UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

DISAGREEMENT AND SKEPTICISM: A GRECOIAN RESPONSE TO THE SKEPTICAL THREAT OF EPISTEMIC SUPERIOR DISAGREEMENT

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

GARY SCOTT OSMUNDSEN
Norman, Oklahoma
2014

DISAGREEMENT AND SKEPTICISM: A GRECOIAN RESPONSE TO THE SKEPTICAL THREAT OF EPISTEMIC SUPERIOR DISAGREEMENT

A DISSERTATION APPROVED FOR THE
DEPARTMENT OF PHILOSOPHY

BY

_____
Dr. Wayne Riggs, Chair

_____
Dr. Neal Judisch

_____
Dr. Martin Montminy

_____
Dr. Linda Zagzebski

_____
Dr. Piers Hale

To my beloved wife Julie, your patience, love, and support cannot be overstated.
I'm so grateful to be able to travel through life with you. There is no one else
I'd want to make this journey with.

To my dear, precious children, Emmafaith, Joel, and Avagrace, you each bring me
more joy, perspective, wonder, and awe than you will ever know.

And to my parents, your constant care, support, and work ethic have contributed to all my successes.
I'm grateful you're here to see and celebrate with me the completion of this journey.

ACKNOWLEDGMENTS

I have been thinking about the problem of disagreement ever since I started graduate school. Reading so much philosophy naturally exposed me to much disagreement. As a result, I often found myself suffering from intellectual dizziness. Part of my interest in pursuing this topic, then, was to stop the room from spinning. The nascent phases of this dissertation started in Wayne Riggs's and Linda Zagzebki's seminars in Social Epistemology, both of which covered the epistemology of disagreement. In writing this dissertation, I have benefited from the comments, suggestions, insights, and questions provided by so many gracious people. I am truly grateful to all my professors, colleagues, friends, and family members who have left their mark on both me and this dissertation. Of all the people to whom I am deeply indebted, four deserve special recognition. First, my friend and committee chair, Wayne Riggs, discussed with me all the different stages and chapters of this dissertation. I am especially indebted to him for his advice, guidance, and insights. Both the quality of my time at the University of Oklahoma and the quality of this dissertation have been enriched by Wayne's jovial spirit and philosophical prowess. Second, my friend and professor, Martin Montminy, was so gracious in offering me his time, talent, and enthusiastic support. Both Wayne and Martin were so instrumental in clarifying the thoughts and writing that went into this dissertation. Third, my friend and professor, Linda Zagzebski, has modeled for me an exemplary philosopher. Her congeniality, largess, and productivity are just a few of her characteristics that I've tried to emulate. Fourth, John Greco's writings and conversations with me have been of great benefit. I take him to be one of the best epistemologists in the field today. If only more philosophers wrote as clearly as John does, many more people would read philosophy and enjoy it! There are still

many people I have benefitted from—indeed, too many to be able to name here all at once. But I'd be remiss if I didn't thank the following people for their comments, suggestions, questions, and much needed support: the late Mabel & Sig Osmundsen, the late James Ricciutti, the late Shyam Patwardhan, Kim Whittington, Ryan & Kara Osmundsen, John & Kelly Kwak, Art & Patricia Hall, Paul Moser, Josh & Sarah Seachris, Bryan Reece, Rusty & Emily Jones, Michael Thune, Nathan L. King, Josh & Yetunde Fogarty, Edwin & Sandra Ware, Sue & Joe Wagner, Jay & Joan Eldridge, Michael Kutepow, and Josh & Mandy Orozco. My greatest debt of gratitude, however, undeniably goes to my wife, Julie, and our children, Emmafaith, Joel, and Avagrace. My children have been a constant source of joy and inspiration as a graduate student. They have kept me grounded and prevented me from spending too much time in philosophical ruminations. Julie's contributions as a wife and mother have not only made the completion of this dissertation possible, she created a home that allowed me to enjoy writing it. I have neither the words nor space to express truly how grateful I am to my wife, Julie. But I'll at least start by appealing to two metaphors: a tug boat and fruitful garden. Julie's patience, composure, strength, and hope provide our family with the kind of vessel that can safely take us through the stormy seas of life by steadily plowing through each crashing wave that comes our way. And Julie's ability to produce love, joy, peace, patience, kindness, gentleness, edifying conversations, amazing meals, warm, festive dinner-parties, wise instruction, skillful budgets, helpful itineraries, and enriching family activities is like having the most productive garden of life—it constantly yields the best produce! (I honestly think if Julie were writing this dissertation, it would have been finished years ago.) It's because of people like her that the cliché "It takes the right spouse to complete a dissertation" exists. Finally, my parents, Gary and Jackie Osmundsen, have constantly, graciously, and lovingly

supported my family and me in too many ways to list. I'm so grateful to see my children get to experience their love, sacrifice, and support, which has nourished me my entire life.

TABLE OF CONTENTS

CHAPTER 5. GRECO'S VIRTUE THEORETIC ACCOUNT OF KNOWLEDGE

CHAPTER 6. HOW EXTERNALISM CAN ACCOUNT FOR DEFEATING EVIDENCE FROM EPISTEMIC SUPERIOR DISAGREEMENT

CHAPTER 7. A GRECOIAN RESPONSE TO THE SKEPTICAL ARGUMENT FROM EPISTEMIC SUPERIORS (SAFES)

ABSTRACT

This dissertation is a response to the skeptical threats and challenges leveled by disagreement. Any plausible response to skepticism should explain what knowledge is and explain why the skeptic's assumptions about what's required for knowledge are false. In this dissertation I assume a virtue theoretic account of knowledge, which is a species of an externalist theory of knowledge. I defend this account of knowledge in the face of two problems I argue any externalist must address. The first problem is whether or not externalists can account for the kind of defeating evidence brought about by disagreement. Put differently: reliabilists need to explain how counterevidence from epistemic superior disagreement fits into a reliabilist scheme given what matters for reliabilists is that a true belief is sourced in a reliable process. The second problem is explaining what is wrong with what I take to be one of the strongest skeptical arguments from disagreement. The scope of this skeptical argument is modest in that one needs to recognize and be aware of an epistemic superior who disagrees with you over the veracity of some proposition. Nevertheless, it's a skeptical argument that avoids what many think are pre-theoretically irrelevant scenarios such as cases involving malevolent demons, brains in vats, Matrixian pods, etc. The pre-theoretical obviousness that one is not in a vat of chemicals doesn't generalize to the more modest claim that one needs to rule out an epistemic superior's belief in theory *t*, which entails the denial of one's belief that p.

Chapter One explicates how peer disagreement provides arguments for modest versions of skepticism. This chapter appeals to some of the leading proponents of this version of skepticism in order to draw together the common threads their skeptical arguments share. Once I identity the common threads shared by proponents of skepticism

from disagreement, I put forth the argument for skepticism from disagreement. Chapters Two and Three argue why the problem of moderate skepticism in Chapter One is a pseudo-problem by arguing why internalism is not the only game in town when it comes to making normative evaluations. Chapter Three continues to explain why the argument for moderate skepticism is a pseudo-problem by arguing how unlikely it is to satisfy the conditions needed for the argument to go through. Chapter Four propounds what I take to be the strongest skeptical argument from disagreement. This chapter shifts the problem of disagreement from internalist's concerns to externalist's ones. The result is a skeptical argument that avoids the problems plaguing the original argument presented in Chapter One. Chapter Five provides Greco's virtue theoretic account of knowledge. This chapter provides the theoretical resources needed to provide a Grecoian response to the challenges leveled against externalists in Chapter Four. Chapter Six explains the role defeaters play in epistemology and what a Grecoian account of defeater looks like. Once an account of defeating evidence is provided, an answer to the question "How can externalism provide an account of defeating evidence that comes from disagreement?" emerges. Chapter Seven provides a response to the new skeptical argument from disagreement laid out in Chapter Four. This chapter appeals to the theoretical resources provided in Chapter Five so as to complete this dissertation's purpose of providing a Grecoian response to skepticism from disagreement.

INTRODUCTION

Skepticism and Disagreement

Life is teeming with people who disagree with each other. We are all too familiar with this phenomenon. Indeed, living in the information age exposes us to so much disagreement that we sometimes walk away feeling intellectually dizzy. Recently, however, the epistemic significance of disagreement has become a very active topic in epistemology.[1] As a result, epistemologists have sought to answer a flurry of questions related to disagreement. Questions like Does rationality demand one or both parties in the disagreement to revise their belief? Does the disagreement provide defeating evidence for one or both disputants' beliefs? When is it epistemically appropriate to stick to one's guns and retain one's degree of confidence in a case of disagreement? Can both disputants rationally refrain from having to revise their confidence in their beliefs—that is, can both disputants stick to their guns in the face of peer disagreement? Can both disputants reasonably disagree and view the other as equally reasonable for sticking to one's guns? What general conclusions can we make from the epistemology of disagreement? Such questions are indeed deserving of our attention. But my aim in this dissertation is to provide a response to the skeptical threats leveled by disagreement. More exactly, this dissertation offers a Grecoian response to two problems disagreement levels against externalist theories of knowledge—specifically, a version of reliabilism.

My Grecoian response will assume the epistemological method employed by John Greco in his approach towards skepticism. There are three main theses that inform this

[1] See for example Audi (2011); Bergmann (2009); Bogardus (2009); Christensen (2007, 2009a, 2009b); Earl Conee (2010); Elga (2007); Elgin (2004); Feldman (2005, 2006, 2007); Fumerton (2010); Goldberg (2009, 2013a, 2013b); Goldman (2010, 2011); Greco (2009a); Kelly (2005b, 2010, 2013), King (2008, 2011, 2013); Lackey (2010a, 2010b, 2013); Matheson (2011); Riggs (2008); Thune (2009, 2010); van Inwagen (2010); Wedgewood (2010); Zagzebski (2007; 2012). Three recent edited volumes devoted to epistemic significance of disagreement are Feldman & Warfield (2010); Christensen & Lackey (2013); and Machuca (2013).

method. They are: (1) Skeptical arguments can be powerful in that it's not obvious why they're not cogent or unsound.[2] And what's powerful about skeptical arguments from disagreement is that they don't appeal to strange or outlandish scenarios involving malevolent demons, brain-in-a-vat scenarios, or Matrixian pods. Because they are based on the fact that real, live people disagree with us, one need not invoke such pre-theoretically implausible scenarios to get the argument off the ground.

(2) Analyzing skeptical arguments is useful because it allows valuable lessons to be discovered thereby advancing progress in philosophy.[3] Skeptical arguments pose strong theoretical problems that motivate us to reconsider our previous pre-theoretical assumptions about the nature of knowledge and evidence. It's true that many externalists don't take skepticism to provide an existential threat; we all seem to get through life just fine despite our awareness of them. Any existential crisis we might experience is short lived at best. But the value of skeptical arguments comes from learning the lessons they provide to our prior assumptions about the nature of knowledge and evidence. We do well, then, to take them serious and allow them to further refine our understanding of knowledge and evidence.

(3) One of the main attractions of externalism—specifically a version of reliabilism—is that it offers a very simple, straightforward response to skepticism.[4] Thus, I will be assuming an externalist approach in this dissertation. I will, of course, interact with and point out problems with internalism. But I don't intend to settle the internalism-externalism debate. As an externalist, I will be assuming that skepticism is false and that

---

[2] John Greco, *Putting Skeptics in Their Place* (Cambridge: Cambridge University Press, 2000), p. 2.

[3] Ibid., p. 3.

[4] Ibid., p. 4.

knowledge is something widespread. For I agree with James Van Cleeve who argues that knowledge is either relatively easy to acquire or virtually impossible.[5] Be that as it may: *I take the skeptical arguments from disagreement seriously because I think they demand a response and can enable externalists to have a better understanding of the nature of knowledge and evidence*. Thus, my Grecoian response will be working within a recent trend in epistemology which is to assume knowledge is widespread but that the epistemic work that needs to be done is to explain what knowledge is and how it's possible to acquire. This approach, then, is at odds with trying to prove one, in fact, has knowledge or trying to vindicate that one has knowledge by satisfying the skeptic's demands.

It's important to point out that my Grecoian response to skepticism employs a method that enjoys a rich tradition that traces back to Thomas Reid and G.E. Moore. This will help the reader know more about some of the things this dissertation will be already assuming. So, in light of the three theses mentioned above, I want to briefly offer three methodological principles informing this dissertation. The three methodological principles that Greco employs, which he credits to both Thomas Reid and G.E. Moore, are the following.[6]

First, Greco holds to a moderate version of foundationalism. This is roughly the idea that justification and knowledge have a certain hierarchical structure. At the base of this structure are sources from which beliefs derive that are sometimes non-inferential and

---

[5] See James Van Cleeve, "Is Knowledge Easy—Or Impossible? Externalism as the Only Alternative to Skepticism," in Stephen Luper, ed., *The Skeptics: Contemporary Essays* (London: Ashgate, 2003).

[6] See G.E. Moore, "Proof of an External World," *Proceedings of the British Academy* 25 (1939): 273-300; Thomas Reid, "Essays on the Intellectual Powers of Man," in Sir W. Hamilton, ed., *The Works of Thomas Reid* (Charlottesville, VA: Lincoln-Rembrandt Publishing, 1993); John Greco, "Reid's Critique of Berkeley and Hume: What's the Big Idea?," *Philosophy and Phenomenological Research* 55, 2 (1995): 279-296; Idem. "How to Reid Moore," *Philosophical Quarterly* 52, 209 (2002): 544-563.

are in no need of further justificatory support.[7] This version of foundationalism is moderate in that such sources are indeed fallible. Unlike classical foundationalism, which says a belief sourced in a basic non-inferential foundation must be infallible, indubitable, and incorrigible, moderate foundationalism says knowledge doesn't require infallibility. Moderate foundationalism, then, gives way to the following methodological principle

M1: One should not try to prove what is not known by proof.[8]

This methodological principle falls out straightforwardly from moderate foundationalism. For if a belief is not sourced in either deductive or inductive reasoning (which are both inferential, belief-forming processes), then providing a proof for a belief that is sourced in a non-inferential belief forming process is to embark upon a justificatory project that appeals to a different justificatory process from whence the targeted belief came. For example, we don't use memory and intuition to form perceptual beliefs. Nor do we use them to prove that perception is reliable. Similarly, we don't use tactile perception to form mathematical beliefs. Nor do we consult visual perception to form memorial beliefs. Thus, we need not justify our memorial beliefs by appealing to perceptual processes. The same goes for reason. Since reason isn't the lone source from which all other beliefs derive, we need not appeal to reason to support beliefs sourced in non-inferential belief forming processes. Appealing to a non-inferential belief's justificatory support by offering an argument in support of it is to engage confabulation (at worst) and offer an alternative, non-original justificatory ground (at best). The true merit behind this principle, then, is that one doesn't need to take the bait of having to provide a proof for a belief that wasn't

[7] Although the kinds of basic sources that make knowledge possible vary amongst foundationalists, some of the well-recognized ones are introspection, memory, perception, testimony, and reasoning—namely, inductive and deductive reasoning.

[8] John Greco, *Achieving Knowledge: A Virtue Theoretic Account of Epistemic Normativity* (Cambridge: Cambridge University Press, 2008), p. 183.

sourced in an inferential belief-forming process like an argument, which the skeptic demands.

Second, Greco endorses a Reidian methodological principle aimed specifically at skepticism. More specifically, Reid's approach to skepticism is that it serves as a heuristic device to better understand one's theory of certain epistemic goods. Indeed, according to Greco, "skeptical arguments are useful and important because they drive progress in philosophy."[9] For example, skepticism motivates one to both refine one's theories of evidence, justification, knowledge, etc., and offer a better understanding of one's account thereof. Skeptical arguments thus draw one's attention to mistaken pre-theoretical assumptions embedded in one's theories. So, rather than take a dismissive response to skepticism, Greco adopts the following Reidian methodological principle.

M2: Rather than trying to prove that external things exist, or that we know that external things exist, one should take a close look at the skeptic's reasons for saying that we do not know.[10]

Notice that this principle doesn't deny skepticism proffers interesting epistemological problems. Nor does it suppose that skeptical arguments cannot offer important epistemological lessons. Indeed, it's because some skeptical arguments are so prima facie compelling that they are deserving of our time and response. And Greco takes Reid and Moore to have done just that: namely, provide "devastating criticisms of [different skeptical arguments] and then show how each [argument] rests on either (a) an assumption that is implausible, (b) an assumption that is clearly false, or (c) a fallacy in reasoning."[11] The true merit of this principle, then, is that instead of trying to satisfy the inordinate demands made

[9] John Greco, *Putting Skeptics in Their Place*, p. 3.

[10] John Greco, *Achieving Knowledge*, p. 184.

[11] Ibid., p. 185.

by skeptics, one embarks upon a more realistic project that can produce much epistemic fruit.

Third, Greco's methodology aims at turning the tables on the skeptic and placing the burden of proof on them. For Greco takes certain sources of knowledge and common sense to have greater authority and reliability than philosophical reasoning. Greco quotes Reid as saying philosophical reasoning and reason, in general, "is like a telescope, which may help a man [or woman] see farther, who hath eyes; but, without eyes, a telescope shews nothing at all."[12] So, the third methodological principle is this.

M3: Common sense has defeasible authority over philosophical theory.[13]

This principle continues to mount evidence in its favor, in addition to the evidence the history of science bears out. For new scientific discoveries, through new innovation and technology, correct philosophical speculation and theories. Indeed, even hundreds of years ago Reid pointed out that the pattern of scientific discoveries "have always tended to refute, but not to confirm, the theories and hypotheses which ingenious men have invented."[14] And thus, by parallel reasoning, it's more reasonable to call in question the arguments put forth by skeptics than it is to call into question beliefs deriving from basic sources like perception and common sense. According to Greco, Reid and Moore have taught him that to think otherwise "is only philosophical arrogance, which is both misplaced and ill-advised."[15]

This dissertation, then, provides a Grecoian response to the skeptical threats leveled by disagreement. It will consist of four projects. The first project aims at entering

[12] Ibid., p. 179.

[13] Ibid., p. 185.

[14] Ibid., p. 186.

[15] Ibid., p. 187.

into the dialectic midstream and dissect the extant arguments for moderate skepticism from disagreement. The second project is to show where the arguments go wrong and why the problem is a pseudo-one. The third project is to offer a more formidable skeptical argument from disagreement that avoids the pseudo-problems plaguing the first. The last project is to see how the new argument from disagreement can allow this dissertation to make progress in epistemology. More specifically, this dissertation will provide a better understanding of a virtue theoretic account of knowledge and defeating evidence. For this dissertation answers two important questions the skeptical argument from disagreement demands an externalist to answer: namely, (1) How can externalism provide an account of defeating evidence that comes from disagreement? And (2) What is wrong with the skeptical argument from disagreement?

Chapter One explicates how peer disagreement provides arguments for modest versions of skepticism. More specifically, I provide and examine the conditions needed to obtain in order for one of the skeptical arguments from disagreement to go through. This chapter appeals to some of the leading proponents of this version of skepticism in order to draw together the common threads their skeptical arguments share. Once I identity the common threads shared by proponents of skepticism from disagreement, I put forth the argument for skepticism from disagreement.

Chapters Two and Three argue that the problem of moderate skepticism from disagreement is a pseudo-problem. Chapter Two shows why the argument for moderate skepticism rests on an implausible assumption: namely, internalism is not the only game in town when it comes to making normative evaluations. This chapter, then, looks at the differences between internalism and externalism and argues why the former is too restrictive to appreciate the inherent complexities in making normative evaluations. Thus,

epistemic progress is made by better understanding normative evaluations as a result of taking skepticism seriously.

Chapter Three continues to argue that the argument for moderate skepticism in Chapter One is a pseudo-problem. The strategy in this chapter is to show that the argument for skepticism demands certain conditions be satisfied. But I argue that such conditions are rarely ever satisfied. This chapter uses skepticism to make progress in better understanding the different accounts of evidence that are on offer and why we shouldn't treat cognitive agents the same way we treat inanimate objects in cases of disagreement.

Chapter Four propounds what I take to be the strongest skeptical argument from disagreement. I take the problem of disagreement to be a problem leveled against knowledge rather than rationality. Shifting the problem of disagreement from internalist's concerns to externalist's ones results in a skeptical threat that avoids the problems plaguing the original argument presented in Chapter One. After this argument is laid out, I specify the kinds of questions an externalist must answer as a result of it.

Chapter Five provides Greco's virtue theoretic account of knowledge. This chapter provides the theoretical resources needed to provide a Grecoian response to the challenges leveled against externalists in Chapter Four. Along the way, this chapter also provides plausible answers to some of the perennial problems in epistemology.

Chapter Six provides an answer to the challenges leveled by the skeptical argument from disagreement. More specifically, this chapter explains the role defeaters play in epistemology and what a Grecoian account of defeater looks like. Once an account of defeating evidence is provided, an answer to the question "How can externalism provide an account of defeating evidence that comes from disagreement?" emerges.

Chapter Seven provides a response to the new skeptical argument from disagreement laid out in Chapter Four. This chapter appeals to the theoretical resources provided in Chapter Five so as to complete this dissertation's purpose of providing a Grecoian response to skepticism from disagreement. This chapter uses Greco's account of testimonial knowledge and understanding to challenge some of the key assumptions embedded in the premises of the skeptical argument from disagreement, in addition to the fallacious reasoning in support of the argument. Again, we see how skepticism allows progress to be made in epistemology by forcing one to better refine and understand their theory of knowledge, evidence, and other epistemic goods.

1

The Skeptical Threats From Peer Disagreement

Whenever two persons make opposite judgments about the same thing, it is certain that a least one of them is mistaken, and neither, it seems, has knowledge. For if the reasoning of one of them were certain and evident, he would be able to lay it before the other in such a way as eventually to convince his intellect as well.

~RENÉ DESCARTES

1. SETTING UP THE PROBLEMS OF PEER DISAGREEMENT

Disagreement between two or more people happens quite often. Just listen long enough to others talk and you'll soon encounter some kind of disagreement. Questions ranging from "Which team will win the World Series?" to "Is determinism compatible with freedom?" invite many contrary answers in response. Life is teeming with people who disagree. Economists, scientists, politicians, philosophers, you name it: listen to their discussions long enough and you'll learn disagreement abounds. This is all too familiar. Recently, however, philosophers have embarked upon inquiring into the epistemic significance of disagreement.

But what, exactly, is motivating this investigation into the epistemology of disagreement? Philosophers are not interested in cases of disagreement that can be easily explained. More specifically, philosophers are not interested in cases involving salient epistemic asymmetries between the disagreeing parties. In other words, having an *epistemic advantage* over one's disputant is not what's motivating the debate in the epistemology of disagreement. For example, if two people disagree about the sum of ten, three-digit numbers, and one uses a calculator while the other does the sum in her head, then explaining the disagreement is fairly simple: one person used a calculator, which is a more *reliable process* for adding numbers; whereas the other person used a less reliable process for

doing a large sum. The use of a calculator provides, then, the kind of asymmetry that helps explain the disagreement by pointing to the fact that one is epistemically better situated than the other. So philosophers are not motivated by cases of disagreement that can easily be explained. Rather, philosophers are interested in cases of disagreement where there are no salient asymmetries to explain the disagreement. Such asymmetries have been called 'symmetry breakers'[16]. A symmetry breaker is that which explains how one person is epistemically privileged or superior to the other disputant. Symmetry breakers can be tokens of reliable processes, superior evidence, superior cognitive abilities, superior epistemic virtues, superior cognitive performances, evidence of the disputant's inferiority, etc.

What motivates the epistemology of disagreement, then, is cases involving people who are both *cognitive* and *evidential* equals, and lack any relevant symmetry breaker to explain their disagreement. It is this kind of disagreement which has sparked a flame of questions ranging from Does rationality demand one or both parties in the disagreement to revise their belief? Does the disagreement provide defeating evidence for one or both disputants' beliefs? When is it epistemically appropriate to stick to one's guns and retain one's degree of confidence in a case of disagreement? Can both disputants rationally refrain from having to revise their confidence in their beliefs—that is, can both disputants stick to their guns in the face of peer disagreement? Can both disputants reasonably disagree and view the other as equally reasonable for sticking to one's guns? What general conclusions can we make from the epistemology of disagreement?

[16] I borrow this term from Jennifer Lackey who borrowed it from Nathan Christiansen. See Jennifer Lackey, "A Justificationist View of Disagreement's Epistemic Significance," in Adrian Haddock, Alan Millar, and Duncan Pritchard, eds., *Social Epistemology* (Oxford: Oxford University Press, 2010), p. 309.

Answers to all these questions go beyond the scope and purpose of this chapter. The purpose of this chapter is to get clear how peer disagreement provides arguments for modest versions of skepticism. In the next section I will explicate the conditions needed in order for one of the skeptical arguments from disagreement to go through. I will appeal to some of the leading proponents of this version of skepticism in order to draw together the common threads their skeptical arguments share. These proponents represent positions in the literature deemed conciliationism, conformism, or the equal weight view. For the sake of brevity, I will hereafter refer to these views as CCE. Once I identity the common threads shared by CCE, I will put forth the argument for skepticism from disagreement that falls out of the CCE view on disagreement.

2. AN OVERVIEW OF CCE

2.1 What Is An Epistemic Peer?

There are a number of responses to the problem of peer disagreement. You can think of these responses as falling somewhere on a spectrum. At one end, there are those who think peer disagreement is not epistemically significant. This view says someone who believes p need not undergo any doxastic revision when faced with a peer who believes not-p. One can retain the same level of confidence in p, despite lacking any independent reason for thinking one's belief is epistemically better, and still be rational in doing so.[17] In the middle of the spectrum, there are those who think *some* cases of peer disagreement rationally require some moderate level of doxastic revision. This cluster of views takes peer disagreement to be epistemically relevant, but doesn't think a significant degree of doxastic

[17] This position has been deemed the 'steadfast view,' the 'non-conformist view,' and the 'no-defeater view.' See David Christensen (2009b); Jennifer Lackey (2010a), and Michael Thune (2010) respectively. Proponents of this view are Gideon Rosen, Thomas Kelly, Ralph Wedgwood, Peter van Inwagen, and Alvin Plantinga.

revision is required (e.g., adopting the other person's belief or withholding belief, etc.) given there are other epistemically relevant factors mitigating the force of peer disagreement.[18] In this dissertation, however, I will be explicitly dealing with those epistemologists who take peer disagreement to require significant doxastic revision. This position, which occupies the logical space at the other end of the spectrum, I have deemed above CCE.[19] CCE says one should compromise with a disputant thereby rationally requiring significant doxastic revision. Rationality demands that all disputants cannot retain the same level of confidence in their belief. Indeed, in Richard Feldman's final analysis, he thinks, "one should give up one's beliefs in light of the sort of disagreement under discussion" which implies a version of skepticism in "that [CCE] denies the existence of reasonable beliefs in a significant range of cases."[20]

CCE gets motivated by considering cases where there is no relevant symmetry breaker that explains the disagreement at hand. So consider a case involving two pocket watches that display two different times. If you believe that both watches are equally reliable and you lack any reason for thinking one to be more accurate than the other, then epistemic symmetry obtains. This means you lack any justificatory basis for believing one watch to be more accurate than the other. Thus, on pain of irrationality, you should significantly adjust your doxastic attitude. Similarly, if you believe your interlocutor is as reliable as you at getting the truth and you lack a symmetry breaker—which gives you a reason to favor your belief—then epistemic symmetry obtains. Accordingly, you are

---

[18] This position has been deemed the 'dynamic view' and is represented by epistemologists such as Jennifer Lackey, Michael Bergmann, and Michael Thune. See Michael Thune (2010).

[19] I deem this position CCE to abbreviate the three names it goes by—conciliationism, conformism, and extra-weight view.

[20] Richard Feldman, "Reasonable Religious Disagreements," in Louis Anthony, ed., *Philosophers Without God* (Oxford: Oxford University Press, 2007), p. 213.

rationally required to compromise and considerably adjust your doxastic attitude by, say, withholding judgment.

But what, exactly, does it mean to say epistemic symmetry obtains? We see how symmetry between two inanimate objects can easily obtain relative to two peoples' epistemic situation regarding information about, say, some inanimate object. For example, there can be parity between two peoples' epistemic situatedness as a result of both parties being privy to the same information. Each party knows that both thermometers were made by the same manufacturer; both were manufactured the same year; both appear undamaged; both have reliable track records; both are placed in the same environment, which is where we want to know the temperature, etc. But what does it mean to say epistemic symmetry obtains between two sentient things like human persons?

To answer this question we must unpack the conditions of *epistemic peerhood* since the kind of epistemic symmetry epistemologists are interested in involves two or more people who are epistemic peers. For 'epistemic peer' is a technical term. It has a stipulative definition in the epistemology of disagreement literature that picks out certain characteristics which are shared between the disagreeing parties. Thomas Kelly describes epistemic peerhood this way:

> Let us say that two individuals are epistemic peers with respect to some question if and only if they satisfy the following two conditions: (i) they are equals with respect to their familiarity with the evidence and arguments which bear on the question, and (ii) they are equals with respect to general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias.[21]

In much the same vein, David Christensen describes epistemic peers as two people who both have "good reason to believe that the other person is one's (at least approximate)

[21] Thomas Kelly, "The Epistemic Significance of Disagreement" in John Hawthorne and Tamar Gendler Szabo (eds.) *Oxford Studies in Epistemology* (Oxford: Oxford University Press, 2005), p.174-175.

equal in terms of exposure to the evidence, intelligence, freedom from bias, etc."[22] Similarly, Catherine Elgin says the two disagreeing parties are epistemic peers when two people "have the same evidence, reasoning abilities, training, and background assumptions."[23] Alvin Goldman says much the same thing when he says epistemic peers are "roughly equal with respect to intelligence, reasoning powers, background information, and so on."[24]

Epistemic peers, then, roughly[25] have the following three characteristics: namely, (i) evidential equality; (ii) cognitive equality; and they have engaged in (iii) full disclosure.[26] These three characteristics are defined as follows. [27]

> (i) *Evidential equality*: A and B are evidential equals relative to the question whether *p* when A and B are equally familiar with the evidence and arguments that bear on the question whether *p*.

---

[22] David Christensen, "Disagreement as Evidence: The Epistemology of Controversy," in *Philosophy Compass* 4/5: 756-7.

[23] Catherine Z. Elgin, "Persistent Disagreement," in Richard Feldman and Ted Warfield, eds., *Disagreement* (Oxford: Oxford University Press, 2010), p.53.

[24] Alvin Goldman, "Epistemic Relativism and Reasonable Disagreement," in Richard Feldman and Ted Warfield, eds., *Disagreement* (Oxford: Oxford University Press, 2010), p.189.

[25] Adam Elga thinks of epistemic peers differently. He thinks epistemic peerhood should be cashed out in terms expected reliability. If two people think of themselves as equally reliable at getting the truth, then both should view the other as an epistemic peer. More specifically, Elga says: "Upon finding out that an advisor disagrees, your probability that you are right should equal your prior conditional probability that you would be right. Prior to what? Prior to your thinking through the disputed issue, and finding out what the advisor thinks of it. Conditional on what? On whatever you have learned about the circumstances of the disagreement." In other words, Elga's way of establishing epistemic peerhood is to think of it in terms of expected reliability. If one thinks the other party is equally reliable at getting the truth, then such a person qualifies as an epistemic peer. Since I have no antecedent, independent reason for thinking I am more likely to be right than my interlocutor—prior to learning that we disagree—I should view that person as my epistemic peer. I don't address Elga's view specifically; I only wish to establish a rough and ready view of what I mean when I use the term epistemic peer. The quotation from Elga is taken from Adam Elga, "Reflection and Disagreement," *Nous* 41 (2007): 490

[26] One might ask Why is full disclosure constitutive of epistemic peerhood? I include it because from a first-person perspective, two or more peers need to share with each other their reasons for and against the targeted proposition so that each party is confident that the other is not ignorant of any relevant information. It's not enough that both parties function equivalently in some domain of inquiry—the parties need to be exposed to all the relevant arguments, reasons, and evidence in order for epistemic peerhood to obtain.

[27] I borrow these three definitions from Jennifer Lackey in "A Justificationists View of Disagreement's Epistemic Significance," p. 302-303.

> (ii) *Cognitive equality*: A and B are cognitive equals relative to the question whether *p* when A and B are equally competent, intelligent, and fair-minded in their assessment of the evidence and arguments that bear on the question whether *p*.
>
> (iii) *Full disclosure*: A and B are in a situation of full disclosure relative to the question whether *p* when A and B have knowingly shared with one another all of their relevant evidence and arguments that bear on the question whether *p*.

The idea, then, behind epistemic symmetry obtaining in a case of peer disagreement can now be explained straightforwardly. Once both parties agree that they are both cognitive and evidential equals, and both have engaged in full disclosure, then they are *epistemic peers*. But *epistemic symmetry* hasn't obtained yet until both parties agree that no symmetry breaker is available. Let's call this (iv) condition the *acknowledgement condition.*[28]

> *(iv) Acknowledgement condition*: A and B must know or justifiedly believe that the previous three conditions are met.[29]

So epistemic symmetry obtains just in case both disagreeing parties acknowledge each other as evidential and cognitive equals; both have engaged in full disclosure with the result that no one is able to avail him- or herself of a relevant symmetry breaker; and both parties justifiedly believe the previous conditions are satisfied.

To be more concrete, let's consider a case that is used to motivate CCE. David Christiansen is a leading proponent of CCE. He thinks one of the best ways to understand the motivation behind CCE is by considering the following case:

---

[28] I'm indebted to Nathan L. King for pointing out this fourth condition in "Disagreement: What's the Problem? *or* A Good Peer is Hard to Find" in *Philosophy and Phenomenological Research* (2011): 1-24.

[29] Nathan L. King, "Disagreement: The Skeptical Arguments from Peerhood and Symmetry," in Diego E. Manchuca, ed., *Disagreement and Skepticism* (New York: Routledge, 2013), p. 200.

> DINNER BILL: 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. How should I react, upon learning of her belief?[30]

This case is supposed to surface the intuition that if both disputants think of each other as epistemic peers and that epistemic symmetry has obtained then, on pain of irrationality, significant doxastic revision is required. The epistemic principle proponents of CCE believe falls out of this epistemic situatedness goes something like this:

Rationality Maintenance (RM): If S believes p, and S* believes not-p, and S acknowledges S* is an epistemic peer, and S lacks a symmetry breaker, then on pain of epistemic irrationality, S should suspend judgment towards p.

Because both parties lack any symmetry breaker—other than the disagreement itself—in thinking one is more reliable at getting the truth than the other, it is question-begging, and thus, irrational to not revise one's doxastic attitude.[31] Virtually all philosophers agree that in

---

[30] David Christiansen, "Epistemology of Disagreement: the Good News," *The Philosophic Review* 116 (2007): 193.

[31] What's more, it's not only irrational, which is already damnable, but Linda Zagzebski says such responses are guilty of epistemic egoism. She has identified three kinds of epistemic egoism: extreme, strong, and weak epistemic egoism: (1) *Extreme epistemic egoism* "maintains that the fact someone else has a belief is never a reason for her to believe it, not even when conjoined with evidence that the other person is reliable. If she finds out that someone else believes *p*, she will demand proof of *p* that she can determine by the use of her own faculties, given her own previous beliefs, but she will never believe anything on testimony." (2) *Strong epistemic egoism* "maintains that she has no obligation to count what another person believes as relevant to her own beliefs, but she may do so if she sees that given what she believes about them, they are likely to serve her desire for the truth, that is, she sees that they are reliable." The difference between extreme and strong forms of epistemic egoism is that the latter is permitted—not obligated—to believe *p* if another person believes *p*; whereas the former says that one is not even permitted to believe *p*—that is, one should never believe *p* in virtue of another person believing *p*. (3) *Weak epistemic egoism* is the attitude that says if I have "evidence that someone else's beliefs reliably serve [my] desire for truth in some domain, [then] not only [am I] rationally permitted, [I rationally should] take their beliefs [as a reason to believe what the other person believes.]" *Epistemic egoism*, then, is contrasted with those epistemic attitudes that entail a presumption of trust in other people's beliefs and cognitive faculties. Such epistemic attitudes have been listed under the umbrella term *epistemic universalism* which says, because S believes *p,* you always have a reason to believe *p,* despite not having

cases like DINNER BILL there should be significant doxastic revision—specifically, withholding judgment until an explanation is discovered. The question that is disputed amongst epistemologists is how far such cases can generalize.

In §§ 2.2 - 2.3, I will look at specific proponents of CCE and lay out the lineaments of their arguments for moderate skepticism via peer disagreement. This will help us see how far proponents of CCE think the skeptical attitude prescribed by DIINNER BILL generalizes to other cases involving disagreement. In §§ 2.4 - 2.5, I will look at two other proponents of CCE who arrive at a similar skeptical conclusion, but take a different route to get there.

2.2 CCE *A La* Feldman

Richard Feldman defends the view that in cases of disagreement that involve epistemic peers who engage in full disclosure, wherein epistemic symmetry obtains, there can be no reasonable disagreement. More specifically, Feldman denies both that (i) epistemic peers can reasonably disagree and (ii) that one epistemic peer can reasonably maintain her own belief, while simultaneously thinking the other disagreeing peer is reasonable in retaining his or her belief after both parties have engaged in full disclosure and epistemic symmetry has obtained.[32] The denial of reasonable disagreement gets motivated by Feldman's endorsement of the Uniqueness Thesis.[33]

---

any evidence in support of S's reliability. In addition, Zagzebski makes a distinction between (5) *strong* and (6) *weak epistemic universalism*. The former says that, even though *p* is defeasible, it is reasonable to believe *p,* given that S believes it. The latter says that, even though *p* is defeasible, you have a reason to believe *p* in virtue of S believing *p*. See Linda Zagzebski, "Ethical And Epistemic Egoism And The Ideal Of Autonomy," *Episteme: A Journal of Social Epistemology* 4:3 (2007), pp. 252-255.

[32] Richard Feldman, "Reasonable Religious Disagreement," p.201.

[33]This kind of principle is endorsed by Feldman (2007); Christensen (2007); and Matheson (2011). Roger White defends the Uniqueness Thesis, too. The specific thesis that he defends reads like this: Uniqueness: If

Uniqueness Thesis (UT): A body of evidence E bearing on proposition p justifies (at most) one proposition out of set of competing candidates or (at most) one doxastic attitude towards p.

A case Feldman uses to capture the intuition behind UT involves two detectives trying to adjudicate whether evidence E incriminates suspects Lefty or Righty. There is strong evidence incriminating Lefty, but there is also strong evidence incriminating Righty. (The evidence rules out that two people committed the crime; so, Lefty and Righty could not have committed the crime together.) It's irrational for either detective to believe in either Lefty's or Righty's guilt because the evidence for Lefty is evidence against Righty and vice versa. So if one detective believes Lefty is guilty, then the inference Righty is not guilty is valid. But Feldman is quick to note that this inference prevents one from also believing that it's reasonable for one's peer to conclude that Righty is guilty, too, because "this combination of beliefs simply does not make sense. […] The only reasonable option is to suspend judgment."[34] In cases involving disputants who are cognitive equals and share the same evidence, UT says there cannot be reasonable disagreement. That is, neither party can rationally retain their belief so they should withhold judgment. So, according to Feldman, if

---

an agent whose total evidence is E is fully rational in taking doxastic attitude D to P, then necessarily, any subject with total evidence E who takes a different attitude to P is less than fully rational. See Roger White, "Evidence Cannot Be Permissive" in Matthias Steup, John Turri, and Ernest Sosa, eds., *Contemporary Debates in Epistemology* Second Edition (Oxford: Blackwell, 2013). It's important to notice how White's version shares a key point with Feldman's: namely, the principle restricts only one doxastic attitude deemed rational given any set of evidence. So, like Feldman, he denies that epistemic peers can reasonably disagree because there is only one rational doxastic attitude towards any total set of evidence E that's shared between two or more peers. Uniqueness is a very strong thesis in that it says given evidence E, there is only one fully rational doxastic attitude towards E that applies to everyone.

[34] Richard Feldman, "Reasonable Religious Disagreements," in Louis M. Antony, ed., *Philosophers Without Gods: Meditations On Atheism And The Secular Life* (Oxford: Oxford University Press, 2007), p. 205. It seems this kind of case can elicit intuitions that are capturing different values (e.g., moral, epistemic, or prudential) than the epistemic one Feldman is assuming. Feldman thinks given such evidence one should suspend judgment in accordance with the epistemic norm that's concerned with forming beliefs that are aimed at getting the truth and avoiding falsehoods. But given the case involves a verdict that may or may not incriminate an innocent person, the stakes are high and other competing norms may be operative in the evaluation of the case. E.g., the moral norm of not convicting the innocent or convicting-verdicts requiring evidence that blocks reasonable doubt may demand an evidential standard that is too strong to govern epistemic norms but appropriate for governing moral norms.

S and S* are both epistemic peers relative to domain D and proposition p, then both S and S* should share the same doxastic attitude towards p. And if both parties disagree, then, on pain of irrationality, they must suspend judgment.[35]

The reasoning behind Feldman's view on peer disagreement takes the form of a dilemma: either epistemic symmetry between epistemic peers obtains or it does not. If it does, then neither party can reasonably disagree nor can they reasonably maintain their belief while also thinking the other party is reasonable in retaining his or her own belief. If epistemic symmetry doesn't obtain, then the process of fully disclosing each other's evidence will vindicate which party's belief is better supported by the relevant evidence. So, either the process of fully disclosing each other's evidence will vindicate one party's belief, or neither party can reasonably disagree nor can they reasonably retain their belief while also thinking the other party is reasonable in retaining his or her own belief.[36]

What makes Feldman's view so friendly to skepticism is his refusal to grant disputants certain symmetry breakers many think can explain the disagreement. For example, perhaps one disputant appeals to the different starting points that are operative in each disputant's assessment of the relevant evidence. Because one disputant believes a certain fundamental, metaphysical claim or an epistemological principle that the other disputant denies, the way the evidence gets evaluated, weighed, and processed can kick out different conclusions. Feldman doesn't accept this explanation because he believes such fundamental starting points can themselves be the targeted proposition both disputants can disclose evidence for and against. The point of bringing up different starting points or

---

[35] Richard Feldman, "Epistemological Puzzles About Disagreement," in Stephen Hetherington, ed., *Epistemology Futures* (Oxford: Oxford University Press, 2006), p. 235.

[36] I am indebted to Micheal Thune for pointing out how Feldman's view is supported by a constructive dilemma. See Michael Thune, *The Epistemology of Disagreement*, unpublished doctoral dissertation thesis, Purdue University, West Lafayette, IN (2008, August), p. 30.

background beliefs just pushes the problem one step back. And, for Feldman, there are no assumptions, beliefs, principles, methods, or worldview left unchallenged. So, in order to reasonably retain your belief in the face of peer disagreement, you better be able to convince your peer that your symmetry breaker not only explains the disagreement, but that it's true. Appealing to pragmatic reasons why you have a certain starting point that explains the disagreement is a woefully inadequate response.

Nor does Feldman accept appeals to private evidence or evidence that can't be disclosed as legitimate candidates for symmetry breakers. For example, appeals to intuitions that p is true, or p's truth status is seemingly obvious, or p's evidence comes from private—non-disclosable—perceptual experiences cannot function as legitimate symmetry breakers. The reason is that, for Feldman, "each [symmetry breaker a disputant] can say in support of his view, the other [disputant] can say something analogous in support of the other view. To stick to one's guns in such a situation is to fail to treat like cases alike."[37] Indeed, Feldman says that retaining belief in such circumstances "is a violation of what I take to be a clear condition of rational belief."[38]

So, according to Feldman, it would be unreasonable and vulnerable to damning bootstrapping objections to think your peer is the inferior one or the one who has the epistemic defect. What's more, appeals to private insights or intuitions are *neutralized* by the disputant having her *own* insight or intuition. Such evidence shouldn't be conveniently weighed in favor of your belief thereby demoting the significance of your peer's piece of evidence. According to Feldman, unless you have reason to think your disputant is the one with the epistemic problem then your reasoning is akin to this: "You have an insight

[37] Richard Feldman, "Epistemological Puzzles about Disagreement," p. 116.

[38] Ibid.

according to which *P* is not true. I have one according to which *P* is true. It's reasonable for me to believe *P* in light of all this because, gosh darn it, *my* insight supports *P*."[39] This, of course, reeks of the bootstrapping objections and epistemic egoism discussed above. Therefore, appealing to non-disclosable evidence can't be a legitimate symmetry breaker because the extant, disclosable evidence shared between the disputants is:

> very similar, and each has evidence about what the other's private evidence supports. […] If I think you do have good evidence for your view, then I admit that there is good evidence for your view, and thus my own beliefs must take this into account. I need a reason to think that you, not me, are making a mistake. The unshared evidence does not help.[40]

It's important to see, then, how Feldman takes the lessons learned from DINNER BILL to generalize far and wide to cases involving disagreement over claims involving philosophy, religion, politics, ethics, economics, etc. It's not that legitimate symmetry breakers are, in principle, not on offer. Clearly, once a calculator is pulled out and all parties agree that the targeted math problem was correctly entered into the calculator a symmetry breaker will obtain. It's just the difficulty in securing one is hard given the work Feldman's Uniqueness Thesis is doing in conjunction with how easily one's private evidence or varying starting points get neutralized as a result of weighing their epistemic import equally.

2.3 CCE *A La* Christensen

Like Feldman, Christensen argues that this account of peer disagreement has the resources to mount a moderate version of skepticism towards many of our significant beliefs. Furthermore, Christensen frames the problem of peer disagreement in terms of a dilemma, as well: either *epistemic symmetry* obtains or it doesn't. If it doesn't, then one will have a

[39] Richard Feldman, "Reasonable Religious Disagreement," p.208.

[40] Ibid., p. 208.

symmetry breaker to explain why one disputant's belief over the other disputant's belief is more reasonable or justified. For example, your symmetry breaker may be that you know that I am not as reliable as you in domain D, relative to proposition p, thereby explaining why it's more reasonable for you to attribute error to me than to yourself. If symmetry does obtain, then neither party can reasonably disagree nor can they reasonably maintain their credence level in the relevant belief while also thinking the other party is reasonable in retaining his or her credence level in the targeted belief.

Christensen employs a different epistemological method than Feldman when dealing with cases of disagreement. More exactly, Christensen employs the following two principles, which he thinks should inform the way one assesses, reacts, and explains peer disagreement:[41]

> (1) I should assess explanations for the disagreement in a way that's independent of my reasoning on the matter under dispute.
>
> (2) To the extent that this sort of assessment provides reason for me to think that the explanation in terms of my own error is as good as that in terms of my friend's error, I should move my belief towards my friend's.

