Phenomenological Research* 88, no. 2 (2014): 314–45.
- Lehrer, Keith, and Carl Wagner. *Rational Consensus in Science and Society a Philosophical and Mathematical Study*, 1981.
- Leslie, Charles. *A Short and Easy Method with the Deists*. Londa: Society for Promoting Christian Knowledge, 1841.
- Levi, Isaac. "Consensus as Shared Agreement and Outcome of Inquiry." *Synthese* 62, no. 1 (1985): 3–11.
- Lipton, Peter. *Inference to the Best Explanation*. 2nd ed. Vol. 102. Routledge/Taylor and Francis Group, 2004.
- . "Is Explanation a Guide to Inference? A Reply to Wesley Salmon." In *Explanation: Theoretical Approaches and Applications*, edited by G. Hon and Sam S. Rakover. Springer, 2003.
- Mackie, John L. *The Miracle of Theism*. New York: Oxford University Press, 1982.
- March, Evita. "Psychopathy, Sadism, Empathy, and the Motivation to Cause Harm: New Evidence Confirms Malevolent Nature of the Internet Troll." *Personality and Individual Differences* 141 (2019): 133–37. <https://doi.org/10.1016/j.paid.2019.01.001>.
- Matheson, Jonathan. "Disagreement and the Ethics of Belief." In *The Future of Social Epistemology: A Collective Vision*, edited by James Collier, 139–48, 2015.
- . "The Case for Rational Uniqueness." *Logic and Episteme* 2, no. 3 (2011): 359–73.
- Matheson, Jonathan, and Bryan Frances. "Disagreement." In *Stanford Encyclopedia of Philosophy*, 2018.
- McGrew, Lydia. "Bayes Factors All the Way: Toward a New View of Coherence and Truth." *Theoria* 82, no. 4 (2016): 329–50.
- . "Undesigned Coincidences and Coherence for an Hypothesis." *Erkenntnis* 85, no. 4 (2020): 801–28.
- McGrew, Timothy. "Confirmation, Heuristics, and Explanatory Reasoning." *British Journal for the Philosophy of Science* 54, no. 4 (2003): 553–67.

- . “Miracles.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2019. Metaphysics Research Lab, Stanford University, 2019.
<https://plato.stanford.edu/archives/spr2019/entries/miracles/>.
- McIntosh, Chad A. “Nontraditional Arguments for Theism.” *Philosophy Compass* 14, no. 5 (2019): 1–14.
- Meierding, Loren. “The Consensus Gentium Argument.” *Faith and Philosophy* 15, no. 3 (1998): 271–97.
- Mill, John Stuart. “Theism.” In *Three Essays on Religion*, edited by Louis J. Matz. Ontario, Canada: Broadview Editions, 2009.
- Miller, Boaz. “The Social Epistemology of Consensus and Dissent.” In *The Routledge Handbook of Social Epistemology*, edited by David Henderson, Peter Graham, Miranda Fricker, and Nikolaj Jang Lee Linding Pedersen, 228–37. Routledge, 2019.
- . “The Social Epistemology of Consensus and Dissent.” In *The Routledge Handbook of Social Epistemology*, edited by David Henderson, Peter Graham, Miranda Fricker, and Nikolaj Jang Lee Linding Pedersen, 228–37. Routledge, 2019.
- . “When Is Consensus Knowledge Based? Distinguishing Shared Knowledge from Mere Agreement.” *Synthese* 190, no. 7 (2013): 1293–1316.
- . “When Is Consensus Knowledge Based? Distinguishing Shared Knowledge from Mere Agreement.” *Synthese* 190, no. 7 (2013): 1293–1316.
- . “Why Knowledge Is the Property of a Community and Possibly None of Its Members.” *Philosophical Quarterly* 65, no. 260 (2015): 417–41.
- Moon, Andrew. “Independence and New Ways to Remain Steadfast in the Face of Disagreement.” *Episteme* 15, no. 1 (2018): 65–79.
- Mulligan, Thomas. “Disagreement, Peerhood, and Three Paradoxes of Conciliationism.” *Synthese* 192, no. 1 (2015): 67–78.
- Nguyen, C. Thi. “Echo Chambers and Epistemic Bubbles.” *Episteme* 17, no. 2 (2020): 141–61.
- NW, 1615 L. St, Suite 800 Washington, and DC 20036 USA 202-419-4300 | Main 202-419-4349 | Fax 202-419-4372 | Media Inquiries. “Americans’ Beliefs about the Nature of God.” *Pew Research Center’s Religion & Public Life Project* (blog), April 25, 2018.
<https://www.pewforum.org/2018/04/25/when-americans-say-they-believe-in-god-what-do-they-mean/>.
- O’Briant, Walter H. “Is There an Argument Consensus Gentium?” *International Journal for Philosophy of Religion* 18, no. 1/2 (1985): 73–79.