It's important to notice how these principles ratchet up the difficulty in providing a legitimate symmetry breaker. Otherwise known as the epistemic principle of independence, the idea behind principle (1) can be stated more formally as follows:

Independence (I): Any symmetry breaker offered to explain why one disputant is epistemically privileged or advantaged over another must do so without employing the very reasons, beliefs, or arguments that originally formed one's disputed belief.[42]

---

[41] Ibid., p.199.

[42] This type of principle is endorsed by Christensen (2007); Elga (2007); and Kornblith (2010).

This principle aims to block off any kind of bootstrapping strategy. The idea is that if I believe you are a cognitive and evidential equal, then I shouldn't appeal to the *token* evaluative performance I employed to bring about my belief. I should, rather, consider the likelihood of my disputant's ability to get the truth of p prior to discovering whether or not we disagree. For if I appeal to my reasons, arguments, and evidence comprising my performance in forming my original belief, then I can easily avail myself to symmetry breakers that allow me to bootstrap my way to demoting you as a peer. I can simply say your performance is inferior to mine or find some fault in your performance.

Independence prevents one from availing him or herself of dogmatically thinking one's token performance is either superior or the disputant's token performance is inferior. More specifically, Independence says I should decide whether or not you are an epistemic peer prior to discovering whether or not we disagree and leave each other's token performances out of it. I should ask myself if there is any reason for thinking you are more prone to err than me relative to the veritic status of the targeted proposition. If I think your epistemic performances are (in general) just as reliable as mine, then I should *not* appeal to the token performance at work forming your belief in order to ferret out a symmetry breaker. If there is a symmetry breaker, it should be appealed to *prior* to each disputants' performances that result in a discovered disagreement. But if I lack any symmetry breakers prior to discovering our disagreement, then I can't try to find one based on your token performance. Nor can I appeal to my token performance as superior since this would be a case of dogmatic reasoning. For either of these responses are vulnerable to bootstrapping objections.[43]

[43]Adam Elga has a similar view and bootstrapping argument in support of a version of CCE. For when you are aware that your epistemic peer disagrees with you, you should view her as just as likely to get things right as your view of yourself. He says, "When you learn of your [epistemic peer's] opposing judgment, you should

Principe (2) is a prescription for how to respond to a case of disagreement with one whom I take to be an epistemic peer prior to the disagreement, according to the conditions behind principle (1) (i.e., independence principle). Thus, unless I have reasons to demote you from the status of epistemic peer prior to discovering our disagreement, then upon learning about our disagreement I should considerably revise my doxastic state.

These principles show that Christensen differs with Feldman about the supervenience base of epistemic symmetry given epistemic peerhood obtains *prior* to the disagreement. But, like Feldman, Christensen also supports the second horn of the dilemma with his own version of the uniqueness thesis: namely, Rational Uniqueness (hereafter RU). RU, like UT, says "there is a unique maximally rational response to a given evidential situation."[44] What's more, both Christensen and Feldman take peer disagreement to provide a subjective rationality defeater[45] to one's belief that requires epistemic peers to *compromise* with each other and withhold judgment or, according to Christensen, split the difference. Despite their methodological differences, both Christensen and Feldman agree that disagreement requires significant doxastic revision and the upshot of DINNER BILL

---

think that the two of you are equally likely to be correct. […] If it were reasonable for you to give your own evaluation extra weight—if it were reasonable to be more than 50% confident that you are right—then you would have gotten some evidence that you are a better evaluator than your friend. But that is absurd." See Adam Elga, "Reflection and Disagreement," *Noûs* 41/3 (2007): 478-502.

[44] Ibid., p.190, n., 5.

[45] I will say more about the different kinds of defeaters there are in Chapter Six. But it will suffice to say, for now, that a defeater is a belief for thinking another belief is false or unjustified. The term 'defeater' admits many different meanings amongst epistemologists. It was first put forth by John Pollock (1986) whose seminal work on defeaters distinguished two kinds: undercutting defeaters and rebutting defeaters. Rebutting defeaters give you a reason or a new belief or experience for thinking your belief is false. Undercutting defeaters, however, defeat beliefs in a different way: namely, by indicating a belief's source is unreliable thereby undercutting the belief's justification. In general, a defeater is a reason to no longer hold a belief. More exactly, a defeater is a doubt, belief, or experience that indicates one's belief is either false, unreliably formed, or irresponsibly retained. Defeaters, then, can defeat a belief's justification in different ways. NB: Goldberg (below in § 2.5) thinks systematic disagreement levels a normative defeater against one's belief, thereby preventing it from being either justified or an instance of knowledge. A normative defeater is a doubt or belief that S *ought* to have, which indicates that one's belief is false or unreliably formed. One may not be aware of the extant defeater, which is why it's called a normative defeater. Once one is made aware of the normative defeater, it then becomes a psychological defeater of one stripe or another.

generalizes far and wide. Indeed, when it comes to some of our most cherished beliefs in the domain of philosophy, religion, morality, and politics, both prescribe moderate skepticism because disagreement is so pervasive and the work their epistemic principles do make it difficult to discover a legitimate symmetry.

In the next two sections, I will look at two alternate routes to skepticism from disagreement. But instead of appealing to epistemic principles that provide a rationality defeater to one's belief vis-à-vis disagreement, these arguments see disagreement as providing an undercutting defeater. This moves the epistemic significance of disagreement from the realm of rationality to the realm of knowledge since systematic disagreement is taken to be an indication that humans are not cognitively equipped at reliably getting the truth in domains wherein disagreement abounds.

2.4 CCE *A La* Kornblith

Disagreement for Hilary Kornblith is indirect evidence for the truth of some proposition p. In the same vein, Feldman makes the point this way: "evidence that there is evidence for P [via an epistemic peer believing it] is evidence for P."[46] Whether one knows the other disputant is just as reliable as one's self at getting the truth or is just as reliable at assessing the disclosed evidence, Kornblith believes epistemic symmetry can obtain either way. So, along with Feldman and Christensen, Kornblith argues that disagreement requires significant doxastic revision. More specifically, Kornblith thinks disagreement with an

[46]Richard Feldman, "Epistemological Puzzles about Disagreement," p. 223.

epistemic peer provides one with an undercutting defeater against one's belief, which requires one to suspend judgment lest one be guilty of epistemic irresponsibility.[47]

Kornblith thinks epistemic symmetry can obtain both prior to the disagreement like Christensen and after both parties have participated in full disclosure like Feldman.[48] But what Kornblith adds to the CCE camp is the epistemic significance of consensus opinion or lack thereof. He thinks consensus opinion about some claim p within the mathematical and scientific community is a reliable indication of the truth of p. But the philosophical community lacks the kind of historical progress mathematical and scientific communities have achieved. Indeed, Kornblith says, and I quote a length:

> [The] sad truth, it seems, is that the history of philosophy does not look remotely like the history of science or mathematics when it comes to the dynamics of consensus among its most esteemed practitioners, and this has a striking bearing on the question of its epistemic credentials. […] The field of philosophy in general […] simply does not have anything like the epistemic standing of the empirical sciences. […] So, much as we all find ourselves forming beliefs about disputed philosophical questions when we immerse ourselves in the arguments, we should acknowledge in quiet moments of reflection that *these views we form are ones that are not epistemically justified*. It would be presumptuous to claim that we are justified [in our philosophical beliefs] as it would in Christiansen's [DINNER BILL] case when we find that our mathematically reliable dinner companion has reached a different conclusion about the division of the check. […] I am forced to conclude, very reluctantly, that the opinions I hold on most philosophical matters—and I have a great many of them—are not epistemically justified.[49]

---

[47] Undercutting defeaters defeat beliefs by indicating a belief's source is unreliable thereby undercutting the belief's justification. Undercutting defeaters provide good reasons for thinking the source of your belief is unreliable, but they—unlike rebutting defeaters—don't necessarily make your belief false.

[48] I don't mean to suggest that Christensen denies epistemic symmetry can obtain after full disclosure. I only intend to contrast Feldman and Christensen this way to highlight their unique contributions to CCE: namely, Feldman's Uniqueness Thesis and how epistemic symmetry supervenes on disagreement persisting after epistemic peers have engaged in full disclosure and, regarding Christensen, how epistemic symmetry can obtain prior to full disclosure given the epistemic principles—like independence—he endorses.

[49] Hilary Kornblith, "Belief In The Face Controversy," in Richard Feldman and Ted Warfield, eds., *Disagreement* (Oxford: Oxford University Press, 2010), pp.45-46 (emphasis mine).

Thus, according to Kornblith, the epistemic upshot of the DINNER BILL case is that it *does* generalize far and wide—especially to the field of philosophy. After all, in his final analysis, he thinks "given the current state of the field[s] [of philosophy, morality, and politics], no one's opinions on these matters […] are epistemically justified. […] And what follows from this, of course, is that a broad skepticism threatens"[50]

Kornblith assumes that the truth of some proposition p is an equal opportunity endeavor amongst the majority of inquirers. If the majority of experts in domain D inquire about p and move closer to a consensus, then this is evidence that such experts are, in fact, reliable at getting the truth in that domain. But because the "experts" in philosophy, morality, politics, etc., fail at achieving the kind of long-standing progress of consensus enjoyed by mathematical and scientific experts, Kornblith believes such "experts" really aren't experts insofar as being reliable gatherers of truth.[51] In other words, according to Kornblith, there are no experts in domain D unless there is a historical precedent of a consensus of opinion amongst such experts. So, although there may be prestigious experts in philosophy and other disciplines as a result of (say) being the recipient of many academic awards, being appointed to a prestigious university, and having many cited, peer-reviewed publications, it doesn't follow that such experts are reliable at getting the truth. The fact that one has acquired academic prestige does not indicate one's academic work is alethically aimed.

---

[50] Ibid., pp. 46-47.

[51] Unlike the discipline of philosophy, Kornblith says, with regards to mathematics and science, that "there is a well-established track record of reliable results issuing from the community, and, although dissenters are sometimes proven right, these communities are sufficiently reliable, and have been so for long enough, that one would always be ill advised to bet on the dissenter in the face of overwhelming majority opinion." See Hilary Kornblith, "Belief In The Face Controversy," in Richard Feldman and Ted Warfield, eds., *Disagreement* (Oxford: Oxford University Press, 2010), p.43.

Kornblith, thus, assumes that the majority of people comprising a discipline are equally suited at reliably getting the truth. If there are a group of people comprising discipline x, then if humans can achieve truths in discipline x, then the majority of people comprising discipline x will eventually converge on those truths. But since there is no majority converging on claims in philosophy, there is, by extension, no philosophical knowledge either.[52]

So rather than taking peer disagreement as something wherein epistemic symmetry either does or doesn't obtain, like Feldman and Christensen, Kornblith takes peer disagreement in philosophy (in general) to imply even stronger skeptical conclusions. Indeed, he takes disagreement in philosophy to be indicative of either the unreliable methods employed by philosophers or the unknowability of the subject matter *tout court*, given our limited epistemic capabilities. More specifically, Kornblith takes peer disagreement to be evidence against the reliability of the methods employed by the participants within any discipline that's saturated with disagreement. So, just as doing large math problems in one's head is an unreliable method evidenced by the pervasive disagreement amongst those employing such methods so, too, is trying to solve philosophical problems with the methods employed by philosophers. The verdict, for Kornblith, is that the upshot of the DINNER BILL case is that it generalizes far and wide. Widespread disagreement calls into question the reliability of philosophers' methods. This provides one with an undercutting defeater against one's philosophical beliefs. Accordingly, philosophical beliefs are unjustified, and thus, one lacks knowledge in such domains.

[52] Hilary Kornblith, "Is Philosophical Knowledge Possible?" in Diego E. Machuca, ed., *Disagreement and Skepticism* (New York: Routledge, 2013), p. 268.

2.5 CCE *A La* Goldberg

Like Kornblith, Sandford Goldberg is not sanguine about the methods employed within the discipline of philosophy nor our ability to get philosophical truths. More specifically, Goldberg argues that given the breadth and depth of disagreement in the field of philosophy, the methods employed by philosophers are suspect and such disagreement creates an epistemically hostile environment. Regarding philosophical methods, take a case involving peer disagreement whereby both disputants have fully disclosed their evidence bearing on p. One disputant believes p; the other believes not-p. Both disputants retain their belief on the basis of their intuitions, which function as a symmetry breaker.[53] But a method that involves appealing to one's intuitions is not a reliable one, according to Goldberg: "For in such a case one and the same method (or set of methods) is used by distinct individuals to reach opposing beliefs. Assuming that no single side of the debate has a preponderance of adherents, the result would be that, no matter which of the various beliefs produced by that method in this dispute is true, the method does not produce a preponderance of truths on the disputed topic."[54] Furthermore, given there is systematic peer disagreement amongst philosophers, even if *a* philosophical method *were* reliable, such disagreement would level an undercutting defeaters against one's belief. Goldberg thus

---

[53] Feldman denies that one's private evidence can be used as a symmetry breaker in order to demote a disputant's status as an epistemic peer. Given each disputant has her own private evidence or insight, such potential symmetry breakers cancel each other out once they are disclosed. Feldman sees "no basis for either [disputant] justifying his own belief simply because the one insight happens to occur inside of him. A point about evidence that plays a role here is this: evidence of evidence is evidence. More carefully, evidence that there is evidence for P is evidence for P. Knowing that the other has an insight provides each of them with evidence" (208). Feldman thus thinks that appealing to private evidence such as an intellectual seeming or intuition or insight is guilty of bootstrapping because one is arbitrarily giving one's private evidence special status and failing to treat like cases alike given each disputant has her own private evidence. See Richard Feldman, "Reasonable Religious Disagreements," in Louis M. Antony, ed., *Philosophers Without Gods: Meditations On Atheism And The Secular Life* (Oxford: Oxford University Press, 2007), p. 208; (Cf. Feldman, 2006).

[54] Sanford Goldberg, "Reliabilism in Philosophy," *Philosophical Studies* vol. 142 (2009): 107.

takes systematic peer disagreement within the field of philosophy to provide strong evidence for skepticism relative to philosophical knowledge.[55]

Goldberg defines systematic peer disagreement (hereafter SPD) as disagreement between two peers over a claim that is disagreed upon in a nonlocal, widespread, and entrenched way.[56] The three ingredients of SPD are (i) Nonlocal Disagreement—that is, disagreement over p, which implies disagreement over many other related topics; (ii) Widespread Disagreement—that is, when both disagreeing parties can appeal to a wide constituency of followers; (iii) Entrenched Disagreement—that is, disagreement that has persisted for some substantial length of time and all parties continue to defend and advance their view in the face of persistent challenges from the other side, which are replied to with evidence and arguments in the defense thereof.[57] SPD, thus, levels an undercutting defeater against one's belief since one's subjective probability relative to the disputed proposition should decrease given the pervasive disagreement is evidence for the disputed belief being a product of unreliable processes and/or methods. This subjective rationality defeater, in the form of an undercutting defeater, provides Goldberg with the resources needed to mount a case for skepticism. Because, again, regardless of whether or not one's controversial philosophical beliefs are reliably formed (and Goldberg argues they probably aren't reliable), such beliefs are not *ultima facie* justified.[58] His skeptical argument from disagreement can be summarized as follows:

---

[55] See Goldberg (2009; 2013a; 2013b).

[56] Sanford Goldberg, "Defending Philosophy in the Face of Systematic Disagreement," in Diego E. Machuca, ed., *Disagreement and Skepticism* (New York: Routledge, 2013), p. 278.

[57] Ibid., pp. 278-279. Local or Isolated disagreement comes from Kornblith (2010) and Elga (2007). A case of disagreement over p that is isolated and not widespread doesn't require suspension of judgment.

[58] It should be noted that Goldberg argues for skepticism with a two-pronged approach: the first prong is to argue for the unreliability of philosophical method. The second prong argues that even if philosophical

(1) In cases in which S believes that p in the face of systematic peer disagreement [SPD] over whether p, there are (undefeated doxastic or normative) defeaters with respect to S's belief that p.
(2) If there are (undefeated doxastic or normative) defeaters with respect to S's belief that p, then S neither knows, nor is doxastically justified in believing, that p.

Therefore,

(3) If p is a proposition regarding which there is systematic peer disagreement [SPD], then if S believes that p, S's belief is neither knowledge nor doxastically justified.[59]

This argument captures Kornblith's reasoning, too. For both think that it is epistemically incorrect to believe that there is knowledge or justified beliefs about a disputed proposition in a domain (esp., philosophy) that has an epistemically hostile environment evidenced by the systematic disagreement pervading it.

3. SKEPTICAL ARGUMENTS FROM CCE

3.1 Feldman and Christensen's Argument For Moderate Skepticism

We have looked at four versions of CCE. I now want to put forth—more formally—a skeptical argument that ties together Feldman's and Christensen's version of CCE. Both Feldman and Christensen believe that in a case of disagreement either epistemic symmetry obtains or it doesn't. If it does, then significant doxastic revision is required. If epistemic symmetry doesn't obtain, then there will be a relevant symmetry breaker that explains which party has the epistemic advantage. The first horn is innocent enough. For it hinges on an epistemic principle virtually all epistemologists agree upon, which I will call epistemic irrationality (hereafter **EI**):[60]

---

method is granted as reliable, systematic disagreement still mounts a subjective rationality defeater in the form of an undercutting defeater. So premise (1) in the argument above gets defended by two different arguments that are independent of each other.

[59] Ibid., p. 279.

[60] This point and principle is similarly made by Michael Thune in "Religious Belief and the Epistemology of Disagreement," in *Philosophy Compass* 5/8 (2010): 713.

> **EI**: If in a case of disagreement about the truth of some proposition p there is *epistemic symmetry* between the interlocutors' beliefs about p, it is epistemically irrational (or intellectually irresponsible) for those interlocutors to reject doxastic revision.

Alternatively: if one lacks a relevant symmetry breaker that can explain the disagreement between two cognitive and evidential equals then, on the pain of irrationality (or epistemic irresponsibility), one should engage in some level of doxastic revision. Clearly, EI is the principle behind the case of conflicting watches or, if you like, conflicting thermometers. If there is no independent reason—other than the conflicting thermometers—that can provide a symmetry breaker, then it's epistemically irresponsible to think your thermometer is the accurate one simply because it's *yours*. Similarly, EI is behind cases of disagreement involving two epistemic peers who disagree and there is no symmetry breaker to explain away the disagreement. This, of course, is what's going on in the DINNER BILL case, which aims to surface the intuition that significant doxastic revision is required when two epistemic peers disagree because they satisfy the antecedent of EI.

We also learned that the antecedent of ES (below) obtains when both parties acknowledge each other as epistemic peers and no relevant symmetry breaker is on offer. So the antecedent of EI (above) obtains when the consequent of ES obtains:

> **ES**: In a case of disagreement over the truth of some proposition p wherein both parties acknowledge each other to be *epistemic peers* and no symmetry breaker is on offer, then *epistemic symmetry* has obtained.

Alternatively: In a case of disagreement over the truth of some proposition p wherein both parties acknowledge each other as evidential and cognitive equals, and both parties have engaged in full disclosure, which results in failing to surface an explanatory symmetry

breaker, then there is no independent reason—other than the disagreement itself—to think one is more reliable at getting the correct answer than the disputant.[61]

Notice how ES can be strengthened once its antecedent explicitly includes the independence principle.[62] Recall independence:

> **Independence (I)**: Any symmetry breaker offered to explain why one disputant is epistemically privileged or advantaged over another must do so without employing the very reasons, beliefs, or arguments that originally formed one's disputed belief.

So ES can be strengthened by reformulating it this way.

> **ES+I**: In a case of disagreement over the truth of some proposition p wherein both parties acknowledge each other to be *epistemic peers* and no symmetry breaker (prior to the discovered disagreement) is on offer, then *epistemic symmetry* has obtained.

This version includes Christensen's independence principle which prevents one from conveniently demoting one's disputant from the status of epistemic peer by availing oneself of a property that is constitutive of one's token belief forming performance and using it as a symmetry breaker. Such bootstrapping strategies get blocked by it.

There is another epistemic principle proponents of CCE appeal to: namely,

> **Rationality Maintenance (RM)**: If S believes P, and S* believes not-P, and acknowledges S* is an epistemic peer, and S lacks a symmetry breaker, then on pain of irrationality, S should suspend judgment towards P.

CCE doesn't have exclusive purchase power on RM since it's endorsed by virtually everyone in the peer disagreement literature. Indeed, it's what gets the significance of peer

---

[61] In trying to generalize the argument for skepticism from CCE, I recognize that Christensen's independence thesis doesn't require both parties to engage in full disclosure subsequent to discovering the disagreement. Nevertheless, something like full disclosure is satisfied prior to discovering the disagreement in order to establish epistemic peerhood, even if *all* the relevant evidence hasn't been disclosed. Christensen would concede that part of what it takes to know another person is one's epistemic peer is knowing that they are familiar with the relevant evidence and arguments relevant to the proposition both parties disagree over.

[62] I thank Nathan King for pointing out this kind of gambit to me. See Nathan L. King, "Disagreement: The Skeptical Arguments from Peerhood and Symmetry," p. 209.

disagreement off the ground. Most of the division is whether cases like DINNER BILL can generalize far and wide.

The Argument From CCE A La Feldman & Christensen

(1) **ES**: In a case of disagreement over the truth of some proposition p wherein both parties acknowledge each other as *epistemic peers*, then *epistemic symmetry* has obtained.
Or **ES+I**: In a case of disagreement over the truth of some proposition p wherein both parties acknowledge each other to be *epistemic peers* and no symmetry breaker (prior to the discovered disagreement) is on offer, then *epistemic symmetry* has obtained
(2) **EI**: If in a case of disagreement about the truth of some proposition p there is *epistemic symmetry*, it is epistemically irrational (or epistemically irresponsible) for each party to not compromise and reject doxastic revision.
Therefore,
(3) **RM**: If S believes P, and S* believes not-P, and acknowledges S* is an epistemic peer, and S lacks a symmetry breaker, then on pain of epistemic irresponsibility, S should suspend judgment towards P.
(4) Epistemic responsibility is a necessary condition for knowledge.
(5) Finding oneself disagreeing with an acknowledged epistemic peer over some proposition p wherein epistemic symmetry obtains happens quite often.
Therefore,
(6) It happens quite often that peer disagreement provides one with a rationality defeater against knowledge.

3.2 Kornblith and Goldberg's Argument For Moderate Skepticism

We looked at Kornblith and Goldberg who argue that any domain that suffers from pervasive or systematic disagreement is grounds for denying any knowledge or justified beliefs in that domain. More specifically, Kornblith and Goldberg argued that since domains like philosophy—unlike mathematics and the hard sciences—suffer from systematic disagreement, we have strong evidence that we are not endowed with reliable methods and/or faculties to get the truth in these domains. (This is what explains the historic and ongoing disagreement.) So, although Feldman and Christensen appeal to epistemic symmetry in order to generate a rationality defeater, Kornblith and Goldberg are

using a different notion of epistemic symmetry: namely, everyone is in the same epistemically damaged boat insofar as getting the truth in domains wherein systematic disagreement abounds. Therefore, I don't assume Kornblith's and Goldberg's arguments to perfectly map onto Feldman's and Christensen's arguments.

But their arguments do converge on disagreement providing one with a subjective rationality defeater. More carefully, Feldman and Christensen argue for skepticism by providing a rationality defeater from disagreement on the basis of epistemic symmetry obtaining; whereas Kornblith and Goldberg argue for skepticism by providing a rationality defeater from disagreement on the basis of philosophical beliefs deriving from unreliable methods and/or sources. Since these CCE proponents all believe epistemic responsibility is a necessary condition for knowledge, they all agree that disagreement provides a subjective rationality defeater to one's belief thereby preventing one from having a justified belief or knowledge. What's more, they both agree that such defeaters get generated—quite often—in domains that lack consensus like philosophy, religion, politics, and ethics. So the argument for moderate skepticism that Goldberg and Kornblith propound, put more formally, goes something like this:

The Argument From CCE A La Kornblith & Goldberg

(1) In cases in which S believes that p in the face of systematic peer disagreement over whether p, there are rationality or normative defeaters with respect to S's belief that p.
(2) If there are rationality or normative defeaters with respect to S's belief that p, then S neither knows, nor is doxastically justified in believing, that p.
Therefore,
(3) If p is a proposition regarding which there is systematic peer disagreement, then if S believes that p, S's belief is neither knowledge nor doxastically justified.

3.3 Summary

My aim in this chapter has been to introduce one of the main problems within the dialectic concerning the epistemology of peer disagreement. More specifically, I provided what the CCE camp takes to be the epistemic significance of peer disagreement: namely, that a moderate version of skepticism follows. The purpose of showcasing these kinds of skeptical threats from disagreement is because I want to respond to them in subsequent chapters. So, in Chapters Two and Three, I will respond to the first argument from Feldman and Christensen. More specifically, I argue why the problem of peer disagreement put forth by Feldman, Christensen, et al. is a pseudo-problem because (a) the antecedent conditions in premise (2) rarely ever obtain and (b) the normative evaluation used to motivate the problem of peer disagreement is an inadequate one.

The argument from Kornblith and Goldberg will not be addressed until Chapter Six. The reason these arguments need to be addressed separately is that I don't think the problem of disagreement Kornblith and Goldberg put forth is a pseudo-problem. Because their argument is not motivated by only internalists' concerns but, also, externalists' ones, it will require its own separate treatment. Indeed, the argument from Kornblith and Goldberg specifically deals with knowledge, which is where I believe the debate should be taking place. So, it will be responded to in Chapter Six which is where I address what I take to be the specific problems disagreement levels against reliabilist theories of knowledge. But first I need to wrestle the problem of peer disagreement away from internalists. That is, I need argue why the problem of peer disagreement—as characterized by an internalist construal of it—is a pseudo-problem. I do this in Chapters Two and Three. We turn now to Chapter Two.

2

Why Epistemic Disagreement Is A Pseudo-Problem: Part 1

Good sense is the best distributed thing in the world: for everyone thinks himself so well endowed with it that even those who are the hardest to please in everything else do not usually desire more of it than they possess. In this it is unlikely that everyone is mistaken. It indicates rather that the power of judging well and of distinguishing the true from the false—which is what we properly call 'good sense' or 'reason'—is naturally equal in all men [and women], and consequently that the diversity of our opinions does not arise because some of us are more reasonable than others but solely because we direct our thoughts in different paths and do not attend to the same things.

~RENÉ DESCARTES

We give the name evidence to whatever is a ground of belief. To believe without evidence is a weakness which every man [and woman] is concerned to avoid, and which every man [and woman] wishes to avoid. Nor is it in [one's] power to believe anything longer than [one] thinks [s/he] has evidence.[…] I confess that, although I have, as I think, a distinct notion of the different kinds of evidence above mentioned [e.g., perception, memory, introspection, consciousness, testimony, axioms, and argument] […] yet I am not able to find any common nature to which they may all be reduced. They seem to me to agree only in this, that they are fitted by Nature to produce belief in the human mind.

~THOMAS REID

In this chapter, I argue that the problem of peer disagreement is a pseudo-problem. There are two arguments I propound that are each independently sufficient to dissolve the skeptical worries leveled by peer disagreement. The first argument I present here in Chapter Two; the second argument I provide next in Chapter Three. The first line of argument exposes how most of the responses to peer disagreement work within an internalist framework. This is problematic because it assumes that the normative status at issue in cases of peer disagreement is one that is abstracted away from externalist concerns and considerations that bear on knowledge. For when we consider cases of peer disagreement, they typically result in evaluations that judge a subject's belief to be either justified or unjustified, reasonable or unreasonable, rational or irrational, responsible or irresponsible. Such normative evaluations, then, attribute to a belief different positive or negative epistemic status. But I argue that proponents of CCE make such normative

evaluations without weighing in externalist intuitions. Indeed, I argue that once externalist concerns do weigh in, the problem of peer disagreement dissolves. I argue that we shouldn't assume internalism (full stop) because it provides an inadequate account of the kind of positive epistemic status at issue in cases of peer disagreement.

There are many different properties that can afford beliefs positive epistemic status of one kind or another. For example, a belief may be justified, rational, formed in an epistemically responsible way, or be subjectively appropriate, etc. These are all candidates for what can satisfy the function of a belief having what I deem *pd-normative status*. In other words, pd-normative status is a functional term that intends to capture (whatever it is) we think our beliefs must (at minimum) possess to have positive epistemic status in cases involving peer disagreement. For we want to know how peer disagreement affects our belief's pd-normative status, which informs us as to whether we should (say) retain our belief or give it up. And although internalist concerns can capture *a* plausible candidate for what constitutes a belief having pd-normative status, it fails to capture many other important pd-normative candidates. Indeed, internalism offers an impoverished account of how to evaluate a belief's pd-normative status because it leaves out relevant candidates that rely upon external factors that also give a belief positive normative status.

1 WHY INTERNALISM IS TOO RESTRICTIVE TO ACCOUNT FOR MANY IMPORTANT NORMATIVE EVALUATIONS

1.1 Internalism vs. Externalism

The debate between internalism and externalism is primarily over what properties, events, actions, abilities, or states of affairs should be evaluated when determining a belief's positive epistemic status. Should the locus of analysis for what constitutes a belief's positive epistemic status be, according to the internalist, that which grounds why one is reasonable?

Or should we, like the externalists, also care about the objective facts about the world and how beliefs form in it—regardless of whether or not we can internally access these facts by reflection alone? Are normative evaluations regarding a belief's positive epistemic status restricted to domains that exclusively deal with subjective appearances or do they extend to domains that take into account objective facts we cannot subjectively access? In short, is a belief's positive epistemic status something internal or external to a subject?

1.1.1 Access Internalism

There are different characterizations of internalism. We do well, then, to look at some popular formulations of it. The first type of internalism is *access internalism*. This view says a belief's epistemic status must meet a justificatory standard comprised of that which is internally accessible to a subject. For example, Robert Audi says, "The central internalist idea about justification [i.e., a belief's positive epistemic status] is that of meeting a certain justificational standard that one can conform to on the basis of a kind of response to accessible elements."[63] Stated more formally:

Access Internalism: What determines whether or not S's belief that p has positive epistemic status is entirely a function of what facts S can internally access via reflection alone.
Alternatively: A belief's epistemic status supervenes upon the states of affairs S can reflectively access via introspection, a priori intuition, or reasoning from truths that are known in such ways.[64]

The idea, then, is that access internalism uses a standard for epistemic evaluation that is solely determined by what properties, facts, and/or states of affairs a subject has direct,

---

[63] Robert Audi, *Epistemology: A Contemporary Introduction to the Theory of Knowledge* (New York: Routledge, 1998), p. 235.

[64] This formulation was inspired by John Greco. See John Greco, *Achieving Knowledge: A Virtue Theoretic Account of Epistemic Normativity* (Cambridge: Cambridge University Press, 2008), p. 48.

internal access to. One of the main motivations for access internalism derives from considering skeptical scenarios involving subjects living in evil demon worlds or as envatted brains. Because we ourselves could possibly be one of these epistemically hapless victims, it's counterintuitive to think that not one of our beliefs—right now!—possess any positive epistemic status whatsoever. Surely some of our beliefs have achieved a positive level of epistemic status despite virtually all of them being patently false. And access internalists think the best way to capture our intuitions after evaluating such skeptical scenarios is to appeal to that which both the deceived subject and non-deceived subject share in common: namely, the states of affairs they both have privileged internal access to through reflection alone. This can be illustrated by considering the following case.

> MY DECEIVED INTERNAL DUPLICATE: There is a molecule per molecule duplicate of my brain in a vat of chemicals. This person is my internal TWIN insofar sharing the same internal, phenomenological perspective of the outside world. Whatever mental states I have internal, reflective access to so, too, does my TWIN. For example, I look down at my arm and see that I have hands. My TWIN does the same thing and his hands *appear* to him the same way phenomenologically as my hands appear to me. The difference, however, is that I *have* hands and my TWIN doesn't—he's a handless brain-in-a-vat. So, when I form the belief that I have hands, my belief is true. When my TWIN forms the belief that he has hands it's false despite sharing the same internal, phenomenological perspective as me.

The internalist thinks this kind of case, as well as others, captures how both my TWIN and me are alike internally, which is the locus of a belief's positive epistemic status. Because we are responsible only for that which is under our control, we are responsible only for what we can internally access. So, a belief's epistemic status is an internal matter because justification is about being epistemically responsible and rational relative to how we form our beliefs. Therefore, if I'm justified in believing that I have hands and my TWIN's internal perspective is indiscriminable from mine then so, too, is my TWIN for believing he has hands. Therefore, given that our internal perspectives are indiscriminable from each

other, neither of us are blameworthy for believing we have hands—that is, both beliefs share the same epistemic status. A principle that falls out of access internalism after considering such skeptical scenarios is this.

ACCESS: If S and S* do not differ in the facts that they are able to know through internal reflection alone then the degree of their belief's epistemic status will not differ.[65]

Such accessible elements, then, will inform the standard by which we evaluate a belief's epistemic status. Therefore, a belief's epistemic status is determined by those states of affairs S is reflectively aware of.

There is another principle that is a more general version of ACCESS: namely, NT.

NT: Indiscriminable situations are normatively identical.[66]

NT is a principle put forth by Timothy Williamson who argues that it captures the driving intuition behind internalist's reactions to the New Evil Demon case. According to this principle, because we and our TWINS share the same phenomenological states, we and our TWINS are in normatively identical situations. In short, to be in an internally indiscriminable situation with respect to your counterpart is sufficient to be in a normatively identical situation with respect to your counterpart. It's important to see that to be in an indiscriminable situation with your counterpart is to be in a situation where both you and your counterpart lack the ability through internal reflection alone to "access" any normative differences

[65] This principle is inspired by one put forth by Duncan Pritchard. See Duncan Pritchard, "Evidentialism, Internalism, Disjunctivism," in Trent Dougherty, ed., *Evidentialism and its Discontents* (Oxford: Oxford University Press, 2011), p. 238

[66] Timothy Williamson, "On Being Justified in One's Head," in Mark Timmons, John Greco, and Alfred R. Mele, eds., *Rationality and the Good: Critical Essays on the Ethics and Epistemology of Robert Audi* (Oxford: Oxford University Press, 2007), p.119.

1.1.2 Mentalism Internalism

Mentalism is a weaker type of internalism. Instead of a belief's positive epistemic status supervening only upon those states of affairs that are reflectively accessible to a subject, the supervenient base now extends to a subject's "occurrent and dispositional mental states, events, and conditions."[67] In other words, "internalism" denotes a theory of justification, but S need not be able to reflectively access *all* such mental states. So the members in the set of mental states deemed "internal"—on which a belief's epistemic status supervenes—need only be reflectively accessible *in part*. There is no explicit requirement for S to have internal access to all the properties that confer onto a belief a positive epistemic status. Stated more formally:

Mentalism Internalism: That which determines whether or not S's belief that p has positive or negative epistemic status is entirely determined by the facts constitutive of S's mental life.
Alternatively: The facts about a belief's epistemic status supervene on the facts about S's mental life.[68]

Notice how one no longer needs to have privileged access to the states of affairs a belief's epistemic status supervenes upon. One may criticize this type of internalism for employing a quixotic interpretation of what it means for something to be internal. For one of the main motivations, again, behind internalism is considering the new evil demon problem which involves evil demon worlds and brain-in-vat scenarios. Internalists intuitively took a belief's epistemic status to be the result of facts shared by both the subject in a bad case (i.e., a subject's brain is in a vat of chemicals being stimulated by mad scientists that bring about false beliefs through manipulative, unreliable processes) and the subject in a good case (i.e.,

---

[67] Earl Conee and Richard Feldman, "Internalism Defended" in Feldman and Conee, eds., *Evidentialism: Essays in Epistemology* (Oxford: Clarendon Press, 2004), p. 56.

[68] The formulation of this principle is inspired by John Greco's characterization of mentalism. See John Greco, *Achieving Knowledge*, p. 54.

a subject's brain is causally connected to the external world delivering true beliefs through reliable processes). Such facts are the ones each subject was able to internally access via reflection alone. But the evaluative principle that falls out of mentalism internalism is broader in scope.

MENT: If S and S* do not differ in their mental states then their beliefs' epistemic status will be the same.[69]

Notice, however, that some versions of MENT are compatible with externalism. Why? Because, according to the externalist, S and S*'s mental states entail more than what both have internal access to. Thus, the externalist denies ACCESS while still retaining a variant of MENT. For example, an externalist may say the following: As long as (a) both S and S* share the same mental states (understood more broadly in scope than ACCESS) and (b) both the etiology of S and S*'s mental states are brought about in the same relevant way, then both S and S*'s mental states will share the same epistemic status. Thus, it's important to see how externalists understand mental states in a way that is much broader in scope than an internalist's construal. For the internalist believes my TWIN and I share the same mental states because we share the same phenomenological appearance of a hand, which is something we can both internally access. Whereas, the conditions that need to be satisfied for S and S* to share the same mental state will require more external facts that go beyond the scope used by internalists.

So, in order to understand what it means for my TWIN and I to share the same "mental states" we must get clear on what "mental states" denote. And the internalist and externalist are going to disagree over the referents. For externalists will happily endorse MENT because they are going to say the term "mental states" denotes factive states that

[69] The formulation of this evaluative principle is inspired by Duncan Pritchard. See Duncan Pritchard, "Evidentialism, Internalism, Disjunctivism," in Trent Dougherty, ed., *Evidentialism and its Discontents* (Oxford: Oxford University Press, 2011), p. 238.

include, but are not limited to, that which is internally accessible. Indeed, externalists say mental states denote both what one has internal, reflective access to *and* the factive states one lacks internal, reflective access to. But the reflectively inaccessible facts still constitute one's mental state. The tent, then, for what constitutes a mental state is bigger for the externalist. This is why most internalists endorse some version of accessibilism, which captures the privileged access one has of one's own mental states. So, it's unclear why mentalism is, strictly speaking, a version of internalism since the debate between internalists and externalists revolves so much around the merits of ACCESS. (Hereafter any reference to MENT, unless otherwise indicated, is to be understood in the restricted internalist sense which includes an ACCESS component.)

Internalists, then, put a primacy on appealing to states of affairs that can be known by a subject without begging any questions. By taking the demands of skeptics seriously, they appeal to mental states that can be internally accessed via reflection so that their epistemic status can be evaluated. Contra externalism, internalists say *de facto* reliability alone is not enough for a belief to achieve a positive epistemic status. So the internalist seeks to *vindicate* why her belief has positive epistemic status in a way that satisfies the skeptic's demands. Moreover, the internalist wants to provide an account of normativity that explains why subjects' beliefs in both the good and bad cases share the same positive epistemic status. Because there are skeptical scenarios involving internal duplicates—one in an epistemically friendly environment and the other in an epistemically unfriendly environment—one's account of epistemic normativity mustn't make one blameworthy for not accessing that which can't be accessed. After all, recall that we, too, could be the hapless victims in bad cases. But surely some of our beliefs have some positive epistemic status even if we were in such an epistemically unfriendly environment. This is why

mentalist-internalists, like Feldman and Conee, say things like this about the New Evil Demon Scenario[70], which is just a restatement of MENT: "If any two possible individuals are exactly alike mentally, then they are alike justificationally [i.e., their beliefs share the same positive epistemic status], e.g., the same beliefs are justified for them to the same extent."[71]

1.1.3 Externalism

To get a sense of how externalism contrasts with internalism and understand its driving intuition, let's consider one of the most prominent examples of it: namely, Alvin Goldman's *reliabilism*. Reliabilism says a belief's positive epistemic status is determined by factors that are reliable at getting the truth. The motivation for reliabilism, then, according to Goldman, results from considering different processes or methods from which beliefs derive. Goldman says beliefs that form from "confused reasoning, wishful thinking, reliance on emotional attachment, mere hunch or guesswork, and hasty generalizations" are

[70] The new evil demon problem is different than the old evil demon problem. The new evil demon problem (Lehrer and Cohen, 1983) is used as an argument against externalism. It aims to provide a counterexample to reliabilism by eliciting an intuition, namely, despite your beliefs externally hooking up to the world in the right sort of way while your TWIN's don't, both you and your TWIN's beliefs are rational. Rationality, then, supervenes on internally accessible mental states—that is, the internal facts shared by you and your TWIN—and not the way such mental states hook up to the external world. *De facto* reliability is insufficient for an adequate account of knowledge. The old evil demon problem traces back to Descartes; it's used as a skeptical hypothesis that allegedly needs to be ruled out before one can truly know. The problem the old evil demon scenario presents is that there is nothing that one can internally access that can distinguish the real world from the deceptive-world manufactured by the evil demon. In other words, my Twin and I access through reflection alone mental states that are phenomenologically indiscernible, yet neither of us can access that which distinguishes the veridical mental states from the ones aimed to deceive. So, the old evil demon problem is trying to find what positive epistemic candidate can be used for distinguishing a belief having positive epistemic status in the good case but is lacking in the bad case. I'm grateful for Wayne Riggs's comments and suggestion about the old and new evil demon problem.

[71] Earl Conee and Richard Feldman, "Internalism Defended," p. 56

all processes or methods that share a common culprit: namely, "*unreliability*."[72] Indeed, he says such processes:

> tend to produce error a large portion of the time. By contrast, which species of belief-forming (or belief-sustaining) processes are intuitively justification-conferring? They include standard perceptual processes, remembering, good reasoning, and introspection. What these processes seem to have in common is reliability: the beliefs they produce are generally true. [A belief's epistemic status, then,] is a function of the reliability of the process or processes that cause it, where (at first approximation) reliability consists in the tendency of a process to produce beliefs that are true rather than false.[73]

It's not that externalists deny that *some* of the states of affairs determining a belief's epistemic status are not internal to the subject. It's just externalists deny that *all* the facts determining a belief's epistemic status are reflectively accessible. This point made by externalists highlights just how strong of a thesis internalism is regarding a belief's epistemic status.

There are, of course, other versions of externalism. Some versions argue that a belief's epistemic status is determined by intellectual abilities or virtues that reliably get the truth. But the ways in which such abilities produce a true belief are not entirely reflectively accessible.[74] And there are still other versions of externalism that say a belief's epistemic status is determined by whether or not S's cognitive faculties are functioning properly. For example, Michael Bergmann says S's belief has positive epistemic status just in case "(i) S does not take [her belief] to be defeated and (ii) the cognitive faculties producing [the targeted belief] are (a) functioning properly, (b) truth-aimed and (c) reliable in the

---

[72] Alvin Goldman, "What Is Justified Belief?" in Paul K. Moser *Empirical Knowledge* (Totowa, NJ: Rowman an Littlefield, 1986), p. 179.

[73] Ibid.

[74] See Earnest Sosa (1991); Zagzebski (1996); and Greco (2000).

environments for which they were 'designed'"[75] Such a version of externalism denies that a subject must have reflective access to all the facts comprising (a) - (c).

It would go beyond the purposes of this chapter to tease out all the species of externalism and how they differ from each other. What is of importance is seeing that what all versions of externalism have in common is a rejection of the accessibility requirement (ACCESS) for determining a belief's epistemic status.

ACCESS: If S and S* do not differ in the facts that they are able to know by internally accessing through reflection alone then the degree of their belief's epistemic status will not differ.

Externalists also reject the idea that phenomenologically indiscriminable experiences implies that two people's beliefs will share the same epistemic status. The reasons why they reject this will be discussed below. And with regards to MENT—that is, if S and S* do not differ in their mental states then they will not differ in the degree of epistemic justification that they have for their beliefs—again, it's unclear why this is exclusively an internalist thesis. For it all depends on whether or not one is, say, an *externalist* regarding the content of mental states. I will discuss content externalism below. But the gist of it is that mental states are fixed by their causal etiology which necessarily depends upon the external world. So, the reason why internalists don't have exclusive claim to MENT is that an externalist that is also a *mental-content-externalist* will gladly endorse MENT.

Contra internalism, externalism is the view that says that what determines a belief's positive epistemic status is *not* entirely a function of the facts S can internally access via reflection alone. It rejects principles like ACCESS. Externalism, then, denies the standard used by internalists for determining a belief's epistemic status. For externalism says a

[75] Michael Bergmann, *Justification without Awareness* (Oxford: Oxford University Press, 2006), p. 133. The first and most influential proponent of this version of externalism is Alvin Plantinga. See Alvin Plantinga, *Warrant and Proper Function* (Oxford: Oxford University Press, 1993).

belief's positive epistemic status is evaluated in terms of whether or not the token process or method or ability that caused the belief is conducive towards getting the truth. And there is no requirement by the subject to have any internal access to such processes, methods, or abilities through reflection alone. Although *some* facts may be accessible, they need not be.

What's important for the externalist is that part of a belief's positive epistemic status is the result of that which has a strong relation to the world; otherwise a belief gets positive epistemic status on the cheap. Indeed, externalists deny that positive epistemic status can be acquired so easily, especially since mental states that are internally accessible can be completely cut off from the external world. Internally accessible facts or evidence, then, is important only if it's a reliable means towards getting the truth through some strong relation—perhaps even a modal one—between one's belief and the world.

1.2 The Inherent Complexity of Epistemic Normative Evaluations

Epistemology is a normative discipline. *When we say a belief has a positive epistemic status we are saying it has a positive normative status—that is, we are saying it is rational, reasonable, responsible, justified, or has some other epistemic good.*[76] When we say people have responsible or justified beliefs, we are making a normative evaluation that contrasts one belief as epistemically *better* than other beliefs that are evaluated as irresponsible or unjustified. Many internalists and most of the ones writing in the area of peer disagreement are interested in capturing a dimension of a belief's epistemic status deemed epistemic reasonableness, rationality, or responsibility, which they take to be an internal matter.[77] This way of thinking about

[76] I grateful for Wayne Riggs for helping me to clarify this point.

[77] It's unfair to paint all internalist by using the same brush. I concede that my descriptions and constraints that are used to characterize internalists will not apply to all of them given all the species of internalism that

epistemic normativity is, again, largely motivated by appealing to cases that involve two or more agents that are alike internally, except one has a true belief and the other has a false belief. Internalists think such cases provide important data that any account of epistemic normativity needs to explain. For the agents in both the good and bad scenarios are acting with equal degrees of epistemic responsibility or share the same justification—that is, their beliefs share the same epistemic status. Why? Because both agents' internal perspectives are indiscriminable. That is, the accessible factors that are internal to S's perspective are indiscriminable from the accessible factors that are internal to S*'s perspective.

There are three theses, then, that capture the motivating intuitions behind internalism.

Principle 1—MENT: If S and S* do not differ in their mental states then they will not differ in the degree of epistemic justification that they have for their beliefs.
Alternatively: The facts about a belief's epistemic status supervene upon S's mental states.

Principle 2—ACCESS: If S and S* do not differ in the facts that they are able to know by reflection alone then the degree of their belief's epistemic status will not differ.[78]
Alternatively: The facts about a belief's epistemic status supervene upon the states of affairs to which S has privileged access—that is, whatever mental states S can reflectively access from S's internal perspective will determine the epistemic status of S's belief.[79]

Principle 3—NT: Indiscriminable situations are normatively identical. [80]

Again, I take NT to be more general than ACCESS and MENT, which are special cases of it. Be that as it may, NT is used to make normative evaluations about a belief's epistemic status. It's an unrestricted claim; so, it can apply to many epistemic candidates depending upon what epistemic good one is interested in a targeted belief satisfying. For example, NT

---

are out there. So, I don't intend to pigeon hole all internalists with the characterizations I use. I do think, however, that these characterizations are playing a vital, operative role in the literature on peer disagreement.

[78] This principle is inspired by Duncan Pritchard but has been slightly amended. See Duncan Pritchard, "Evidentialism, Internalism, Disjunctivism,"p. 238.

[79] This principle was inspired by what Greco deems "privileged access internalism." See John Greco, *Achieving Knowledge* (Cambridge: Cambridge University Press, 2010), p. 63.

[80] Timothy Williamson, "On Being Justified in One's Head," in Mark Timmons, John Greco, and Alfred R. Mele, eds., *Rationality and the Good: Critical Essays on the Ethics and Epistemology of Robert Audi* (Oxford: Oxford University Press, 2007), p.119.

may be true with respect to, say, blameworthiness. My TWIN and I are blameless insofar as forming the beliefs that we have hands. But NT, and by extension ACCESS and MENT, are false once we widen the scope of normative concerns beyond the normative dimension of blameworthiness.

The remainder of this section, then, aims to argue why principles 1-3 prevent one from acquiring an adequate account of normativity in cases involving disagreement. The kind of epistemic status I argue that we should care about in cases of disagreement takes into account externalist concerns that elude principles 1-3. Beliefs that do satisfy both internally and externally relevant facts relative to cases involving disagreement will have positive *pd-normative status*. My argument against the adequacy of principles 1-3 is threefold. First, I argue that if content externalism is true, then ACCESS and its variant NT fail to capture normatively relevant evaluations. Second, I build upon the significance of content externalism by arguing that NT and ACCESS are false as a result of there being *bona fide* normative differences NT and ACCESS would have us deny. Third, principles 1-3 prevent one from making normative evaluations that pick out other relevant normative values that give beliefs positive pd-normative status.[81]

1.2.1 Content Externalism And Its Relevance to Epistemic Evaluations

First, many externalist and philosophers (in general) believe that the content of one's mental states are dependent upon and partly constituted by the external environment. That is, the content of one's belief are not *merely* the result of the goings-on inside one's head but, rather, partly constituted by external world that partly caused it.[82] For example,

[81] I'm grateful for Wayne Riggs who helped me better refine my argument for why principles 1-3 prevent one from acquiring an adequate account of normativity.

consider a case involving my TWIN and me. My TWIN lives in "Twin World," which is a world just like mine.[83] The only difference is that the clear, liquid stuff we drink, shower with, and freezes below 32° F is composed of a different molecular structure than the liquid stuff in my world. In TWIN's world, the liquid stuff is composed of XYZ instead of H_2O. Now consider a case where my TWIN was transported to my world. If I asked him What are these cubes in my ice tray made out of? and he answered "water," would his belief be true? Furthermore, and more importantly, would we share the same belief were we to perceive *those* cubes in my ice tray as frozen cubes of water? Before answering this remember that both my TWIN's perception of the ice cubes and my perception of the ice cubes are phenomenally indiscriminable. So, we can't see the cubes' molecular structure.

It seems that even if our perceptual experiences of the cubes are indiscriminable, intuitively we are not in identical epistemic states. Why? Answer: my belief is true and my TWIN's is false given our beliefs are about different referents: namely, H_2O and XYZ, respectively. The answer to both questions, then, seems to be intuitively no. Indeed, my belief that the cubes are made from water is true and my TWIN's belief is false because we are referring to two different things. For my belief is *about* that which is composed of a molecular structure H_2O and my TWIN's belief is *about* that which is composed of molecular structure XYZ.[84] Furthermore, even our English term 'water' refers to different

---

[82] There is a distinction between mental states sharing either "narrow" content or "wide" content. Narrow content in circumscribed by only the goings-on inside one's head. Jaegwon Kim puts it this way: "[T]he content of an intentional state is narrow just in case it supervenes on the internal-intrinsic properties of the subject who is in that state, and that it is wide otherwise." Wide content, however, is individuated by "reference to conditions external to the believer." See Jaegwon Kim, *Philosophy of Mind*, 2nd ed., Cambridge, MA: Westview, 2006), p. 254-255.

[83] This kind of case involving "Twin Earth" was originally given by Hilary Putnam in his paper "The Meaning of 'Meaning'" in *Minnesota Studies in the Philosophy of Science* 7 (1975): 131-193. This is also the paper where the distinction between narrow and wide content is made.

[84] Logicians would say my term 'water' has a different extension than it does in Twin Earth. Were I to be transported to Twin Earth and discover the liquid stuff there is not H_2O but, rather, XYZ, then I would have

liquid stuff. Therefore, the intuition for why my TWIN and I have different beliefs seems best explained by the concept water denoting different things as a result of its etiology tracing back to two different worlds and being partly constituted by such external inputs. [85]

The upshot, then, is that content externalism explains why we intuitively take my TWIN and me to have different beliefs and why their truth conditions turn on external environmental factors that lie outside the goings-on inside our heads. Content externalism says a beliefs' mental content *can't* be circumscribed to that which supervenes solely upon one's internal mental states. Otherwise, both my TWIN and I would share the same beliefs regardless of which world's clear liquid stuff we perceived. But that seems counterintuitive because there is a truth to the matter about which of our beliefs is true. (Whether perceiving fool's gold is phenomenologically indiscriminable from perceiving atomic number 79 won't convince a jeweler the weight of your fool's gold is of equal value to the weight of atomic number 79.) And the truth conditions depend upon both the mental content of our beliefs and the external facts—outside our skin—that brought about and partly constitutes such perceptual beliefs. This is why content externalists say that if S and S*'s beliefs are truly the same, then the truth status of their beliefs turns on the content of their beliefs, which necessarily depends upon being related to the external world in the right sort of way and partly constituted by such external facts.

---

to refer to it by a new term, say, 'twater' or 'fools water.' The extension of twater would refer to the molecular structure XYZ.

[85] This example comes from a classical argument that was first given by Saul Kripke for content externalism. Kripke argued that the referent for proper names and natural kind terms causally supervene upon the states of affairs in the external world in the right sort of way. Later, Hilary Putnam argued for a similar thesis by using a thought experiment involving "Twin Earth," which aimed to show that the appearance of water was not sufficient to establish the reality of water since chemical constituents XYZ and H_2O can appear phenomenologically in indiscriminable ways. Put simply: our beliefs are not the same given their mental contents are different. The amount of time it would take for my TWIN's use of the term "water" to accurately refer to H_2O is a puzzle that I will not attempt to answer given the difficulty and consequence of taking us too far afield from the purposes of this chapter. See Saul Kripke, *Naming and Necessity* (Cambridge, Mass: Harvard University Press, 1972).

Mental content, then, seems to intuitively entail both a relation to the external world that partly constitutes the content of the mental state in question. The content of my belief that this liquid stuff is water is not the result of a supervenient thesis that says the content of my belief entirely supervenes on my brain's physical-psychological states. Individuating beliefs by their physical-psychological states alone fails to explain the intuition we have for why my TWIN's belief about water in Twin Earth is different than my belief about water in this world. A better explanation is that beliefs get individuated by their mental content which is partly constituted by the external facts comprising one's environment, which figure into the causal etiology of the targeted belief. So, if the chemical molecules of the liquid stuff in my world is composed of H_2O and in my TWIN's world is composed of XYZ, then we can't share the same belief—period! Our beliefs lack the same "wide" mental content even though they may, perhaps, share the same "narrow" content given we share the same phenomenology. Therefore, being phenomenological duplicates isn't sufficient to establish the same mental content comprising our beliefs.

Let's take stock: the epistemic status of my TWIN's belief (in the brain-in-vat scenario) is different than mine because our beliefs don't share the same mental content given their causal etiology is necessarily dependent upon two different external worlds.[86]

Second, we can now build upon the significance of content externalism to straightforwardly show the falsity of NT. Again, ACCESS provides the content for a normative evaluation principle used by internalists. Williamson says internalists are committed to a principle he deems Normative Tolerance (NT), which is a variant of ACCESS.

NT: indiscriminable situations are normatively identical.

[86] This is why epistemic externalists endorse a version of MENT that includes *all* the facts comprising one's mental state—even the facts that aren't internally accessible and lie outside one's head.

Indiscriminable situations can mean either MENT (in the restricted internalist sense) or ACCESS. So, if ACCESS or MENT are true, then S and S* are in normatively equivalent situations. Therefore, S and S* should receive the same normative evaluation. In other words, any two cases we evaluate are normatively identical just in case any difference between (say) my TWIN and me is not internally accessible to us in the relevant way. And a situation is normatively identical—that is, should be normatively evaluated the same way—just in case ACCESS and MENT are true. NT, then, is supposed to capture why internalists are motivated to give my TWIN and me the same normative evaluation. For the internalist thinks the following kind of normative principle about indiscriminate situations falls out of cases like MY DECEIVED INTERNAL DUPLICATE.

INDISCRIMINABLE: If the phenomenological experiences had by S and S* are indiscriminable then S's and S*'s beliefs will share the same positive epistemic status.

It's important to notice that INDESCRIMINABLE assumes something like ACCESS and a *restricted* notion of MENT is true. Thus, because internalists endorse INDISCRIMINABLE they are motivated to evaluate my TWIN and me the same way since what we are both accountable for is determined by what we can both internally access (ACCESS). And since we are not in control of the external environment, we can't blame my TWIN for believing falsely that he has hands for he believes, according to the internalist, in the *same* thing as me.

But externalists can, for good reasons, deny this. My TWIN and I didn't form the *same* beliefs. So, we can be evaluated differently. The internalist, however, wants to argue that my TWIN and I still believe the same thing, regardless of the implications of content externalism. So, if I'm not blameworthy for forming the belief that I have hands, then neither should my TWIN be blameworthy for believing likewise. This is why

indiscriminable situations are treated by internalists as normatively identical. Furthermore, the internalist will argue that if I'm not blameworthy for not being able to control my external environment, then my TWIN can't be blameworthy for believing falsely since he can't control his external environment either.

Regarding NT, I think Williamson is correct when he says my TWIN and I are not in normatively identical situations.[87] Given NT is an unrestricted claim, it may get right the fact that my TWIN and I are blameless. But it's false we are normatively identical. The reason why is that one essential feature of identity is that it's symmetrical thereby making normative identity symmetrical, as well.[88] But indiscriminable situations don't necessarily share identical properties. For example, I know that I have hands, yet my TWIN doesn't know he has hands because he's a handless-brain-in-a-vat. Indiscriminable situations, then, are not normatively identical. Why? Because identity is symmetrical yet I know I have hands given the relation I stand to my perceptual evidence; but my TWIN doesn't stand in the same evidential relation to his hands as I do to mine. Therefore, indiscriminable situations are not always normatively identical. Therefore, NT is a false principle. Granted, we are both *blameless* in believing we have hands. So, we are in the same normative boat relative to blameworthiness. But this is only one kind of normative evaluation, which leads to the third problem with internalist normative evaluations: namely, blameworthiness alone

[87] Timothy Williamson, "On Being Justified in One's Head," p. 120.

[88] Identity has three characteristics: it's transitive, reflexive, and symmetrical. Regarding transitivity, if X is identical to Y and Y is identical to Z, then X = Z. Regarding identity having reflexivity, if some relation R is reflexive, then for all X, X stands in R to itself. For example, consider the relation of 'being the same size as' as opposed to 'being larger than'. Only the former is a reflexive relation because X can be the same size as itself. But X can't be larger than itself at the same time and in the same relation without losing identity since it wouldn't be the same thing. Identity is symmetrical because if some relation R is symmetrical, then X stands in relation R to Y just in case Y stands in relation R to X. Put differently: for any X and Y, if X is identical to Y, then Y is identical to X. For example, if Gary Osmundsen is identical to Jackie Garofalo's son, then Jackie Garofalo's son is identical to Gary Osmundsen.

is inadequate to account for all the normatively relevant dimensions of a belief's positive epistemic status.

1.2.2 Blameworthiness Doesn't Exhaust Normative Values

Third, let us consider a case that shows why evaluating for blameworthiness alone is not enough to account for all the relevant facts needed for an adequate account of a belief having a positive pd-normative status. The first way to get at the different kinds of normative evaluations there are is to consider the following kind of moral case.[89]

> BRAIN SURGEONS: Linda and Lorna are both surgeons scheduled to remove malignant tumors lodged deep in their patients' brains. They both have scrupulously perused over their patient's medical history and all relevant information, including MRI scans, CAT scans, the best guidelines on how to perform the surgery, etc. Both Linda and Lorna are well rested, prepared, in good working condition, and have taken all necessary steps to ensure the best possible outcome. Both surgeons skillfully perform the surgery the best they know how with all the relevant available information they have to work with. Unfortunately, despite both Linda and Lorna preparing and performing in virtually identical ways, Lorna loses her patient while Linda saves hers.

For starters, let's be clear: it's obvious that Lorna is *not* blameworthy for losing the patient. She has done all that is expected of her. So, she need not feel shame, regret, or the need to apologize. She may, perhaps, feel regret and be obligated to explain to the patient's family what happened. But Lorna has what Timothy Williamson calls a "cast-iron excuse" for losing the patient, and thus, is exempt from any *blame* for the patient's death.[90] Regarding Linda, there is no need for her to have a cast-iron excuse given she doesn't need to be excused for anything: her patient survived.

---

[89] This case was inspired by a similar one put forth by Robert Audi. See Robert Audi, "An Internalist Theory of Normative Grounds," in *Philosophical Topics* 29 (2001): 19-46.

[90] Timothy Williamson, "On Being Justified in One's Head," p.116.

But just because both Linda's and Lorna's actions are blameless, it doesn't follow that there is not a normative difference between the two.[91] There are other normative values different evaluations can pick out. For example, we can say the outcome of Linda's surgery was in fact *better* than Lorna's. Regardless of what kind of luck you attribute the difference between Linda's and Lorna's cases to, the fact is, Lorna's patient died as a result of her actions.[92] Suppose Lorna decided not to operate. Perhaps she thought the surgery

---

[91] I am indebted to Timothy Williamson for both making this point and fleshing out some of the other normative evaluations on offer. See Timothy Williamson, "On Being Justified in One's Head," p. 116.

[92] The different kinds of luck I have in mind come from Thomas Nagel who describes four kinds of moral luck: namely, resultant luck, constitutive luck, circumstantial luck, and antecedent luck. First, Nagel appeals to Kant's example of a truck-driver who failed to get a routine brake maintenance check-up for his truck. As a result, the truck's brake system failed thus preventing the truck from running over a child. Or consider Nagel's case involving two drunk drivers who acted identically insofar as driving home drunk and swerving onto the sidewalk. The only difference between the two is that one driver had the bad luck of hitting a pedestrian who happened to be on the sidewalk where he swerved. Resulted luck, then, captures the luck involved in the way factors outside one' control change the way one's actions or projects turn out. Second, constitutive luck captures the luck involved in how one is constituted—that is, one's inclinations, temperament, capacities, and dispositions, which are constitutive of one's character. But such traits are sometimes the result of our natural endowment or determined by the people we learned to emulate while growing up. Either way, the vices we get for free or acquire are the result of factors outside our control. Moral luck, then, seems to undermine our moral evaluations of blaming others for their vices and praising others for their virtues given such character traits were the result of factors outside our control. Nagel says "Yet people are morally condemned for such [vices], and esteemed for [virtues which are equally both] beyond control of [one's will to compensate for them in action]: [for we morally evaluate people] for what they are *like*" (p. 299). Third, circumstantial luck captures the circumstances we find ourselves in, which determines the kind of actions we can make, as well as influence our choices and actions. E.g., situationism says experimental social-psychological studies suggest behavior is influenced more by trivial environmental factors rather, than, robust character traits. Nagel points out how the people of Germany, under the Nazi regime, had to face moral decisions most people never had to face given their geographic circumstances. Yet we morally judge the German citizens for sins of omission or commission (or both) that many other people would have been guilty of had they been in such circumstances. Yet we don't morally judge others for what they would have done; we morally judge others for what they do or didn't do, yet their circumstances—where they are geographically—contribute to what one actually does. Lastly, Nagel uses antecedent luck to capture the problem of free will and determinism. If one's antecedents in the form of constitutive luck (e.g., one's temperament) or circumstantial luck (e.g., one's present whereabouts) or both, determine one's character or one's geographic location, then what one is like and what one does is due to factors beyond one's control. Nagel asks, "how can one be responsible [for who they are and what they do] even for the stripped-down acts of the will itself, if *they* are the product of antecedent circumstances outside of the will's control? (p, 300). In other words, if all there is is event-event causation and no agent-causation, then even pure acts of will are governed by involuntary desires, temperaments, and beliefs constitutive of our character that we have no control over. Given the universe is determined (or even if it's indetermined), we are still properties of it and our internal view of being a substance with the power to bring things about is false. We are aggregates of properties or property-things. So our internalist view of ourselves having the ability to do otherwise is false. This informs Nagel's error theory for why we think the control principle is true: namely, we are morally assessable only to the extent that what we are assessed for depends on factors under our control and we are the kind of things that can, on a conceptual level, be legitimately evaluated morally. And this tension between (i) our internal views of ourselves as substances having agent-causation that informs our pre-theoretical

was too risky moments prior to making the first incision. Williamson says, "Of course, the surgeon followed the best available guidelines in deciding to intervene, and operated with state-of-the-art skill and technology. Nevertheless, it turns out that it would have been better for the patient if the operation had not been performed."[93] For despite Lorna having an excuse from being culpable of any blame whatsoever, there is still logical space to make the following kinds of normative evaluations.

First, as mentioned above, the decision to operate, despite Lorna's decision being blameless, turns out, unfortunately, to be *wrong*. It's wrong in the sense that the patient would still be living had Lorna decided not to perform the surgery. Second, Lorna would have been *more praiseworthy* had the surgery resulted in saving the patient's life. Alternatively: Lorna is *less praiseworthy* given the surgery resulted in the patient's death rather than her survival. Third, we can say that despite Lorna's excused failure to save the patient, it doesn't follow that excusability is normatively equivalent to *success*.[94] Linda's actions led to an achievement whereas Lorna's didn't. Indeed, Lorna failed to acquire Linda's type of achievement. Fourth, again, despite Lorna's iron-clad excuse, we can still say that Linda *succeeded* in fulfilling her *obligation* to do no harm to the patient whereas Lorna didn't succeed. Not only did Linda achieve success in removing the malignant tumor, she also didn't cause her patient any significant harm unlike Lorna. Fifth, we can say that Linda

---

intuitions on morality and its incompatibility with (ii) our external view of ourselves as property-things with only event-event causation needs to be taken on the chin and lived with. Our conception of the nature of morality, then, needs to be revised—our counter-intuitive and visceral responses to it notwithstanding. See Thomas Nagel, *Mortal Questions* (Cambridge: Cambridge University Press, 1979), Chpt. Three.

[93] Timothy Williamson, "On Being Justified in One's Head," p. 116. Of course, many of the pertinent details of this case have been left out. It shouldn't be too hard, however, filling in the kind the details which present a scenario where a cost benefit analysis of the pros and cons of operating and not operating were negligible. Even if there was a potential benefit to getting the surgery; the surgery wasn't necessary to save the patient's life.

[94] Ibid., p. 117,

deserves more *praise* than Lorna for successfully acquiring such an achievement and fulfilling her obligation.

One might reasonably take umbrage here. One might say, perhaps, that such normative evaluations veer away from features that Linda and Lorna have internal access to, which makes such evaluations unfair. Because Lorna can only act on the limited basis of information she has access to, her ironclad excuse covers a multitude of negative normative evaluations, including the ones mentioned above. Furthermore, such evaluations are made without acknowledging the role luck plays in Lorna's operation relative to the outcome in Linda's operation. Indeed, most share the intuition that the control principle is true—that is, we are morally assessable only to the extent that what we are assessed for depends on factors under our control.[95] And the death of Lorna's patient *is* the result of factors outside her control. Fair enough.

But this only strengthens the point I want to make: namely, the "epistemic net," so to speak, is larger than what principles 1-3 allow for it to be.[96] There are other epistemic goods and values that lie beyond the scope of such principles. Granted, there are some ways Linda and Lorna are in the same normative boat as a result of NT, especially when it's a version of ACCESS. But there are other epistemic values that are relevant to other kinds of epistemic evaluations. And trying to carve out a domain for the epistemic (that is completely accessible to the subject's internal perspective) fails to take into account other epistemically relevant facts.

After all, it fails to take seriously how an externalist accounts for and utilizes facts about reality that are not always accessible to the subject. More exactly: there is a difference

---

[95] Dana K. Nelkin, "Moral Luck," *The Stanford Encyclopedia of Philosophy* (Summer 2008 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/entries/moral-luck/

[96] I'm grateful to Wayne Riggs for helping clarify this point.

between (i) all the facts of the case and their normative relevance and (ii) the subject's position to internally access these facts. And any adequate treatment of the normative domain can't carve out space which leaves out these other important dimensions. Because the human epistemic condition is such that we only have limited access to a set of appearances that don't exhaustively capture all the normatively relevant facts, we do well to learn to how to live with this fact about our cognitive predicament just as we have in other domains of inquiry. Indeed, assuming something like ACCESS can account for *all* the normatively relevant facts in cases involving *peer disagreement* is not only intellectually dishonest but, in addition, fails to be consistent with how we apply the appearance-reality distinction to virtually every other domain.[97] Thus, a robust, realist view of the normative domain must appreciate and deal with the fact that all normative facts are not perfectly accessible from an internalist perspective.[98] Therefore, subjects that are in indiscriminable situations from their internal point of view are not in normatively identical situations.

1.3 Why CCE Employs An Internalist Account Of Normative Evaluations To Cases of Peer Disagreement

Let's take stock of what we've learned about internalism and see how its motivations apply to CCE. So consider an extreme case that showcases the epistemically relevant features proponents of CCE fixate on. The following case involves conflicting thermometers which illustrates how two disputants are indiscriminable in the ways I previously discussed above; for appealing to disagreeing thermometers is about as explicit as we can get relative to the

[97] I'm indebted to Timothy Williamson for helping me see these points and how inconsistent our methodology can be when it comes to the normative domain.

[98] This is why the principle Williamson deemed normative tolerance is false. Recall normative tolerance says indiscriminable situations are normatively identical. But because identity is reflexive, then normative identity is reflexive. However, indiscriminability is not reflexive because there are relevant normative facts that go beyond what is internally accessible. I know I have hands, but my TWIN doesn't. Furthermore, you can have indiscriminate situations as illustrated in the BRAIN SURGEONS case, yet there are still normative differences between the two surgeons.

kind of indiscriminability that motivates the kind of epistemic symmetry believed to be shared between two epistemic peers that disagree and have fully disclosed to each other their evidence.

> DISAGREEING THERMOMETERS: You and I are a bit odd in that we both carry in our pockets thermometers like most people carry cell phones. We both take care of our thermometers and frequently consult each other's thermometer from time to time for fun and to test for accuracy. We each take each other's thermometer to be reliable on the basis of their track records. Today, like many times in the past, we ask each other what the other's thermometer indicates the ambient temperature to be. So we take out our thermometers, place them in virtually the same space, and allow them to respond to the environmental inputs. After sufficient time has lapsed to get an accurate reading, we consult each other's thermometers and discover that they conflict. Your thermometer reads 73° F; whereas mine reads 68°F.

When we consider how to respond to this sort of case, withholding judgment seems to be the intuitively correct one. Until we discover which thermometer is defective or discover that, perhaps, both are defective, given our extant evidence, we should withhold judgment and remain agnostic about the room's exact temperature. This is the driving intuition behind proponents of CCE. For in cases like this digging my heels in and believing the temperature is 68°F just because *my* thermometer says so reeks of epistemic irresponsibility. Such a display of epistemic egoism is surely epistemically irresponsible.

But if this prescription seems intuitively right for conflicting thermometers, then why not for cases involving disagreement between epistemic peers? Recall back to Chapter One where I laid out the three characteristics of two people functioning as epistemic peers: namely, (i) evidential equality; (ii) cognitive equality; and both disputants having engaged in (iii) full disclosure.

> (i) *Evidential equality*: S and S* are evidential equals relative to the question whether *p* when S and S* are equally familiar with the evidence and arguments that bear on the question whether *p*.

> (ii) *Cognitive equality*: S and S* are cognitive equals relative to the question whether *p* when S and S* are equally competent, intelligent, and fair-minded in their assessment of the evidence and arguments that bear on the question whether *p*.
>
> (iii) *Full disclosure*: S and S* are in a situation of full disclosure relative to the question whether *p* when A and B have knowingly shared with one another all of their relevant evidence and arguments that bear on the question whether *p*.

Notice how we can apply these characteristics to disputants consulting their thermometers. Regarding evidential equality, both my disputant and I can access each other's thermometer to see what they indicate the temperature to be. Furthermore, both thermometers are exposed to the same environmental inputs given we see them both occupying the same region of space. Both disputants are cognitive equals insofar as we both take each other to be equally competent, intelligent, fair-minded, and the like. In short, we both take each other to be reliable at getting the truth—specifically, in this case, the room's temperature. And both disputants have fully disclosed their evidence—that is, both disputants have allowed the other to read the other's thermometer. Proponents of CCE think that what seems clear from DISAGREEING THERMOMETERS generalizes to other cases of disagreement far and wide. Thus, it's easy to see why epistemologists are so attracted to the CCE position given it offers an intuitively correct model that provides straightforward verdicts on how to respond to disagreement.

In the remainder of this section I argue that this kind of model breaks down. More specifically, the problem is the intuitive evaluation from DISAGREEING THERMOMETERS is merely about blameworthiness. But as I will argue the normative status that we're interested in is *not* only about blameworthiness.

1.3.1 ACCESS and NT Work Well For Thermometers But Not For Cognitive Agents

Access internalism is motivated by the New Evil Demon Problem, which captures a dimension of epistemic symmetry shared by both subjects in the good and bad cases. Similarly, CCE gets motivated because it's believed that both disputants share the same dimension of epistemic symmetry. Notice, then, how the same kind of intuitive evaluation is employed in both the New Evil Demon case and DISAGREEING THERMOMETERS. The subjects' beliefs in both the former and latter cases are believed to share the same kind of epistemic status as a result of principles like ACCESS or NT *governing* the evaluation. In the former case, both subject's beliefs are not epistemically blameworthy. And in the latter case, both subject's beliefs will be blameworthy if they don't withhold judgment. But the evaluation that both cases prompt is one of blameworthiness. And this is a key piece of datum that ACCESS and its variant NT capture which govern such normative evaluations. So whether the intuitive evaluation is that one has committed an epistemic sin of omission (if we don't withhold judgment in cases of disagreement) or an epistemic sin of commission (when we stick to our guns in cases of disagreement), CCE evaluations are capturing two sides of the same coin of blameworthiness. And such a coin is forged by internalists using ACCESS, NT, or (restricted) versions of MENT when making normative evaluations.

It's important now to draw from the lessons we learned in (§1.2) above in order to see how internalist evaluations are too constrained to pick out some of the *other* normative facts that are relevant to peer disagreement. There are many other ways to evaluate others epistemically other than for blameworthiness. For example, just as my TWIN and I share the same appearance of the epistemically relevant evidence captured by ACCESS, proponents of CCE similarly take epistemic peers to share the same epistemically relevant

evidence. And just as my TWIN and I are taken to be in evidentially indiscriminable situations so, too, are both disputants taken to be in evidentially indiscriminable situations. Such evidential indiscriminability comes from testimony once both disputants fully disclose to each other their evidence. But the upshot is the same: both disputants are in evidentially indiscriminable situations thereby making them both normatively identical.

Fixating on blameworthiness is what makes the verdict for cases like DISAGREEING THERMOMETERS so attractive.[99] Both you and I can access the same evidence for the room's temperature by consulting each other's thermometer. We are now evidentially indiscriminable. Similarly, the proponent of CCE asks why treat cases of disagreement between epistemic peers who have fully disclosed to each other their evidence any differently? Recall the case involving the dinner bill from Chapter One. This case aims to show how evidential indiscriminability easily obtains thus making the disputants in normatively identical situations thereby preventing both contrary beliefs from having positive pd-normative status.

> DINNER BILL: 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 desert, or drank more of the wine. I do the math in my head and become highly confident that our shares are $43 each. Meanwhile, my friend does the math in her head and becomes highly confident that our shares are $45 each. How should I react, upon learning of her belief? [100]

---

[99] Another way of thinking about the thermometer model is to imagine the case involving us somehow lodging our thermometers into our brains so that it becomes integrated with the rest of our cognitive system. Our beliefs, then, would be sourced in our thermometers. And we both would be in evidentially indiscriminable situations as a result of it. The same verdict of withholding judgment should apply. Were we to do otherwise and, say, retain our beliefs, then we would be epistemically blameworthy. I thank Thomas Kelly for sharing this point with me in our correspondence via email.

[100] David Christiansen, "Epistemology of Disagreement: the Good News," in *The Philosophic Review,* 116 (2007): 193.

Notice how we intuitively take evidential indiscriminability to exist between the disagreeing peers as a result of full disclosure. Because we can't access where the epistemic defect is or evidential difference lies, we are in normatively identical situations. To do anything but withhold judgment would be epistemically blameworthy. Neither belief has positive pd-normative status. So, just as the epistemically responsible thing to do in DISAGREEING THERMOMETERS is to withhold judgment so, too, is it epistemically responsible to withhold judgment in DINNER BILL. But the worry is that proponents of CCE think cases like these can generalize far and wide.

There are three problems, however, with how proponents of CCE think evidential indiscriminability is shared between disputants, which is necessary for a principle like NT to go through. First, proponents of CCE allow ACCESS to govern their evaluations. But as we saw above, ACCESS can only inform normative evaluations concerned with blameworthiness. A domain of normativity that only evaluates for blameworthiness is, however, anorexic. As we saw above, it is inadequate because it fails to capture other normative evaluations besides blameworthiness. Indeed, the evidence may, in fact, support one disputant's belief and not the other's in which case the former did a *better* job of weighing and processing the evidence than the latter. The disputant who got the truth gets *more praise* for this achievement; *succeeded* in getting the truth; and fulfilled the epistemic *obligation* to get the truth. Furthermore, we do well not to evaluate a disputant as blameworthy for believing what she arguably knows. Proponents of CCE would have us avoid charges of blameworthiness at the price of giving up what we already know. This is, indeed, a strange result. So, unless the disagreement itself is an undefeated-defeater, most interesting cases of disagreement will involve facts that can explain the epistemic inequality between disputants. But either way, we are already invoking externalist's concerns.

Therefore, utilizing a principle like ACCESS allows the psychological evidence of disagreement alone to virtually swamp all the other kinds of evidence in play which seems misguided. (And, again, it's not that normative evaluations shouldn't be made based on what a disputant has internal access to; it's just what a disputant has access to *should not be restricted to what is internally accessible.*)

Second, evidential indiscriminability depends upon a conception of evidence that overlooks the epistemic import of the totality of a disputant's beliefs, knowledge, starting points, and epistemic goals. Such things all contribute—at some level and degree—towards weighing the targeted proposition on evidence E that's getting fully disclosed. That is to say, the scope of relevant evidence in support of the targeted proposition will extend beyond the species of evidence that's operative when both disputants disclose to each other their evidence. One's background beliefs, knowledge, or worldview will ineluctably color the species of evidence that gets disclosed. To suggest otherwise is psychologically implausible. For each disputant will have her own subjective probability of how strong each *disclosed* piece of evidence should be weighed given her other background beliefs, knowledge, starting points, and epistemic goals. Indeed, disputants who share the same disclosed evidence may nevertheless *rationally* (in the epistemic sense) assess it differently based on their involuntary dispositions towards certain risks involved in believing, disbelieving, or withholding judgment towards the disputed proposition.[101]

Because there are different theories of evidence, there will be different species of evidence that influence the weighting of the disclosed evidence. This is why the

[101] For a paper that argues against the implausibility of assuming an objectively correct way to weight the probability of p on evidence e in all cases to establish a notion of epistemic rationality see Wayne D. Riggs, "Epistemic Risk and Relativism," in *Acta Analytica* 23 (2008): 1-8.

Uniqueness thesis rarely applies to cases of disagreement yet is crucial for the CCE position to go through.

Uniqueness: If an agent whose total evidence is E is fully rational in taking doxastic attitude D to P, then necessarily, any subject with total evidence E who takes a different attitude to P is less than fully rational.

Notice this thesis requires a species of evidence that is universally accepted. But there isn't any such notion of evidence. Yet, such universally accepted species of evidence is required to even get evidential indiscriminability off the ground, which proponents of CCE need in order to motivate the kind of normative evaluation they make relative to a belief's pd-normative status. And the best shot at satisfying Uniqueness is to stipulate a notion of evidence that satisfies ACCESS, which ignores the epistemic import a disputant's beliefs, knowledge, starting points, and epistemic goals will have on the disclosed evidence.

What's interesting is that proponents of CCE permit psychological evidence to be a member of the set of disclosed evidence. The disagreement itself, then, is additional evidence shared by both disputants. This is where Feldman's view of "evidence of evidence is evidence" kicks in. Recall Feldman says,"[E]vidence [i.e., disagreement—specifically, a disputant believing p while you believe not-p] that there is evidence for P is evidence for P."[102] That is, the fact that a disputant still disagrees is evidence that there is additional evidence for not-p. It seems, then, that ACCESS would permit one disputant to disclose to another disputant a piece of private evidence such as an intellectual seeming or intuition. Such evidence, perhaps, could be grounds for discovering an evidential difference that would no longer make two disputants evidentially indiscriminable. But proponents of CCE block this move and retain evidential indiscriminability by either neutralizing or preventing such evidence to weigh in by condemning it as an act of blatant bootstrapping. They

[102] See Richard Feldman, "Reasonable Religious Disagreements," Louis M. Antony (ed.) *Philosophers Without Gods: Meditations On Atheism And The Secular Life* (Oxford: Oxford University Press, 2007), p. 208.

reason that one is arbitrarily giving one's private evidence special status and failing to treat like cases alike given each disputant has her own private evidence. Appeals to private evidence, then, cancel each other out or, at least, neutralizes its evidential import thereby preserving evidential indiscriminability.

Setting bootstrapping objections aside, the ease with which one can neutralize a disputant's intellectual seeming or intuition, on the one hand, and one's background beliefs, knowledge, starting points, and epistemic goals from influencing how p on E gets weighed, on the other hand, is particularly worrisome. Indeed, I take this to be one of the major flaws with CCE. For it assumes that such background items can be neutralized from influencing the disclosed evidence E. Such background items, however, will ineluctably prevent two disputants from reaching evidential indiscernibility. (This, of course, is another reason why ACCESS is too restrictive when it governs normative evaluations.) And trying to cancel out the epistemic import such background items have on the evidence by saying the other disputant has them, too—so evidential indiscriminability is preserved—overlooks how psychologically implausible such an account is. For it suggests such items will not get factored in at the conscious or subconscious level, or some combination of both. It's not clear to me that one can even have the kind epistemic wherewithal to control how one's background beliefs, knowledge, starting points, and epistemic goals will weigh whatever disclosed evidence E turns out to be after full disclosure.

Third, the ACCESS thesis isn't sufficient to establish evidential indiscriminability once we extend the competing kinds of relevant evidence to that which (say) externalists endorse. Put differently: ACCESS depends upon a universally accepted view of evidence which, unfortunately, simply doesn't exist. In Chapter Three, I develop this reason further by showcasing all the different notions of evidence that elude ACCESS. The upshot is that

evidential indiscriminability rarely ever obtains thereby preventing each disputant from being normatively identical (full stop).

1.4 Summary

I grant there are some cases where such universally accepted kinds of evidence are, perhaps, operative. Again, this is why cases like DISAGREEING THERMOMETERS initially make CCE an attractive position. For if two disputants each had, say, their own thermometers that shared all the same properties except for the fact that they indicate different temperatures, then there is nothing preventing the two disputants from accessing and fully disclosing all the relevant evidence E. E, then, will be the kind of universally accessible evidence that allows ACCESS to obtain.[103] This is why we think both disputants are in evidentially indiscriminable situations; neither has access to some evidential difference or epistemic defect that explains the variance. It's assumed, then, that both disputants are just as likely to be in error were they to believe their thermometer's reading and retain their belief in the face of disagreement. It follows that since both disputants are in evidentially indiscriminate situations they are both normatively identical, which is why Uniqueness can't be violated. The the same kind of "universal" evidence seems to be, perhaps, in play in *some* cases like DISAGREEING THERMOMETERS and DINNER BILL.[104] But this kind of discursive evidence will not capture *all* the relevant evidence that's

[103] But even this case of disagreement depends on the relation between the evidence and targeted proposition to be abstracted away from background beliefs, knowledge, and epistemic goals. But if one's prior beliefs, knowledge, and epistemic goals do weigh in to the probability calculation in some nomologically necessary way, then it's an inaccurate way of describing and evaluating the case. In other words, Tom Kelly points out how many proponents of CCE assume the evidential support relation evidence E has on p is a two-place relation (e.g., p/E) rather than a three-place relation (e.g., p/E relative to background Z). See Thomas Kelly, "Evidence Can Be Permissive," in Matthias Steup, John Turri, and Ernest Sosa, eds., *Contemporary Debates in Epistemology*, 2nd ed., (Oxford: Blackwell, 2013).

[104] For example, in DINNER BILL, the numbers on the bill will be akin to a universally accepted kind of evidence needed for ACCESS to obtain.

operative in many other cases we're concerned about. So, even if we can find cases of disagreement where, perhaps, something like a universally accepted kind of evidence is the only operative and relevant kind, it doesn't follow that such a notion of evidence will be the only game in town in many other cases of disagreement, especially ones where the disputant's error or defect is hard to discover. This will result in the "normative net" expanding to include different epistemic values so as to allow a disputant's belief to have positive pd-normative status once a different kind of normative evaluation is made.

To sum up: proponents of CCE assume there is a universally accepted kind of evidence that allows the epistemology of disagreement to be operational insofar as prescribing how disputants should respond to peer disagreement so as to avoid being epistemically blameworthy. But because such notions of evidence are an impossible ideal, evaluating for blameworthiness by using principles like NT, ACCESS, or NT is misguided. If one of our epistemic values is getting the truth, then we do well utilizing *all* of our evidence, beliefs, and knowledge. Therefore, issues involving a belief's pd-normative status depend upon normative evaluations that acknowledge the appearance-reality distinction and other externalist concerns. But then such evaluations can't be governed by principles like ACCESS, NT, and (restricted forms of) MENT. Therefore, a belief's pd-normative status and the kind of normative evaluations epistemologists should make is not orthogonal from externalist concerns. And any normative prescriptions that requires one to give up what one arguably already knows in order to avoid being epistemically "blameworthy" or preserve "rationality" or avoid "irrationality" or "intellectual irresponsibility" is in serious trouble. Therefore, externalism is motivated. The epistemic significance of disagreement, then, shifts to a new challenge for externalism: namely, externalists need to provide an account of defeating evidence that's generated by peer

disagreement. More specifically, externalists need to account for the kinds of potential defeaters peer disagreement can level against knowledge. I will address this problem later in Chapter Four. But for now, I move on to Chapter Three wherein I argue that the normative domain carved out by CCE is not only impoverished, but is vulnerable to some damning objections.

3

Why Epistemic Disagreement Is A Pseudo-Problem: Part 2

'How am I to prove now that 'Here's one hand and here's another?' I do not believe I can do it. In order to do it, I should need to prove for one thing, as Descartes pointed out, that I am not now dreaming. But how can I prove that I am not? […] I have conclusive evidence that I am awake; but that is a very different thing from being able to prove it.

~G.E. MOORE

Whatever evidence is, one is not always in a position to know what one has of it. […] [F]or although we could show that one's evidence in the good case had the same [phenomenal] properties appropriate to the [phenomenal evidence in the] bad case, we could not show that one's [total] evidence in the good case as one's [total] evidence in the bad case. Indeed, we could not show that one was always in a position to know which properties of evidence were appropriate to one's own case.

~TIMOTHY WILLIAMSON

In this chapter, I argue that the problem of peer disagreement is a pseudo-problem for another independent reason than the one previously given in Chapter Two. This argument is more general in scope. For it demonstrates why one shouldn't assume the internalist perspective is the correct one to adopt for all epistemic evaluations of cases involving "peer" disagreement. For although an internalist evaluation along the lines of blameworthiness is one kind of normative evaluation, it is not the only one; moreover, the kind of conditions needed to establish epistemic peerhood, which is required to get the intuitive blameworthiness-evaluation to go through, is rarely ever satisfied. The latter is damning, indeed, because the internalist model depends upon both disputants being evidential equals in order to capture intuitive evaluations that fixate on blameworthiness. The argument in section two highlights the fact that there are a great number of possible asymmetries between two or more disputants, which makes it very unlikely that any two such disputants would ever actually meet the demands for being "epistemic peers." As a result, the conditions needed to satisfy the kind of epistemic equality motivating the

skeptical threat are rare. Hence, the skeptical threat raised by actual peer disagreement is mild at best.

1. WHY INTERNALISM IS NOT THE ONLY GAME IN TOWN

In this dissertation I'm not assuming the truth of internalism. Rather, I explore the possibilities an externalist approach to peer disagreement opens; moreover, I aim to show why an externalist approach has real advantages over internalist ones.[105] In this section I will reformulate Greco's argument against knowledge- and justification-internalism into one that is directed more specifically at CCE. More specifically, I will argue that the normative resources CCE employs to govern their normative evaluations is severely impoverished. This is because the normative evaluations made by proponents of CCE are governed by principles that are too restrictive. Accordingly, CCE's prescribed epistemic responses lack any normative property that we *ultimately* care about. Even more specifically, I argue that any account of a belief's pd-normative status that excludes externalist concerns and its relation to knowledge is uninteresting (at best) and false (at worse).

The following arguments aim not only to show that justification and knowledge are not best understood within an internalist framework but also generalizes to concerns relative to a belief's pd-normative status. For one can form a belief in a vicious way and forget the relevant details for how one formed the belief *negligently*. Such facts, however, are not only epistemically relevant, but will elude any internalist normative evaluation. Indeed, as I will argue below, the kind of constraints governing internalist normative evaluations prevent one from making evaluations that discriminate between, say, beliefs formed from

[105] Trying to establish the truth of externalism over internalism relative to a belief's epistemic status in such a way that all internalists will be convinced is not only naively ambitious, but arrogant. I don't intend to settle this debate once and for all in a section or two of this dissertation. This is why I will be assuming the truth of externalism regarding what constitutes a belief's pd-normative status. I still think, however, I'm required to motivate my reasons for moving the discussion about a belief's pd-normative status vis-à-vis disagreement into a domain that allows externalist concerns and matters involving knowledge to weigh in.

apparent memories rather than *real* memories. Furthermore, internalist evaluations fail to capture the difference between a belief formed in conjunction with *having* good reasons versus a belief formed *from* good reasons. Nor can internalist evaluations account for epistemically relevant facts relative to both bare memorial beliefs and beliefs formed by sub-personal processes. But such facts about a belief's etiology *are* epistemically relevant to informing our normative evaluations. So, a belief's epistemic status shouldn't be abstracted away from such externalist concerns and restricted to a normative domain governed by internalist's constraints.

1.1 Greco's General Argument Against Internalism

Recall back to the BRAIN SURGEONS case. This case involved normative evaluations concerning Linda's and Lorna's performances. Some evaluations were made from a subjective point of view; others we made from an objective point of view. From the objective point of view, I was concerned with whether or not the two surgeons' performances were successful. From the subjective point of view, I was concerned with whether the two parties were blameworthy for the outcomes. Similarly, when we evaluate the beliefs of two peers who disagree, we can either evaluate them from an objective point of view or a subjective one. The subjective point of view is concerned about questions regarding whether or not a subject's belief has positive pd-normative status. For example, what is the rational response or reasonable response or responsible response when facing the evidence of a peer who disagrees with you? But the objective point of view is going to be concerned about which peer's belief is true. Thus it will be concerned with other related matters such as whether or not the belief was reliably formed or whose evidence is objectively accurate.

The problem is that normative evaluations from a subjective point of view are either uninteresting or false if they're abstracted away from certain external facts. For example, consider the relationship the notions reasonableness and responsibility have with the notions of praise and blame. When someone pushes me into you thus causing your drink to spill, I am not blameworthy for staining your shirt. Why? Because I'm not responsible!—the person who aggressively dashed to the desert table, impolitely bumping into others along the way, is responsible—he's blameworthy! The same relationship holds between epistemic responsibility and praise and blame. For evaluations concerning whether or not one's belief is an epistemically responsible one are akin to evaluations concerning whether or not one is epistemically blameworthy for forming the belief. Greco says the upshot, then, is this: "whether a person is epistemically blameworthy for holding some belief is partly a function of the person's prior behavior: if S's reasons for believing b are the result of *prior negligence*, then S is not now blameless in believing b [regardless of whether or not S remembers being negligent]."[106] Here's a case he takes to flesh out this point.

> DEAN MARTIN: Maria believes that Dean Martin is Italian. She believes this because she seems to remember clearly that it is so, and she presently has no reason for doubting her belief. But suppose also that Maria first came to this belief carelessly and irresponsibly (although she has now forgotten this). Many years ago, she formed her belief on the basis of testimony from her mother, who believes that all good singers are Italian. At the time Maria knew her mother was an unreliable source in these matters, and she realized that it was not rational to accept her mother's testimony.