- Palmira, Michele. "Expert Deference about the Epistemic and Its Metaepistemological Significance." *Canadian Journal of Philosophy* 50, no. 4 (2020): 524–38.
- Pasnau, Robert. "Disagreement and the Value of Self-Trust." *Philosophical Studies* 172, no. 9 (2015): 2315–39. <https://doi.org/10.1007/s11098-014-0413-x>.
- Peter, Fabienne. "Epistemic Self-Trust and Doxastic Disagreements." *Erkenntnis* 84, no. 6 (2019): 1189–1205.
- . "Epistemic Self-Trust and Doxastic Disagreements." *Erkenntnis* 84, no. 6 (2019): 1189–1205.
- "Philosophy in Canada - Department of Philosophy - Arts and Science - University of Saskatchewan." Accessed January 11, 2021. <https://artsandscience.usask.ca/philosophy/departement/Resources/philosophy-in-canada.php#WhoStudiesPhilosophy>.
- "Population Clock." Accessed January 2, 2021. <https://www.census.gov/popclock/>.
- Porter, Anne Marie, Jack Pold, and Susan White. "The State of Philosophy Departments in Four-Year Colleges and Universities (2017)." American Academy of Arts & Sciences, 2018. https://www.amacad.org/sites/default/files/media/document/2020-05/HDS3_Philosophy%20Profile.pdf.
- Priest, Maura. "Inferior Disagreement." *Acta Analytica* 31, no. 3 (2016): 263–83.
- Raleigh, Thomas. "Another Argument Against Uniqueness." *Philosophical Quarterly* 67, no. 267 (2017): 327–46.
- "Report on the Canadian Philosophical Association Equity Survey 2018." The Canadian Philosophical Association. Accessed January 11, 2021. [Report on the Canadian Philosophical Association Equity Survey 2018](#).
- Rescher, Nicholas. *Pluralism: Against the Demand for Consensus*. Oxford University Press, 1993.
- Rey, Georges. "Does Anyone Really Believe in God?" In *The Experience of Philosophy*, edited by Daniel Kolak and Raymond Martin. Oxford: Oxford University Press, 2006.
- Roche, William A. "Coherentism, Truth, and Witness Agreement." *Acta Analytica* 25, no. 2 (2010): 243–57.
- Rosa, Luis. "Justification and the Uniqueness Thesis." *Logos and Episteme*, no. 4 (2012): 571–77.
- . "Justification and the Uniqueness Thesis Again." *Logos and Episteme* 7, no. 1 (2016): 95–100.

- Schechter, Joshua. "Rational Self-Doubt and the Failure of Closure." *Philosophical Studies* 163, no. 2 (2013): 428–52.
- Schoenfield, Miriam. "Permissivism and the Value of Rationality: A Challenge to the Uniqueness Thesis." *Philosophy and Phenomenological Research* 99, no. 2 (2019): 286–97.
- Schwartz, W.A., J.B. Cobb, F.J. Beckwith, G.R. McDermott, J. Walls, J. Cumming, D.W. Shenk, et al. *Do Christians, Muslims, and Jews Worship the Same God?: Four Views*. Counterpoints: Bible and Theology. Zondervan Academic, 2019.
<https://books.google.com/books?id=cVGNDwAAQBAJ>.
- "SDA." Accessed January 11, 2021. <https://sda.bodleian.ox.ac.uk/sda/#!/about>.
- Shogenji, Tomoji. "Why Does Coherence Appear Truth-Conducive?" *Synthese* 157, no. 3 (2007): 361–72.
- Sliwa, Paulina, and Sophie Horowitz. "Respecting All the Evidence." *Philosophical Studies* 172, no. 11 (2015): 2835–58.
- Stegenga, Jacob. "Robustness, Discordance, and Relevance." *Philosophy of Science* 76, no. 5 (2009): 650–61.
- Steup, Matthias. "Is Epistemic Circularity Bad?" *Res Philosophica* 90, no. 2 (2013): 215–35.
- Sullivan, Meghan. "Semantics for Blasphemy." In *Oxford Studies in Philosophy of Religion, Vol. IV*, edited by Jonathan L. Kvanvig. Oxford University Press, 2012.
- Surowiecki, James. *The Wisdom of Crowds*. New York: Anchor Books, 2005.
- Thagard, Paul. "The Best Explanation: Criteria for Theory Choice." *Journal of Philosophy* 75, no. 2 (1978): 76–92.
- Preliminary Survey results. "The PhilPapers Surveys," 2009.
<https://philpapers.org/surveys/results.pl>.
- Tucker, Aviezer. "The Epistemic Significance of Consensus." *Inquiry: An Interdisciplinary Journal of Philosophy* 46, no. 4 (2003): 501–21.
- . "The Inferences of Common Causes Reduced to Common Origins." *Studies in History and Philosophy of Science Part A* 81 (2020): 105–15.
- White, Roger. "Epistemic Permissiveness." *Philosophical Perspectives* 19, no. 1 (2005): 445–59.
- Whiting, Daniel. "Whither Higher-Order Evidence?" In *Higher-Order Evidence: New Essays*, edited by Asbjørn Steglich-Petersen and Mattias Skipper. Oxford University Press, forthcoming.
- Wilson, Alastair. "Disagreement, Equal Weight and Commutativity." *Philosophical Studies* 149, no. 3 (2010): 321–26.

Zagzebski, Linda. "Epistemic Authority." *Epistemology and Philosophy of Science* 53, no. 3 (2017): 92–107.

———. *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief*. Oup Usa, 2012.

———. "Epistemic Self-Trust and the Consensus Gentium Argument." In *Evidence and Religious Belief*, edited by Kelly James Clark and Raymond J. VanArragon. Oxford University Press, 2011.

———. "Epistemic Self-Trust and the Consensus Gentium Argument." In *Evidence and Religious Belief*, edited by Kelly James Clark and Raymond J. VanArragon. Oxford University Press, 2011.

Zuckerman, Phil. "Atheism: Contemporary Numbers and Patterns." In *The Cambridge Companion to Atheism*, edited by Michael Martin. New York, NY: Cambridge University Press, 2006.