The point of this illustration, according to Greco, is that it should be obvious that despite Mary's failing memory, she is still blameworthy for forming this belief given her prior negligence. It's also important to see that Mary's prior negligence still factors into the normative evaluation despite such negligence being absent from what she can internally

[106] John Greco, "Is Justification Internal?: Justification Is not Internal," in Matthias Steup, ed., *Contemporary Debates in Epistemology* (Oxford: Blackwell, 2005), p. 261(emphasis mine).

access. In other words, just because I forgot about doing a blameworthy act in the past—thereby preventing me from accessing it from my present, subjective perspective—it doesn't follow that I am no longer responsible for having done it.[107]

There is another relevant distinction Greco makes between apparent and real memories. This distinction provides additional reasons for why the internal states one has reflective access to are not sufficient to inform the kind of interesting and accurate normative evaluations we care about. He thinks the following cases capture this distinction.[108]

> APPARENT VS. REAL MEMORY: Case 1. Sofia believes that the [Phillies] won the World Series the year she was born, because she remembers that her parents have told her this. Indeed they have told her this, reliably reporting an important fact about her early childhood. Case 2. Sofia believes that the [Phillies] won the World Series the year she was born, because she believes she remembers that her parents have told her this. In fact, they have never told her this, although she has a "false memory" that they have.

This distinction points out how the facts Sofia has internal access to are the same in both cases. The upshot is that we cannot internally access all the relevant facts that comprise both a belief's etiology and the kind of objective relation our evidence or source of our belief has to the world. But such facts *are* interesting and relevant when making accurate normative evaluations.

---

[107] This raises metaphysical concerns over the nature of personal identity, responsibility, and justice. Pursuing arguments as to whether or not people should be punished later in life for crimes they don't even remember committing takes us beyond the scope of this section of the dissertation. Furthermore, it surfaces, again, the messiness of trying to carve out a normative domain where internal and external concerns don't conflict. This is a problem not just for the nature of normativity, but an implication of our limited epistemic condition where we have access to only a related subset of facts in contrast to the set of facts that are modally tied, somehow, to the word. Be that as it may, it is still an interesting and related topic deserving of careful thought and reflection.

[108] John Greco, *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity* (Cambridge: Cambridge University Press, 2010), p. 52.

In addition to the role prior negligence plays in making normative evaluations, Greco also considers the implications of the following distinction in order to show why normative evaluations require incorporating externalist concerns. The distinction is this: there is a difference between (a) *having* good reasons for what one believes, and (b) believing *for* good reasons. So consider Greco's case involving a wishful thinker.[109]

> WISHFUL THINKER: Charlie is a wishful thinker and believes that he is about to arrive at his destination on time. He has good reasons for believing this, including his memory of train schedules, maps, the correct time at departure and at various stops, etc. However, none of these things is behind his belief—he does not believe what he does because he has these reasons. Rather, it is his wishful thinking that causes his belief. Accordingly, he would believe that he is about to arrive on time even if he were not.

From this case, it should be intuitive that Charlie's belief is not praiseworthy. Rather, because his belief is the result of merely having good reasons, it's still deserving of epistemic blame. What's more, this case points out that a belief's etiology matters insofar as what considerations inform our normative evaluations. More exactly, facts about a belief's *etiology* and whether the belief was formed *from* good reasons loom large in how we make accurate normative evaluations. These facts are external concerns that go beyond the scope of what one has subjective, internal access to, which raises significant problems for carving out a normative domain that is exclusively internal to one's reflective perspective alone.

1.2 The General Argument's Bearing on CCE

So far I am trying to argue that externalist concerns do and should factor in any notion of a belief's epistemic status (in general) and into how it's applied to cases involving peer disagreement involving a belief's pd-normative status (in particular). In this section I intend

[109] John Greco, "Is Justification Internal?: Justification Is not Internal," p. 261.

to show that any account of a belief's pd-normative status that abstracts away from external concerns is implausible. I begin by considering a case of disagreement.

> YIELDING TO CCE: You and I both begin our day with a tall cup of coffee at our favorite café. There is one thing we aim to achieve today and that's deciding whether or not all the evidence we've been collecting for the last three months supports the truth of hypothesis h. We begin by disclosing to each other all the reasons, arguments, and evidence we have that bears on h. We each take the time to summarize our findings and its bearing on h. Fortunately, we learn that we both agree that our shared evidence supports h. So we both believe that h is true given our evidence, yet you're slightly more confident than me. Your credence level is .9 and mine is .7. Because we both subscribe to CCE's prescriptions, we recognize that on the pains of irrationality, we should split the difference and both adjust our credence levels to .8.

According to CCE, we have done the epistemically correct thing; both our beliefs now enjoy a positive pd-normative status. The decision to split the difference, after discovering ourselves in a case of peer disagreement, is to be praised as personifying epistemic responsibility. In short, we are exempt from any epistemic blame given our belief's both share positive pd-normative status.

Notice, however, that the normative evaluation made turns on an evaluation that's governed by what is internally accessible via reflection alone. But given what has already been said above, it's clear that this evaluation can miss some epistemically relevant facts. Indeed, it can be closed off to facts that can turn this seemingly correct response into an epistemically blameworthy one. The reason why is this: beliefs may satisfy the blameless condition yet still be epistemically deficient in a normatively significant way. It seems, then, internalist evaluations attribute a positive pd-status to beliefs in an overly permissive way—accounting for blamelessness notwithstanding.

Internalist evaluations afford beliefs a positive pd-status in an overly permissive way because both of us could be guilty of (i) *prior negligence* and/or (ii) not reasoning *from* good reasons and/or (iii) employing *apparent* memories (not *real* ones) in our reasoning

processes. Yet our beliefs could still be afforded a positive pd-normative status (even strong justification), which is why internally accessible evidence alone can sometimes be far too removed from the objective relationship it has with the world.

Put differently: if our credence levels of .9 and .7 are the result of fallacious reasoning—at some level—then our beliefs have acquired a positive pd-normative status on the cheap. CCE, then, has the result of prescribing two or more people who have formed their beliefs from unreliable processes to bootstrap their way up the pd-normative status ladder.[110]

This account of pd-normative status is an implausible one because it fails to take into consideration externalist concerns. When we evaluate what the epistemically correct thing to do is in YIELDING TO CCE, we need to answer the following kinds of questions that require externalist answers. For example, regarding the evidence we have, (i) is the evidence set true?; (ii) is the evidence objectively probable?; (iii) is the evidence reliably formed?; (iv) is the evidence responsibly formed?; (v) is the evidence leaving out important information?; (vi) are our beliefs objectively probable on evidence *e*?; and (vii) are our beliefs formed *from* our evidence or just in accordance with it in some lucky way?[111] Answers to these questions require us to appeal to external facts. Thus, the normative domain must extend beyond the narrowly carved territory set out by proponents of CCE

[110] Here we see how a belief's pd-normative status, within an internalist framework, can turn out to be capturing only a thin notion of reasonableness. Such reasonableness can be appealed to in defense of one's belief. But it can't handle charges that go beyond some of the basic requirements for a mere, rational belief. Indeed, because reasonableness is pretty thin stuff already, it's vulnerable to objections that point out that one may be reasonable in believing epistemically defective things if it's effective at securing one's desired ends, despite those ends being epistemically defective.

[111] These kinds of questions that are concerned about externalist matters are put forth by Greco in "Is Justification Internal?: Justification Is not Internal," p. 268.

who allow principles like NT, ACCESS, or NT to govern their evaluations thereby preventing externalist concerns and intuitions to weigh in.[112]

In light of these kinds of questions that capture externalist concerns, it's worth pointing out another related flaw in the normative evaluations made my CCE. This particular flaw has recently been brought out by Thomas Kelly. He says CCE divides the kind of evidence at work in cases of disagreement into psychological evidence and non-psychological evidence. Some cases, at one end of the spectrum, solely rely upon psychological evidence (e.g., a case wherein one lacks access to any non-psychological evidence and only knows what her peers believe regarding the truth value of some proposition p). At the other end of the spectrum, there are cases where one only has access to non-psychological evidence (e.g., a case where one neither has formed a belief nor knows what her peer believes). Then there are cases like YIELDING TO CCE that fall somewhere in the middle of the spectrum. This case involves the set of non-psychological evidence *e* that we disclosed to each other. And, in addition to *e*, there are two pieces of psychological evidences: namely, my belief that *h*/*e* is true with .7 confidence and your belief that *h*/*e* is true with .9 confidence.

According to CCE, we should split the difference and adjust our levels of confidence to .8. But notice that it's the psychological evidence that is doing all the work after we've discovered our disagreement and disclosed our evidence to each other. But why, given the kinds of externalist concerns put forth above, should we allow the psychological evidence to swamp the significance of the non-psychological evidence? It's as

[112] Recall from Chapter Two these three principles which govern internalist normative evaluations: ACCESS: If S and S* do not differ in the facts that they are able to know through internal reflection alone then the degree of their belief's epistemic status will not differ; NT: Indiscriminable situations are normatively identical; MENT: If S and S* do not differ in their mental states then their beliefs' epistemic status will be the same.

if all the interesting and accurate external concerns completely fall out of the picture and everything hinges on our degrees of confidence, which gets abstracted away from not only *e*, but also *h*'s objective probability on *e*. Furthermore, if the psychological evidence swamps the non-psychological evidence, then we set ourselves up for bootstrapping objections. For if we ignore that one or both of us can be guilty of (i) *prior negligence* and/or (ii) not reasoning *from* good reasons and/or (iii) employing *apparent* memories—not *real* ones—in our reasoning processes, then we get to afford our beliefs a positive pd-normative status (even justification) despite our evidence being too far removed from an objective relationship to the world. Normative evaluations that ignore these externalist concerns allow uninteresting and even false epistemic evaluations to be made. The motivation for expanding "the normative net" to include more epistemic goods that derive from externalist concerns is thus motivated.

1.3 Two More Reasons Why CCE's Normative Evaluations Are Inadequate

Recall back to Chapter Two, § 1.3, where one of the criteria for epistemic peerhood is *evidential equality*. This condition says that S and S* are evidential equals relative to the question *p* when S and S* are equally familiar with the evidence and arguments that bear on the question whether *p*. This notion of both S and S* being "equally familiar with the evidence" not only assumes that all the evidence informing one's one belief can be disclosed, which Chapter Two, § 1.2, showed to be implausible. But it also assumes dialectical evidence is the only game in town insofar as (a) what constitutes a belief having a positive pd-normative status and (b) what relates a belief to the world in an externally strong way required for knowledge. In this section I argue why such assumptions are false.

In § 1.3.1, I argue that bare memorial beliefs are intuitive candidates for knowledge, yet lack any internally accessible evidence. And in §1.3.2, I argue that there are beliefs

sourced in sub-personal processes that also can't be consciously accessed via reflection alone. The upshot is this: if some bare memorial beliefs and beliefs formed by sub-personal process are genuine instances of knowledge, then CCE normative evaluations are not only uninteresting, but false. For CCE evaluations will sometimes abstract a belief's pd-normative status from its relation to knowledge. But why think that knowledge is orthogonal to making evaluations of what is normatively appropriate? Surely, if there was ever anything that is epistemically correct to believe it is that which you know. Fixating on dialectical evidence and ignoring external factors related to knowledge is going to leave one's normative evaluation of cases involving disagreement impoverished.

Let's turn now to these two arguments which show why the kind of dialectical conception of evidence CCE works with is insufficient to account for the kind of external facts that are relevant to whether or not a belief has positive pd-normative status.

1.3.1 Bare Memorial Knowledge

Alvin Goldman says internalism is in trouble when trying to account for bare factual memory because many of our memories lack the kind of phenomenology internalists need to appeal to in order to establish a belief having positive epistemic status.[113] It's one thing to say the experience of an eidetic memory provides non-doxastic states that can be internally accessed thereby providing one's belief with justification; it's another thing to say memory can still play this kind of epistemic role once we consider a memorial belief from long ago. For we have many memorial beliefs that simply lack or are too isolated from the kind of *internal justifiers* that can be internally accessed the way a recent ostensible memory can. After all, we may not even remember how we acquired the memorial belief. Of course, as Greco acknowledges, some memorial beliefs "might be accompanied by a characteristic

[113] Alvin Goldman, "Internalism Exposed" *Journal of Philosophy* 96 (1999): 271-93.

sort of feeling—a kind of confidence, or perhaps a kind of attraction. But it is a stretch to think that these states are evidence—to think that these states 'support' my memory beliefs, or that my beliefs are evidentially based on such states."[114] According to Greco, it is psychologically implausible to suggest that all instances of memorial knowledge come with a colorfully rich phenomenology. Indeed, some instances of memorial knowledge may come without any phenomenology at all. One may believe that one's memory is a reliable source for acquiring true beliefs. But its actual features that make it reliable are not accessible from an internalist perspective. For example, consider the following case from Goldman.

> ICHABOD'S BARE MEMORIAL BELIEF: Years ago Ichabod formed a belief in proposition *q* by acquiring it in an entirely justified fashion. He had excellent evidence for believing it at that time (whether it was inferential or non-inferential evidence). After ten years passed, however, Ichabod has forgotten all of this evidence and not acquired any new evidence, either favorable or unfavorable. However, he continues to believe *q* strongly. Whenever he thinks about *q*, he (mentally) affirms its truth without hesitation. At noon today Ichabod's belief in *q* is still present, stored in his mind, although he is not actively thinking about it.[115]

Goldman wants to say that even if Ichabod has no other beliefs that can epistemically support *q*, he is still justified in believing it. Even when he retrieves *q* from memory and he lacks any phenomenology or additional memory related to *q,* Ichabod's belief still possesses a positive epistemic status. To say otherwise is counterintuitive. Furthermore, Goldman thinks that "if we refuse to grant justifiedness to beliefs of this sort, which derive from preservative memory, there will be serious skeptical ramifications: people will fail to know a

---

[114] John Greco, "Evidentialism about Knowledge," in Trent Dougherty, ed., *Evidentialism and Its Discontents* (Oxford: Oxford University, 2011), p. 170.

[115] Alvin Goldman, "Toward a Synthesis of Reliabilism and Evidentialism? Or: Evidentialism's Troubles, Reliabilism's Rescue Package," in *Evidentialism and its Discontents*, ed., Trent Dougherty (Oxford: Oxford University Press, 2011), p 260.

great many things that common sense credits them with knowing."[116] Of course, internalists could jettison the demand that Ichabod's memorial belief that *q* needs to "fit" or be supported by the original evidence that grounded the belief. But such synchronic, time-sliced evaluations that appeal to one's occurrent mental states is an essential tenant of internalism. Abandoning it is tantamount to taking an externalist position.

Therefore, paradigm cases of memorial knowledge lack the kind of internal justifiers an internalist perspective requires in order to afford a belief positive epistemic status. Of course one can bite the bullet and say that in such cases involving bare memorial knowledge one knows that p but the belief in p lacks positive epistemic status. That is, such a belief is unjustified or is held in an epistemically irresponsible way. This response, however, has the implausible result of affording one knowledge that p without one's belief that p having a positive epistemic status. Put differently: one's belief that p lacks positive epistemic status thereby leaving one epistemically blameworthy despite the fact that one *knows* that p. As a result, in order to avoid charges of epistemic blameworthiness one will have to give up what one already knows! But if it takes giving up what one already knows in order to acquire a positive epistemic status, then so much the worse for carving out a normative domain that is restricted to internalist concerns.

1.3.2 Current Models in Cognitive Science

"Recent empirical studies make it doubtful that paradigm cases of knowledge, such as perceptual knowledge, memory knowledge, and inductive knowledge, can be understood entirely in terms of person-level representation states, as evidence is understood to be,"[117] according to Greco. This is because the reliability of such belief forming processes may

[116] Ibid.

[117] Ibid.

depend upon sub-personal inputs and processes that can't be captured by an internalist construal of evidence. It's an empirical question as to whether human cognition is governed by person-level representations or seemings. Thus, we do well to leave open the possibility that knowledge is not always governed by rules which inform us of what kind of inferences we ought to make given the internal properties—deemed evidence—we can internally access via reflection alone. The fact there are phenomena such as blindsight and other types of non-conscious perception, motivates the plausibility of making logical space for information that can be received from the eye and delivered to the brain without being consciously experienced.[118] Indeed, Greco points to studies that argue that non-conscious perception is a part of normal cognitive functioning.[119] But whether this phenomenon can be established as fact or fiction—to everyone's satisfaction—is not important. What is important is that according to our best cognitive science, "the reliability of our cognitive systems might not depend just on the facts about [dialectic] evidence."[120] It seems not only possible, then, but plausible that perception involves the input, processing, and output of information that is not entirely consciously accessible.[121]

What's more, connectionist models of cognitive processing suggest there are units of information that escape our conscious detection. Some connectionist models suggest

---

[118] Blightsight is a documented phenomenon. It involves subjects who are blind in some part of their visual field. Despite this deficiency, Greco summarizes the studies by saying subjects can still "discriminate size, shape, location and/or orientation of objects in the blind part of the field despite the absence of any conscious visual experience." Ibid., p. 35. See L. Weiskrantz, *Blindsight: A Case Study and Implications* (Oxford: Clarendon Press, 1986).

[119] See Paul Lewicki, Thomas Hill, and Mara Czyzewska, "Nonconcious Acquisition of Information," *American Psychologist* 47, 6 (1992): 796-801.

[120] John Greco, "Evidentialism about Knowledge," p.171.

[121] Jack Lyons points to cases involving sensationless perception, which involve blind peoples' ability to detect obstacles like walls and chairs through a subtle capacity to echolocate. Such a species of perception is taken by Lyons to be a source from which justified beliefs can derive. See Jack Lyons, *Perception and Basic Beliefs: Zombies, Modules, and the Problems of the External World* (Oxford: Oxford University Press, 2009).

there is a level of information processing that takes place by various kinds of interactions between the sub-systems operating within a larger system. Such interactions between the sub-systems are said to be governed by strict physical laws. This implies that the reliability of our perception depends upon the interactions amongst sub-systems working within a larger system that cannot be internally accessed, much less understood as a kind of evidence that can be disclosed to one's peer. We do well to heed an insight from Timothy Williamson: "If agents have access to some information to which they lack internal access [e.g., the uptake of information from impersonal-processes], what is the point of assessing their beliefs and actions as if they were based only on [a dialectical conception of evidence or] information to which the agents have internal access? Why not assess their beliefs and actions in relation to *all* the information to which they have access?"[122]

Therefore, the kind of evidence CCE works with cannot account for the fact that we may know many things on the basis of certain kinds of cognitive functioning that can't be accessed and disclosed from an internalist perspective.[123] The restrictions CCE places on what constitutes the kind of access an agent has to information and processes that reliably, at some level, connects them to their environment results in an inadequate account of how to evaluate a belief's pd-normative status. The internalist perspective assumed by the argument from CCE gives the verdict that one can respond to disagreement in an epistemically incorrect way for retaining a belief that can be an instance of knowledge. This is because CCE is constrained by internalist resources which can't account for cognitive processes that are not inferential ones and don't depend upon a dialectical conception of evidence. Thus, externalism regarding epistemic normativity's relation to knowledge is

[122] Timothy Williamson, "On Being Justified in One's Head," p. 118 (emphasis mine).

[123] A related objection against simple reliabilism's claim that knowledge is a true belief deriving from a reliable process will be addressed in Chapter Six.

motivated and objections that I'm begging the question against internalist notions of normativity are misguided. Because CCE makes pd-normative status supervene entirely upon facts about one's internally accessible evidence, it fails to capture a plausible account of pd-normative status.

Let's take stock: in this section, I argued why internalism can't be the default position employed when evaluating cases of disagreement by giving a general argument against internalism. This argument highlighted how a belief's epistemic status is inextricably tied to external concerns. Things like (a) prior negligence, or (b) a belief formed from apparent memories rather than real memories, or (c) a belief formed in conjunction with *having* good reasons versus a belief formed *from* good reasons, or (d) bare memorial beliefs that lack any internally accessible evidence, or (e) beliefs formed by sub-personal processes are all relevant to informing our normative evaluations. Yet, (a)-(e) are all external concerns. So a belief's epistemic status shouldn't be abstracted away from such externalist concerns and restricted to a normative domain governed by internalist's constraints.

This all bears on an important question related to CCE's position on peer disagreement. How can one respond to disagreement in an epistemically incorrect way as a result of not significantly undergoing doxastic revision when one's belief is arguably already an instance of knowledge? This point surfaces what I take to be CCE's misguided approach on the epistemic significance of peer disagreement and how it generates a pseudo-problem that ignores the real problem: namely, 'Does disagreement function as a defeater against knowledge?' One can respond in an epistemically incorrect way once a belief is retained vis-à-vis peer disagreement only if the disagreement provides counterevidence that functions as a defeater against one's claim to know. The epistemic

interest in peer disagreement, then, should be how to understand when disagreement provides defeating evidence to one's account of knowledge.

2. WHY THE ANTECEDENT CONDITIONS FOR EPISTEMIC PEERHOOD RARELY EVER OBTAIN

Consider a case involving two conflicting thermometers. First, we are both standing on your front porch, in the shade. But your thermometer indicates the temperature is 76° Fahrenheit, while mine says it's 80° Fahrenheit. So there is something the thermometers "disagree" about: namely, the outside temperature. Call this the disagreement condition.[124] Second, because we're holding our thermometers next to each other—thus, virtually occupying the same shaded space on your front porch—the thermometers are receiving the same environmental inputs. Furthermore, we can both access the readings of each other's thermometers. Call this the same evidence condition. Third, both thermometers are the same kind of mercury thermometers. That is, they are both the same brand, model, age, and neither one shows signs of damage. Up to this point, they both have reliable track-records and so are believed to accurately indicate the temperature. Call this the reliability condition. Fourth, both of us agree that our thermometers conflict with each other, share the same "evidence," and have reliable track-records. Call this the acknowledgement condition.

According to CCE, a case such as this aptly illustrates why, epistemically, both of us should adjust our doxastic attitudes. We should withhold judgment as to whose thermometer is the more accurate one. After all, given there is no reason to deem one more reliable than the other, we should withhold judgment until we can verify the

[124] I'm indebted to Nathan King for spelling this condition and following ones out. See Nathan L. King, "Disagreement: The Skeptical Arguments from Peerhood and Symmetry," in *Disagreement and Skepticism* (New York, NY: Routledge, 2013), pp. 199-200.

thermometers' accuracy through some reliable metric. Or we find reason to demote one of the thermometers in light of discovering why one of the conditions above was not satisfied. But because these conditions have been satisfied and there are no known symmetry breakers on offer to explain why one or both thermometers are inaccurate, we ought to be skeptical and withhold judgment relative to what is the exact outside temperature. Sticking to one's guns without a symmetry breaker is not the epistemically correct response. Having a skeptical attitude, then, towards the veracity of the thermometers is obvious.

It's clear how a case like this involving two conflicting thermometers can easily satisfy the disagreement, evidence, reliability, and acknowledgment conditions, and therefore, motivate a localized, moderate kind of skepticism. But can these conditions be easily satisfied when it comes to conflicting agents? In this section I answer 'No.' More specifically, I argue that these conditions are rarely ever satisfied given there are a number of symmetry breakers on offer one can avail oneself of without acting in an epistemically incorrect way.[125]

[125] Again, a symmetry breaker is something that indicates one's epistemic situatedness is either superior or inferior relative to the disputant's epistemic situatedness. In other words, a symmetry breaker is some kind of indicator that breaks the symmetry that needs to obtain in the disagreement, evidence, reliability, and acknowledgment conditions. For example, it may later turn out that the original disagreement was the result of an equivocation over how the disputants were using the term "moral goodness." One was, say, using it to capture the kind of acts an agent was responsible for, which were constitutive of the agent's moral goodness; the other was using it to refer to an agent's moral worth, which required a modal account of how an agent would behave in a vast array of circumstances across a range of close possible worlds. Here the symmetry breaks down in the disagreement condition. Or the symmetry could break down for the reliability condition. I learn that my thermometer is damaged; it lost some of its mercury, which explains why it's unreliable. Nathan Christensen first coined the term 'symmetry breaker' and I originally borrowed it from Jennifer Lackey. See Jennifer Lackey, "A Justificationist View of Disagreement's Epistemic Significance," in *Social Epistemology* (Oxford: Oxford University Press, 2010), p., 309.

2.1 Symmetry Breaker #1: Disclosable Evidence (Dialectical Evidence) vs. Non-Disclosable Evidence (Non-dialectical Evidence)

The problem with peer disagreement (in general) and skeptical arguments from peer disagreement (in particular) is that in order for the evidence condition to be satisfied, both disputants must "share" or "disclose" with each other all the relevant evidence bearing on the disputed proposition. But the only way this approach can be carried out is by assuming or privileging disclosable or dialectic evidence over other kinds of evidence. So the skeptical threat leveled by peer disagreement relies upon a stringent conception of evidence in order to make satisfying the evidence condition even possible. This strategy, however, ironically overlooks the significance of disagreement over whether all evidence is ontologically propositional or non-propositional, or both. [126] Accounting for one's justificatory reasons for a belief, I argue, must go beyond appeals to mere dialectical conceptions of evidence. For constraining a belief's justificatory support to only the kinds of evidence that can be publically shareable is too impoverished. The reason why is because there are other kinds of evidence that agents have privileged access to, yet it can't be disclosed without losing epistemic properties. That is, there are other epistemically relevant facts that distinguish true beliefs and false beliefs from instances of knowledge that cannot always be explained by appealing to a dialectical conception of evidence. Furthermore, there are relevant kinds of evidence that can go undetected, yet still be epistemically relevant to how one's belief is justified. For example, experiences, intuitions, intellectual seemings, a belief's etiology, one's relation to a disputed belief's source (including testimonial sources), personal information, the monitoring of defeaters, dispositions,

[126] For example, Timothy Williamson (2000), Peter Unger (1975), and Donald Davidson (1986) takes all evidence to be propositional whereas Alvin Plantinga (1993b), Paul Moser (1989), Alan Miller (1991) and Earl Conee and Richard Feldman (2004) take evidence to be both propositional and non-propositional. I am indebted to Thomas Kelly for carving up the logical space relative to the kind of positions philosophers have taken regarding ontological candidates for evidence. See Thomas Kelly, "Evidence: Fundamental Concepts and the Phenomenal Conception," in *Philosophy Compass* 3/5 (2008): 933-955.

epistemic norms or norm-evidence, are all candidates for evidence that can't, in principle, be shared with a disputant. For trying to share all the epistemically relevant facts the aforementioned kinds of evidence have is either beyond our ken or awaits further empirical support.[127] Therefore, trying to satisfy the evidence condition in a way that parallels two conflicting thermometers is an unrealizable condition in most cases. The following subsections aim to flesh this point out in more detail.

2.1.1 Non-Transferability of All Testimony's Evidence

If there are genuine cases of testimonial knowledge that are not reducible to some other species of inductive knowledge, then proponents of CCE can't endorse it. In other words, if *some* cases of testimony are *sui generis*, then the kind of constraints governing CCE's conception of evidence will prevent one from endorsing genuine instances of testimonial knowledge. The trouble with CCE, then, is that non-reductionism relative to testimony is a strong possibility with a long tradition in support of it. So, if non-reductionism is true, then CCE will not be able to account for the kind of evidence that support some instances of testimonial knowledge.

And there are good, intuitive reasons to think that testimony can, sometimes, be an originating source of knowledge and not some species of inductive knowledge. For example, when my spouse, Julie—whom I have closely known for over ten years and know to be a trusted source for true information—discloses to me (in confidence) that she has cancer, my belief that she has cancer is a non-inferential one. This testimonial knowledge I gain of her condition is not reducible to a species of inductive knowledge; rather, this kind of testimony is an originating source for knowledge. I don't employ an argument that

[127] I leave this loss of epistemic significance somewhat vague in order to accommodate as much internalists and externalist accounts of evidence.

includes premises about her reliability and other relevant information, and then make the inference that she has been diagnosed with cancer. I don't synchronically need good inductive reasons before I can be intellectually responsible in believing her; nor do I need good inductive reasons in order to make an inference that she has been diagnosed with cancer before I can acquire knowledge of her condition.[128]

Thus, it seems there are cases where something epistemically *sui generis* is going on between the speaker and hearer given the kind of trusted relationship in play. Indeed, we see such epistemically significant relationships that result in the acquisition of testimonial knowledge through non-inferential processes all the time. For example, take cases involving the transfer of information between parents and children, and teachers and their students. A child asks her mom where the milk is and knows that it's in the fridge merely on the mother's say-so. So, too, when a child knows what the capital of New Jersey is as a result of being told so by her elementary school teacher. In both cases, the mother and teacher speak, and the hearers gain testimonial knowledge by way of a non-inferential process. Such cases intuitively capture the non-reductionist thesis that testimony can *sometimes* function as an originating source of knowledge. But CCE requires us to not only be reductionists regarding testimony, but use a dialectal conception of evidence to support the testimonial claim that's a part of our evidence set.

So, one reason why it's so hard to satisfy the evidence condition is that when it comes to testimony, there is a difference between (a) receiving testimonial evidence and (b)

[128] I am happy to concede to the reductionist that some kind of positive reasons are needed to first establish a speaker as a reliable source for testimonial knowledge. But after I have verified this speaker as a reliable one, the belief forming process shifts from an inferential one to, in some if not most cases, a non-inferential one. Just as *a priori* knowledge—assuming there is such a thing—sometimes first requires *a posteriori* knowledge, similarly, testimonial knowledge from person A may first require reductionist accounts of testimonial knowledge first before becoming regular instances of non-reductionist accounts of testimonial knowledge of the non-inferential stripe relative to the kind of operative belief forming process at work.

merely hearing about testimony.[129] It's very hard for two disputants to *always* share the same epistemic situatedness relative to a testimonial source. That is, there's good evidence to suggest that there is some kind of special relationship between a speaker and hearer that allows for testimonial knowledge to be a distinct kind of knowledge, as is argued by non-reductionists.[130]

Now consider the different epistemic situatedness between my spouse and me, on the one hand, and my spouse to a complete stranger, on the other, in the following case.[131]

> PRINCE WILLIAM DID WHAT!?: Suppose I'm sitting next to a complete stranger. We are both in New York City waiting for the next train to Philadelphia to arrive. My spouse Julie arrives at the train station a little later. Once she finds me, the first thing she tells me is how today at work Prince William—the Duke of Edinburgh—verbally abused her at the shareholders meeting today. The stranger next to me thinks Julie is out of her mind: for he just saw Prince William on television a couple hours ago walking into a bed and breakfast with his wife Kate in Dublin, Ireland, as reported by a CNN news anchor. The facts of the case, however, are these: Prince William's whereabouts are being strategically concealed so as to foil an assignation plot; the man and woman the stranger saw on CNN's report were government employed decoys; Julie is the CEO of a software company and held a meeting today amongst shareholders which included Prince William; Prince William did indeed verbally assault Julie.

Notice how that even though the stranger and I both heard Julie's testimony, only I *received* the full weight of testimonial evidence on offer whereas the stranger merely *heard* a piece of testimony. Indeed, my belief that my spouse was the victim of verbal assault came about non-inferentially given my unique relation to her, which includes a history of me knowing her to be a trustworthy, reliable, and always truthful in her assertions, esp., regarding

---

[129] I'm indebted to John Greco for point out the epistemic significance of this distinction. See John Greco, "Religious Knowledge In The Context Of Conflicting Testimony," in *Proceedings Of The American Catholic Philosophical Association*, vol. 83 (2009): 74.

[130] For a brief description of the reductionist/non-reductionist distinction relative to testimony see Chpt. Seven, § 1.1. Proponents of non-reductionist views of testimony are Audi (2006); Coady (1992); Goldberg (2006); Graham (2006); Plantinga (1993b); Reid (1983); Sosa (2006); Weiner (2003); Williamson (2000).

[131] This kind of case is inspired by Laurence Bonjour's clairvoyant cases. See Laurence Bonjour, "Externalist Theories of Empirical Knowledge," *Midwest Studies in Philosophy* 5 (1980): 53-73.

matters such as these. She said it and I believed it. Now I have testimonial knowledge of it. So, *if* testimony *sometimes* requires a privileged relationship between the speaker and the hearer, by virtue of which the epistemic status of the testimonial claim gets transferred, then that's not something the speaker can, then, just share with anybody. And however the details of this testimonial model gets fleshed out, it will eclipse any dialectical or discursive conception of evidence.

Now notice that even if the stranger believed truly that Prince William verbally abused Julie, as a result of Julie's testimony, he is not going to know it. For the conflicting news reports serve as potential defeaters or normative defeaters that need to be ruled out. Otherwise, the stranger violates some kind of subjective rationality condition for knowledge. And even if we tweaked the case so that it wasn't a stranger who witnessed this exchange between Julie and me, but rather, a colleague of mine, the outcome will probably be the same. My colleague John knows Julie and me fairly well. He has seen both Julie's and my interactions with other colleagues and students. He takes us to be normal people just trying to live a flourishing life. But given John lacks the kind of relationship with Julie that I have, he lacks the special kind of epistemic relationship that is operative between us, which enables me to arguably have testimonial knowledge through a non-inferential belief. John is, of course, reasonable for withholding judgment about the particulars of the case. After all, it seems he would be irrational if he were to believe Julie given the evidence he has from the CNN report and lack of inductive reasons for thinking Julie is as trusting, honest, and reliable as I know her to be. Therefore, the evidence condition is so hard to satisfy when it comes to testimonial evidence. Once there is relevant—especially pertinent—evidence deriving from testimonial sources that involve special relationships between the speaker and hearer, there is going to be a loss of epistemic properties during

the exchange. Why? Answer: because merely hearing testimony doesn't share epistemic parity with receiving testimony through a relationship such as this.[132]

There is another reason why testimony makes the evidence condition so hard to satisfy: even if both disputants share the same kind of special relationship with the speaker as noted above, the speaker still needs to choose whom she is going to disclose her information to. Greco has recently called attention to the models of *interpersonal perception* and *interpersonal testimony* which I argue sharpens the problem of satisfying the evidence condition. Perception and testimony can take on an interpersonal dimension as when I perceive that my spouse is upset and when my spouse chooses to disclose to me that she has cancer. This interpersonal dimension involves a key ingredient: namely, *self-disclosure* which, according to Greco, happens when "someone discloses something about themselves by either allowing themselves to be perceived in a particular way, or by telling about themselves."[133] Self-disclosure, then, can be mediated through interpersonal perception or interpersonal testimony. The former is when my spouse *allows* herself to be vulnerable around me and chooses to disclose her suppressed feelings and emotions to *me* about the diagnosis. She chooses not to be perceived this way in front of our children. The latter is when she chooses to disclose information about the diagnosis and prognosis to me through testimony. The upshot is that interpersonal perception and testimony involves the speaker *selecting* whom they will disclose such information to. So even if both disputants

---

[132] One could also argue that the reliability condition is compromised when it comes to disclosing testimonial evidence to a disputant who lacks the kind of relationship originally shared between the speaker and hearer. If my disputant did have the kind of relationship with the speaker that I have, then the disputant could receive all the epistemic import constituting the testimony. But because such a relationship is lacking, the disputant can't reliably receive everything that's epistemically significant. Indeed, it may not be psychological possible to even form the belief. But if the testimony is true, and the source is reliable, then it seems the defect can be located somewhere in the hearer's performance.

[133] John Greco, "Friendly Theism," in *Religious Tolerance Through Humility*, eds., James Kraft and David Basinger (Burlington, VT: Ashgate, 2008), p. 56.

share the kind of the relationship that enables non-inferential beliefs to form into instances of testimonial knowledge, it doesn't follow that one disputant disclosing such testimonial evidence to the other disputant will result in evidential parity. *Unless* the original speaker chooses to also disclose that same information to the disputant, evidential parity will likely fail.

2.1.2 Personal Information

Unlike technicians who *can* inspect the thermometers' mechanisms for damage, humans can't open up the hood of their disputant's skull and see if there is any damage (at least in the cases of disagreement we're considering). The best one can do is introspect one's internal states, consult memory, and monitor for epistemic vices in an effort to uncover any reasons for why one might not be epistemically performing well. For example, I can access whether or not I am sick, sleep deprived, starving, depressed, physically or mentally exhausted, under the influence of narcotics, or acting in some kind of epistemically blameworthy manner. But my access to my disputant's cognitive condition is far less accessible. Jennifer Lackey calls this epistemically relevant information personal information.[134] Therefore, there is an asymmetry between the accessibility one has of one's own epistemic performance and the disputant's epistemic performance on any given occasion.

Richard Fumerton, for example, seems to take an even stronger view of the significance of personal information than Lackey. Because Fumerton thinks it's reasonable to conclude that other philosophers he takes to be his cognitive equals suffer from biases that corrupt their reasoning, he thinks he is justified in thinking his disputant is more likely suffering from such epistemic defects. Indeed, Fumerton thinks he is justified in thinking

[134] See Jennifer Lackey, "A Justificationist View of Disagreement's Epistemic Significance," p. 310.

he is less likely to suffer from the epistemic defects caused by certain biases because he knows how *he* reasons better than he knows how other people reason; he is no privy to the kind of information he only has about his self. Thus, this egocentric perspective, according to Fumerton, is not only inescapable but provides one with access to personal information that we lack about our disputants.[135] This is another kind of symmetry breaker one may avail oneself of. And it is a kind of evidence that is needed to truly satisfy the conditions for two disputants to share cognitive equality.

2.1.3 Dispositions As Evidence

Some epistemologists take dispositions to be evidence.[136] According to Greco, all evidence ought to be understood in terms of "contingent belief-forming dispositions of cognitive agents."[137] Indeed, Greco thinks that one of the most important insights reliabilism has to offer is that all evidence reduces to contingent belief-forming dispositions.[138] To say evidence is best explained as contingent belief-forming dispositions is to deny that the propositional content of one's beliefs are the *relata* within defeat relations. Rather than appealing to the propositional content of the relevant beliefs supplying the basal conditions for defeat relations to supervene upon, one's contingent belief-forming disposition is the locus of defeating evidence. Dispositions not only provide defeating evidence against one's belief that p, but they also ground positive evidence for p. This view of evidence assumes our cognitive systems are modular in that some modules that produce beliefs (e.g., perception) will override beliefs produced by other modules (e.g., theoretical reasoning

---

[135] See Richard Fumerton, "You Can't Trust A Philosopher," in Richard Feldman and Ted A. Warfield, eds., *Disagreement* (Oxford: Oxford University Press, 2010), p. 105-106.

[136] I take up this theme is great detail in Chapter 6, § 2.

[137] John Greco, *Achieving Knowledge*, p. 162.

[138] Ibid.

processes). So, for example, when one believes that an argument against the reality of the external world is sound, because one lacks the disposition to deny the existence of the external world, the disposition to believe in the external world is evidence against the argument's soundness. The defeater system's outputs are governed and realized by the agent's disposition to believe, disbelieve, withhold judgment or reduce confidence or strengthen confidence in one's belief.

The merits of this account of evidence are not important for our present purposes. Despite this account of evidence being coherent, plausible, and held by expert epistemologists, what's important is this: if such an account of evidence is true, then *not* all the relevant evidence can be disclosed to others. In the same way I can't share with you the *quale* of my feeling of hurtfulness, I can't share with you my disposition nor its degree of strength to either believe or disbelieve or adjust my credence level. Therefore, if evidence is a belief-forming disposition, then it's not the kind of thing that can be disclosed to another. Of course, one can testify to another that one has a particular disposition. But this piece of testimony can't transfer the epistemic significance of a belief-forming disposition to a hearer. Abstracting away one's disposition from how it fits within one's overall cognitive system through a piece of testimony will result in an evidential asymmetry. Therefore, dispositions are another kind of symmetry breaker one can avail oneself of that would prevent the evidence condition from being satisfied.

2.1.4 The Misguidedness of Privileging Certain Epistemic Norms Over Others

Scottish philosopher Thomas Reid called into question the way skeptics act like imperialists when they force all other faculties to bow down to the faculty of reason. Since *a priori* intuition, perception, memory, introspection, and reasoning were all bequeathed to us from

the same source—whether that source be naturalistic or not—they are all, according to Reid, in the same epistemic boat. To put the point in Reid's words:

> Reason, says the sceptic, is the only judge of truth, and you ought to throw off every opinion and every belief that is not grounded in reason. Why, sir, should I believe the faculty of reason more than that of perception?—they came both from the same shop, and were made by the same artist; and if he pits one piece of false ware into my hands, what should hinder him from putting another?[139]

The problem, then, with skeptical arguments from disagreement with epistemic peers is that the same kind of chauvinism is displayed. Since proponents of CCE stipulate that the correct epistemic norm is the one that places a primacy on evidence that comes from disclosable arguments and reasons, they overlook the importance of how one's evidence should dictate the appropriate epistemic norm to employ. CCE proponents, then, privilege certain epistemic norms, which dictate what kinds of evidence can be shared. But this is to make the mistake of putting the cart in front of the horse: for one's evidence should dictate the correct epistemic norm, not the other way around. For example, something like the following principle is employed when adjudicating the reasonableness of retaining or withholding judgment of one's belief vis-à-vis peer disagreement.

> ***Dialectical Symmetry + Independent Principle***: If S believes P and T believes ~P, and S and T are aware of their disagreement, and neither S nor T has an argument that would compel every rational person to assent to her position regarding P, *nor any evidence that is independent of the disagreement and the original grounds for her attitude regarding P*, then neither S nor T is justified in her attitude toward P, and both should suspend judgment.[140]

---

[139] See Thomas Reid, "An Inquiry into the Human Mind and the Principles of Common Sense," in *Philosophical Works*, ed. H.M. Bracken, 2 vols. (Hildesheim: George Olms, 1983), p. 185. The same principle applies if the story of how we acquired our faculties is a naturalistic one wherein such faculties were selected for given their conduciveness to adapt and survive.

[140] This principle was formulated by Nathan King. See Nathan L. King, "Disagreement: The Skeptical Arguments from Peerhood and Symmetry," p. 209. This principle includes another principle frequently endorsed by many proponents of CCE, namely, the independence principle. This principle says "In evaluating the epistemic credential of another's expressed belief about P, in order to determine how (or whether) to modify my own belief about P, I should do so in a way that doesn't rely on the reasoning behind

CCE proponents demand their disputants to provide discursive evidence, but in addition to that, they also prevent them from appealing to the justification that's a part of the targeted belief's etiology.[141] This second conjunct in the antecedent is known as the independence principle.[142] So not only is the evidence condition restricted to discursive kinds of evidence, one can't appeal to the kind of evidence that originally grounded the belief. But such a principle suffers from countless counterexamples. For example, consider Alvin Plantinga's notable missing letter case.

> MISSING LETTER: I am applying to the National Endowment for the Humanities for a fellowship; I write a letter to a colleague, trying to bribe him to write the Endowment a glowing letter on my behalf; he indignantly refuses and sends the letter to my chairman. The letter disappears from the chairman's office under mysterious circumstances. I have a motive for stealing it; I have the opportunity to do so; and I have been known to do such things in the past. Furthermore an extremely reliable member of the department claims to have seen me furtively entering the chairman's office at about the time when the letter must have been stolen. The evidence against me is very strong; my colleagues reproach me for such underhanded behavior and treat me with evident distaste. The facts of the matter, however, are that I didn't steal the letter

---

my initial belief that P." See David Christensen, "Disagreement as Evidence: The Epistemology of Controversy," in *Philosophy Compass* 4: 756-67.

[141] Feldman even wants all foundational evidence to be disclosed and argued for leaving no starting points left unchallenged. For example, he says "It is difficult to accept the claim that the starting points are beyond rational scrutiny. […] The fact that one's own starting point has, let us grant, some justification in isolation hardly suffices to defend the view that it retains that justification once one realizes that other people, otherwise as capable as oneself, have a different stating point with as much 'objective' initial credibility as one's own." Feldman goes on to say that sticking to one's starting point and using it as a symmetry breaker is arbitrary if one recognizes a competing starting point has as much justification in favor of it. Adding that, "The problem is that once these starting points are brought out into the open, they are every bit as open to rational scrutiny as anything else." This, of course, assumes we can provide all the relevant evidence and processes involved that support our starting points. Furthermore, it invites all kinds of regress problems once we conflate nondiscursive evidence with discursive evidence or evidence with argument. (I will address this problem late in Chapter Seven.) See Richard Feldman, "Epistemological Puzzles about Disagreement," in Stephen Hetherington, ed., *Epistemology Futures* (Oxford: Clarendon Press, 2006), p. 224-225.

[142] This principle, recall, was first mentioned and expounded upon in chapter 1. A version of the independence principle is this: In evaluating the epistemic credentials of another's expressed belief about P, in order to determine how (or whether) to modify my own belief about P, I should do so in a way that doesn't rely on the reasoning behind my initial belief that P. See David Christensen, "Disagreement as Evidence: The Epistemology of Controversy," in *Philosophy Compass* 4: 758

> and in fact spent the entire afternoon in question on a solitary walk in the woods; furthermore I clearly remember spending that afternoon walking in the woods.[143]

Here is a case where the faculty of memory is the original ground for one's attitude towards the disputed proposition: namely, I'm innocent of stealing the letter. It's intuitively obvious that my colleagues' testimony and use of arguments for my guilt should not be assessed independently of my memorial beliefs, which would violate the Dialectical Symmetry + Independent Principle. The upshot, then, is this: I am perfectly reasonable in retaining my belief in my innocence despite (a) lacking an argument that convinces my colleagues and (b) appealing to properties comprising my belief's etiology—namely, it deriving from memory—despite these properties being a part of my initial reason for why I believe I'm innocent. Indeed, such an eidetic memory of walking through the woods, replete with additional memories of hearing the crisp, autumn leaves crackling under my feet while gazing at all the beautiful wildflowers, is sufficient for not taking a skeptical attitude of suspending judgment relative to my innocence. Just because there is strong testimonial evidence and arguments that suggest otherwise and I've appealed to the original reasons constitutive of my belief, I'm not epistemically blameworthy despite violating the Dialectical Symmetry + Independent Principle.

In other cases, perception will be the original ground for my belief. But demanding that one can't appeal to it because it's not independent of how the belief came about is epistemically inappropriate. Consider a case involving disagreement over whether determinism prevents one from being able to do otherwise. One disputant appeals to the

---

[143] Alvin Plantinga, "Intellectual Sophistication and Basic Belief in God," *Faith and Philosophy* 3 (1986): 306-12. Admittedly, there does seem to be some point at which testimony will trump memory. I think this is another case of ontological vagueness insofar as identifying the exact point at which one more additional testifier guarantees you will no longer trust your memory and yield to testimony. Just as we can't tell which hair, after being plucked, causes one to become bald, I don't think there is a specific number of testifiers it would take for one to no longer trust one's memory and admit to being delusional.

faculty of perception to justify her claim that she knows, in *some* cases, she both can and could have done otherwise. She can just see that she has this ability to exert her will to do A or not-A, despite not being able to formulate it into an argument. The other disputant puts forth a theoretical argument. Given the evidence in support of determinism, she argues that we lack the ability to do something other than what we in fact do. Here, too, it seems one is perfectly reasonable in retaining one's belief in having the ability to do otherwise despite (a) lacking an argument that convinces the disputant and (b) relying upon evidence that is not independent of the belief's original etiology. *If* perception delivers a reliably true belief that one can do otherwise, it doesn't follow that one must only appeal to discursive kinds of evidence in order to avoid epistemic blame.

The point of this section is to further flesh out the importance of recognizing the difference between discursive and non-discursive evidence and how it's epistemically inappropriate to demand that one's disputant not appeal to the kinds of reasons that originally grounded the disputed belief. What's more, the assumption that all relevant evidence can be disclosed in a way that generates epistemic symmetry relative to the evidence condition is implausible (at best) and false (at worst). Such an approach conflates (a) the conditions for *defending* one's undefeated evidence for the targeted belief (e.g., providing an argument which relies upon theoretical reasoning) with (b) the conditions that *ground* one's having undefeated evidence for a targeted belief (e.g., having a true, reliably produced perceptual belief or memorial belief).

2.1.5 Intuitions As Evidence

Intuitions are no stranger to the controversy that surrounds the nature of evidence. Virtually all philosophers are willing to concede that Edmund Gettier provided strong counterexamples to the JTB theory of knowledge as evidenced by the intuitions his cases

elicited.[144] Many times disagreement over a philosophical topic bottoms out at the level of intuition, which many philosophers take to be evidence in favor of their position. But not all—indeed, some philosophers take intuitions to be nothing more than either a belief masquerading as evidence or an insidious way of calling a mere opinion something more authoritative than a belief.[145] Intuitions, the reasoning goes, merely describe the psychological state the subject is in, not anything objective about the subject matter that's up for debate. And employing them in one's philosophical method is merely a convenient way to support that which they strongly believe, which is begging the question (at best) or wishful thinking (at worst).

It goes beyond the scope of this chapter to weigh in on the epistemic import of intuitions, the strength of evidence they may or may not provide, and their legitimacy in philosophical method. The point I want to raise, however, is that philosophers do, in fact, use and recognize intuitions as pieces of evidence, in addition to playing a vital role in philosophical method. Recall Chisholm's statement regarding the relation between intuitions and philosophical method: "[Employing intuitions into our philosophical method] may seem [like] the wrong place to start. But where else *could* we start?"[146] That is to say, if our philosophical method doesn't involve providing adequate theories that best explain our intuitions after evaluating cases, then what other kinds of evidence are we

---

[144] See Edmund Gettier, "Is True Justified Belief Knowledge?" in *Analysis*, 23 (1963): 121-123.

[145] Peter van Inwagen, "Materialism and the psychological-continuity account of personal identity" in J. Tomberlin. ed., *Philosophical Perspectives, 11: Mind, Causation and Word* (Oxford: Blackwell, 1997), p.,309. David Lewis says intuitions are nothing more than one's opinion. See David Lewis, *Philosophical Papers*, vol. 1 (Oxford: Oxford University Press, 1983), p.x. Although not much hangs on this, I take opinions to be beliefs that people don't want to have to defend (at worst) or a belief held with extremely low confidence. Beliefs are propositional attitudes we typically take to be defensible at some level. Opinions, however, are propositional attitudes we find ourselves inclined to believe but do not want to have to defend. Opinions, then, serve a practical role in today's societal norms by allowing one from ever having to enter into the dock and defend their belief. I adopted this view of opinions from Linda Zagzebski who expressed this position through correspondence.

[146] Roderick Chisholm, *The Theory of Knowledge*, 2nd ed. (Englewood Cliffs, NJ: Prentice Hall, 1977), p. 16.

looking for to justify our theories? Analytic philosophy takes intuition to be a kind of experience that provides the relevant datum that needs to be explained through valid reasoning. Indeed, many philosophers prefer theories that make sense of their intuitions rather, than, preferring theories that entail the falsity of them. Intuitions, then, play a significant methodological role for many philosophers as a way to judge the adequacies of theories competing for their endorsement. Therefore, as long as intuitions are still considered by many philosophers as evidence, they are going to be (a) another kind of symmetry breaker a disputant can avail oneself of and (b) function as a type of evidence that can't be disclosed to another without losing its epistemic bite thereby frustrating the evidential condition from being satisfied.

2.2 Symmetry Breaker #2: The Incommunicability of a Belief's Etiology

Another problem with obtaining the kind of symmetry needed to satisfy the evidence condition comes from the late William Alston.[147] Alston made the point that we often conflate the difference between (a) being justified in believing that p and (b) being able to justify to another one's believing that p. For example, my son Joel is justified in believing that there is milk in the fridge merely upon my say so. But it's a confusion to assume that in order for Joel's testimonial belief to be justified, he must also provide an argument in support of it. Further, suppose Joel was autistic and was incapable of articulating his reasons in support of his testimonial belief or even one of his perceptual beliefs. His inabilities to articulate his reasons, evidence, or arguments in support of his respective beliefs don't undercut the justification of his beliefs nor the reasonableness of retaining it. This point ferrets out a false assumption in the way skeptical threats from disagreement get

[147] William Alston, "Level-Confusions in Epistemology," in *Midwest Studies in Philosophy*, vol. 5, no. 1 (1980): 135-150.

motivated: namely, it assumes one can retrieve enough relevant evidence, reasons, or justification from memory and articulately "share" it with another. Such a demand is not only inappropriate, it fails to recognize the inscrutable subtleties involved in a belief's etiology. But even though such subtleties cannot be disclosed, they are still epistemically relevant insofar as deciding what the epistemically correct response is vis-à-vis peer disagreement.

For consider the diachronic nature of a belief's etiology before it becomes fully formed, according to Ernest Sosa:

> A belief forms in us over time through the subtle influence of diverse sources. Some are testimonial, others perceptual, others inferential, and so on. The belief might owe importantly to the believer's upbringing, or to later influence by his community. […] [Because the full grounding for a belief] is thus complex, and temporally extended with the aid of memory, it lies beyond our present view. […] And so it is, I submit, for nearly the whole of one's body of beliefs. The idea that we can always or even spot our operative 'evidence' for examination is a myth.[148]

The idea, then, is that a belief's etiology has so many kinds of different evidentially relevant influences. Some, perhaps, come from sub-personal process that will elude our conscious detection. Others come from process that extend over such a long period of time that articulating how the disputed belief was formed by citing its "evidence" is overly ambitious (at best) and impossible (at worst). So, if we can't even consciously detect all the relevant evidence and processes constituting our beliefs' etiology, then surely we can't disclose our "evidence" for the disputed belief to our disputant.

After all, consider the possibility that each person's defeater system performs, among other things, the function of monitoring for defeaters. This process presumably takes place at both the conscious- and sub-conscious level (at least until a potential defeater

[148] Ernest Sosa, "The Epistemology of Disagreement," in Alan Haddock, Adrian Millar, and Duncan Pritchard, eds., *Social Epistemology* (Oxford: Oxford University Press, 2010), p. 290-291.

lights up). But this kind of epistemic influence on beliefs can't be "shared" with our disputant. Furthermore, even if we could consciously detect all the justificatory support for a belief, it doesn't follow that even the best of us are going to be able to articulate it into words and propositions—let alone strong arguments—to support it. Therefore, the kinds of synchronic, sub-conscious, undetectable, ineffable, justificatory support enjoyed by one's belief does not fit well with the assumptions on what counts as evidence undergirding the evidence condition.

2.3 Symmetry Breaker #3: Pragmatic Constraints Over The Amount of Time Needed to Disclose and Evaluate All the Relevant Evidence

This next kind of symmetry breaker calls into question the intellectual tenacity, conscientiousness, and stamina of the disputants. Do we really think disputants are going to put in the time, effort, and energies required to exhaustively examine all the total evidence a disputant has in support of a belief? Granted some disputed beliefs rest upon a small set of evidence and examining such evidence is doable within a reasonable amount of time. But once we consider domains wherein the subject matter is both broad in scope and can be plumbed to great depths, the difficulty of adequately assessing the total evidence seems insurmountable given our limited time and resources we can actually devote to the disagreement. For example, consider the following case put forth by Nathan King. King takes a real life case of disagreement between David Lewis and Peter van Inwagen over the nature of human freedom to ask the sober question: Do philosophers put forth the amount of time, effort, and energy into scrupulously and conscientiously assessing all the total evidence disclosed by the disputant?

> SCRUPULOUS PHILOSOPHERS: Peter and David are professional philosophers of the first rank. Both exhibit a wide range of intellectual skill and virtues, and exhibit them to a great extent. Peter and David have acknowledged that they exemplify these virtues more or less equally. These

> philosophers are, in a word, *scrupulous* with respect to their assessment of evidence. But while many philosophers are scrupulous in the way they assess evidence, Peter and David are scrupulous in another way. Namely, they take great care to *share* their evidence with respect to the claim that <Genuine human freedom requires indeterminism>. In conversation with one another, they proceed slowly and carefully so as to ensure mutual understanding. They read the same books and journal articles, and attend the same conferences, thereby acquainting themselves with the same arguments. They carefully log time spent reading and thinking about human freedom, and end up with matching logs. Yet after all this, Peter believes that genuine freedom must be indeterminist in character, while David denies this.[149]

The point of this case is to illustrate the difficulty of satisfying, again, the evidence condition between disputants. What's more, King is restricting the kind of permissible evidence to a dialectical account wherein all relevant evidence is propositional and can, in fact, be disclosed. Unlike non-dialectical accounts of evidence that include everything epistemically relevant to a belief's etiology, including, but not limited to, intuitions, perceptual experiences, rational insights, or intellectual seemings, etc., this case assumes both disputants' contrary beliefs are supported by only dialectical evidence. Still, the amount of time and conscientious attention exercised by Lewis and van Inwagen perusing over each other's disclosed evidence is laudable because the achievement is so rare. The upshot, then, is this: even if we restrict the kinds of evidence we deem legitimate to that which is disclosable, namely, a dialectical conception of evidence, the likelihood of the evidence condition obtaining in many cases is going to be inscrutable (at best) and low (at worst).

[149] Nathan L. King, "Disagreement: What's the Problem? Or A Good Peer is Hard to Find" in *Philosophy and Phenomenological Research* (2011): 8 (emphasis his).

2.4 Symmetry Breaker #4: Goldman's Incommensurable E-Systems and Their Influence on the Interpretation of Evidence

One of the biggest differences between comparing two conflicting thermometers with two disputants is the former set presumably shares parity relative to the internal mechanisms at work during the uptake of information from the external environment. Conflicting thermometers motivates skeptical concerns about the reliability of one's thermometer because unless there is some explanation for the conflict, both thermometers function the same way and appear indistinguishable from each other. There is no detectable defect between the two. But as Alvin Goldman has pointed out there are "many different communities, cultures, social networks, and so on [that] endorse different epistemic systems (E-Systems)—that is, different sets of norms, standards, or principles for forming beliefs and other doxastic states."[150] Put simply: disputants are very likely to employ different epistemic norms that piggyback off of one's inherited worldview that colors the disclosed evidence shared between disputants. Of course there will be much overlap between people and communities relative to what counts as a valid inference and what sources of belief are more reliable (e.g., perception, memory, introspection, a priori intuition, testimony) than others (e.g., wishful thinking, hasty generalizations, and reading tea leaves). The kinds of norms used, however, in deciding what proposition is best supported by evidence can not only differ, but result in both disputants' beliefs being reasonable. Thus, Goldman thinks there can be reasonable disagreement as a result of the conflicting beliefs deriving from different epistemic norms and such beliefs suffering from no epistemic defect.

[150] Alvin Goldman, "Epistemic Relativism and Reasonable Disagreement," in *Disagreement*, eds., Richard Feldman & Ted A. Warfield (Oxford: Oxford University Press, 2010), p. 187.

Just recall back to the MISSING LETTER case above. Suppose we tweak the case a bit so that it involves me being tried by a jury. The jury hears about the eye witness testimonies from trusted sources; they hear that I have been convicted of doing similar things in the past; they know I have a motive for stealing the letter; etc. Based upon this preponderance of evidence, which rules out any reasonable doubt, the jury convicts me. More exactly, based upon the circumstantial evidence, which includes testimonial evidence, motive, and evidence that I have a track-record of such crimes, the jury employed the epistemic norm of weighing this kind of evidence more heavily than my testimonial evidence to the contrary. But I'm not surprised that my testimonial evidence is not convincing to the jurors. Such evidence can't be disclosed to another without the loss of epistemic properties. The rich phenomenology of, say, my eidetic memory of not committing the crime is not shareable. Therefore, I'm not surprised by the outcome. Indeed, I probably would have formed the same belief were I a member of the jury.

So, on the one hand, I don't blame the jury for convicting me. They did not respond in an epistemically incorrect way, even if their belief is false. On the other hand, neither did I respond in an epistemically incorrect way by employing the epistemic norm of weighing my memorial evidence more heavily than the circumstantial evidence against me.[151] After all, in this case, the source of my belief that I'm innocent is much more reliable than the evidence against me: I vividly remember not stealing it and, what's more, I vividly remember not even being in the same state where the theft took place!

[151] The question of whether the epistemic norm of relying upon memorial evidence should always trump testimonial evidence is an interesting one. Presumably, there will be cases where the correct epistemic norm will be to weigh testimonial evidence more heavily than memorial evidence. If there were a large number of my friends and family members—people I know to be paragons of trust, reliability, and truthfulness—who all testified to have seen me take the letter, then I would no longer place my memorial beliefs over their testimony. The number of people needed to do this is vague and would resemble Sorites paradoxes. But there is no question that such testimonial counterevidence can provide one with an undercutting defeater for one's memorial beliefs.

2.5 Summary

Let's take stock: the purpose of this chapter was to argue that the problem of peer disagreement is a pseudo-problem. More specifically, the kind of skeptical threats generated from peer disagreement require an internalist normative evaluation that fixates on blameworthiness. Internalism, however, is not the only way to govern our normative evaluations. So, in section one I argued why internalism can't be the default position employed when evaluating cases of disagreement by giving a general argument against internalism. This argument highlighted how a belief's epistemic status is inextricably tied to external concerns. Things like (a) prior negligence, or (b) a belief formed from apparent memories rather than real memories, or (c) a belief formed in conjunction with *having* good reasons versus a belief formed *from* good reasons, or (d) bare memorial beliefs that lack any internally accessible evidence, or (e) beliefs formed by sub-personal processes are all relevant to informing our normative evaluations. Yet, (a)-(e) are all external concerns. So a belief's epistemic status shouldn't be abstracted away from such externalist concerns and restricted to a normative domain governed by internalist's constraints.

In section two I argued that the skeptical threat form peer disagreement can get traction only if the evidential condition for both disputants can be satisfied. Despite cases like DISAGREEING THERMOMETERS and DINNER BILL illustrating how the evidential condition can seemingly be satisfied, such models fail to generalize far and wide. Indeed, I argued that the evidence condition is rarely ever satisfied for the following reasons. First, I argued that the idea that there is no universally agreed upon notion of evidence. So the evidence condition is restricted to a stipulated definition of evidence which amounts to that which can be disclosed through propositions, reasons, and arguments via testimony. Such a "disclosable," "discursive," or "dialectical" construal of

evidence fails to take into account the epistemic significance of the following features about views on evidence that can't be disclosed in a way that satisfies the evidence condition: (i) a hearer's evidential relation to a speaker whereby the hearer can *receive* a piece of testimony verses merely *hearing* a piece of testimony; (ii) personal information—that is, one is only privy to his or her *own* cognitive performance in a way that can't be accessed in one's disputant. For example, only you have privileged access to monitor for and be aware of certain epistemic defects, intellectual vices, or poor health conditions which can negatively affect your cognitive performance; (iii) some take dispositions to be a kind of evidence which can't be disclosed to another without loss of epistemic properties; (iv) certain epistemic norms should and do govern belief formation in a way that leaves both disputants in asymmetrical evidential conditions; (v) some take intuitions to be a species of evidence yet intuitions can't be disclosed without a loss of epistemic properties; (vi) some disputed beliefs or beliefs that support the targeted disputed-belief have etiologies that elude our conscious detection and, as Sosa argues, it is a myth to suggest we can retrieve and disclose all the operative evidence in a belief's etiology; (vii) even if evidence was restricted to the dialectical kind, there are severe pragmatic constraints on the amount of time it would take for disputants to seriously and conscientiously peruse all the relevant evidence that each disputant discloses to each other; and (viii) each disputant has an inherited worldview that entails different epistemic norms that color or undermine the disclosed evidence. This kind of evidence which Goldman calls "norm evidence" is thus ignored and isn't the kind of evidence that can be disclosed.

Both sections point out the false assumption that motivates the case for skepticism from disagreement: namely, all relevant evidence can be "disclosed" between the disputants in a way that plausibly satisfies the evidence condition needed for epistemic symmetry to be

shared between disputants. But despite having argued for why the problem of peer disagreement is a pseudo problem, I do think it can be recast as a skeptical argument against knowledge. So, in the next chapter of this dissertation we turn to what I take to be a more formidable argument from disagreement and lay out the problems that deserve an externalist response.

4

The Real Problem: Disagreement With Epistemic Superiors

[C]onsidering how many diverse opinions learned men [and women] may maintain on a single question—even though it is impossible for more than one to be true—I held as well-nigh false everything that was merely probable.

~RENÉ DESCARTES

After motivating why peer disagreement is a pseudo-problem in Chapters Two and Three, in this chapter I recast the skeptical argument from peer disagreement into a more formidable one. For this argument takes into account both internalist and externalist concerns. This new argument levels a potential defeater against knowledge, which is where the debate over the significance of peer disagreement should have taken place from the beginning. The reason why is this: concerns over what constitutes normative appropriateness are not orthogonal from externalist concerns (in general) and knowledge (in particular). But before I introduce this argument, I must first return to the argument put forth by Goldberg and Kornblith, which I introduced in Chapter One. This argument does not provide a pseudo-problem because it *does* take into account externalist concerns, which is why it needs to be treated independently of Feldman's and Christensen's argument. § 1, then, will address the remaining argument from peer disagreement. After I show what I take to be wrong with their argument, in § 2 I put forth a more formidable one. In § 3 I show how reliabilism already suffers from an acute problem of accounting for counterevidence in a way that preserves our pre-theoretical intuitions about counterevidence. But with the argument for skepticism I put forth in § 2, it seems reliabilism faces an even greater challenge: namely, not only must it reflect how defeating evidence is accounted for, but it must also provide an answer to the skeptical threat put

forth by this argument. Unless reliabilism can account for how this skeptical threat doesn't provide knowledge defeating evidence to those who face such disagreement, then many of us don't know a lot of commonsensical things, in additional to our most cherished beliefs.

1. WHY GOLDBERG'S AND KORNBLITH'S ASSUMPTION THAT DISAGREEMENT INDICATES UNRELIABILITY IS TOO STRONG

Recall back to Chapter One where we looked at Kornblith's and Goldberg's argument for skepticism from disagreement. There they argued that any domain that suffers from pervasive or systematic disagreement is grounds for denying any knowledge or justified beliefs in that domain. More specifically, Kornblith and Goldberg argued that since domains like philosophy—unlike mathematics and the hard sciences—suffer from systematic disagreement[152], we have strong evidence that we are not endowed with reliable methods and/or faculties to get the truth in these domains. Indeed, they both agree that systematic peer disagreement generates a defeater in any domain that lacks consensus like philosophy, religion, politics, and value theory. Their argument for moderate skepticism is as follows:

The Argument From CCE A La Kornblith & Goldberg

(1) Reliability is a necessary condition for epistemic justification.
(2) Epistemic justification is a necessary condition for knowledge.
(3) Systematic peer disagreement over whether p is true indicates a lack of consensus.
(4) Lack of consensus is an indication of unreliability.

[152] Goldberg defines systematic peer disagreement (hereafter SPD) as disagreement between two peers over a claim that is disagreed upon in a nonlocal, widespread, and entrenched way. The three ingredients of SPD are (i) Nonlocal Disagreement—that is, disagreement over p, which implies disagreement over many other related topics; (ii) Widespread Disagreement—that is, when both disagreeing parties can appeal to a wide constituency of followers; (iii) Entrenched Disagreement—that is, disagreement that has persisted for some substantial length of time and all parties continue to defend and advance their view in the face of persistent challenges from the other side, which are replied to with evidence and arguments in the defense thereof. SPD, thus, levels an undercutting defeater against one's belief since one's subjective probability relative to the disputed proposition should decrease given the pervasive disagreement is evidence for the disputed belief being a product of unreliable processes and/or methods.

(5) In cases in which S believes that p in the face of systematic peer disagreement over whether p, there are rationality or normative defeaters with respect to S's belief that p.
(6) If there are rationality or normative defeaters with respect to S's belief that p, then S neither knows, nor is doxastically justified in believing, that p.

Therefore,

(7) If p is a proposition regarding which there is systematic peer disagreement, then if S believes that p, S's belief is neither knowledge nor doxastically justified.[153]

I grant all the premises of this argument are true except for premise (4), which I think everything seems to turn on. So for the argument to go through, everything hangs on the merits of premise (4). In what follows, I will argue why premise (4) is false.

1.1 Consensus is Too Strong of a Metric for Determining Reliability In Difficult Domains

The first thing I want to challenge is an implausible assumption undergirding premise (4): namely, lack of consensus is an indication of unreliability. Why think a thing like that? More specifically, why think there should be consensus in a domain that is very difficult? Getting the truth in challenging domains sometimes requires hard work, extended periods of time, and the cultivation and exercise of the right epistemic virtues; moreover, it requires a willingness to know the truth regardless of the cost. Furthermore, it seems some epistemic peers may be epistemically gifted better than other epistemic peers yet still be considered epistemic peers in some broad sense of the concept. For example, there are many professional athletes that play sport X, but some are still athletically gifted better than other professional athletes.

I am a good perceiver just as my friend—the radiologist—is a good perceiver. After all, my perceptual abilities are better than average; I have 20-15 vision. But I can't see

[153] Sanford Goldberg, "Defending Philosophy in the Face of Systematic Disagreement," in Diego E. Machuca, ed., *Disagreement and Skepticism* (New York: Routledge, 2013), p. 279.

subtle, hairline fractures in x-rays or diagnose tumors or ligament tears by looking at CAT scans like a skillful radiologist. Radiologists can make finer discriminations than others. And even amongst radiologists, there is going to be better perceivers than others. Not every radiologist is going to be a member of the set of *best* radiologists or most reliable radiologist. But they are still radiologists and can discriminate some things better than most people can. So, it's a fact about our cognitive abilities that some people have better ones than others, which is why we recognize some professional athletes and radiologists as better than others. So, I don't see why one should accept a principle that says if a cognitive ability or epistemic method *m* is reliable at getting the truth in domain *d*, then there will be a consensus among the community of participants pursuing truths in domain *d*. For if knowledge is an achievement, then getting the truth because of an ability is not going to be easy in domains where the subject matter is hard. Thus, it's no surprise consensus is lacking in some domains.

Or consider this way of thinking about it. There are only a small set of people who make it to the professional level in sports. So why assume when it comes to philosophy or other intellectual domains, humans are equally equipped to get the truth? More carefully: why assume that if any person can achieve cognitive contact with reality in one specific domain, then virtually everyone else has that same ability? Thus, I don't see why one should accept a principle that says *if a cognitive ability or epistemic method is reliable at getting the truth in domain d, then there will be a consensus among the community of participants pursuing truths in domain d*. Again, if knowledge is an achievement, then getting the truth because of an ability is not going to be easy in domains where it is hard. Granted, it would be nice if certain methods enjoyed the kind of consensus shared by others. But I'm not willing to grant that

methods which don't share the same degree of consensus as other methods therefore lack the kind of reliability needed for one to avoid rationality or normative defeaters.

Consider this analogy as a way to drive home the point. In a normal environment, all professional baseball players can hit a beach ball that is slowly tossed to them in a normal environment and circumstances from twenty feet away, 100% of the time. Such an achievement is easy to do. The fact that everyone can do it, indeed, indicates how reliable they are achieving such things. But not many players can achieve base hits more than 30% of the time. Indeed, out of the (roughly) twenty-thousand major league baseball players that have a lifetime batting average of .300 or higher, there are fewer than 200.[154] Because this kind of achievement is difficult, we should expect a small percentage of achievers. Similarly, once the domain moves from that which is epistemically easy to that which is hard, there is going to be pervasive disagreement. But this doesn't indicate that our cognitive abilities and methods are unreliable *tout court* any more than inferring from the fact that only one percent of all major league baseball players have a lifetime batting average of .300 or higher that some major league players are not reliable at hitting .300 *tout court*. It's just some cognitive achievements are very difficult while others are easier. The former achievements will sit in company with many who disagree; the latter will sit in company with many who agree. So, we shouldn't be surprised there is disagreement in domains like philosophy, religion, politics, and value theory given the subject matter is so difficult.[155]

---

[154] See Baseball Almanac, http://www.baseball-almanac.com/hitting/hibavg1.shtml

[155] Indeed, philosophers like Collin McGinn et al. don't think evolution has bequeathed to us faculties that are equipped to do philosophy, esp., in areas like philosophy of mind. More specifically, McGinn holds to 'transcendental naturalism,' which is the view that the brains evolution has bequeathed to us are, in principle, ill-equipped to solve many philosophical and scientific problems. See Colin McGinn, *Problems in Philosophy: The Limits of Inquiry* (Oxford: Blackwell, 1993).

This is not to say that math and science are not difficult subjects. Indeed, they can be tremendously difficult. But much of the subject matter is publically accessible unlike much of the subject matter in philosophy. God, moral properties, mental properties, intuitions, a priori knowledge, etc. are controversial entities that if they exist can't be publically disclosed, and thus, verified like we can with proofs in math. Nor can such entities, if they exist, be publically accessible, domesticated, and disclosed for others to verify as can be done in the hard sciences. Therefore, I only mean to say that humans may, in general, be better equipped to get at truths in science and math than truths in philosophy because that which is publically accessible can benefit from collaborative thinking and achieve levels of consensus that don't easily generalize to domains like philosophy.

Therefore, assuming that the aforementioned domains should fit the pattern of consensus we see in fields like math and science in order to determine reliability is too strong of a standard and overlooks the ontological differences within the subject matter, assuming such entities exist. Furthermore, because we lack an objective, external way—let alone a consensus—of determining what degree of reliability a philosophical method has that employs, say, intuitions or, say, perception of abstract entities, we can't say it's objectively unreliable. That is, we can't say definitively that intuitions or certain methods of belief formation fall below 50% reliability.

1.2 Consensus is Relative to Periods of Time or Current Paradigms

Second, premise (4) assumes consensus to be a sufficient condition for reliability. Goldberg and Kornblith, for example, appeal to the kind of consensus enjoyed by science. They use this data to argue for the reliability of science's methods. But I want to argue that it takes too rosy a view of science that overlooks the history of science and even its current state.

New discoveries, new theories, new instruments, and new paradigms have caused consensus opinion to shift and it will most likely shift again, esp., in light of the motivation to make general relativity compatible with quantum mechanics. Moving from geocentricism to heliocentricism is but one familiar example. But unless one is willing to adopt a coherence view of truth and its related problems, we must acknowledge that consensus—though a strong indicator of justification—is not sufficient for getting the truth.

In addition, because there is no neutral procedure for settling cases of disagreement over conflicting intuitions between two or more disputants within (say) philosophy, it doesn't follow that the sciences are immune from the same problem. For example, the hard sciences, which Goldberg and Kornblith take to be employing reliable methods as evidenced by the amount of consensus they enjoy, depends upon (among other things) the method of perception. But the method of perception itself can't be neutrally verified as reliable without relying upon it thereby begging the question. So why think scientific beliefs that currently enjoy a consensus is an indication of a type of reliability that is sufficient to get the truth? For there can certainly be disagreement over such "consensus" beliefs sourced in basic perception. To appreciate these possibilities consider the following case.

> Suppose that there were a diversity of sense perceptual doxastic practices as diverse as [the belief outputs sourced in intuitions] are in fact. Suppose that in certain cultures there were a well established "Cartesian" practice of seeing what is visually perceived as an indefinitely extended medium that is more or less concentrated at various points, rather than, as in our "Aristotelian" practice, as made up of more or less discrete objects scattered about in space. In other cultures we find a "Whiteheadian" [perceptual method] to be equally socially established; here the visual field is seen as made up of momentary events growing out of each other in a continuous process. Let's further suppose that each of these practices serves its practitioners equally well in their dealings with the environment. We may even suppose that each group has developed physical science, in its

> own terms, to about as high a pitch as the others. But suppose further that, in this imagined situation, we are as firmly wedded to our "Aristotelian" [perceptual method] as we are in fact. The Cartesian and Whitheadian *ausländer* seem utterly outlandish to us, and we find it difficult to take seriously the idea that they may be telling it like it is. Nevertheless, we can find no neutral grounds on which to argue effectively for the greater accuracy of our way of doing it.[156]

Notice that such a case does not involve an object that cognizers believe appears to them a certain way and then they allow their metaphysical theories to inform them of the true nature of the object. Rather, the Cartesian sees the object *as* a continuous medium, which is taken to be an accurate description of reality. Similarly, the Whiteheadian believes the appearance of the object *as* a flux of momentary events accurately describes reality. Because each community shares equal success in predictability and advancing science, there is nothing epistemically irresponsible or irrational about each scientist's beliefs, despite the extant disagreement over the nature of material objects. (And scientists today may, indeed, disagree over the ontological properties had by objects they perceive and/or disagree that the "objects" perceived really are individuated substances. Here the alleged consensus is the result of extant practices and conventions for how the scientific community have agreed to talk about their findings and leaving aside their metaphysical analysis of "objects" they perceive.

Furthermore, despite the Aristotelians, Cartesians, and Whitheadians lacking noncircular reasons for the reliability of their method for forming perceptual beliefs, it's not irrational for them to continue forming beliefs the way they do. One of the three may be employing the most reliable method, but there is no objectively neutral, noncircular way to demonstrate it as such. Alston's envisaged case, then, captures why a lack of consensus is not sufficient to establish a method as unreliable. Accordingly, a lack of consensus is not

[156] William P. Alston. *Perceiving God: The Epistemology of Religious Experience* (London: Cornell University Press, 1991), pp. 273-274.

sufficient to provide a subjective rationality defeater for beliefs sourced in philosophical method full stop. The fact there is disagreement over the outputs of intuitions is, for example, no reason to reject its reliability than it is to reject the reliability of perceptual belief outputs because people disagree over the ontological status of the "objects" perceived.

1.3 Intuitions Aren't Fixed and Can Possibly Become More Accurate Over Time

One might think that because philosophy relies so heavily upon intuitions, that philosophical *method* is unreliable because of all the conflicting philosophical intuitions employed by philosophers. Many philosophers take intuitions to be a method of sorts that aid us in getting the truth. But because intuitions produce contrary outputs relative to other philosophers' intuitions, one might argue that they can't be taken as a reliable. In other words, it's as if intuitions are like a calculator everyone has been issued; and when we consult our own calculator to answer a specific problem we learn we each have contrary answers. Thus, we rightly conclude that our calculators are not reliable since they aren't converging on the same number. This kind of argument from analogy may be used in support of premise (4). But this assumes intuitions are fixed, unalterable, pieces of innate cognitive equipment that have been bequeathed to us in a "one size fits all" version. But surely intuitions can be refined or calibrated through specialized training in order to be more accurate thereby improving one's cognitive system over time and, perhaps, accompanied by some good luck. Indeed, Goldberg himself says as much by granting that some peoples' intuitions may track truth better than others.[157] For if intuitions are but only

[157] Sanford Goldberg, "Reliabilism in Philosophy," *Philosophical Studies* vol. 142 (2009): 107. Goldberg doesn't argue that intuitions are unreliable full stop. He thinks there may be some philosophers whose intuitions are highly reliable just like the mathematician Ramanujen has highly reliable mathematical intuitions. What Goldberg does argue, however, is that philosophical intuitions can't be checked against some proof unlike

one operative faculty working within one's overall cognitive system, then they are going to be influenced, monitored, and governed by countless contributing factors.

Furthermore, intuitions would, then, be integrated with one's defeater system—working within the overall cognitive system—for better or for worse. The more cognitively integrated one's intuitions are with one's defeater system (thereby making intuitions sensitive to counterevidence) and overall cognitive system, the more stable and, perhaps, reliable the outputs. The less integrated one's intuitions, the less stable and unreliable the outputs will be. So, if something like this view of intuitions is even close to being true, then we should not be surprised the outputs of one's intuitions will, at times, conflict with another's intuitions. Indeed, after learning about all the potential symmetry breakers that can vitiate the evidence condition in Chapter Three, it seems one's intuitions will be uniquely influenced to one degree or another by all the varying kinds of operative evidence in one's cognitive system influencing the targeted intuition. Therefore, the significance of intuitions producing contrary outputs should only be an indication of unreliability were the conflicting outputs coming from the *same* intuition. In other words, if one were to evaluate, say, the same Gettier-style counterexample over a period of time and one's intuition on whether or not S knows were in a constant state of flux thereby undermining itself, then, intuitions would be an unreliable component to philosophical method.

---

mathematical intuitions which can be checked. One can certainly grant this distinction yet still argue that Goldberg is comparing apples with oranges. In other words, just because mathematical intuitions can be checked against some proof, it doesn't follow that philosophical intuitions can't be trusted because we lack such a proof to verifying their accuracy. Indeed, this is but only one way to illustrate our cognitive predicament—namely, that we can't provide a non-question-begging argument for the reliability of our cognitive faculties without eventually using the very faculty that is being challenged as unreliable. Furthermore, the subject matter of math and philosophy may very well be radically different. At an ontological level, intuitions may have formed from previous events involving perception. As section 1.4 below illustrates, there may exists moral properties that are discovered via perception and thus informs our subsequent intuitions relative to considered cases. If that's true, then such intuitions can't, in principle, be checked by a proof the way mathematical intuitions can because moral properties are discovered via perception and not deduced from a proof.

1.4 Some Intuitions Possibly Track Truth Better Than Others

Perhaps some peoples' intuitions are more reliable than others. For example, some people share the same moral intuitions which cause them to make the same moral judgments after evaluating cases. Others, who lack those intuitions, make contrary judgments. Notice, however, that neither disputant has a noncircular proof for the reliability of their intuitions. More importantly: epistemologists not only lack a noncircular proof for the reliability of our faculties (in general), it's hard to conceive of what one would even look like.[158] Given this: why think one needs a noncircular proof for the reliability of one's intuitions in order to justify one's beliefs that derive from them? To put the point differently, just because one lacks a noncircular proof for the reliability of the method from which one's belief derives, it doesn't follow that one is irrational or has a normative defeater undercutting one's subjective justification. For if such a proof was required, then virtually all of our beliefs are irrationally held since even perceptual beliefs can't be vindicated as reliable from a neutral standpoint without begging the question.

Granted, we do well by appealing to the superior track record satellite images have in predicting the weather over the method of appealing to groundhogs and aching joints.[159] But the problem of lacking a neutral way of determining which intuitions are more reliable at producing true beliefs than others is not restricted to philosophy. Quantum physicists disagree over the determinacy or indeterminacy of quantum events; psychoanalysts disagree with behaviorists over what constitutes legitimate behavioral data; biologist disagree over whether or not genes encode information about phenotypic traits; etc. But it's too strong

---

[158] One might consider how even an omniscient person could provide such a proof. This person presumably knows all true propositions and doesn't believe any false ones. But other than the authority of this person's testimony, it's hard to see how a noncircular proof can be provided, despite knowing all true propositions.

[159] I borrow this example from William Alston. See William P. Alston, *Perceiving God*, p. 271.

to suggest such disputants are irrational for believing what they do in these domains despite lacking a non-question begging argument for why their method is more reliable than their disputants. Nor do they seem to suffer from a normative defeater or subjective rationality defeater against their belief that they know such things. Such disagreement may provide a partial-defeater (at best) or a potential-defeater against one's belief. But a full subjective rationality defeater does not seem to be the right way to normatively evaluate cases of disagreement. To get at this put a little differently, let us consider a familiar case of moral disagreement.

> TORTURING CHIDLREN: Jamie thoroughly enjoys torturing children in ways I will protect the reader from hearing. Each episode is the same: Jamie spends days torturing the new innocent child that has been recently abducted in front of the child's parents before eventually murdering, first, the child, and then the parents, just so the parents can see their children in agony before they experience their own death. Jamie is so elusive and knows how to avoid getting captured by others, including the legal authorities. Jamie thereby continues terrorizing a community that is stricken with grief and paralyzed by fear of the next attack.

Suppose a non-philosopher or layperson was asked to evaluate the moral status of Jamie's actions. Further, suppose such a person was asked to morally evaluate this case alongside two redoubtable professional philosophers who specialize in value theory. One is an ethical egoist; the other is a moral nihilist.[160] The ethical egoist believes Jamie's actions are not morally wrong because Jamie's self-interests are being satisfied. Of course, the ethical egoist may accede to having the intuition or sentiment that such an act is "wrong." But the ethical egoist also has some kind error theory that undercuts the reliability of the intuition's ability to track truth. Furthermore, the ethical egoist places a greater value on satisfying one's self-interests. So, the error theory and primacy of the value of self-interest can, in principle, provide justification for such acts of torture.

[160] This case was inspired by Michael Bergman. See Michael Bergman, "Rational Disagreement After Full Disclosure," in *Episteme* (2009): 345-346.

Regarding the moral nihilist, s/he believes Jamie's actions are not morally wrong because such an evaluation is based on a false assumption: namely, objective moral values exist. Since there are no objective moral values, there are no morally right or wrong actions—period. So the moral nihilist may share an intuition or sentiment that such acts are "bad." But this intuition or sentiment is a mere appearance that doesn't accurately capture anything about reality—specifically, that objective moral values exist.

Now suppose that each disputant disclosed all the relevant reasons, evidence, and arguments for both their ethical theories and evaluation of the TORTURING CHILD case. Suppose further that the only ground for the layperson's belief that Jamie's actions are wrong is an intuitive insight—that is, s/he just "sees" that Jamie's actions are wrong in a basic, non-inferential way. The other two disputants may, perhaps, share this intuition; nevertheless, they both have an error theory for why their intuition doesn't track truth.

Now the upshot: the fact there is disagreement between the disputants at the level of conflicting intuitions doesn't imply that the layperson, or anyone else for that matter, is being epistemically irresponsible or irrational for retaining belief in the face of disagreement. Nor does it demonstrate that intuitions are an unreliable method for evaluating the case. That is, the disagreement itself is not enough evidence to conclude intuitions are an unreliable method to get the truth. Therefore, disagreement that is sourced in a philosophical method that employs intuitions doesn't provide a subjective rationality defeater to one's belief that Jamie's actions are wrong; moreover, disagreement alone is not enough to undercut that such a belief is an instance of knowledge.[161]

[161] Indeed, I agree with Bergman who says, "Instead of being moved to doubt the reliability of our own beliefs on the topic, we are moved to feel badly for the ones with whom we disagree and to be glad that we are fortunate enough not to lack the [intuition] we have or to have the misleading apparent insights that they have." See Bergman, "Rational Disagreement After Full Disclosure," p. 346.

In light of these reasons, it seems the problem of philosophical method (in general) and intuition (in particular) lacking a consensus amongst its practitioners is not enough to undercut the justification and potential knowledge such methods can deliver. The fact that one cannot provide an objectively neutral, noncircular reason for the reliability of philosophical method and intuition is not a problem unique to philosophical method. Indeed, it is a problem that plagues all basic sources of justification and knowledge. Even perception, which is the method used by scientists, is not immune to this problem, despite the kind of consensus some of the hard sciences enjoy. Just ruminate over the typical case involving eyewitnesses testifying to the facts about a car accident. There will inevitably be discrepancies once we compare each independent report. So, unless we are willing to relegate sources like perception and memory to the bin of other unreliable sources like wishful-thinking and Ouija boards, then inter- or external-inconsistency alone is not sufficient to undercut reliability. If it were, we wouldn't be left with any reliable sources.[162] Goldberg and Kornblith do well, then, to heed Alston's advice, which traces back to Reid, when he says, "[A]n absolute ban on inconsistency in output will deprive us of all sources of justification, [so] it is clearly the better part of wisdom to require only that the source not generate 'persistent and massive [internal] inconsistency.'"[163] I think once we recognize

[162] I'm not denying that perception used in easy domains like distinguishing middle-sized objects from other ones is more reliable than intuitions used in difficult domains like philosophy, morality, and politics. What I am pointing out, among other things, is that we lack a metric to determine what degree of reliability a method must meet in order to be epistemically responsible for employing it. The degree of reliability perception must meet is going to be much higher than that of a person we deem reliable at getting base hits. Having the ability to perceive a vehicle bearing down on you as you cross the road at only a 30% success rate is not going to meet any standard for reliability. But getting base hits 30% of the time in Major League Baseball is a degree of success taken to manifest reliability. And any Major League coach would be rational in putting such a hitter at the top of their lineup. But because we lack a neural, non-circular way of verifying the reliability of many belief forming sources we use in an array of domains, we lack the data and resources needed to come up with a definitive answer to how reliable some methods in certain domains need to be in order to be rationally employed. I borrow this point from William Alston in *Perceiving God*, p. 238.

[163] William P. Alston, *Perceiving God*, p. 235. The scope of persistent and massive inconsistency seems more relevant to inter- and intra-faculty coherence, first. And then intra-social coherence among people working

one of the goals of philosophy is to get the truth in domains that are by their very nature difficult and are, in principle, resistant to a neutral, noncircular way of demonstrating the truth to others, the benefits of convincing everyone—thereby achieving consensus—will be seen as misplaced. As a result, we will see disagreement as something we should expect.

2. THE ARGUMENT FOR SKEPTICISM FROM EPISTEMIC SUPERIOR DISAGREEMENT

2.1 Shifting From the Pseudo-Problem to the Real Problem

There are two reasons why peer disagreement, as represented by CCE, is a pseudo-problem. First, the rareness of epistemic symmetry obtaining undercuts the skeptical threat argument from peer disagreement. Second, CCE's account of a belief's pd-normative status is false because it carves out a normative domain for dealing with peer disagreement that abstracts away from external concerns and knowledge. The etiology of a belief matters when making normative evaluations. Furthermore, cases involving bare memorial knowledge and beliefs produced by sub-personal processes that can't be internally accessed provide additional reasons for why going internal is not the only game in town for evaluating a belief's pd-normative status.

In this section I argue that disagreement does present a real epistemological problem. But the way to frame the problem shifts, first, from peer disagreement to expert disagreement. And, second, the notion of a belief having pd-normative status involves whether or not disagreement provides a subjective rationality *defeater* against knowledge. One is epistemically incorrect for not undergoing some level of doxastic revision when

---

within a broadly similar inherited worldview. Less relevant is the kind of inconsistency existing at the level of inter-social coherence given one's inherited worldview replete with different epistemic norms and what Goldman called 'norm evidence' in Chapter Three helps explain the disagreement between two communities or traditions.

disagreement levels an undefeated defeater against one's subjective justification, which is a condition for knowledge.

The reason why the epistemic significance of disagreement requires shifting from peer disagreement to expert disagreement is this. In Chapter Three we learned how rare the conditions are for establishing epistemic peerhood. The number of relevant symmetry breakers on offer mitigates the force of the argument from CCE given it needs epistemic symmetry to motivate it. But what does seem to obtain quite often is finding a recognized epistemic superior who disagrees with you. Despite the conditions for epistemic symmetry rarely obtaining, the conditions for satisfying epistemic superiority can easily be satisfied. I take this to be the real problem posed by disagreement from a perceived expert or epistemic superior: namely, when does disagreement with a perceived cognitive superior provide one with a defeater against knowledge? But first we need to get clear on what is an epistemic superior. In the next section I provide a list of conditions that afford one the status of being an epistemic superior relative to another person.

2.2 Who Are My Epistemic Superiors?

When I use the notion of an epistemic superior, I simply mean someone who has both more intellectual prowess and has spent considerably more time thinking hard about some given topic. More exactly, some S is my epistemic superior regarding p if and only if S satisfies the following seven conditions:[164]

---

[164] I indebted to Bryan Frances for both inspiring this kind of argument for skepticism from epistemic superiors and providing some of the necessary conditions constitutive of an epistemic superior. See Bryan Frances, "The Reflective Epistemic Renegade," *Philosophy and Phenomenological Research* vol. 81, no. 2 (2010): 419-463. Such necessary and jointly sufficient conditions may seem too strong for one to be your epistemic superior. But I am looking for something strong in order to achieve a level of reliability that surpasses one's own. For if I knew that I exceeded a disputant in one of the criteria listed in (a) – (g), then I would have a symmetry breaker on offer that would prevent the disputant from having the kind of reliability I take epistemic superiors initially to possess. I say initially because I do work in Chapter Seven that provides reasons for thinking the title 'epistemic superior' doesn't entail more reliable at getting the truth in all cases.

(a) S is generally more informed than I am on the topics involving p.
(b) S has more raw intelligence than I do.
(c) S has thought about and investigated whether p is true longer and in more depth than I have.
(d) S has thought about and investigated the topics surrounding p longer and in more depth than I have.
(e) S is more intellectually careful than I am.
(f) S is less relevantly biased than I am.
(g) S has a significantly greater publication and citation record—that is, S is recognized in the field as an expert given the number of peer reviewed publications S has written and the number of times S's work has been cited.

So, for example, when it comes to the topic of, say, metaphysics, it's easy to *recognize* that Saul Kripke, Patricia Churchland, David Lewis, John Hawthorne, David Chalmers, and the like meet these conditions, and thus, are my epistemic superiors. Indeed, they are my epistemic superiors in other domains as well. To make the point more technically precise, such philosophers can "kick my philosophical ass."[165]

The fact that there are many people who are our epistemic superiors is no surprise. This explains why we take our cars to mechanics, our bodies to physicians, and our computers to technicians when we can't figure out how to solve our respective problems. But I want to emphasis that we are conscientiously aware and recognize that such people exist. Indeed, whenever we read their work, hear them speak, or engage in conversation with them, we conscientiously recognize there are epistemic superiors with whom we disagree. So the question I want to explore is this: does a case of disagreement with a recognized epistemic superior provide one with defeating counterevidence against one's claim to know? More specifically, can epistemic superior disagreement mount a skeptical threat to many things we claim to know, including some of our cherished beliefs? In the next section I develop an argument for skepticism towards knowledge from epistemic superior disagreement.

[165] I borrow this hilarious phrase from Bryan Frances. See ibid., p. 419.

2.3 The Skeptical Argument From Epistemic Superiors (SAFES)

The following argument need not be restricted to a version that utilizes the role relevant possibilities play in providing potential defeaters. But because I am assuming a virtue theoretic account of knowledge in this dissertation, I formulate the argument in a way that provides serious challenges to externalism (in general) and a virtue theoretic account of knowledge (in particular).

In Chapter Three we saw that there are six kinds of symmetry breakers—in most cases of peer disagreement—that can and should be appealed to in order to undercut the case for epistemic peerhood obtaining. Because one or more of these symmetry breakers can often be appealed to, the threat of CCE's argument for moderate skepticism can be deflected. But the epistemic significance of peer disagreement—specifically, the skeptical threat it puts forth—is still a problem. For even though the antecedent conditions for epistemic peerhood rarely obtain, what does seem to obtain more often is epistemic superior disagreement. We are much more confident there are people who are our epistemic superiors and disagree with us than we are at trying to establish epistemic symmetry through satisfying implausible evidential conditions for epistemic peerhood. I argue that the skeptical argument from CCE can be recast as an argument from disagreeing *epistemic superiors*, rather than disagreeing "*epistemic peers*."

The skeptical argument from epistemic superiors (hereafter SAFES) is traditional in form. It says that in order to know something, one must rule out relevant alternative possibilities to one's belief that p. The kind of possibilities, of course, will be the fact that many epistemic superiors believe the denial of what one believes or believe some theory *t* that entails the denial of one's belief that p. So, for example, consider a graduate student who believes p. This student presents a paper defending p to her colleagues, which includes

the department faculty. The department faculty is admittedly her epistemic superiors. She learns that the department faculty believes some theory *t* that entails the denial of p. Moreover, she understands how p is inconsistent with *t*; more exactly, she sees how *t* entails the denial of p. She also learns that the faculty members are not persuaded by her argument in defense of p, and thus, haven't reduced their confidence in the denial of p one scintilla. So, unless she can rule out this possible alternative, namely, that theory *t* is true, it seems she will respond in an epistemically incorrect way were she to retain her belief in p given the expert disagreement provides counterevidence that needs to be ruled out in order to know p. Therefore, unless she can provide a defeater-defeater against *t*, she cannot know p given her belief lacks a positive pd-normative status. Here is the argument expressed more formally.

SAFES

(1) If …
   a. S is aware that X is an epistemic superior, and
   b. S is aware that X believes *t*, and
   c. S understands that (if *t*, then not-p), and
   d. S can't rule out *t*,

   …then S doesn't know p.

(2) (a) S is aware that X is an epistemic superior, (b) S is aware that X believes *t*, (c) S understands that (if t, then not-p), and (d) S can't rule out *t*.

   Therefore,

(3) S doesn't know that p.

This reconstructed argument is plausible in that it doesn't appeal to what many think are pre-theoretically irrelevant scenarios such as cases involving malevolent demons, brains in vats, Matrixian pods, etc. The pre-theoretical obviousness that one is not in a vat of chemicals doesn't generalize to the more modest claim that one needs to rule out an

expert's belief in theory *t*, which entails the denial of p. This kind of argument for skepticism is what Bryan Frances has recently called a *live*[166] one—that is, there are real live cognitive experts who disagree with us over many things. E.g., there are epistemic superiors that say there is no such thing as pain or that there is no such thing as free will or that there are no objective moral values or that there is no color in the world or that there is no such thing as consciousness or that there are no character traits and so on.[167]

2.3.1 Premise (1) of SAFES

How plausible are the premises? The first thing to say about SAFES is that satisfying the conditions for an epistemic superior are easier than satisfying the conditions in the epistemic principle behind ES[168] in the CCE argument. This is because it is much easier to recognize there are cognitive superiors in many different domains that we are not epistemically strong in.

Regarding premise (1), first, it is intuitively plausible. Just as proponents of CCE appeal to inanimate objects like conflicting thermometers to motivate their case for

---

[166] Bryan Frances, "When a Skeptical Hypothesis Live," *Noûs* 39, 4 (2005): 559-595. Frances uses 'live' to denote the number of error theories that are currently held by respected members of both the philosophical and scientific community. Because one can encounter these live hypotheses by interacting with community members the threat of skepticism is live and not a logical possibility that abstracts away from one's environment. Similarly, I think the existence of epistemic superiors who disagree with us also presents a live hypothesis that functions as a defeater against knowledge that needs to be ruled out with some kind of defeater-defeater.

[167] Those who have endorsed error theories regarding beliefs, which is to say that beliefs are not a mental state that captures one's mental assent towards a proposition or statement, are Patricia Churchland (1986); Paul Churchland (1989); Stephen Stich (1983); Daniel Dennett (1978); Paul Feyerabend (1963); Richard Rorty (1970); Willard Quine (1960; 1985). Those who have endorsed an error theory for pain are Bertrand Russel; Broad; Gilbert Ryle; David Lewis; and David Armstrong. Those who endorse an error theory for color are Galileo; Larry Hardin; Paul Bogossian; Daid Velleman; Emmett Holmon; and Frank Jackson. The scientists who have done likewise are Semir Zeki; Stephen Palmer; Werner Backhous; and Randolf Menzel. Those who have endorse an error theory regarding character traits are Gilbert Harman (1999; 2000); and John Doris (1998, 2002). I'm indebted to Bryan Frances for putting together this list of various error theorists. See Bran Francis, "Live Skeptical Hypotheses" John Greco, ed., *The Oxford Handbook of Skepticism* (Oxford: Oxford University Press, 2008), p. 230.

[168] ES: In a case of disagreement over the truth of some proposition p wherein both parties acknowledge each other to be epistemic peers and no symmetry breaker is on offer, then epistemic symmetry has obtained.

disagreement requiring significant doxastic revision, SAFES is even stronger. The reason is that CCE appeals to analogies like two conflicting thermometers whereas SAFES appeals to a thermometer one believes is more reliable than his or her own thermometer. The driving intuition behind SAFES is that you don't know the temperature by relying upon your thermometer when it conflicts with a thermometer deemed more reliable than yours. Such counterevidence provides a potential defeater that first needs to be ruled out in order to know.

Second, premise (1) is similar to plausible closure principles. Closure principles try to capture how knowledge is closed under some function. The function is typically some variation of deduction—specifically, 'competent' deduction. For if S deduces ~p from *t* by employing an invalid inference or from an unreliable process like wishful thinking or by not retaining knowledge of p throughout the inference, then arguably S doesn't know. But if S knows that *t*, and competently deduces ~p from *t* and believes *t* entails ~p, then S knows that ~p. Further conditions can arguably be added and much ink has been printed in doing so.[169] But the point is that there seems to be plausible closure principles at work that support the idea that knowledge is closed under competent deduction.

For example, if I know that I'm presently typing this sentence at home on my computer, and I know that typing on my computer entails I'm not at the golf course, then I know that I'm not presently playing golf.[170] A more popular expression of this principle, which aims to capture how knowledge is closed under entailment, goes like this. If I don't know that I'm not a handless brain in a vat, then I don't know that I have hands. For to

[169] For further references on how epistemic closure principles have become more precise and additional conditions added to the antecedent condition see David and Warfield (2008); Harman (1986); Williamson (2000); Kvanvig (2006).

[170] Say I know the implications of, say, the axiom of localization whereby objects like me can't exist in different spatial locations at the same time.

know that I have hands while at the same time not knowing that I'm not a handless brain in a vat is to believe in what Keith DeRose calls an "abominable conjunction."[171] Abominable conjunctions look like this.

(i) I know that I am the offspring of Jacqueline Garofalo and Gary Osmundsen, but I don't know that I'm not the result of the first replicating cell synthetically created in a lab.
(ii) I know that my body is not encapsulated in a Matrixian pod, but I don't know that this body I perceive and touch is not the result of electrical stimuli programmed by aliens into my brain, which is located in a Matrixian pod.
(iii) I know that I'm presently writing my dissertation, but I don't know that I'm not presently on the golf course putting for birdie.

To deny plausible epistemic closure principles is to epistemically embrace the absurdity of these conjunctions. But premise one is even more plausible than epistemic closure principles for the following reason.

There are theories of knowledge, for example, that say certain possibilities can properly be ignored.[172] Either some possibilities are deemed pre-theoretically implausible because they are wild and outlandish (e.g., brain in vat scenarios) or some possibilities are not relevant given the standards for knowledge change according to one's context, *a la* contextualism, when we make knowledge attributions. For example, regarding contextualism, the possibility that the bank is closed tomorrow in a situation where you need to deposit a check sometime within the next two weeks so the check you just wrote doesn't bounce is considered a low stakes situation. Contrast that situation with one where you need to deposit a check by tomorrow or else the check you just wrote will bounce thus

[171] Keith DeRose, "Solving the Skeptical Problem," *Philosophical Review* 104 (1995): 1-52. Granted there are some who deny that knowledge is closed under entailment (e.g., Robert Nozick (1981:206) and Fred Dretske (2003:113; 2005:19-20), but debating the veracity of this principle will take this dissertation too far afield from it focus. So I will assume that there is a plausible formulation of the epistemic closure principle that is true.

[172] For example Davis Lewis says that S knows that p just in case S's evidence eliminates every possibility in which not-p—Psst!—except for those possibilities that we are properly ignoring. Lewis provides a list of rules for adjudicating which possibilities can be ignored. But possibilities the subject believes or ruminates on cannot be ignored. See David Lewis, "Elusive Knowledge," *Australasian Journal of Philosophy* 74 (1996): 549-67.

causing you all sorts of problems, including losing your job. It's popular now to say we can attribute you with knowing the bank is open tomorrow in the former, low-stakes case. But we can't attribute you with knowing in the latter, high-stakes case. The reason is that in the high stakes case the possibility that the bank is closed can't be properly ignored given the severe consequences of being wrong.[173]

But the possibility that you are wrong in light of the fact that there are epistemic superiors who disagree with you and you are *aware* that they do—that is, you have either read their work or directly heard from them—prevents you from ignoring this alternative possibility. Indeed, this alternative possibility is counterevidence that seems to level a potential defeater against your belief that needs to be ruled out. The idea that this potential defeater for your belief is irrelevant or can be properly ignored seems to violate your belief's pd-normative status. It's not as though you have a belief and aren't aware that your cognitive superiors disagree with you. You *are* aware of the disagreement! Of course, the scope of this disagreement's liveness is restricted to only those who in fact encounter the disagreement from epistemic superiors. So the threat of SAFES is restricted to those who meet the following five conditions.[174]

(a) You believe that p.
(b) You believe that theory *t* entails the falsity of p and your epistemic superior believes *t*.
(c) You understand how *t* entails the denial of p.
(d) You have encountered an epistemic superior who believes *t*, and thus, disagrees with you. So the alternative possibility that the epistemic superior is right and you're wrong is a live alternative possibility.

---

[173] Proponents of contextualism are Cohen (1988; 1998); Lewis (1996); DeRose (1995). Contextualism is a thesis about how the term 'knows' functions in sentences that are used when making knowledge attributions. It is not a thesis about how the conditions for knowledge per se or the conditions for when one knows change according to context. Knowledge attributions, then, are interested in different things than traditional epistemology, which typically is concerned with what is knowledge, warrant, and justification.

[174] I'm indebted to Bryan Frances for picking out some of the features that constitute 'liveness' and how those features inform us regarding what is a relevant alternative that needs to be ruled out before one can claim to know. See Brayn Frances, "Live Skeptical Hypotheses," p. 239.

(e) You recognize that disagreeing with an epistemic superior is live potential defeater for your claim to know p.

So the scope of SAFES is limited; it doesn't level a potential defeater unless one comes in contact with an epistemic superior who disagrees with you. Nevertheless, because you can't rule out the epistemic superior's belief or theory that entails the denial of your belief, you don't know. Indeed, all the non-epistemic-superiors fail to be in a position to know the epistemic superiors are wrong.

2.3.2 (1.A) and (1.B) of SAFES

Premise 1, clause (a), is uncontroversial and can be easily satisfied. It only requires one perceives that there is another well-functioning, healthy person who satisfies the seven conditions for being an epistemic superior, as laid out in § 2.2 above; moreover, one believes such a perceived person disagrees with the targeted proposition one believes to be true. As long as one recognizes there is a person who is one's epistemic superior and this person believes some theory *t* that entails ~p, then one is aware of a case of disagreement with an epistemic superior.

2.3.3 (1.C) of SAFES

Premise 1, clause (c), is one of the conjuncts comprising the antecedent conditions in premise (1) that needs to be satisfied. It also aims to alleviate any worries that the operative epistemic closure-like principle in premise (1) is not vulnerable to unrefined versions of epistemic closure. For example, to say S knows that p entails q is not to say that the deductive function that's operative is some implausible one such as:

If S knows that p, and *if p entails q*, then S knows that q.

Only an omniscient person can deduce every logical consequence of what is known. This is not what's going in the kind of cases envisaged by SAFES. For example, once the closure

principle is stated this way, it follows that from knowing one mathematical axiom then one can also know all the mathematical theorems that fall out of that axiom.[175] Or some make a distinction between knowing and being in a position to know. So although S knows that p, and knows that p entails q, then S is only in a position to know that q. Others want to say that although one knows that p and knows that p entails q, one may still fail to know q because this person simply can't form the belief that q. And if one can't believe that q, then one can't know q, assuming knowledge requires, among other things, a belief.[176] These are interesting debates but would take us too far afield from the purpose of this dissertation. The point of clause (c) in premise 1 is to make clear that S understands through competent deduction that if theory *t* is true, then S's understands that p is false.

2.3.4 (1.D) of SAFES

Premise (1), clause (d) is another one of the conjuncts comprising the antecedent conditions in premise (1). It is supported by the following reasons. First, it indicates that S can't rule out the plausible, live alternative possibility: namely, that S's epistemic superior might be right and S wrong given what she knows. Clause (d) assumes something like the following plausible epistemic principle is operative in SAFES.

A person knows that p on the basis of evidence *e*, only if *e* rules out all *relevant* alternative possibilities to p.[177]

A case of disagreement with an epistemic superior, then, is akin to epistemic principles governing scientific inquiry. If a scientist thinks hypothesis h best explains evidence set *e*,

---

[175] I borrow this example from Peter Baumann, "Epistemic Closure," in Sven Bernecker and Duncan Pritchard, eds., *The Routledge Companion to Epistemology* (New York: Routledge, 2011), p. 597.

[176] Gilbert Harman raises this objection in *Change in View: Principles in Reasoning* (Cambridge, MA: MIT Press, 1986), p. 11-12.

[177] I'm indebted to John Greco for the formulation of this epistemic principle. See John Greco, *Putting Skeptics in Their Place*, p. 57.

but is also aware that there are competing hypotheses h_1, h_2, …, h_n that are also consistent with evidence *e*, then her belief cannot retain the kind of epistemic status needed for knowledge.[178] Because a set of evidence *e* for some hypothesis h doesn't entail h, it's ill-equipped to rule out other competing hypotheses that also plausibly explain the datum. Such hypotheses are relevant and function as potential defeaters against one's knowledge if they are not ruled out or, as it were, at least neutralized with a defeater-defeater.

Second, clause (d) captures why Laurence Bonjour's clairvoyant doesn't know. Below we will look at the case in more detail. But the point of the case is to provide a counterexample to reliabilist theories of knowledge. A person that has, say, a clairvoyant faculty that reliably provides one with true beliefs is not sufficient for knowledge. For one can encounter strong counterevidence that intuitively provides one with either a rebutting or undercutting defeater. So, if one has a clairvoyant faculty that reliably produces true beliefs about, say, the president's whereabouts, it doesn't follow that one knows where the president is. Why? Answer: because one can also be exposed to strong counterevidence like media reports and video footage of the president being elsewhere. This evidence is misleading. But since it's such strong counterevidence that can't be ruled out, we intuitively make the verdict that the clairvoyant doesn't know the president's whereabouts. And were the clairvoyant not to rule out the counterevidence and continue believing what the clairvoyant faculty produced, we would say the clairvoyant didn't respond in an epistemically correct way. And since knowledge seems to require a subjective justification condition, we would say the clairvoyant doesn't know, despite having a true belief.

Similarly, one needs to rule out the counterevidence that comes from cognitive superior disagreement. One knows a cognitive superior believes some theory t which

[178] I'm indebted to John Greco for pointing out this kind of epistemic principle. See John Greco, Ibid.

entails the denial of one's belief. So, one needs to rule out this counterevidence or fail to satisfy a subjective justification condition for knowledge. For a true belief sourced in a reliable process is not sufficient for knowledge.

Third, depending on the details of some of the cases involving disagreement with an epistemic superior, clause (d) can be supported by Humean arguments that conclude there is no way of justifying the following assumptions without arguing in a circle[179]: (i) the way things appear is a reliable indication of the way things are and (ii) past observations are a reliable indication of future states of affairs. Because there is no necessary relation between (i) appearances and reality and (ii) past observations and future states of affairs, one's evidence set *e* neither entails nor makes probable—in a non-circular way—one's belief that ~p is true or some competing hypothesis h or theory *t* is false.[180] So, for

---

[179] David Hume, *An Enquiry concerning Human Understanding and concerning the Principles of Morals*, 3 ed. L.A. Selby-Bigge, 3rd ed. (Oxford: Oxford University Press, 1975).

[180] According to Hume, all of our beliefs about the external world depend upon an unjustifiable assumption: namely, (A) that the way things appear to us is a reliable indication of the way things, in fact, are. But this assumption can't be justified because attempts to do so result in either circularity or denying the contingency between our appearances of the world and the way the world, in fact, is. Stated more formally, the assumption can't be justified because:

1. All my beliefs about the external world depend for their evidence on both a) the way things appear to me, and b) an assumption (A) that the way things appear to me is a reliable indication of the way things are.
2. The assumption (A) in question is itself a belief about the external world.
   Therefore,
3. The assumption (A) depends upon itself for its evidence. (1,2)
4. Beliefs [that] depend upon themselves for their evidence can't be justified.
   Therefore,
5. The assumption (A) in question can't be justified. (3,4)

This argument shows the kind of circularity involved when trying to justify the assumption (A) the way other contingent claims are justified, namely, by appealing to the way things appear. But saying assumption (A) is justified on the basis that the way (A) appears to me is a reliable indication that (A) is true is viciously circular. And because appearances either do or do not accurately reflect reality, one can't justify the assumption by appealing to *a priori* reflection. Whether or not appearances do accurately reflect reality isn't something that can be known through *a priori* reflection given such facts are contingent and not necessary. Let's place this discussion in Hume's broader argument, which says there is no necessary relation between sensory appearances and perceptual beliefs. It goes as follows.

1. All my beliefs about the external world depend for their evidence on both a) the ways things appear to me, and b) an assumption (A)—the way things appear to me is a reliable indication of the way things really are.
2. But assumption (A) can't be justified.
   Therefore,

example, in a case where some epistemic superior, say, Daniel Dennett, believes there is no self that believes anything, one can't appeal to the way things appear to justify that's the way reality, in fact, is.

It's hard, then, to see where SAFES goes wrong. Dismissive responses that are typically employed against skeptical threats cannot be used. Nor can one dismiss this as an irrelevant possibility given the *liveness* of it being a plausible alternative possibility. Recall that the argument from CCE tries to trade on the idea that two epistemic peers function like two conflicting thermometers—perceived as indiscernible insofar as their reliability and exposure to evidence is concerned. But here the analogy would be more like you recognizing that your thermometer is less reliable than some other thermometer that displays a conflicting temperature relative to yours. Such a conflicting thermometer is not judged as equal, but rather, better—more reliable!—than your thermometer. As a result, the argument is powerful in that its scope is moderately broad thereby threatening many of our cherished beliefs. But it also calls into question some of our beliefs and assumptions about the nature of knowledge, warrant, and evidence. *Therefore, SAFES demands a response that says something substantive about the nature of knowledge and evidence.*

---

3. All my beliefs about the external world depend for their evidence on an unjustifiable assumption (A). (1,2)
4. Beliefs that depend for their evidence on an unjustifiable assumption do not count as knowledge. Therefore,
5. None of my beliefs about the external world count as knowledge. I don't know anything about the external world. (3,4).

I am indebted to John Greco for reformulating Hume's arguments into their most persuasive forms, which I have borrowed from in this footnote. See John Greco, "External World Skepticism," pp. 638-639.

3. SAFES' CHALLENGE TO RELIABILISM AND ITS THREAT TO KNOWLEDGE

3.1 Why SAFES Presents a Relevant Possibility That Needs To Be Ruled Out

The problem of epistemic superior disagreement, then, provides a relevant, potential defeater that one's belief is false, and thus, needs to somehow be ruled out. But what does it mean to say there is a possibility that first needs to be ruled out in order to know? To answer this question we need an account of what makes a possibility a relevant one regarding what does and does not need to be ruled out in order for one to know. More specifically, we need an account that harmonizes with our pre-theoretical intuitions towards possibilities that do and don't need to be ruled out in order to know.

In Greco's previous work he provides an account of what a relevant possibility is insofar as what type of skeptical hypotheses didn't need to be ruled out in order to know. This account was originally aimed at explaining why, say, brain-in-vat possibilities were not relevant possibilities that needed to be ruled out in order to know. Greco employed "possible worlds" semantics of modal logic to explain why our pre-theoretical intuitions towards different skeptical hypotheses needn't be ruled out because they were only possibilities in "far away" possible worlds—that is, worlds that are much further away from the facts constituting the actual world we find ourselves in. (Recall Greco's modal account of an ability, which implies that if one has an ability, then one has a disposition to achieve certain ends across a range of close possible worlds. And since worlds where one is the victim of an evil demon or a brain-in-vat are far away possible worlds, such skeptical hypotheses are irrelevant ones.)

So what is Greco's account of relevant possibilities? More exactly: what possibilities always need to be ruled out in order to know and which one's don't? Greco's principled answer comes only after consulting our intuitive judgments on particular cases. He appeals

to the epistemic significance of Tom-Grabit and Barn Façade cases to flesh out his account of relevant possibilities. For both kinds of cases deal with the relevant possibilities that intuitively need to be ruled out in order to know. For example, Tom-Grabit cases involve Tom and his identical twin brother, Jack. In some cases S knows about Tom's twin brother and in other cases S lacks such knowledge. The idea is that if you saw Tom take out a library book, but you didn't know that (a) Tom had a twin brother and (b) Tom's twin brother Jack was also in the library, then you wouldn't know Tom grabbed the book unless you ruled out the possibility that it might have been Jack. For if Jack grabbed the book then you would be tricked into thinking it was Tom. Similarly, there are the familiar barn façade cases: you wouldn't know that the object you were looking at was a barn unless you could rule out that it was a barn façade given the environment you perceived the barn in is peppered with barn facades and only a few genuine ones.

Once we change the environment, however, our intuitions change. In cases where Tom doesn't have a twin brother and cases where counties lack barn facades, one need not rule out the possibility that one is seeing Jack or a barn façade, respectively. But these intuitive judgments need to be expanded to nuanced versions of Tom-Grabit and Barn Façade cases.

Suppose Tom doesn't have a twin brother, Jack, or, say, Cleveland County lacks barn facades; nevertheless, S *believes* otherwise. So, even if Tom doesn't have a twin brother and such a world is, in fact, a far off possible world, the fact that S *believes* Tom has a twin brother, Jack, is sufficient to undercut S's justification for believing Tom took the book. Thus, S doesn't know Tom took the book unless the possibility that Jack grabbed the book is ruled out.

Another worry is when we describe the case this way. Tom doesn't have a twin brother, Jack, and S doesn't believe Tom has a twin brother. But S *ought* to believe that Jack has a twin brother given all the strong (albeit misleading) evidence S is aware of. Just like the case involving all the misleading evidence Bonjour's clairvoyant is aware of. Were one to continue believing p despite the kind of operative counterevidence one is aware, we would intuitively say one doesn't know because one ought to have undergone some type of doxastic revision. Let's consider Bonjour's case which illustrates this point:

> Case I. Samantha believes herself to have the power of clairvoyance, though she has no reasons for or against this belief. One day she comes to believe, for no apparent reason, that the President is in New York City. She maintains this belief, appealing to her alleged clairvoyant power, even though she is at the same time aware of a massive amount of apparently cogent evidence, consisting of news reports, press releases, allegedly live television pictures, etc., indicating that the President is at that time in Washington, D.C. Now the President is in fact in New York City, the evidence to the contrary being part of a massive official hoax mounted in the face of an assassination threat. Moreover, Samantha does in fact have completely reliable clairvoyant power, under the conditions that were then satisfied, and her belief about the President did result from the operation of that power.[181]

Even though all the strong evidence (albeit misleading) points to the president's whereabouts being in Washing D.C., the clairvoyant intuitively doesn't know—despite having a reliably formed true belief — that the president is in New York City. Why? Answer: The clairvoyant hasn't successfully ruled out all the news reports saying otherwise.

What these normative evaluations surface, according to Greco, is this: an adequate account of what constitutes a relevant possibility that needs to be ruled out in order for S to know depends on whether the relevant possibility q is, in fact, true; and (a) q is true in some close possible world; or (b) S believes that q is true; or (c) S ought to believe that q is likely to be true. Here is what Greco takes to be the sufficient conditions constitutive of a

[181] Laurence Bonjour, "Externalist Theories of Empirical Knowledge," *Midwest Studies in Philosophy* 5 (1980): 53-73.

relevant possibility (RP) that needs to be ruled out in order to know spelled out more formally:

(RP) q is a relevant possibility with respect to S's knowing that p is true if and only if
- i. If q is true, then S does not know that p is true, and
- ii. Either (a) in some close possible world q is true, or
  (b) S believes that q is likely to be true, or
  (c) S ought to believe that q is likely to be true.

What's so powerful about epistemic superior disagreement is that it's a relevant possibility that can satisfy conditions (i) and (iib) and (iic)—like the clairvoyant case above. But it can, sometimes, also satisfy condition (iia). Indeed, one can think of the clairvoyant case as involving epistemic superiors since the news reporters are presumably more reliable at getting such truths than Samantha. For example, take a case of superior disagreement involving philosophers who are better than you at doing analytic philosophy. Even if the majority of professional philosophers are not, say, eliminativists today, eliminativism is still a view that could be the majority view in some nearby possible world. Once we tell a story of how eliminativism could have easily become the dominate view if (i) certain realist philosophers were to choose some other vocation or (ii) eliminativism was adopted as orthodoxy amongst all the leading influential philosophers of the day, or some other plausible alternative, we can see how easily such a possible world is not a far off one like worlds involving brain-in-vat scenarios. So the problem of epistemic superior disagreement is a relevant possibility that needs to be ruled out in order for one to know.

3.2 The Challenge SAFES Levels Against Externalism

No reliabilist worth her salt denies that counterevidence (in general) and an argument like SAFES (in particular) can defeat a reliably formed belief. Indeed, reliabilists have had to deal with Laurence Bonjour's nasty counter examples that intuitively weaken a belief's epistemic status—it being reliably formed notwithstanding. More specifically, Bonjour's

cases involve subjects who possess the power of clairvoyance thereby providing examples where (a) one's belief that p is, in fact, true and the result of a reliable process; (b) the counter evidence against one's belief that p is not the result of a reliable processes; and (c) S's belief that p intuitively lacks a positive epistemic status. Indeed, most share the pre-theoretical intuition regarding counterevidence that despite Samantha's belief being both true and the result of a reliable process, she has clear counter-evidence that defeats her belief having positive epistemic status. Similarly, it seems that were one's belief sourced in a reliable processes, once one comes in contact with an epistemic superior who disagrees, the possibility that one is wrong now becomes a live, relevant possibility that needs to be ruled out. Indeed, such disagreement seems to level a defeater against the epistemic status of one's belief.

So the challenge I take SAFES to level against externalism is exactly this: *How can reliabilism account for this kind of defeating evidence when it places all the emphasis for knowledge on a belief's provenance being either reliable or unreliable?* Indeed, Richard Feldman indicts externalist theories of knowledge which, of course, includes reliabilism, for lacking an explanation for how to deal with defeaters, despite the belief having a reliable provenance. He says:

> [E]xternalists […] must face the fact that there are puzzles about how to deal with cases in which one is faced with the prospect of respecting the evidence [of disagreement with an epistemic peer.][…][I]t is one thing to say that we can have knowledge or justified belief without having information about the sources of our beliefs [so says the externalist]. It is quite another to say that our knowledge survives the acquisition of evidence that our reasons are not so good or that our processes are not so reliable [once we face disagreement with an epistemic peer].[182]

Feldman thus challenges externalist accounts of knowledge to develop their theories in such a way that defeaters are properly dealt with and reflected within the theory. Evidence of epistemic superiors disagreeing with us is counterevidence that can defeat a reliably

[182] Richard Feldman, "Respecting the Evidence," *Philosophical Perspectives* 19 (2005): 107.

formed belief just as all the counter evidence Samantha encounters can defeat her reliably formed true belief. Because SAFES provides an argument from epistemic superior disagreement that can defeat one's belief, an externalist theory of knowledge needs to give an account of evidence that can account for this. Indeed, an externalist theory of knowledge needs to give an account of defeating evidence from epistemic superior disagreement that preserves our pre-theoretical intuitions about counterevidence.

After all, reliabilism already suffers from an acute problem regarding how it should account for counterevidence in a way that preserves our pre-theoretical intuitions about counterevidence. But with the formulation of SAFES, it seems reliabilism faces an even greater challenge. For not only must it reflect how defeating evidence is accounted for, but it must also provide an answer to the skeptical threat put forth by SAFES. Unless reliabilism can account for the theoretical problems leveled by SAFES against the nature of knowledge and evidence, then externalism cannot provide an adequate response to this kind of skepticism.

In the Chapter Six, I offer a Grecoian response to SAFES. This response aims to achieve three things. First, it will provide an account of defeating evidence that meets the challenge leveled by Feldman, et al. who say reliabilism can't give an account of defeating evidence in an externalist way (in general) and to counterevidence in the form of superior disagreement (in particular). In Chapter Seven, I will argue how knowledge is possible vis-à-vis epistemic superior disagreement. More specifically, I explain how one can retain one's belief vis-à-vis this kind of counterevidence without being epistemically irresponsible or subjectively unjustified. Furthermore, I explain how disagreement between two or more parties can be the rational response whereby no one is guilty of responding in an epistemically incorrect way. But before we proceed to Chapters Six and Seven, we must

first look at Greco's account of knowledge, which will be laid out in Chapter Five. Chapter Five, then, sets the stage for how a Grecoian response to SAFES will draw from the theoretical resources provided by Greco's virtue-theoretic account of knowledge.

5

Greco's Virtue Theoretic Account of Knowledge

[A]nd if any action is well performed when it is performed in accordance with the appropriate excellence: if this is the case, then the human good turns out to be activity of the soul in accordance with virtue….
~ARISTOTLE

1. VIRTUE EPISTEMOLOGY

The purpose of this section is lay out the main lineaments of virtue epistemology. I will explain some of the main differences between virtue epistemology and other traditional approaches. I will then focus my interests on a particular version. This version represents a cluster of views that share the credit thesis, which says because knowledge is a success from an intellectual virtue, one can't be attributed with knowledge if one doesn't deserve credit for getting the truth. This will then set the stage for introducing Greco's own particular virtue-theoretic account of knowledge.

1.1 ***What is a Virtue Theoretic Account of Knowledge?***

John Greco, a leading proponent of virtue epistemology, says "The term 'virtue epistemology' was first introduced into the literature by Ernest Sosa, who argued that epistemology might benefit by adopting an approach analogous to virtue theory in ethics."[183] And ever since Ernest Sosa's seminal paper "The Raft and the Pyramid" first introduced the term "virtue epistemology" (hereafter VE), VE has burgeoned into a formidable epistemological enterprise. Epistemology has always been a normative

[183] John Greco, "Virtue Epistemology" in Jonathan Dancy, Ernest Sosa, and Matthias Steup (eds.) *Blackwell Companion to Philosophy: A Companion to Epistemology 2nd Edition* (Oxford: Blackwell Publishing, 2010): 81. The seminal work where Sosa defends these claims is in Ernest Sosa, "The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge," in *Midwest Studies in Philosophy* (1980): 3-25.

discipline.[184] But what Sosa did differently was allow virtue theory to inform the kind of normativity epistemologists cared about; virtue theory offered its own account of normativity to epistemology. Virtue ethics offered epistemologists a new way to account for epistemic normativity and provided new resources in an effort to answer some perennial problems in epistemology. One of the things VE borrowed from virtue ethicists is the idea that when we attribute knowledge or justification or some other epistemic property to someone, we are making a normative judgment. That is, when we say S1 has knowledge and S2 doesn't, we are making a value judgment because we are saying S1's evidence, justification, belief, etc., is better than S2's.

This new approach to epistemology differed from traditional approaches by employing a thesis about the direction of analysis. Traditional approaches to epistemology made the properties of beliefs and evidential relations between beliefs most fundamental to an analysis of knowledge. For knowledge and justification are taken to be different evaluations of a belief.[185] VE, however, uses the virtue ethicist's approach. This approach

---

[184] John Greco says virtue epistemology's commitment to being a normative discipline entails at least two things: "First, it signals opposition to Quine's radical suggestion in "Epistemology Naturalized" that philosophers should abandon questions about what's reasonable to believe, and should restrict themselves to questions about cognitive psychology instead. […] Second, it implies that epistemologists should focus their efforts on understanding epistemic norms, value and evaluation." Greco takes this to be a defining feature of the field and is why it's at the epicenter of the recent "value turn" in epistemology." See John Greco, "Virtue Epistemology" in *The Stanford Encyclopedia of Philosophy* (Winter 2011 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/entries/epistemology-virtue//>. The "value turn" Greco quotes refers to a paper written by Wayne Riggs, namely, "The Value Turn in Epistemology" in V. Hendricks and D. Pritchard, eds., *New Waves in Epistemology* (Aldershot: Ashgate, 2006).

[185] Heather Battaly characterizes some of the leading analysis of justification and knowledge that make properties of belief fundamental as *belief-based epistemology*, which takes a different approach to *person-based epistemology* or *virtue epistemology*, which makes properties of persons fundamental. Regarding justification, she says "leading rival analyses of justification have claimed that beliefs are justified when they are (for instance): in accordance with one's epistemic obligations (Chisholm), supported by one's evidence (Conee and Feldman), produced by a reliable process (Goldman, *Epistemology and Cognition*), or based on adequate grounds (Alston)" (p.640). Regarding leading rival analyses of knowledge that comprise the traditional approach, she cites those who say "knowledge is undefeated justified true belief (Lehrer and Paxson); that knowledge requires one's belief to track the truth (Nozick); or that the standards for knowledge shift with changes in context (Cohen)" (p.640). See Heather Battaly, "Virtue Epistemology," in *Philosophy Compass* 3/4 (2008):639-663.

analyzes epistemic properties like knowledge and justification in terms of properties of the person. Properties of the person—not the properties of beliefs, which abstract away from the properties of a person—are, then, taken to be most fundamental, according to virtue epistemologists. Evidence, justification, and knowledge thus get explained in terms of intellectual virtues which are properties of a person and not a belief. This is why virtue theories are considered person-based, rather than belief-based.[186]

For example, John Greco says "virtues are [not, according to traditional approaches, mere] dispositions to believe in accordance with good evidence [rather, a] justified belief [is] one that manifests an intellectual virtue."[187] Evidence, justification, knowledge, and other normative notions get explained, then, in terms of the intellectual virtues possessed by agents and manifested in their performance in acquiring a belief. Regarding knowledge, VE determines whether one knows by evaluating how one acquired a true belief; more specifically, it analyzes whether a true belief resulted from vices such as wishful thinking, hastiness, bias, or rather from virtues like perception, conscientiousness, sobriety, or open-mindedness.

Virtue epistemologists differ over what constitutes a virtue. For example, *Virtue Responsiblists* take intellectual virtues to be particular character traits or skills like intellectual-courage, tenacity, attentiveness, open-mindedness, carefulness, conscientiousness, humility, responsibility, etc. that enable one who desires the truth to be efficaciously equipped to obtain it.[188] *Virtue Reliabilists*, on the other hand, take intellectual virtues to be akin to

---

[186] See Linda Zagzebski, *Virtues of the Mind* (Cambridge: Cambridge University Press, 1996), pp.1-29.

[187] John Greco, "Virtue Epistemology."

[188]"Virtue responsibilism" is a term Guy Axtell first used in a paper he wrote entitled "Recent Work in Virtue Epistemology," *American Philosophical Quarterly*, 34 (1997): 1-27. Epistemologists who endorse the responsibilist characterization of the intellectual virtues are Loraine Code (1987), James Montmarquet (1993), and Linda Zagzebski (1996).

"psychological mechanisms"[189]—that is, reliable cognitive abilities or powers such as memory, perception, and introspection.[190]

Despite the differences amongst virtue epistemologists, there are two things that unite them: (i) epistemology is a normative discipline and (ii) the locus of analysis should be on people—specifically, their cognitive powers, skills, abilities, and character traits—rather than on a belief and its properties to understand certain epistemic goods.

1.2 A Virtue Theoretic Account of Knowledge

Now I want to sketch a specific version of a virtue theoretic account of knowledge. It is one that I will be assuming throughout the dissertation. This particular account of

---

[189]Ernest Sosa, "The Place of Truth in Epistemology," in *Intellectual Virtue,* ed. Linda Zagzebski and Michael DePaul (Oxford: Oxford University Press, 2007), p. 163. Epistemologists who endorse this characterization of the intellectual virtues are Ernest Sosa (1991; 2003); Alvin Goldman (1992); Jonathan Kvanvig (1992); Alvin Plantinga (1993); and John Greco (2000). One reason for this division is that John Greco and others don't think character virtues are an essential or a necessary condition for knowledge. Greco says, "it is plausible that [character] virtues are sometimes needed to turn mere faculties into excellences." That is to say, character virtues may be necessary for making our cognitive abilities more reliable. But it's wrong to think that character virtues are necessary for paradigmatic instances of knowledge acquired through simple perception. Be that as it may, Greco still sees character virtues as essential for a complete epistemology insofar explaining their contribution towards securing other epistemic goods other than knowledge and how they can be intrinsically valuable thus worth pursuing. See John Greco, *Achieving Knowledge* (Cambridge: Cambridge University Press, 2010), p.10.

[190]"Virtue reliabilism" is a term Lorraine Code coined in "Toward a Responsibilist Epistemology," *Philosophy and Phenomenological Research,* 45 (1984): 29-50. I am indebted to Jason Baehr for bringing this and the origin of the term "virtue responsibilism" to my attention in "Character, Reliability, And Virtue Epistemology," *The Philosophical Quarterly,* 56 (2006): 193. It goes beyond the scope of this dissertation to give a taxonomy of all the representatives and ways virtue epistemology has been carved up. But it will suffice to say that in addition to the different ways intellectual virtues are characterized, there are those who put a primacy on the intellectual virtues insofar as how they should comprise the loci of epistemological analysis. But they do not take the intellectual virtues as fundamental and, furthermore, deny that the intellectual virtues can be used to develop a theory of knowledge and justification or explain the latter in terms of the former. Heather Battaly has deemed such epistemologists as virtue epistemologists of the anti-theory stripe in light their denial that intellectual virtues can solve some of the perennial problems in epistemology. Battaly divides the anti-theorists into two camps: virtue eliminativists and virtue expansionists. The former argue that epistemologists should abandon the project of trying to provide a theory of knowledge and justification. Instead, epistemologists should focus their efforts on exploring the intellectual virtues and give them a rigorous analysis. The latter agrees with virtue eliminativists that intellectual virtues should be further explored and analyzed, but still think an analysis of justification and knowledge should be pursued despite the intellectual virtues lacking the ability to provide a systematic account. Jonathon Kvanvig (1992) is considered a virtue eliminativist. Christopher Hookway (2003), Mirada Fricker (2007), Robert Roberts and Jay Wood (2007) are examples of virtue expansionists. See Heather Battaly, "Virtue Epistemology," p. 642-643.

knowledge, whose home lies within the virtue epistemology tradition, is called a credit theory of knowledge. According to *credit theories of knowledge* (hereafter CTK), knowledge is conceived of as a kind of achievement from ability. The thought is that knowledge is a true belief that can be attributed to those features constitutive of how the believer successfully obtained it unluckily. Knowledge, then, is evaluated by employing an explanatory strategy—specifically, one that puts a primacy on explaining *how* the agent successfully obtained a true belief. For a belief that is acquired as a result of luck or chance as opposed to getting the truth as a result of one's powers, skills, and/or abilities is not deserving of the kind of credit believed to be necessary for knowledge.[191]

There are different nuanced versions of CTK. Below are some putative statements, constitutive of the theory, from leading proponents:

---

[191] This view of knowledge has a lot going for it given its explanatory power and scope in providing plausible solutions to (i) the Gettier problem; (ii) the value problem; (iii) why luck is such a threat to knowledge; and (iv) the nature of knowledge. Some of these solutions will be addressed later on the in chapter. But nothing hangs on their success or failure relative to the purposes of this dissertation (in general) and this chapter (in particular). I'm indebted to John Greco for compiling the following list of articles that showcase how a virtue theoretic account of knowledge can provide cogent answers to perennial epistemological problems. Regarding CTK's answer to the Gettier problem see Linda Zagzebski, *The Virtues of the Mind* (Oxford: Oxford University Press, 1996), 293-299; idem, "The Inescapability of Gettier Problems," *Philosophical Quarterly 44,* no. 174 (1994): 65-73; John Greco, "Knowledge as Credit for True Belief," in *Intellectual Virtue,* ed. Linda Zagzebski and Michael DePaul (Oxford: Oxford University Press, 2007), 127-132. Regarding the Value Problem see Linda Zagzebski, *Virtues Of The Mind* (Oxford: Oxford University Press, 1996), 260-261; Ibid., "The Search for the Source of the Epistemic Good," *Metaphilosophy*, 34 (2003): 12-28; Wayne D. Riggs, "Reliability and the Value of Knowledge," *Philosophy and Phenomenological Research* 64 (2002): 92-94; Ernest Sosa, *A Virtue Epistemology* (Oxford: Oxford University Press, 2007), 70-91; John Greco, "Knowledge as Credit for True Belief," 133-134. Regarding luck's deleterious effects to knowledge see Wayne Riggs, "Why epistemologists are so down on their luck," *Synthese* (2007): 329-344; Idem, "Luck, Knowledge, and Control," in *Epistemic Value*, eds. A. Haddock, A. Millar, & D. H. Pritchard (Oxford: Oxford University Press, 2009); Idem., "What are the "chances" of being justified?," *The Monist* 81.3 (July 1998): 452-472; Duncan Pritchard, *Epistemic Luck* (Oxford: Oxford University Press, 2007), 123-177. Regarding the nature of knowledge the sources cited above which address the Gettier problem, the value problem, and the relation between luck and knowledge are already constitutive parts to this answer. But there is also CTK's strategy of employing a normative epistemological strategy so as to analyze the features associated with knowledge attribution. And by couching knowledge in terms of an achievement through a successful from an ability, as Greco remarks: "knowledge attributions can be understood as credit attributions: when we say that someone knows something, we credit them for getting it right. […] [T]he sort of crediting and valuing associated with success from ability (or excellence, or virtue) is ubiquitous in human life. It is instanced in the moral realm, the athletic, the artistic, and many more. In virtually any arena where there is human excellence or ability, there is a normative practice that attaches to it. [CTK] makes knowledge and epistemic evaluation another instance of that more general, familiar sort of normativity" quoted in John Greco, "The Nature of Ability and the Purpose of Knowledge," in *The Metaphysics of Epistemology,* ed. Ernest Sosa (Malden, MA: Blackwell Publishing, 2007), p. 57-58.

Sosa: "[K]knowledge is true belief out of intellectual virtue, belief that turns out right by reason of the virtue and not just by coincidence."[192] Elsewhere Sosa says one of the differences between knowledge and true belief is that the former explains "the good that attaches to an epistemic action credible to the agent, who brings about the good [of believing truly] for himself, and is more than just the recipient of blind epistemic luck."[193]

Riggs: "Credit theories of knowledge hold that S knows that p only if being right about p in this instance is attributable to S as a cognitive agent."[194] Elsewhere Riggs has drawn attention to the necessary ability and anti-luck conditions needed for S to know p by saying, "No matter […] how epistemically virtuous you are, if you arrive at a true belief in some way that is not (sufficiently) due to your epistemic abilities, you have *achieved* nothing. Something nice has happened to you, but it is not something you have brought about or deserve credit for. Only when you deserve credit for some outcome does the outcome count as an *achievement* of yours."[195]

Greco: "S knows that p *if and only if* S believes the truth (with respect to p) because S's belief that p is produced by intellectual ability."[196] Greco, too, draws attention to the ability and anti-luck conditions for knowing p when he says the following: "When we say that S knows p, we imply that it is not just an accident that S believes the truth with respect to p. On the contrary, we mean to say that S gets things right with respect to p because S has recognized in an appropriate way or perceived things accurately, or remembered things well, etc. We mean to say that getting it right can be put down to S's own abilities, rather than to dumb luck, or blind chance, or something else."[197]

Zagzebski: "Knowledge is a state of cognitive contact with reality arising out of acts of intellectual virtue."[198] Elsewhere Zagzebski has called alleged instances of knowledge without her motivation component—that is, a desire for the truth and not some other belief-state motivated by, say, wishful thinking "a matter of luck in a certain sense of luck.

---

[192] Ernest Sosa, *Knowledge in Perspective* (Cambridge: Cambridge University Press, 1991), p. 277.

[193] Idem, "The Place of Truth in Epistemology," in *Intellectual Virtue* eds. Linda Zagzebski and Michael DePaul (Oxford: Oxford University Press, 2007), p. 173.

[194] Wayne Riggs, "Two Problems of Easy Credit," *Synthese* 169, 1 (2009):201-216.

[195]Idem, "Reliability and the Value of Knowledge," *Philosophy and Phenomenological Research* 64 (2002): 95 (emphasis in original).

[196] John Greco, *Achieving Knowledge* (Cambridge: Cambridge University Press, 2010), p. 71.

[197] Idem, "Knowledge as Credit for True Belief," in *Intellectual Virtue* eds. Linda Zagzebski and Michael DePaul (Oxford: Oxford University Press, 2007), p. 116.

[198] Linda Zagzebski, *Virtues of the Mind,* 270. How Zagzebski specifically defines an intellectual virtue is the following: "*An act of intellectual virtue A* is an act that arises from the motivational component of *A*, is something a person with virtue *A* would (probably) do in the circumstances, is successful in achieving the end of the *A* motivation, and is such that the agent acquires a true belief (cognitive contact with reality) through these features of the act (p.270)."

[…] [That is,] from a certain point of view it *is* an accident that I got the truth. I don't get credit for getting the truth."[199]

What these statements point out, as they relate to CTK, is this: instances of knowledge must be able to explain how S obtained a true belief as result of some power, skill, and/or ability in a way that avoids the deleterious effects of luck or accidentality.[200] In other words, CTK takes knowledge to be a matter of whether or not S's abilities are sufficient to explain why S truly believes p in a way that deserves credit. So in order for there to be instances of knowledge, S must have acquired a true belief *because* S's belief is the result of intellectual virtue and not because of luck or accident. What has fallen out of this is the "credit thesis," which says because knowledge is a success from an intellectual virtue, one can't be attributed with knowledge if one doesn't deserve credit for getting the truth. So if one acquired a true belief on account of luck or accident, then we can't attribute one's success to one's abilities in a way that merits the kind of credit attribution necessary for achieving knowledge.

2. THE GRECOIAN ACCOUNT OF KNOWLEDGE

At this point I want to further develop some of the main lineaments of Greco's virtue theoretic account of knowledge that will be relevant to this dissertation. But first some provisos: Greco denies his account provides a conceptual analysis of knowledge. Traditionally, analytic philosophers (in general) and epistemologists (in particular) analyze complex concepts by breaking them down into simpler conceptual parts. Greco denies

---

[199] Idem, "Intellectual Motivation and the Good of Truth," in *Intellectual Virtue* eds. Linda Zagzebski and Michael DePaul (Oxford: Oxford University Press, 2007), p. 151.

[200] For an insightful paper on the differences between 'luck' and 'accidentality' and how they function as two common culprits at undermining instances of knowledge see Wayne Riggs, "What are the "chances" of being justified?," *Monist* 81.3 (July 1998): 452. The paper also shows how these species of 'chance events' are what many competing theories are trying to locate and capture in their negative projects' work in developing notions of epistemic justification.

trying to do either this or providing an ontological analysis. Nor is he trying to give necessary and jointly sufficient, informative conditions for knowledge. Rather, Greco says his account of knowledge is "informative in a straightforward way sense: it provides insight into what knowledge is by identifying it as an instance of a more general, familiar kind."[201] The "familiar kind" Greco is, of course, referring to is the idea of an *achievement*. Knowledge is a kind of achievement thus positioning itself within the broader normative domain. Because we are so familiar with the domain of achievement, Greco believes by reflecting on our thinking about it, we will gain insight and understanding of what is knowledge.

2.1 Six Themes Informing Greco's Account of Knowledge

The following themes play an important role in Greco's theory that knowledge is a success from an ability (hereafter called KSA). It's important to introduce the reader to them since they will be employed later in Chapters Six and Seven when I offer a Grecoian response to different skeptical threats and related problems. The first three themes play a broader role insofar as providing a theoretical framework to place his account within. The latter three are recent trends in epistemology that can stand independently of Greco's account of knowledge. But he wants these three trends to be consistent with his account and inform it.

2.1.1 Theme 1: Epistemology is a Normative Discipline

The first theme is that epistemology is a normative discipline. Grasping the nature of epistemic normativity, however, is as slippery as an eel; it admits many different meanings, which makes it hard to pin down. But Greco wants to employ some of the most uncontroversial claims regarding normativity. For example, the claim that knowledge

[201] John Greco, *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity* (Cambridge: Cambridge University Press, 2010), p. 4.

attributions are value judgments—given they imply one person's judgment is preferable to another's mere opinion—is uncontroversial and virtually universally accepted by epistemologists. Because epistemology (in general) and knowledge (in particular) have normative dimensions, Greco takes epistemology's central task to "provide an account of the normativity involved."[202]

2.1.2 Theme 2: KSA Has A Tradition Stemming Back to Plato

The second theme related to Greco's account of knowledge is that it has a home within traditional epistemology which traces back to Plato. He makes a distinction between two different kinds of traditional epistemological projects. The first project, he claims, can be traced back to challenges leveled by Pyrrhonian skepticism, which focuses on the question "What do we know?" in order to establish knowledge. Such Pyrrhonian challenges are twofold: the first challenge is to provide a fully general account of knowledge that employs a non-question begging argument for the conclusion that we do indeed have knowledge. It demands a *vindication* of knowledge. Greco thus calls this project The Project of Vindication (hereafter PV). For the kind of vindication it demands must be fully general insofar as accounting for all the knowledge we claim we have. And it must be non-circular by providing a non-question begging argument that vindicates our knowledge claims. But it must do so in a manner that does not presuppose any knowledge to begin with.

The second project can be traced back to Plato's *Theatetus* wherein an answer to the question "What is knowledge?" is pursued. Greco calls this The Project of Explanation (hereafter PE). For it's concerned with explaining what knowledge is; how it's possible; and tries to explain the difference between knowing and not knowing.

---

[202] Ibid., p. 4.

Greco sees PV making demands that are ill conceived (at best) and incoherent (at worst). For he says, PV is, in the language of proof, "asking that we give a proof but without employing any premises. […] [And] why should anyone think [PV] is possible or desirable?"[203] For example, why think we should concede to the skeptic's demands of vindicating that we have perceptual knowledge without relying on our perceptual faculties to do so? Or why concede to the demands of the skeptic who wants us to vindicate that we have knowledge of the external world without relying upon any knowledge of the external world?[204] This is why Greco says twentieth century epistemologists "clearly privileged the project of explanation over the project of vindication […]."[205] Furthermore, Greco adds, "there is an emerging consensus that [PV] is somehow flawed or misguided, and that the proper task of epistemology is [PE]."[206]

Greco offers an example to further illustrate this point. He asks us to consider the following skeptical argument.[207]

(1) Knowledge of the world is possible only if sensory perception is a reliable way of forming true beliefs about the world.
(2) But sensory perception is not very reliable—it leads us to make mistakes about the world all the time.
Therefore,
(3) Knowledge of the word is impossible.

---

[203] John Greco, "Religious Knowledge in the Context of Conflicting Testimony," in *Proceedings of The American Catholic Philosophical Association* 83, (2009): 65.

[204] I borrow these examples from Greco, "Religious Knowledge in the Context of Conflicting Testimony," p.66. Similarly, Paul K. Moser says such demands suffer from a kind of incoherence. He likens the skeptic's demand to requiring one to stand somewhere while not being allowed to stand anywhere. For, necessarily, if all sources of knowledge are under question—that is, the fully general requirement—and in need of verification, then "none will be non-question-begging [which is] analytically true, and importantly analytically true." See Paul K. Moser, "Skepticism Undone?" in John Greco, ed., *Ernest Sosa and His Critics* (Oxford: Blackwell, 2004), p. 141.

[205] John Greco, "Epistemology" in Constantin V. Bounds, ed., *Columbia Companion to Twentieth-Century Philosophies* (New York: Columbia University Press, 2007), p. 172.

[206] John Greco, *Achieving Knowledge*, p. 5.

[207] John Greco, "Epistemology," p. 173.

Greco says one obvious place to show where this argument is unsound is by demonstrating the falsity of premise (2). But in order to do so one would have to rely upon sensory perception which would beg the question, and thus, violate the skeptic's demands of non-circularity. In this case, it would be faculty circularity. That is, the skeptic would indict the use of perception to vindicate the reliability of perception. Such a demand, however, is incoherent, because "the only way we could know that premise 2 is false is by relying upon our knowledge of the world—it is only by knowing the world that we know that sensory perception is (in general) reliable. But it is precisely knowledge of the world that is at issue in the argument, and so showing that premise 2 is false seems impossible."[208] And if the skeptic required an argument for the reliability of one's faculties in general, the demand would be even more incoherent. For one cannot employ any other faculties than those one has already been bequeathed. Indeed, this goes for any cognizer, in any possible world. The fact that one can think only by means in which one, in fact, does think is the condition any cognizer will be in.[209]

This is one of the main reasons Greco thinks the project of epistemology should be PE and not PV. The task should be to explain how knowledge is possible. What's more, PE should provide understanding of "what knowledge is, and [provide] an account of epistemic normativity…."[210] Such an approach removes some of the gross ways of begging the question that's dialectically inappropriate. And if one demands an answer to how we know knowledge is possible we concede such an answer has not been provided. But this

---

[208] Ibid.

[209] Greco credits this point to Keith DeRose which was made in a paper presented at Fordham University .

[210] John Greco, *Achieving Knowledge*, p. 5. Greco thinks an account of knowledge should do even more. In addition to explaining the nature of knowledge, he thinks the project of explanation should also explain our concept of knowledge or term "knowledge" and its cognates, as well as how the language of epistemic evaluations (e.g., knowledge attributions) function both semantically and pragmatically.

should not worry us given it only admits that "the project of explanation is not the project of vindication."[211] But Greco clearly wants to give the skeptic her due. For he thinks skepticism is a heuristic device that sheds light on the nature of knowledge and evidence. PE, then, should explain where skeptical arguments go wrong even if skepticism is assumed at the outset to be false.[212] Indeed, Greco's book *Putting Skeptics In Their Place* indicts those who don't take skepticism seriously and exposes the fallacious reasoning behind dismissive responses towards it.[213]

---

[211] Ibid. Sometimes the vindication project assumes knowledge must be transparent and thus one needs to fulfill a subjective condition whereby one can show the skeptic how one knows. Greco argues in *Putting Skeptics in Their Place* that there is a common subjective condition placed upon knowledge to the effect that knowing requires knowing that one knows thus implying knowledge is transparent. But such a transparency principle: namely, 'Knowledge entails knowing that one knows' is either too demanding or borders incoherency, according to Greco for the following reasons. First, the principle that S knows only if S knows that S knows (i.e., Kp only if KKp) leads to total skepticism. One can't even know that one exists if this is a condition for knowledge. For I could know I exist only if I know that I know I exist. But the necessary condition for knowing must itself also meet the demands of this principle. So in order for me to know that I know I exist, I must know that I know that I know that I exist (i.e., KKp only if KKKp). Greco points out, then, that "Kp implies KKp, which in turn implies KKKp, which implies KKKKp, which implies KKKKKp , and so forth" (p.183). Second, the principle can be applied to one's cognitive faculties. The demand is now this: I know that p only if I know that my belief is reliably formed. Put differently: I know that p only if I know p´—namely, p is reliably formed. But to know p´ I must also know p´´—namely, p´ is reliably formed, and we're off and running again. Third, another version of the principle could go like this. One can have knowledge only if one knows that the conditions for knowledge have been fulfilled, say, conditions X, Y, and Z. But if conditions X, Y, and Z are sufficient for knowing p, then nothing more needs to be added in order to know p, right? Wrong! Because the transparency principle requires one to know that conditions X, Y, and Z have been fulfilled. But if X, Y, and Z are sufficient for knowing p, then this further requirement is not necessary. This takes us back to the first reason why assuming one needs to know that one knows is a misguided assumption. Why? Because knowing an infinite number of increasingly complex propositions is beyond our ken and total skepticism ensues from such a demand. Greco is not to first to make these kinds of points. William Alston's "Level Confusions in Epistemology" in *Midwest Studies in Philosophy* 5 (1980):135-150 makes similar points and offers extended arguments why the transparency principle is a misguided assumption for a theory of knowledge to have to shoulder.

[212] Greco thinks part of this assumption is supported by the idea that knowledge is either easy or impossible. Greco thinks James Van Cleve and others have cogently argued that knowledge is either virtually impossible, as the skeptic claims, or fairly easy and widespread, as common sense dictates. Part of the project of explanation, then, is to explain how it is that knowledge can be easy and widespread. Greco, however, does concede that there are some sharp teeth in the bite of the objection that knowledge can't be easy and treats this objection as a legitimate one that deserves a response. Greco thinks epistemologists should heed to the wisdom of James Van Cleve. See James Van Cleve, "Is Knowledge Easy—Or Impossible? Externalism as the Only Alternative to Skepticism" in Steven Luper, ed., *The Skeptics: Contemporary Essays* (Burlington, VT: Ashgate), p.45-59.

[213] The three main theses he defends in that book are: (1) that a number of historically prominent skeptical arguments make no obvious mistake and therefore cannot be easily dismissed; (2) that the analysis of skeptical arguments is philosophically useful and important and should therefore have a central place in the

2.1.3 Theme Three: Knowledge and Understanding Are Distinct Epistemic Goods

The third related theme that governs Greco's theory of knowledge is this: "epistemology benefits from a distinction between knowledge and understanding."[214] Greco's account of understanding builds off of Aristotle's account of *episteme*, which requires knowing the cause of something. To know how or why something is the case Aristotle provided four kinds of causes: formal-, material-, efficient-, and final-cause. Why is there a wooden chair in a classroom? Formal cause: because the chair's blueprint has those dimensions; material cause: because wood is the material the chair is made from; efficient cause: because the carpenter cut, shaped, glued, and nailed pieces of wood together; final cause: because the carpenter wanted to build something for students to sit on. Understanding why there is a chair in this classroom, then, requires knowing how this chair depends upon these four causes in a systematic way that sees how they all fit together. So, understanding, for Greco, "involves 'grasping,' 'appreciating,' or knowing causal relations taken in a broad sense: i.e. the sort of relations that ground explanation."[215] The epistemic good of understanding is,

---

methodology of philosophy, particularly in the methodology of epistemology; and (3) that taking skeptical arguments seriously requires us to adopt an externalist, reliabilist epistemology—more specifically, an "agent reliabilism" that is positioned in an externalist version of virtue epistemology. See John Greco, *Putting Skeptics in Their Place: The Nature of Skeptical Arguments And Their Role In Philosophical Inquiry* (Cambridge: Cambridge University Press, 2000), p. 1.

[214] John Greco, *Achieving Knowledge*, p. 7. He also thinks wisdom is a distinct epistemic good, as well.

[215] Ibid., p.9. In addition to preserving Aristotle's points about causal explanations, Greco's account aims at preserving points made by Jonathan Kvanvig (2003) and Wayne Riggs (2003) where understanding involves "internal grasping" of how different causal relations and pieces of information "hang together." Greco may have just as well included Lorraine Code's and Linda Zagzebski's points about understanding. Code says that understanding is having the ability to see the "interconnections" and "significance" of known facts in a way that places them in a "pattern, or a whole structure," which elicits causal explanations. Zagzebski says understanding enables one to not only know p, but rather, explain how p is "part of […] a system or network of truths…." Because seeing the "interconnections" and explaining how an atomic truth is part of a "system" requires appealing to causes, it seems Greco wants to preserve these insights about understanding, too. See Lorraine Code, *Epistemic Responsibility* (Hanover: University Press of New England, 1987), pp. 149-150. Linda Zagzebski, *Virtues of the Mind* (Cambridge: Cambridge University Press, 1996), p. 49.

then, contrasted with knowledge since one can know *that* a particular event, E, occurred, but lack knowledge of *why* or *how* E occurred, which requires understanding.

Greco, however, builds upon Aristotle's account *episteme* and offers his own neo-Aristotelian account of understanding by updating Aristotle's in two ways.[216] First, he recognizes that what Aristotle's four causes have in common is that they cite different dependency relations. He, then, takes Aristotle's four dependency relations—replaces them with modern concepts (e.g., final causes are teleological dependency relations and formal causes are replaced with essential dependency relations)—and includes more kinds of dependency relationships (e.g., mereological-, logical-, mathematical-, conceptual-, and different kinds of supervenient-relationships). Second, he emphasizes that unlike knowledge, understanding "consists in *systematic* knowledge of dependency relationships."[217] Such systematic knowledge will come in degrees relative to knowing how many different kinds of dependency relationships there are and which ones are fundamental.

These amendments allows for there to be different objects one can understand. If understanding X is being able to locate it and see it within a web of different modally strong dependency relations, then understanding X involves knowing both the dependency relationships and their *relata*. From this, Greco makes a further distinction: namely, the objects of the dependency relations may be of real objects, representations, or both. For example, S can understand X in one of three ways.

> (a) S understands the "real" dependency relations [and their *relata*] in the world comprising a system that X is located within (e.g., ecosystems, economies, a machine, or historical event).

[216] This account is put forth in John Greco, "Episteme: Knowledge and Understanding," in Kevin Timpe and Craig Boyd, eds., *Virtues and Their Vices*. Oxford: Oxford University Press, 2014.

[217] Ibid., p. 7 (draft pagination).

(b) S understands a representation [replete with its own dependency relations and *relata*] of a real system in the world (e.g., a theory, narrative, model, or a set of equations).
(c) S understands the relations between a real system and a representation (e.g., the relations between a diagram and a machine that it represents or the relations between a narrative of a historical event and the actual event it represents).[218]

Let's consider the following two cases Greco uses to illustrate how one can have stronger and weaker degrees of understanding depending upon which objects function as the *relata* in one's systematic knowledge of dependency relations.

> Case 1. Jill knows what the ideal gas laws says (i.e. she knows relevant facts about the representation), Jill knows that the ideal gas law is an idealization of how actual gases behave in the world (i.e. she knows relevant facts about the representation-world relation), and Jill knows that actual gases behave so as to approximate the ideal gas law (i.e. she knows relevant facts about the world).
>
> Case 2. Jack knows what the ideal gas law says, but does not know that it is supposed to be an idealization. Accordingly, Jack knows relevant facts about the representation, but he does not know relevant facts about the representation-world relation, and he does not know relevant

[218] These examples are taken from Greco with slight amendments made. See John Greco, "Episteme," p. 8 (draft pagination). Greco thinks the literature on understanding sometimes conflates these distinctions or confuses them. E.g., Catherine Elgin (2004) objects that understanding entails truth because in order to understanding or make progress in understanding something one need only be "true enough" rather than factually true or strictly true. E.g., she points to cases where Copernicus's theory is superior in its understanding of planetary movement within our solar system to Ptolemaic ones; and likewise with Kepler's theory being more accurate than Copernicus's. But each theorist demonstrated a degree of understanding yet their theories were not literally true, and thus, understanding doesn't entail truth, unlike knowledge. Greco believes understanding does entail truth and it's clear how once these distinctions are made we can see that Elgin and others are conflating truths about the model with truths about the relation between the model and the world. Or one's conflating truths about the relationship between real objects in the world and truths about the relations comprising the representation of the world. Or one's conflating truths about one's representation of the model and its relation to the world. Such distinctions also handle objections that understanding unlike knowledge doesn't have to be immune from luck. Jonathan Kvanvig (2003) thinks one can understand something as a result of luck as evidenced when one reads a history book on why the Comanche Indians dominated the southern plains given their superior weapons. The history book was, indeed, historically accurate, but it was acquired luckily. That is, the book laid in company with many other inaccurate history books on why the Comanche Indians dominated the southern plains. So the reader of the accurate book acquired his true beliefs luckily which prevents him from having knowledge. Be that as it may, such a person still has understanding. Again, according to Greco, these distinctions amongst different kinds of understanding allow us to explain the case by saying the lucky person only understands the historical narrative given in the book. But he doesn't know that the narrative accurately captures the historical events because he doesn't know whether the narrative is true. So it's not lucky that he understands the narrative proffered by the author, which is all Kvanvig's point establishes. But he doesn't know that the narrative (which he understands) corresponds to actual events in history because he doesn't know whether the narrative is true or not.

> facts about the world (for example, that actual gases behave only so as to approximate the ideal gas law).[219]

It's clear Jill has a greater degree of understanding gas laws than Jack. Indeed, Jill understands more than just (b) the *representation* of the ideal gas law, which is all Jack understands, because she understands both (c) the relation between the law-statement and the world, and (a) how gas molecules behave in the world. That is, Jill knows that the representation of the law is an idealization of how gas molecules behave and she also knows that the representation of the law doesn't correspond to how gas molecules actually behave. Jack, however, both misunderstands (c) the relation between the law and the world, and (a) how gas molecules actually behave.

Greco takes understanding to consist in a "systematic knowledge of dependency relations, where dependency relations can be of various sorts, including "real" relations between parts of the world, conceptual and logical relations between parts of theory, and semantic relations between theory and world."[220] By distinguishing knowledge from understanding, the concept of knowledge is free from having to do the insurmountable job of accounting for so many different kinds of epistemic goods. More specifically, it provides room to accommodate some of the "deep-seated intuitions motivating internalism, evidentialism, and coherentism, while rejecting them as intuitions about knowledge per se."[221]

---

[219] I borrow these cases from John Greco, "Episteme: Knowledge and Understanding," p. 10-11 (draft pagination).

[220] Ibid., p. 8 (draft pagination).

[221] Ibid., p. 8.

2.1.4 Theme Four: Good Methodology Asks What Function The Concept of Knowledge Performs?

The fourth theme that figures into Greco's account of knowledge is a methodological one. Greco thinks that when we do epistemology, we should, among other things, be asking questions about the concept of knowledge. More specifically, Greco thinks epistemologists should heed to Edward Craig's proposal about epistemic methodology: namely, epistemologists ought to be asking questions related to what role the concept of knowledge plays in our conceptual-linguistic economy.[222] For example, "What is the *point* of knowledge?"—that is, why do we have the concept of knowledge to begin with? "What is the *purpose* of knowledge?"—that is, what purpose does the concept of knowledge serve? Greco thinks answers to these kinds of questions can provide the resources needed to explain both the nature and value of knowledge. Thus, he finds himself in agreement with Edward Craig's engineering point: namely, the purpose of a concept is to constrain the kind of content that gets built into it because *content follows function*.[223]

So what is the function of the concept 'knowledge'? Greco takes seriously Craig's argument that the primary point and purpose of knowledge is to flag down both good information and good sources of information so we can employ it into our practical reasoning. More exactly: "The concept of knowledge serves to govern the production and flow of actionable information, or information that can be used in action and practical reasoning, within a community of information sharers. Since we need information we can act upon, the concept of knowledge is used to pick out the relevant sources of information needed. Indeed, since we are social, information-dependent creatures in daily need of

---

[222] Edward Craig, *Knowledge and the State of Nature* (Oxford: Oxford University Press, 1990).

[223] John Greco, "Testimonial Knowledge and Flow of Information" in Greco and Henderson, eds., *Epistemic Evaluation* (Oxford: Oxford University Press, *forthcoming*), p. 8 (draft pagination).

actionable information, we need the concept of knowledge to perform this function."[224] And if such a concept didn't already exist then, according to Craig, we would have to invent it.[225]

2.1.5 Theme Five: Knowledge and Practical Reasoning Are Intimately Related

This fifth theme is closely related to the fourth. Greco wants his account of knowledge to be consistent with the following thesis: knowledge has a dimension that is captured by the idea that it is the norm of practical reasoning. Thus, *ceteris paribus*, one should act only on what one knows.[226] According to Greco, knowledge and action are intimately related thereby making knowledge intimately related to practical reasoning.

2.1.6 Theme Six: An Adequate Theory of Knowledge Should Deal With Knowledge Attributions

The last theme that plays a role in Greco's account of knowledge is his position on contextualism—specifically, epistemic contextualism. Epistemic contextualism is an epistemological theory only because it's concerned with sentences that attribute epistemological words like 'knows' and 'knowledge' to others. Greco has recently developed his own version of contextualism; for he thinks knowledge attributions are somehow sensitive to one's context.[227] That is to say, he thinks it's relatively uncontroversial that our "disposition towards making and accepting knowledge attributions

---

[224] Ibid.

[225] John Greco, "Religious Knowledge in the Context of Conflicting Testimony," p. 69.

[226] This thesis is defended by John Hawthorne in *Knowledge and Lotteries* (Oxford: Oxford University Press, 2004). Greco is sympathetic to it and wants his account of knowledge to at least make room for it and be consistent with it.

[227] See John Greco, "What's Wrong with Contextualism" in *The Philosophical Quarterly*, 58 (2008):416-436.

are, as a matter of fact, influenced by features of practical context."[228] His is a version of semantic (attributor) contextualism insofar as it borrows a point represented by subject-sensitive invariantism.[229] The point borrowed is this: the relevant pragmatic interests and purposes of the attributor may be the same as that of the subject—that is, the *practical reasoner* at hand. The practical environment that the reasoner's interests and purposes are in will determine, not the justificatory standards but, rather, the causal explanation for how the subject's abilities either saliently explain or not explain why a true belief is acquired in the right sort of way necessary for receiving a knowledge attribution. But the practical environment the subject's interests and purposes are in can also be the same as that of the attributor, which is where he differs from subject-invariantists and why Greco calls his version interest-dependent, subject-sensitive contextualism.

Greco, thus, wants to accommodate the invariantists' point made by Stanley and Hawthorne that the subject's interests and purposes (not the attributor's), relative to one's practical environment, should govern the justificatory standards for knowledge attributions by making sure the attributor is sensitive to it.[230] For example, he thinks if *we* are the subjects and considering whether the plane will land in Chicago, then it makes sense to consider what a testifier knows relative to *our* interests and purposes. But if we are

---

[228] John Greco, *Achieving Knowledge*, p. 72. On this point he is in agreement with Stewart Cohen (1988); Keith DeRose (1995); and David Lewis (1996).But he differs with their overall project—specifically, how they think contextualism is a plausible response to skepticism.

[229] Invariantism is the view that there is only one standard of justification at work in determining whether or not one should be attributed with knowing that p. Subject-sensitive invariantism, says the subject's interests and purposes that are operative in the subject's context govern the justificatory standards for knowledge attributions. Unlike invariantism, contextualism says the justificatory standards for knowing do shift depending on the attributor's context. It's not that the standards for knowledge in and of itself change but, rather, the standards for how the term 'knows' is used can change. For according to contextualism, what is sufficient evidence or sufficient reliability to know in one context may not be sufficient in another context. There will be a continuum of justificatory standards for how the term 'knows' is ascribed to the attributor's context.

[230] See John Hawthorn, *Knowledge and Lotteries* (Oxford: Oxford University Press, 2004) and Jason Stanley, *Knowledge and Practical Interests* (Oxford: Oxford University Press, 2005).

considering whether someone else, say, Smith (who now plays the role of the subject) knows the plane will land in Chicago, then it makes sense to consider whether Smith knows relative to *his* interests and purposes. So the subject governs the justificatory standards, but the subject can also shift to the attributer herself.

Sometimes the justificatory standard for knowledge will be determined by the attributor's interests and practical reasoning environment; other times it will be the subject's interests and practical reasoning environment. For Greco takes knowledge attributions to be credit attributions. Thus, contextualism does work for Greco because knowledge attributions involve causal explanations. And causal explanations require being sensitive to contextual features. Therefore, knowledge attributions require contextual semantics so as to provide a causal explanation for how the subject's abilities contributed to the success in the right sort of way necessary for attributing the subject with knowledge.[231]

---

[231] Greco parts company with most contextualists, however, given he denies contextualism is successful at responding to skepticism. Contextualism is partly motivated by its ability to capture our intuitions about ordinary knowledge claims, like I know the external world exists, and our intuitions about skeptical arguments to the contrary, namely, that we don't know an external world exists. Contextualism explains these conflicting intuitions by ascribing an error theory regarding the implicit shift in justificatory standards once our context changes. That is, in normal everyday contexts we know lots of things. But once the context shifts to the "philosophy classroom" and we engage the skeptic, the standards for knowing can shift so high that we no longer can claim to know lots of things. The shifting standard for knowing explains why skeptical threats seem so cogent. Greco, however, thinks contextualism concedes too much to the skeptic. What's more, he thinks there is a more stable standard for knowledge than the shifting one contextualists employ to save ordinary knowledge attributions from skepticism. Because the interests and purposes operative in the reasoner's practical environment determine the justificatory standards, there will be an increased general stability and not a widely shifting one for the following reasons. First, because the concept of knowledge performs the function of flagging down good sources and testifiers of information, knowledge is something that can be acquired by many people without having to satisfy overly demanding conditions. As a result, knowledge is somewhat ubiquitous, and thus, there's downward pressure on what the standards for knowledge are. Second, the kind of balance Greco is envisaging comes as a result of the upward pressure placed upon the standards of knowledge via the demands placed on knowledge for practical reasoning. In order for me to be attribute S as a "knower" I must "be confident that the standards by which S counts as knowing in her context are at least as high as my practical reasoning requires" (p. 116). The upshot is that because knowledge performs a certain function in our social linguistic economy, practical reasoning, and social environment, the standards for knowledge will vary across different contexts comprising our practical environment. But the standards will neither vary "widely nor wildly" across such contexts and environments because of the upward and downward pressures at play. See John Greco, "Achieving Knowledge," p. 116.

2.2 Greco's Subjective Justification Condition

Greco's conditions for a belief having a positive epistemic status are something like this: S's believing that p is licensed by the dispositions that S manifests when trying to believe the truth, where "trying to believe the truth" signals a default mode of thinking conscientiously in an effort to get the truth, which is opposed to believing what is convenient or chic. (Greco doesn't want the subjective condition to imply one needs some intentional attitude towards one's belief or the source from which that belief derived; for we don't typically have such a perspective in instances of knowing.) More exactly, a belief has positive epistemic status when the following condition obtains.

VJ: A belief p is subjectively justified for a person S if and only if S's believing p is grounded in the cognitive dispositions that S manifests when S is thinking conscientiously.[232]

Greco qualifies this definition in the following ways. First, "thinking conscientiously" captures a distinction between thinking honestly as opposed to thinking that aims at securing psychological stability, or thinking that aims to get attention from others, or thinking that aims to maintain a belief despite facing normative defeaters, which are defeating evidence one ought to have but fails to manifest because of, say, pigheadedness or some other intellectual vice.

Second, a justified belief does not reduce to a mere conscientious belief. The latter can obtain without striking the balance between trying too hard to get the truth and not trying hard enough. By not trying hard enough to get the truth one's affective states can possibly prevent one from reliably getting the truth as when a father allows his love for his son to suppress the weight of evidence indicating his son's use of drugs. But by trying too hard, one can stymie one's ability to perform well. Just consider how athletes perform

[232] John Greco, *Putting Skeptics in Their Place* (Cambridge: Cambridge University Press, 2000), p.190.

poorly when trying too hard (e.g., a golfer grips the club too tight and swings too hard, which disrupts ideal speed, tempo, rhythm, timing, and fluency).

Third, "cognitive dispositions" signals that the stable properties manifested in conscientious thinking are constitutive of one's cognitive character and thus are a part of the agent. This idea fits well within a virtue epistemology since intellectual virtues are taken to be stable dispositions that are constitutive of one's cognitive character thereby being internally related to the agent.

Last, even though a justified belief is grounded in one's cognitive character, one need not have beliefs about one's character in order to be justified.[233] This makes the condition more psychologically plausible given we rarely have a perspective about many of our beliefs nor a perspective on our beliefs' sources. Just as a good athlete need not have a perspective or beliefs about her abilities in order to manifest her dispositions that reliably achieve the aim of playing well. We see this by recognizing that some of the best coaches cannot perform well because they only possess true beliefs and an accurate perspective on the matter at hand. But they lack the dispositions to perform the athletic feat they can teach well. (Just consider the overweight gymnastic coach who can coach a girl to win an Olympic gold medal, yet can't even do a cartwheel. Or consider the athlete that can hit a homerun every eight at bats yet can't articulate, let alone explain, how he does this to someone who wants to learn from him. Such a person would be a terrible coach.)

VJ is a condition that is intended to capture an epistemically normative requirement for knowledge. It aims to satisfy the pre-theoretical intuition we have: namely, that one is responsible for how one carries out one's epistemic conduct and that *de facto* reliability is

[233] I will say more about objections leveled against accounts of justification that requires a perspective on one's reliability later in chapters 6 & 7.

necessary, but not sufficient for knowledge.[234] In Greco's recent work, he replaces "subjectively justified" with "epistemically responsible." He defines the knowledge-relevant responsibility or, as I've been deeming it in this dissertation, namely, when a belief has positive epistemic status, as follows:

Resp. S's belief that p is epistemically responsible if and only if S's believing that p is properly motivated; [in other words] if and only if S's believing that p results from intellectual dispositions that S manifests when S is motivated to believe the truth.[235]

2.3 Greco's Answer to 'What is the Nature of Knowledge?'

According to Greco, in order to answer the question "What is the nature of knowledge?" one's answer must also provide an answer to an additional question: namely, "Why is knowledge valuable?" He thus endorses both Linda Zagzebski's and Jonathon Kvanvig's point that an adequate account of knowledge must also explain why knowledge is valuable.[236] I will first give Greco's answer to what is the nature of knowledge and then explain how this answer explains why knowledge is valuable.

Greco claims that knowledge is a success from ability and not the result of luck or accidentality.[237] In cases of knowledge, S knows because there is a causal explanation for how S acquired a true belief in a credit deserving way. S's abilities explain why S acquired a

---

[234] I will say more about what motivates this intuition in Chapter Six. But it will suffice to say that since Greco holds to a version of reliabilism—more specifically, agent reliabilism—there are cases that surface a worry that objective reliability alone is not adequate to account for instances of knowing. The following cases bear this out: Carl Ginet's Fake Barn Facades (1975); Plantinga's Serendipitous Brain Lesion (1993); Laurence Bonjour's Clairvoyant, Norman (1980).

[235] John Greco, *Achieving Knowledge*, p. 43.

[236] Linda Zagzebski, *Virtues of the Mind* (Cambridge: Cambridge University Press, 1996). Jonathon Kvanvig, *The Value of Knowledge and the Pursuit of Understanding* (Cambridge: Cambridge University Press, 2003).

[237] What is known as the *incompatibility thesis* dates back to Plato and basically says that epistemic luck is incompatible with knowledge. Because true beliefs can be acquired through guessing or wish fulfillment or reading tea leaves, epistemologists have reached a virtual consensus that knowledge requires more than mere true belief. Justification was believed to be the difference between true belief and knowledge until Edmund Gettier's 1963 paper showed otherwise as explained further below in note 52.

true belief in a way that avoids the deleterious effects luck or accident have on knowledge by undermining it. More formally:

S knows that p *if and only if*
1. *p* is true;
2. S believes that *p*; and
3. S believes the truth because S's belief is produced by intellectual ability.

The thesis that falls out from this is that knowledge is a success from ability (KSA) and is formally stated as follows:

KSA. S knows that p if and only if S believes the truth (with respect to p) because S's belief that p is produced by intellectual ability.[238]

Greco distances himself from other virtue epistemologists who also endorse the credit thesis by attempting to give a principled account of the because-of relation between one's true belief and the sources from which it derives. Linda Zagzebski, for example, is content on leaving the because-of relation primitive, and thus, sees no need in trying to unpack it. Of course, all credit theorists agree that in cases of knowing, S's intellectual virtues explain why S acquired a true belief as opposed to dumb luck or chance or some other explanation. It's because S reasoned well or remembered accurately that merits knowledge-relevant credit. To put the point differently: it's all fine and well when S acquires a truth belief that is acquired in an epistemically responsible and reliable way: but it's far *better* if S acquires a true belief *because* it's responsibly and reliably formed. According to Greco, "this marks the difference between virtuous belief and belief from virtue."[239]

[238] Greco, *Achieving Knowledge*, p. 71.

[239] Greco, *Achieving Knowledge*, p. 44.

Now I will showcase how Greco's account of knowledge handles some familiar cases in the epistemological literature and how it handles some notable objections. I will first give a Gettier-style case:[240]

> RODDY: Roddy is a farmer. One day he is looking into a field near-by and clearly sees something that looks just like a sheep. Consequently, he forms a belief that there is a sheep in the field. Moreover, this belief is true, in that there is a sheep in the field in question. However, what Roddy is looking at is not a sheep, but rather a big hairy dog that looks just like a sheep and which is obscuring from view the sheep standing just behind.[241]

---

[240] Gettier-style cases derive from Edmund Gettier's seminal paper "Is Justified True Belief Knowledge?" in *Analysis* 23 (1963): 121-123, which became one of the most famous papers in 20th century philosophy. Its inception sent many epistemologists into a tailspin. In that paper, Gettier successfully presented counter examples to the tripartite account of knowledge, which was, at the time, the most widespread account of knowledge. The three conditions of this tripartite view of knowledge were (1) the truth condition: If S knows that p, then p must be true. After all, one can't know something if it's not even true. (2) the belief condition: If S knows that p, then S must believe that p. After all, if S knows p, then it better be S who believes that p. There first two conditions are still recognized as necessary by virtually all epistemologists as necessary for any plausible theory of knowledge. The third condition, which Gettier showed was not the third necessary and jointly sufficient condition needed to explain instances of knowing is this. (3) the justification condition: If S knows that p, then S's belief that p must be justified. This was supposed to explain the difference between knowing and having a mere true belief or a lucky belief or something else. But Gettier showed that one can believe a proposition that is true and justified, yet fail to be an instance of knowledge. Later, Linda Zagzebski showed that all Gettier-style cases have a similarly structured recipe in her paper "The Inescapability of Gettier problems" in *Philosophical Quarterly* 44 (1994): 65-73. This two-step recipe involves constructing a case wherein a stroke of bad luck (e.g., one's justification for the relevant belief that p is such that by itself would result in a false belief) gets cancelled out by a stroke of good luck (e.g., despite the justification one has for believing p by itself would result in a false belief that p, p turns out to be true). The stroke of good luck makes the belief turn out true, but for reasons that are disconnected from one's justification. Zagzebski diagnosis the problem Gettier-style counterexamples surface against different accounts of knowledge as this: any account of knowledge that assumes the nature of knowledge can be divided up into constitutive parts, namely, a true belief plus some other condition and also assume that a false belief plus that condition is possible, then there will always be enough space between the third condition and the true belief to insert a Gettier-style counterexample. She suggests closing the gap between the true belief and third condition such that the third condition *guarantees* a true belief. More specifically, she thinks the most plausible way to close the gap is not by employing some version of fallibilism but, rather, the third condition entailing truth. Truth turns out to be a component of the third condition because the third condition guarantees it. More exactly, knowledge is a belief in which the believer gets the truth because of her good epistemic behavior. She says, "The relation *A because of B* is the key element in the approach to avoiding Gettier problems I have endorsed." Zagzebski endorses the following definition: Knowledge is belief in which the believer gets to the truth because she acts in an epistemically conscientious way. She thinks this definition avoids the Gettier problem and answers why knowledge is more valuable than true belief. See Linda Zagzebski, *Virtues of the Mind* (Cambridge: Cambridge University Press, 1996), p.259-339 and *On Epistemology* (Belmont, CA: Wadsworth, 2009), pp.114-129.

[241] I borrowed this version of the case from Duncan Pritchard, *Knowledge* (New York: Palgrave Macmillan, 2009), p. 11. This kind of case was first introduced by Roderick Chisholm in *Theory of Knowledge*, 2nd. ed. (Englewood Cliffs, NJ: Prentice-Hall, 1977), p. 105.

RODDY offers a case where one has a true belief and believes as a result of intellectual ability. But Roddy doesn't believe the truth *because* of his ability. He believes the truth because there happens to be—luckily—a sheep behind the dog that enables his belief to be true. Greco proposes that in genuine instances of knowledge, one believes the truth because S believes from an intellectual ability or power. What one believes truly from an intellectual ability will explain why a true belief was acquired. So how does Greco go about unpacking the because-of relation?

Greco proposes the because-of relation between a true belief and its source should showcase the *salient* part or, perhaps, the *most* salient part of what explains why a true belief was acquired. But what determines or governs that which is deemed casually salient or explanatorily salient? What are the mechanisms at work, as it were, that govern what gets picked out as explanatorily salient? Greco borrows from the work of Joel Feinberg—specifically, the pragmatics involved in assigning blame to others—in order to provide an answer.[242] Fineberg cites two major factors that govern explanatory salience: (i) that which is abnormal in the case and (ii) what our interests and purposes we take into the inquiry. Regarding the latter, Greco says our causal explanations typically pick out only one part of the set of causal conditions that bring about the effect of interest. So, for example, he says that when "we say that poor lending practices caused the current crises in the housing market, our explanation cites only one part of a complicated causal story" or "we cite a botched defensive play as the reason for losing a game" or "we cite drunk driving as the

[242] Greco says he reconstructs Feinberg's account of blaming others, which is presented in three different papers by Feinberg. See Joel Feinberg, "Problematic Responsibility in Law and Morals," "Action and Responsibility," and "Causing Voluntary Actions," in *Doing and Deserving: Essays in the Theory of Responsibility* (Princeton: Princeton University Press, 1970). Greco first develops this principled rational for carving Gettier-style cases this way in "Knowledge as Credit for True Belief" in Michael DePaul and Linda Zagzebski (eds.) *Intellectual Virtue: Perspectives from Ethics and Epistemology* (Oxford: Oxford University Press, 2007), pp. 116-123.

cause of a car crash." [243] But these explanations are the result of our interests and purposes and not the necessary and jointly sufficient conditions that comprise a very complicated story about how these events were brought about. Indeed, given our limited cognitive resources, we are typically epistemically closed to an exhaustive explanation. Causation is tough stuff to understand. Regarding the former, the cause of a fire in a welding shop, for example, is not going to be explained by appealing to the spark that initially started the blaze since sparks flying around are quite normal for this type of environment. It will rather be something abnormal—for instance, someone forgetting to remove a gas container that was set down temporarily in order to attend to a fallen co-worker.[244]

Greco acknowledges he lacks a precise way of understanding the rules governing explanatory salience. This makes his account of knowledge a non-maximally specific and informative one. But such a consequence should be expected given the nature of causal explanations. *A fortiori*, deciding what is appropriate to cite in our causal explanation language is complicated and only partially understood. Be that as it may, he thinks two points can be made regarding the mechanisms at work that contribute to what gets picked out as explanatorily salient. Moreover, these two points are relatively uncontroversial and can explain a gamut of Gettier-style cases:

(1) Explanatory salience is partially a function of our interest and purposes.
(2) Explanatory salience is partially a function of what is abnormal or unusual.

These two points, according to Greco, provide the resources needed to explain virtually all Gettier-style cases.[245] His explanation is as follows and I quote at length:

---

[243] Greco, *Achieving Knowledge*, p. 74.

[244] John Greco, "Knowledge as Credit for True Belief," p. 127.

[245] There is another thing worth points out about what all Gettier cases share in common: namely, the Zagzebski-formula which involves a double-luck structure is responsible for generating virtually all Gettier cases. To see this, first add to your case an example of a subject who has a warranted or highly justified false

> Given our interests and purposes as information-sharing beings, our intellectual abilities have a default salience in explanations of true belief. In Gettier cases, this default salience is trumped by something abnormal in the way that S gets a true belief. In effect, Gettier cases involve something akin to a deviant causal chain. […] In cases of knowledge, S believes the truth because S believes from intellectual ability—S's believing the truth is explained by S's believing from ability. But the success of this explanation requires more than that ability is involved. It requires that S's ability has an appropriate level of explanatory salience. Such salience is there by default in normal cases, owing to our interests and purposes as information-sharing beings in need of reliable informants. But default salience is trumped by abnormality manifested in Gettier cases. Specifically, it is trumped by the abnormality manifested in the way that S ends up with a true belief.[246]

So, returning to RODDY, what is explanatorily salient is the deviant causal chain[247] between the ability of perceiving a sheep, which happens to be a sheep-dog, and the true belief that there is a sheep in the field.

---

belief that's false as result of bad luck. Second, neutralize the effects of the bad luck by inserting an element of good luck that subsequently affords the subject a true belief. More specifically, insert a stroke of good luck in such a way that it's independently related to the relation the false belief has with strong justification. The finished product is a true belief that is warranted or justified, but the justification is not connected to the true belief in the right kind of way to elicit the institution that the subject knows. E.g., consider the following case Zagzebski uses to illustrate this point. Consider a virologist who has tons of evidence indicating her patient suffers from virus X. Envisage the amount of evidence needed to satisfy any reasonable standard for knowing a patient has virus X. The virologist has, say, blood samples indicating certain protein concentrations, visual indicators deriving from the patient's skin conditions, the patient has a fever and is nauseous, the patient reports certain kinds of pain in certain locations, etc., which all point to the conditions accompanying having virus X. This is the first ingredient involving the insertion of bad luck resulting in the virologist having a justified false belief because, in fact, the indicators are deriving from an unknown virus the patient has contracted, viz., virus Y. Now we insert an ingredient of good luck. Let's say prior to the virologist forming the belief that her patient has virus X, virus X does, in fact, germinate in the patient's body, but has not caused any identifiable indications of its existence to the virologist. So, now when the virologist forms the belief that her patient has virus X, she has a justified, true belief. But we intuitively recognize that the virologist doesn't know because her true belief is a result of luck. See Linda Zagzebski, "The Inescapability of Gettier Problems," in *The Philosophical Quarterly* 44, no. 174 (1994): 65-73.

[246] Ibid., p.75.

[247] Without having to get into the contrasting realism and ant-realism approaches between Aristotle and Hume, respectively, relative to causal processes, by a causal deviant chain I simply mean and take Greco to mean there is a chain of events between an intention or disposition that causes a chain of events that result in satisfying one's intention or what a disposition reliably brings about, but does so through a chain of events that is abnormal or deviates from the statistically normal chains of events that derive from the targeted disposition. Roderick Chisholm provided one of the first examples of a causal deviant chain. This case involved intentions. More specifically, it involved a man who intended to kill his uncle. But the chain of events that resulted in his uncle dying as a result of the man's actions didn't come about as the man originally intended. What happened was that the man was so worked up prior to driving to his uncle's house with the intention of killing him that he drove his car erratically and too fast. This resulted in him swerving the car

Now Greco wants us to consider a similar, yet different kind of case that epistemologists have gotten a lot of mileage out of since it was first introduced by Carl Ginet and used by Alvin Goldman.[248] Because this kind of case can't be dealt with the same way Greco deals with RODDY, Greco thinks a different explanatory strategy must be employed than the one above. The case goes like this.

> BARNEY: Barney is driving through the county and happens to look out of the window into a field. In doing so, he gets to have a good look at a barn-shaped object, whereupon he forms the belief that there is [a] barn in the field. This belief is true, since what he is looking at really is a barn. Unbeknownst to Barney, however, he is presently in 'barn façade county' where every other object that looks like a barn is actually a convincing fake. Had Barney looked at one of the fake barns, then he would not have noticed the difference. Quite by chance, however, Barney just happened to look at the one real barn in the vicinity.

The approach employed in RODDY, Greco admits, has limitations and can't be employed here for BARNEY. For one thing, there is no deviant causal chain that can be picked out. So the strategy Greco employs is one that appeals to the nature of ability.

So what, according to Greco, does it mean to have an ability? For starters, he says an ability is a stable disposition to achieve some result under appropriate conditions.[249] But he is particularly concerned about two kinds of cases. The first kind of case is where an ability achieves some result due to good luck. The second kind of case is when one has an

---

into a pedestrian on the sidewalk. The irony is that the pedestrian happened to be his uncle and died, thus satisfying his intention. But it came about in the wrong way—the death was a result of a chain of events that deviated from the intended plan on how he was going to kill his uncle. Analogously, a dispositional belief deriving from perception may deliver a true belief, but the causal chain between the true belief and perceptual ability may have come about in a way that deviates from the normal processes involved in an ability's capacity to get the truth relative to a hospitable, epistemically cooperative environment. See Roderick Chisholm, "Freedom and Action" in Keith Lehrer (ed.) *Freedom and Determinism* (New York: Random House, 1966), pp. 28-44.

[248] I borrow this case from Duncan Pritchard in *Knowledge* (New York: Palgrave macmillan, 2009), p. 12 who borrows it from Alvin Goldman in "Discrimination and Perceptual Knowledge," in *Journal of Philosophy* 73 (1976): 771-91.

[249] Greco, *Putting Skeptic in Their Place*, p.211.

ability but, due to an inordinate amount of bad luck, the ability never achieves the relevant result, which it would achieve in more normal circumstances.

To be more concrete, Greco likes to refer to his favorite Yankee baseball player who has the ability to hit baseballs—Derek Jeter. The idea, then, is to give an account of an ability wherein Jeter's ability to hit baseballs is insulated from both good and bad luck.[250]

---

[250] Epistemologists have identified different kinds of luck that are epistemically relevant. Some of their characterizations seem to overlap. Here are some examples of epistemically relevant luck floating about in the literature. Mylan Engel Jr. (1992) has identified two kinds of luck that can potentially undermine knowledge: *evidential luck (EL)* and *veritic luck (VL)*. But Engel thinks only veritic luck can destroy knowledge. Evidential luck has to do with how one acquires the kind of evidence that results in a true belief. But this doesn't undermine knowledge as is shown when we consider a case involving a masked bank robber whose mask accidentally falls off his face. Here the bank teller sees the robber is Tom—the bank's manager! The perceptual evidence is tainted by luck insofar as how it was acquired, but this kind of luck doesn't undermine knowledge. The teller knows the manager is the robber through a reliable perceptual belief.

Veritic luck, however, does undermine knowledge, according to Engel. Consider any Gettier-style case involving a deviant causal chain between one's evidence for p and p. E.g., in Chisholm's sheepdog case, the perceiver, Roddy, has perceptual evidence of a sheep that is connected to the proposition 'There is a sheep in the field' in a causally deviant way. That is, Roddy unluckily perceives what appears to be a sheep but is, in fact, a sheepdog. But such bad luck is (using the Zagzebski formula) neutralized by another stroke of good luck, namely, a sheep standing behind the sheepdog. So when he forms the belief that 'There is a sheep in the field,' he has a justified, true belief. But the connection between the perceptual evidence and proposition is tainted by luck expressed in the form of a deviant casual chain.

Duncan Pritchard (2003) identifies a different kind of knowledge-undermining-luck than Engel's veritic luck: namely, a *modal account* of veritic luck (MVL). MVL is different from VL in that the former is concerned with the relevant *method of belief formation* and proposition believed, whereas VL is concerned with *S's evidence* and the proposition believed. Pritchard's MVL is modal because he invokes counterfactual possible world scenarios he thinks captures luck. For example, he says if an event is lucky then it is an event that occurs in the actual world but doesn't occur in a wide class of the nearest possible worlds where the relevant initial conditions for the targeted event are the same as the actual world. Pritchard adds a second condition for luck: namely, if an event is lucky then it is an event that is significant to the agent concerned.

Wayne Riggs (2009) identifies a different kind of knowledge-undermining luck from Pritchard and Engel, namely, *agential luck*. Agential luck is unlike Pritchard's modal account of luck because it emphasis things like "agent control," "inadvertentness or lack of intention," and "probabilistic unlikelihood," rather than invoking modality and possible world scenarios. Riggs thinks the debate between internalism and externalism can be understood in terms of each camp recognizing different kinds of knowledge-undermining luck. E.g., internalists capture *subjective luck* where one's belief is the result of subjective luck to the extent that one's belief is not responsibly formed. More specifically, the agent lacked the intention to form a belief that derives from a desire or motivation to get the true as in cases involving one forming a belief out of wishful thinking. And externalists capture *veritic luck* where one's belief is the result of veritic luck to the extent that getting the truth was not the result of one's abilities or was not the result of process under one's control or was the result of unreliable belief forming processes. So any belief that forms which lacks an intention to get the truth or is primarily the result of explanatory factors that lie outside one's abilities or control or is the result of unreliable belief forming processes is undermined by luck's influence on how the true belief was acquired.

There is also such a thing as *environmental bad luck*. This kind of luck occurs when the subject is in an unfriendly epistemic environment, but nevertheless forms a true belief. For example, Goldman's barn façade country is a case involving an environment where it looks like the country is peppered with barns, but only a few are real barns. The rest of them are barn facades. So if one were unknowingly driving through barn façade country and looked out their window and formed the perceptual belief that there is a barn, which

On the one hand, even if one is always successful at hitting baseballs, it doesn't follow one has the ability to hit baseballs *per se* given the success can be attributed to mindboggling good luck. For example, there are cases where, miraculously, every time you happen to swing at a baseball—while closing your eyes!—you hit it. But your success is the result of good luck and not ability. On the other hand, one may have the ability to hit baseballs but, because of a lifetime string of bad luck, is never successful at manifesting the ability. Sickness, injury, severe weather, you name it—bad luck can prevent an ability from manifesting its capacity to deliver a high rate of success. In order to extricate the targeted ability from luck's tentacles, Greco thinks an ability must not only have success but also modal success. The formal account of an ability Greco thinks avoids the deleterious effects of both good and bad luck is this.

Ability: S has an ability to achieve result R in conditions C if and only if, across the range of close possible worlds where S is in C, S achieves R in C with a high rate of success.[251]

What falls out of this account is that abilities are dispositional properties. More specifically: "to say that S has the ability to achieve result R is to say that S has a disposition or tendency to achieve R across some range of relevantly close worlds."[252] Greco believes this account weeds out strings of good or bad luck that prevent the kind of success a bone fide ability will bring about. For consider the possibility that in some nearby possible world you

---

happened to be one of the few real barns, one would have a true belief. But it would be a belief that is undermined by environmental luck.

Lastly, there is what Prichard calls *reflective epistemic luck*. This kind of luck captures the undermining luck involved in internalistic, reflective knowledge. E.g., one may, through introspection alone, believe one is not a brain in a vat. But if this belief just happens to be true in the actual world in which it was formed, but not in most of the nearby possible worlds, then it is lucky. In other words, if the belief was formed by reflection alone that one is not a brain in a vat, but false in mostly all the nearby possible worlds, then we cannot have reflective knowledge, esp., the kind reflective knowledge that falsifies the skeptics' radical hypotheses. For an overview of how epistemic luck is incompatible with knowledge see Josh Orozco, "Epistemic Luck" in *Philosophy Compass* 6/1 (2011):11-21.

[251] Greco, *Putting Skeptic in Their Place*, p.212.

[252] Greco, *Achieving Knowledge*. P. 77.

can flip a fair coin one hundred times and it lands heads every toss. This is obviously an example of an amazing string of luck you're experiencing.[253] But across the range of nearby possible worlds, your tosses will on average land heads 50 percent of the time. A modal account of ability, then, does the work of offsetting strings of both bad and good luck.

When applying this modal account to Jeter's ability to hit baseballs, we learn that there may be a nearby possible world wherein Jeter hits a pathetic average of .000 for the month of April. So, if you started following Jeter's career as baseball player at the start of this month, you'd think he was a terrible batter. But across a range of close possible worlds, Jeter hits over .300 because the string of bad luck he experienced in April will be offset. Indeed, even if you continued watching him throughout the season into October, you'd see his average climb back up to around .300. So it's not Jeter's actual success in one plate appearance or in one week or one month that matters. What matters is Jeter's ability to hit baseballs with a high rate of success across a range of close possible worlds.[254]

---

[253] I borrow this example from Greco, *Putting Skeptics in Their Place*, p. 212.

[254] On an unrelated matter, I think once a modal account of abilities is employed with respect to virtues the kind of studies situationists appeal to in order to undermine the existence of virtues turn out to be less significant. For example, in *Lack of Character* (Cambridge: Cambridge University Press, 2002), John Doris argues that virtue ethicists who appeal to character traits to explain why people consistently act moral or do the right thing make claims that are incompatible with empirical studies conducted by experimental social psychologists. More specifically, Doris argues that the research conducted by social scientists favor situationism insofar as suggesting that morally irrelevant factors comprising a subject's synchronic situation swamp any other influence bearing on how the subject morally behaves. This conclusion is contrary to virtue ethicists' claims that virtues or character-traits are causally efficacious in successfully bringing about moral actions. E.g., a study showed that 87% of the people who found a dime in a phone booth subsequently helped a stranger who dropped her bags in contrast to only 4% who didn't find a dime and helped. Doris appeals to other similarly manufactured studies to infer that such non-moral environmental factors comprising one's situation are more determinate of how one will morally respond to a situation than the efficacy of possessing virtue, which gives reason to be skeptical towards their existence.

A criticism I have towards Doris's use of these empirical studies is this. Attributing to an agent the possession or lack of a virtue on the basis of a one-time performance is misguided. It would be akin to evaluating whether or not Ted Williams possessed the virtue of hitting Major League pitching by watching one plate appearance. But even if one were to watch the whole game, or a whole week of games, or even a month of games, it would still be an inadequate sample size for determining whether or not one should attribute Ted Williams with the ability to hit Major League pitching. Such attributions would require evaluating Williams' performance for at least three months to rule out instances of bad and good luck that skew the numbers. Therefore, evaluating whether or not one has a particular virtue on the basis of a single time-sliced evaluation is an inadequate method and fails to take into consideration the agent's reasons for not

Greco has built upon this account of ability by also adding an environment component alongside the condition-component. For abilities, according to Greco, are tied to both conditions and an environment. We don't say Jeter lacks the ability to hit baseballs because the conditions for playing baseball have been compromised. That is, we don't say Jeter lacks the ability to hit baseballs because he can't do so in the dark or in a hail storm or during an earthquake or with his hands tied behind his back. Although Greco thinks the use of conditions and environment eliminate some of the vagueness behind his account of an ability, he admits that our ideas of conditions and environment can overlap. In response, he says "we can think of 'environments' as sets of relatively *stable* circumstances and 'conditions' as sets of *shifting* circumstances within an environment."[255]

With the inclusion of the environment component, we are now in a position to draw all these threads together and state the structure of an ability, according to Greco:

S has an ability A(R/C) relative to environment E = Across the set of relevantly close worlds W where S is in C and in E, S has a high rate of success in achieving R.[256]

There are four dimensions in play, then, when stating the conditions needed for one to have an ability. The first three dimensions remain constant: namely, (i) the laws of nature; (ii) the set of conditions associated with the ability in question; and (iii) S's constitution, which Greco defines as "those characteristics of S that, given the actual laws of nature, are

---

responding in a way that is considered to be the appropriate moral behavior. If virtues are abilities that have a modal nature, then its success can't be measured by some (if any) of the methods used by social scientists. If one has a virtue, then one will be successful at achieving some result across a set of relevantly close worlds where the appropriate context and environment allow the virtue to manifest itself.

[255] Greco, *Achieving Knowledge*. p. 77 (emphasis mine).

[256] Ibid. While Sosa (1999) adds a virtue-theoretic condition to a safety condition, Greco argues that a safety condition for knowledge falls out of a virtue theoretic account of knowledge (2003). In other words, Sosa thinks we must understand safety in light of how a belief gets formed by intellectual virtues which grounds a belief being safe. But Greco thinks safety falls out of seeing how abilities (in general) and intellectual abilities (in particular) need to be understood modally. E.g., to say that 'In relevantly close possible worlds, if S believes that p because of a perceptual ability, then p is true,' entails that most formulations of a safety condition will be satisfied.

relevant to whether S's actions result in producing R when S is in conditions C [relative to environment E]."[257] Dimension (iv) is the conditions within an environment that are not constant but, rather, vary in minor degrees (e.g., the trajectory of the ball pitched to Jeter may be slower or faster or curve to the left or right etc.).

This general account of ability provides the structure for what it means to have a cognitive ability. When applied to cognitive abilities, R is believing true propositions and not believing false propositions. Greco thinks one of the virtues of this account is its ability to distinguish cognitive abilities like vision, memory, and reliable reasoning from non-abilities like wishful thinking, dreaming, hasty generalizations and the like. Beliefs sourced in, say, wishful thinking do not enjoy a high rate of success of being true. And even if a belief that p sourced in wishful thinking came out true, wishful thinking would not enjoy a high rate of success in both the actual world and across close possible worlds. Greco's account of a cognitive ability stated more formally is this:

S has a cognitive ability A (R/C,F) with respect to proposition p and relative to environment E if and only if there is a field of propositions F and a set of conditions C such that

i. p is in F, and
ii. Across the set of relevantly close possible worlds W where S is in C and in E, S has a high rate of success with respect to believing correctly about propositions in F.

Let's revisit the BARNEY case. Because abilities are always understood in relation to an environment, Greco thinks Barney lacks the ability to distinguish fake barns from real ones. Put differently: relative to Barney's environment—that is, Barn Façade County—he lacks the ability to discriminate barns from barn façades. Greco says this is not an *ad hoc*

[257] Greco, *Putting Skeptics in Their Place*, p. 215. Such a definition of one's constitution allows us to decide what worlds are close to the actual world relative to whether S has the targeted ability or not. If S has the same constitution in world W as the actual world and the laws of nature remain the same in both worlds, then the possible world we are considering is a close one.

response given we already pre-theoretically think of abilities this way, and thus, can generalize this approach to cases such as these.[258] What's more, Greco thinks that knowledge attributions are always context-dependent relative to the practical reasoning context involved by the attributor or the subject of the attribution or some third party. He indicts cases like BARNEY and others as being too *under-described*. This results in an inability to attend to the relevant features of the case in question, which demonstrates a methodological flaw.

For example, consider a case involving a government employee and us. It's our job to count the number of fake and real barns on each property for the purpose of determining tax rates. Suppose our co-worker was new and didn't know he was in Barn Façade Country. Furthermore, he hasn't received the training on how to discriminate fake barns from real ones. If this rookie viewed a barn from a couple hundred yards away and recorded it as a real barn, then we wouldn't attribute him with knowing that the building he saw is a barn. Given our practical reasoning context (and his, too, for this matter), he is not a good source of information regarding the ability to discriminate fakes barns from real

[258] Greco does think there is a generality problem that can be leveled at his response. More specifically: Depending on how narrowly or how widely we specify S's environment E, world W, conditions C, field of propositions F, and success R, we will get varying success rates relative to S believing the relevant true proposition. This problem becomes particularly acute once one recognizes that his account of knowledge as a success from ability requires knowledge attributions that depend upon a number of different parameters, which is exactly the same problem any reliabilist account of knowledge faces. Greco's answer to the generality problem involves marrying a thesis developed by Edward Craig with Mark Heller's insights on the generality problem. Edward Craig (1990) defends the thesis that one of the main characteristics about the concept of knowledge is that it's a concept we use to flag down trustworthy sources of information to use in our practical reasoning. Because humans are social and depend heavily upon others for information, humans had to invent the concept of knowledge in order to establish a system whereby we can efficiently and effectively come into contact with information we can trust to act upon. Greco thinks something like this is broadly true regarding the concept of knowledge. Mark Heller (1995) argues that reliabilists err when they accept the demands embedded in the generality problem. That is, reliabilist do not need to provide a fixed principle for specifying the relevant levels of generality because context determines the correct level of generality and the appropriate context will depend upon the practical reasoning context of the attributor, subject, or some third party. The relevant parameters, according to Greco, are "set by the interests and purposes that are operative in the relevant practical reasoning context. So, for example, if we are trying to decide what we should do, the parameters are set by our practical reasoning concerns. If we are trying to decide what S should do, the parameters are set by her practical reasoning concerns, etc." (Greco, 2010:79).

ones. But suppose the practical reasoning context changed. Let's say it involved a farmer who knew they were on a farm with a real, functioning barn with no barn façades on the property and ourselves; our task is to bring a cow back to the farmer's barn. Given our practical reasoning context (and his, too, for this matter), Greco wants to say we can attribute the farmer as knowing the building we are walking the cow towards is a barn. The farmer is good source of information relative to our practical reasoning interests and thus provides a good source of information that we can act upon.[259]

What I have done so far is canvass the main lineaments of Greco's account of knowledge and show how it handles some of the cases that have vexed previous accounts of knowledge. Although it is not my intention, given the purposes of this dissertation, to defend Greco's account of knowledge, it will be helpful to further understand his account by considering some objections leveled against it and how Greco responds. The first objection comes from Duncan Pritchard.

2.4 Objection to Greco's Virtue Theoretic Account of Knowledge

2.4.1 Some Cognitive Achievements Are Not Instances of Knowledge

Duncan Pritchard challenges Greco's thesis that knowledge is a success from ability by trying to show how one can perform a cognitive achievement from ability yet fail to acquire knowledge.[260] Pritchard grants that Greco's account does well at explaining why certain kinds of successes fail to be an achievement. For example, Pritchard employs a Gettier-Style case involving an archer whose shot experiences the double-luck structure in the form of countervailing winds. A gust of wind pushes the arrow's intended trajectory to

[259] This case is borrowed from Greco, *Achieving Knowledge*, p. 80.

[260] See Duncan Pritchard, "The Value of Knowledge," in Edward N. Zalta, ed., *The Standford Encyclopedia of Philosophy* (Fall 2008), available at: http://plato.standford.edu/archives/fall2008/entries/knowledge-value/.

the left (bad luck) but is counterbalanced by a gust of wind that pushes it back (good luck) to an accurate trajectory. The arrow hits the bull's eye! But the archer's success was lucky as a result of the intervening gusts of wind. Pritchard says Greco's account of knowledge fares well at explaining why the archer's success fails to be an achievement given the explanatory salience of the countervailing gusts of wind. Indeed, Greco's account does well at explaining why certain successes are not achievements when it comes to Gettier-style cases that involve intervening luck.

But Pritchard indicts Greco's account of knowledge for failing to be able to account for *environmental* luck—its ability to account for *intervening* luck notwithstanding. So if one can provide a case wherein one achieves a success from ability yet the success is still tainted by luck, then Greco's account fails; for then knowledge would seem to be something more than a mere cognitive achievement. The case Pritchard envisages is one involving another archer.

> ARCHIE: Archie, the archer, decides to practice at his craft of archery at a new indoor target range. Archie randomly picks out a bully's eye he intends to hit amidst a set of targets that are scattered about. He draws his arrow back, releases, and successfully hits the bull's eye as a result of his ability to shoot a bow well. But there is a twist to the case. Instead of intervening luck manifested in countervailing winds, there is environment luck manifested in the set of targets Archie gets to choose from; more specifically, there is only one target without an invisible force-field encasing it and Archie randomly chose to shoot at it. Luckily, this results in the arrow hitting the target as opposed to missing the target due to the arrow ricocheting off of it.[261]

Pritchard's point is that:

> [L]uck of this [environmental] sort does not seem to undermine the thesis that Archie's success is a genuine achievement. Indeed, we would still ascribe an achievement to Archie in this case even despite the luck involved. It is, after all, *because of* his skill that he is successful, even

[261]This case is borrowed from Duncan Pritchard, "Knowledge and Final Value," in Duncan Pritchard, Alan Millar, and Adrian Haddock, eds., *The Nature and Value of Knowledge: Three Investigations* (Oxford: Oxford University Press, 2010), p. 35.

> though he could very easily have not been successful in this case. That is, his success is still primarily creditable to his archery abilities, even despite the luck involved in that success.[262]

Pritchard, thus, thinks he has provided an argument from analogy that shows Greco's thesis is still compatible with, at least, one kind of knowledge undermining luck. So given the incompatibility thesis (i.e., knowledge is incompatible with luck), Greco's thesis that knowledge is a success from ability seems forced to say similar cases involving environmental luck are still achievements provided that the success is because of one's abilities. The problem is that achievements—as construed by Greco—are, then, compatible with luck—specifically, environmental luck. Indeed, the structure of Archie is, according to Pritchard, structurally analogous to cases like BARNEY that involve environmental luck, as manifested in Barn Façade County. The upshot, then, is this: since Archie's athletic achievement is undermined by luck so, too, is Barney's cognitive achievement—that is, Barney succeeds at getting the truth through ability which, according to Greco is an achievement, yet lacks knowledge.

Greco reconstructs Pritchard's argument from analogy as follows:

1. The Archie case (involving force-fields around archery targets) is analogous to the Ginet-Goldman bard façade case in all the relevant respects.
2. The Archie case is a case of success from ability.
   Therefore,
3. The barn façade case is a case of success from ability.
4. There is no knowledge in the bard façade case.
   Therefore,
5. The barn façade case is a case of success from ability (i.e., cognitive achievement) without knowledge.[263]

With what has been said thus far, relative to the BARNEY case and the nature of ability, you may already be anticipating Greco's response. The cases are disanalogous because

---

[262] Ibid, (emphasis Pritchard's, which is referring to Greco's account of explanatory salience).

[263] John Greco, *Achieving Knowledge*, p. 88.

Barney lacks the ability to discriminate between fake barns and real ones. As pointed out above, Barney lacks knowledge in the BARNEY case because his success is not from ability. Relative to Barney's environment, he lacks the ability to have a high rate of success of distinguishing barns from barn façades. But Archie does have the ability to hit the bull's eye; his success is from ability, and thus, is an achievement. So Greco is going to the deny premise 1.

But there is still more to say about Greco's account of ability and the ARCHIE case. First, Pritchard conflates two different abilities that are operative in the ARCHIE case. There is (i) the ability to reliably hit bull's eyes with a bow and arrow—that is, archery-skills; and (ii) the ability to reliably discriminate between targets encased by force-fields from targets that are not—that is, a particular perceptual skill.[264] This distinction can be highlighted once we consider a new sport Greco calls Archery*.[265] Like the ARCHIE case, Archie is trying to hit bull's eyes with his bow. But this sport involves *first* picking out which target is not encased by a force-field and then shooting at the target. So in order to receive the kind of credit needed for achieving success at Archery*, Archie needs to successfully pick out the right target *and* successfully hit the bull's eye. Notice that this requires two distinct abilities. Lacking either ability, even if successful, will not afford Archie the kind of credit needed for achieving success at Archery*. Success must come from both the relevant perceptual ability and athletic ability.

2.4.2 Objection to KSA: Some Instances of Knowledge Are Not Achievements

We have just seen how Greco's account of knowledge handles the objection made by Pritchard, namely, that knowledge involves more than just a cognitive achievement. But

[264] Ibid.

[265] Ibid.

there is another notable epistemologist who indicts Greco's account because knowledge, according to Jennifer Lackey, sometimes doesn't even require a cognitive achievement. The kind of case that is supposed to motivate this worry involves testimonial knowledge.

Lackey thinks Greco cannot account for testimonial knowledge. More specifically, Lackey thinks testimonial knowledge (in general) and, especially, knowledge acquired from expert testimony (in particular) are cases where the hearer can acquire knowledge but fail to be deserving of any credit since the hearer's abilities are swamped by the speaker's abilities. The case she uses goes like this.

> MORRIS: Having just arrived at the train station in Chicago, Morris wished to obtain directions to the Sears Tower. He looks around, approaches the first adult passer-by that he sees, and asks how to get to his desired destination. The passer-by, who happens to be a Chicago resident who knows the city extraordinarily well, provides Morris with impeccable directions to the Sears Tower by telling him that it is located two blocks east of the train station. Morris unhesitatingly forms the corresponding true belief [P].[266]

Lackey uses this garden-variety case that is "nearly universally accepted" as an instance of testimonial knowledge. For she thinks "cases involving testimony provide the clearest and most compelling counterexamples to [Greco's thesis that knowledge is a success from ability]."[267] The thought is that Morris's abilities are not the most salient features that explain how he obtained a true belief. Rather, what explains why Morris truly believes that p has everything of epistemic relevance to do with the testifier. So if there is any intellectual credit that should be afforded to someone, it should go to the testifier. For it is the testifier's "experience with and knowledge of the city of Chicago that explains why Morris ended up with a true belief rather than a false belief. […] Thus, though it is plausible to say that Morris acquired knowledge from the passer-by, there seems to be no substantive sense

[266] Jennifer Lackey, "Why We Don't Deserve Credit for Everything We Know," *Synthese* (2007): 352.

[267] Ibid., p. 357.

in which Morris deserves credit for holding the true belief that he does."[268] This case, then, generalizes to other cases involving knowledge via testimony thereby showing how one can have knowledge without deserving any credit. Therefore, Lackey thinks Greco's account is unable to provide a necessary condition for testimonial knowledge.

The two main charges Lackey levels against Greco are, then, two:

(1) **Hearer Lacks Most Salient Casual Factor (HLS):** Hearers of testimony (in general) and Morris (in particular) should not be afforded the kind of credit Greco disperses to hearers because hearer's powers, skills, and/or abilities are not the *most salient* part of the total set of causal factors that give rise to the hearer acquiring knowledge.
(2) **Testifier Possesses Most Salient Casual Factor (TPS)**: The epistemic relevance of the hearer's contribution to obtaining a true belief is inconsequential to the testifier's contribution. Therefore, it is the speaker's powers, skills, and/or abilities that are the *most salient* part of the total causal factors that explain why the hearer obtained the truth. Therefore, if anyone deserves credit, it is the testifier who is deserving of it, not the hearer.

The first thing that can be said (and has been said) in response to Lackey's charges is that the MORRIS case is underdescribed.[269] We don't know whether Morris exercised his abilities or not. Did Morris use any discriminatory abilities while receiving the information from the passerby or did he blithely believe the testimony the way a small infant believes what mommy or daddy tells her?

Second, Greco thinks both HLS and TPS dissolve once testimonial knowledge is given a virtue theoretic account. This requires not carving up the epistemology of testimony into either reductionism or non-reductionism; it requires recognizing there is more logical space available for a virtue theoretic account.[270] Because reductionism places

---

[268] Ibid. p. 352.

[269] Wayne Riggs, "Two Problems of Easy Credit," *Synthese* 169: 201-216. John Greco, "The Value Problem" in Sven Bernecker and Duncan Pritchard, eds., *The Routledge Companion to Epistemology* (New York: Routledge, 2011), 219-231.

[270] Reductionism and non-reductionism are positions held relative to the *reducibility* of testimonial knowledge—that is, is testimonial knowledge justifiedly basic or not? In other words, (a) is testimonial

all the epistemic emphasis on the hearer (specifically, evaluating whether or not the hearer has reduced a piece of testimony to either perception, memory, or inductive inference in order to justifiedly believe the testimony on offer) and non-reductionism places all the epistemic emphasis on the speaker (specifically, how testimony is a reliable, basic source of knowledge in company with perception, memory, and induction inference), neither take into account the phenomenon of being a reliable *receiver* of testimony.[271] Greco thinks a virtue-theoretic account of epistemology places emphasis on the hearer; more specifically, evaluations will focus on whether or not the hearer has reliable capacities during the uptake process of the testimonial exchange to discriminate reliable sources of testimony from unreliable ones. Once this approach is employed, the MORRIS case should be evaluated in

---

knowledge reducible to inductive knowledge that relies upon perception, memory, and reason or is testimony a basic source of knowledge? Non-reductionism says testimony is a basic source of knowledge, which is on an epistemic par with perception, memory, reasoning processes, etc. Like other basic sources of knowledge, in the absence of any relevant defeaters, hearers are justified in believing what they are told on the basis of the speaker's testimony (e.g., Reid, Burge, Weiner, and Audi), and if the testimony is true, then testimonial knowledge is acquired. Reductionism, however, says testimonial knowledge reduces to perception, memory, and inductive inference. But there are two kinds of reductionism. *Global reductionism* reduces testimonial knowledge *in general* to perception, memory, and inductive inference. So for a hearer to be justified in believing a speaker's testimony that p, the hearer will need to possess non-testimonially based positive reasons for the reliability of testimony *in general*. *Local reductionism* reduces *each instance of testimonial knowledge* to perception, memory, and inductive inference. So for a hearer to acquire testimonial knowledge that p, the hearer will need to possess non-testimonially based positive reasons for the reliability of *each particular instance* of testimony. For testimony is not a basic source of knowledge but only a means through which a hearer becomes acquainted with a proposition that needs inductively based positive reasons—from non-testimonial sources—to support it. Such inductively based positive reasons can be (i) the kind of report given is reliable; (ii) the context in which the report is given is reliable; and (iii) the discriminatory abilities of hearer can reliably indicate the trustworthiness of the speaker. All these positive reasons, then, contribute towards providing an accumulative case for thinking the speaker's testimony that p is true. Local reductionism is the most common form of reductionism. It says S has testimonial knowledge IFF S believes p, p is true, there are no relevant undefeated-defeaters, and the hearer has non-testimonially based positive reasons in support of p (e.g., Hume, Fricker, Adler, Van Cleve). For a great overview of the landscape comprising the epistemology of testimony see *The Epistemology of Testimony,* Jennifer Lackey and Ernest Sosa (eds.) (Oxford: Oxford University Press, 2006) and Jennifer Lackey's *Learning From Words: Testimony as a Source of Knowledge* (Oxford: Oxford University Press, 2008), esp., chpt.5. I'm indebted to Lackey's work on testimony and relied upon heavily in distinguishing reductionism from non-reductionism in this footnote, in addition to the global and local distinction within reductionism.

[271] I'm indebted to John Greco for pointing out the fruitful implications a virtue-theoretic approach contributes to the epistemology of testimony by shifting the focus of reliability off the testifier (reductionism) and testimony (anti-reductionism) and onto the receiver of testimony (virtue theoretic account of testimonial knowledge). See John Greco, "Recent Work in Testimonial Knowledge" *American Philosophical Quarterly* (2010) and "Testimonial Knowledge and the Flow of Information" *Epistemic Evaluation: Point and Purpose in Epistemology*, John Greco and Dave Henderson, eds., (Oxford: Oxford University Press, *forthcoming*).

terms of whether or not Morris manifested the ability to discriminate a reliable testifier from an unreliable one. And if Morris does manifest this ability, then he can be credited with forming a true belief and achieving knowledge.

But more still needs to be said given Lackey's objection deals with the *amount* of credit someone like Morris should be afforded. Even if we grant that Morris manifests an ability to discriminate reliable testifiers from unreliable ones, HLS and TPS basically says the testifier's contribution swamps Morris's contribution towards getting the truth. So the credit thesis fails to account for testimonial knowledge. Greco thinks the problem Lackey raises through HLS and TPS is a result of a misguided methodology. He says Lackey is "*first* deciding whether [Morris] knows, and *then* deciding (on that basis) whether [Morris's] contribution is 'important enough' in the case in question."[272] A better method, according to Greco, is to "draw analogies to non-epistemic cases where our intuitions are both firm and uncontroversial. The second thing we can do is give a principled account of the analogies."[273]

Taking heed to Greco's advice, we'll consider a non-epistemic case involving the great sport of ice hockey.[274] In hockey, players employ their powers, skills, and/or abilities in an effort to successfully score goals. Just as there are all kinds of ways to score goals there are all kinds of ways players get rewarded with points that bolster their statistics depending upon how they contributed to the process involved in scoring a goal.

For example, some goals are scored without the help of a teammate's pass. It's the job of the score card evaluator to decide whether the goal was, as they say, unassisted. But

---

[272] Greco, *Achieving Knowledge*, p. 82 (emphasis is Greco's).

[273] Ibid.

[274] My use of a case involving the sport of hockey was inspired by John Greco's use of soccer analogies. See John Greco, *Achieving Knowledge* (Cambridge: Cambridge University Press, 2010), pp.82-83.

some goals are assisted. Sometimes the success of a goal is attributed to one or more teammates who helped assist the goal-scorer with the necessary passes that contributed to the goal. Such passes are rewarded by the score card evaluator by giving the passer an assist, which gets added to the player's career statistics. Assisted goals, then, cannot be attributed solely to the goal-scorer; the success of the goal is just as much attributable to the abilities of the players who set the goal-scorer up, as it were, with the necessary passes that culminated in a goal.

Indeed, there are some goals that are *almost* entirely the result of the pass made to the goal-scorer. Sometimes the goal-scorer's abilities seem trivial in comparison to abilities manifested in the teammate's pass. For example, a teammate who skates through four defenseman before feathering the puck between one opponent's legs and over another opponent's stick so that it lands right on the tape of goal-scorer's stick! Such a pass is extremely difficult whereas the goal is, by comparison, easy. (The goal-scorer only needed to stay put so as to allow the puck to ricochet of the blade of her stick.) Be that as it may, the score card evaluators still award the goal-scorer with a point despite the disproportionate abilities of the passer.

But now consider a fourth way in which goals are scored. Some goals are scored by using the backside of the opposing team's goalie the same way a basketball player uses the backboard to make a bank shot. But in this scenario the goalie (thankfully) is not awarded with scoring a goal against his team. And the reason why is that the goalie did not *intend* to score a goal nor were his abilities appropriately involved in the right sort of way. The success of the goal, then, cannot be attributed to the abilities of the goalie because they were either not involved or not involved in the right sort way. But a goal-scorer who receives an amazing pass *is* involved in the right sort of way. The success of the goal can

thus be attributed to her abilities which are a necessary part of the causal story for how a goal was achieved.

From this non-epistemic case, it seems uncontroversial that the recipient of an amazing pass still gets credit for the achievement of scoring a goal. However, were the puck to ricochet off her head, our intuitions are firm in refusing to credit her with a goal just as they are in refusing to credit the opposing goalie with a goal for allowing an opponent to bank a shot off her back. Greco says the principled explanation we can derive from this kind of analogy is this: "[C]redit for success, gained in cooperation with others, is not swamped by the able performance of others. It is not even swamped by the outstanding performance of others. So long as one's own effort and abilities are appropriately involved, one deserves credit for the success in question."[275] It's not as though the recipient of the awesome pass scored the goal as result of luck. Rather, the recipient of the pass participated in a *collaborative achievement* of scoring a goal. And those who contributed were accordingly awarded points in the form of assists or a goal. (The score card evaluator doesn't say, "Sorry!—your abilities didn't contribute enough to the goal so we're not crediting you with a goal.")You can, if you want, say the recipient of the pass was lucky. But this kind of luck doesn't undermine the achievement; rather, it makes that kind of achievement possible. Analogously, Morris's contribution is not swamped by the testifier's contribution given his own effort and abilities were appropriately involved for purpose of getting the truth. Morris's abilities and the testifier's abilities that contributed in Morris getting the truth constitute what Greco call's a *social achievement*.[276] Morris may be lucky for being in an environment where a reliable testifier happened to be present. But

[275] Ibid, p. 83

[276] Greco called such instances of testimonial knowledge a social achievement in correspondence with me.

that sort of luck doesn't undermine the social achievement; it, like the hockey pass, enables it.

2.5 How KSA Answers The Value Problem

Recall that Greco's account of knowledge works within the project of explanation, which historically traces back to Plato's *Theatetus* wherein the question of "What is knowledge?" is pursued. This approach contrasts with the project of vindication, which historically traces back to Pyrrhonian skepticism wherein answers to the question "What do we know?" are pursued. But despite Greco's answer to what knowledge is, he takes heed to Kvanvig's point that one cannot provide an answer to the "What is knowledge?" question without also answering the question "Why is knowledge valuable?" Kvanvig thinks there is a interdependency relation between the answers given to the nature question and value question.[277] This is because any adequate theory of knowledge should be able to answer both questions given they are both at the heart of the project of explanation.[278] Explaining why knowledge is valuable is to also, in part, explain what is the nature of knowledge and vice versa. Thus, if an answer to one question can't shed light on the other question, then this is a mark of inadequacy against one's theory of knowledge. For an inadequate answer to, say, the nature question will also cash out an inadequate answer to the value question and vice versa. Indeed, Zagzebski thinks it's a necessary condition for any adequate theory of knowledge to explain why its valuable.[279]

---

[277] Jonathon Kvanvig, *The Value of Knowledge And The Pursuit of Understanding* (Cambridge: Cambridge University Press, 2003).

[278] Linda Zagzebski calls the problem of what makes knowledge *better* than true belief the value problem. See Linda Zagzebski, "From reliabilism to virtue epistemology," Guy Axtell, ed., *Knowledge, Belief and Character: Readings in Virtue Epistemology* (Lanham, MD: Roman & Littlefield). Zagzebski says Michael DePaul proposed another version of the value back in 1993. See Michael DePaul, *Balance and Refinement: Beyond Coherence Methods of Moral Inquiry* (New York: Routledge, 1993).

The trouble with answering "What is the value of knowledge?" is getting clear on what the question is asking, exactly. Traditionally, the value problem traces back to Plato's *Meno* where Socrates suggests to Meno that knowledge is of no greater value than true belief. Socrates notes that either true belief or knowledge will get you to Larissa. So, there is no greater benefit afforded to the one who has knowledge because: "[As] long as he has the right [belief] about that which the other has knowledge, he will not be a worse guide than the one who knows, as he has true [belief], though not knowledge."[280] And therein lies *a* problem: our intuitions pull us towards believing knowledge is more valuable than mere true belief, but we are vexed over how to explain the source of this additional value. Let's call this the Meno problem.

The Meno Problem: Why is knowledge more valuable than mere true belief?[281]

There are, however, four more questions relative to the value of knowledge. I'll start by listing a more general question and then cite the remaining three, which are more specific.[282]

The General Value Problem: Why (how, in what way) is knowledge valuable? **The Secondary Value Problem**: Why is knowledge more valuable than any of its proper parts? Alternatively: Why is knowledge more valuable than that which falls short of knowledge?[283]

---

[279] Linda Zagzebski, *Virtues of The Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge* (Cambridge: Cambridge University Press, 1996).

[280] Plato, "Meno" in *Plato: Complete Works,* ed. John M. Cooper and D.S. Hutchinson (Indianapolis, IN: Hackett Publishing Company, 1997), p. 895.

[281] Ibid.

[282] I borrow the formulations of these questions from John Greco, whom traces the secondary and tertiary problem back to Kvanvig. Recently, Greco has introduced the quaternary problem. See John Greco, "The Value Problem," p. 221.

[283] This question was first put forth by Jonathon Kvanvig. It assumes some kind of JTB theory of knowledge, plus whatever further condition is added to account for Gettier-style counter-examples. See Jonathon Kvanvig, *The Value of Knowledge And The Pursuit of Understanding*, p. 107; 112. The idea of distinguishing the primary value problem (Meno problem) from the 'secondary value problem' and 'tertiary problem' and providing the nomenclature of the value problem is attributed to Duncan Pritchard. See Duncan Pritchard, 'Recent Work on Epistemic Value', *American Philosophical Quarterly*, 44 (2007b): 85–110.

The Tertiary Value Problem: Why is knowledge more valuable than the value of all its parts taken together?[284]

The Quaternary Value Problem: Is knowledge more valuable than anything in the neighborhood?[285]

Here is how Greco's thinks his theory of knowledge can meet the demands of these questions centering around the value problem.

Recall that Greco says knowledge is a kind of success from ability. More exactly, when S knows that p, S's having a true belief that p is to be explained by the fact that S has exercised some cognitive ability, such as reliable perception, or reliable memory, or sound reasoning. This kind of achievement is to be contrasted with those successes that can be explained as a result of luck, accident, or chance.[286] So, in cases of knowledge, S has a true belief *because* S's belief is brought about by an ability. According to Greco, this captures Aristotle's distinction in the *Nichomachean Ethics*: namely, (a) achieving some end by luck or accident, and (b) achieving the end through the exercising of one's abilities or virtues. Achievements that are brought about by way of the latter are both intrinsically valuable and

---

[284] Duncan Pritchard coined the terms 'secondary value problem' and 'tertiary value problem' in Duncan Pritchard, 'Recent Work on Epistemic Value', *American Philosophical Quarterly*, 44 (2007b): 85–110. Greco traces this question back, however, to Kvanvig; although he's not as confident attributing this problem to Kvanvig as he is the Secondary Value Problem. Be that as it may, Greco still considers it a problem that can be inferred from Kvanvig's writing and despite the possibility that Kvanvig doesn't explicitly endorse it, he wants to list it in order to show how his theory of knowledge can provide an answer to a problem that even makes inappropriate demands. See Jonathon Kvanvig, *The Value of Knowledge And The Pursuit of Understanding*, p. xiv; xv-xvi; p., 116. Greco thinks these two questions (i.e., The Secondary and Tertiary Value Problems) originally put forth by Kvanvig are far too removed from the pre-theoretical intuition that first motivated the value problem: namely, What makes knowledge more valuable than true belief? The Secondary Value Problem makes an inordinate demand and the demand made by the Tertiary Value Problem makes even a stronger one. Both of these questions lack plausible presuppositions. That is, they lack any support from our pre-theoretical convictions. Therefore, according Greco, such demands are inappropriate because they are too strong for an adequate theory of knowledge to have to answer.

[285] This question is a new addition Greco has recently formulated to highlight a difference between knowledge, understanding, and wisdom. He takes these as separate epistemic goods that have their own value but are not equal in value. See John Greco, "The Value Problem," p. 221; 225.

[286] This basic kind of solution is not unique to Greco. It has also been championed by Wayne Riggs and Ernest Sosa. See Wayne Riggs, "Reliability and the Value of Knowledge," *Philosophy and Phenomenological Research* 64 (2002): 79-96; Ernest Sosa, *A Virtue Epistemology: Apt Belief and Reflective Knowledge,* vol. 1 (Oxford: Oxford University Press, 2007).

constitutive of human flourishing. Such achievements are more valuable than mere lucky successes and are valuable for their own sake—that is, they have final value.[287] So the successful exercise of an intellectual virtue, or one's understanding, or one's wisdom is both intrinsically valuable and constitutive of human flourishing. From this, answers to the value problem questions above fall out of it straight forwardly.

The answer to The Meno Problem and The General Problem is this. Knowledge has value over and above the practical value of true belief because: "knowledge is a kind of success from ability, and in general success from ability is both intrinsically valuable and constitutive of human flourishing, which is also intrinsically valuable. Moreover, both success from ability and human flourishing have 'final' value, or value as ends in themselves, independently of any instrumental value that they might also have."[288] This answer can also be given to The Secondary Value Problem. If knowledge is a certain kind of success from ability, it is more valuable than any subset of its constituents because "virtuously produced true belief is more valuable than both true belief that is not virtuous and virtuous belief that is not true."[289]

This solution answers even The Tertiary Value Problem. Knowledge is more valuable than the value of all its parts taken together because "success from ability is more

---

[287] By intrinsic value, Greco simply means that x's value is not derived from that which x is externally related to, which would be a case of extrinsic value. E.g., Princess Dianna's dress has extrinsic value because its value derives from it being owned and worn by Princess Dianna. The dress's value is parasitic on Princess Dianna. Princess Dianna, then, has intrinsic value because she is valuable because of herself. Final value is to be contrasted with instrumental value. X has final value only if its value is an end in and of itself and is not valuable because it's a means to some other end. So, e.g., human flourishing has final value, regardless of its relation to having instrumental value in achieving some end. Money, by contrast, has instrumental value insofar as it is a means to acquire some end. But successes from ability are constitutive of human flourishing. Therefore, they have final value independently of whatever instrumental value they have. See Michael Zimmerman, "Intrinsic vs. Extrinsic Value" in Edward N. Zalta, ed., *The Standford Encyclopedia of Philosophy* (2010), available at: http://plato.stanford.edu/entries/value-intrinsic-extrinsic/#WhaExtVal

[288] John Greco, *Achieving Knowledge*, p. 99.

[289] Ibid.

valuable than an act that is both successful and from ability, but not successful *because* from ability."[290] The idea here is that even if Scott were to win an Olympic gold medal in the 100 meters competition in London as a result of his ability to run, it would not follow that such an achievement would be as valuable if Scott beat Husain Bolt and some of the other best sprinters in the world. In other words, if Scott won the race *because* the other nations' best sprinters sat out of the event or, even worse, Scott won *because* the other Olympians threw the race in order to collect money from a bet, then such a success from ability has inferior value in comparison to winning the race against the best sprinters in the world.[291] Greco says the structure of the running cases are analogous to the structure of Gettier cases in that both are successes from ability but not successes *because* of the ability. For there is an important distinction between (a) a belief's being true and formed from an ability and (b) a belief's being true *because* of an ability.

Regarding The Quaternary Question, which asks "Is knowledge more valuable than anything in the neighborhood?" Greco answers with a clarion, "No!" Because he holds to an Aristotelian account of understanding and wisdom, they both have more final value than knowledge. Since understanding is having knowledge of causes and wisdom is understanding what are the most important, valuable things, they are both achievements, and thus, have final value. But their final value is superior. Indeed, Greco thinks "wisdom has more final value than mere understanding and understanding has more final value than mere knowledge. That is, in each case the former is more valuable *for its own sake* than the latter."[292]

---

[290] Ibid., (emphasis his).

[291] This case was inspired by the one Greco used to illustrate his point. See Ibid., p. 99.

[292] John Greco, "The Value Problem," p. 225.

2.6 Summary: The Benefits of KSA

Because this theoretic account of knowledge is simple, elegant, and consistent with independently verifiable facts, it provides a powerful answer to the question 'What is knowledge?" in a way that can be consistent and unified with other phenomena we take to be true. The benefits of this virtue theoretic account of knowledge are these. First, it captures our pre-theoretical intuitions when we evaluate cases to determine whether one has genuine knowledge or not. Two, it captures the widespread assumption that knowledge must be immune from certain kinds of luck—specifically, the kind of luck that prevents one from being attributed credit for getting the truth in a non-accidental way. (Points one and two do work in providing an answer to the Gettier problem.) Third, it helps us make sense out of knowledge attributions by understanding knowledge as a species of a general kind of human achievement. Lastly, it provides a cogent answer to problems dating back to Plato: namely, Why is knowledge more valuable than mere true belief? And although most virtue epistemologists broadly agree with this definition of knowledge and the kind of work it can do, there is division over how to understand what an intellectual virtue is, consequently resulting in different theories once the details are fleshed out.

6

How Externalism Can Account For Defeating Evidence From Epistemic Superior Disagreement

We are born under a necessity of trusting to our reasoning and judging powers; and a real belief of their being fallacious cannot be maintained for any considerable time by the greatest skeptic, because it is doing violence to our constitution.

~THOMAS REID

This chapter provides an account of defeating evidence. More specifically, this chapter provides a response to the acute challenge leveled against reliabilism: namely, How can a reliably formed true belief still be defeated by strong counterevidence? Even more precise, reliabilists need to explain how counterevidence from epistemic superior disagreement fits into a reliabilist scheme given what matters for reliabilists is that a true belief is sourced in a reliable process.

1. GRECO'S ACCOUNT OF DEFEATING EVIDENCE

The purpose of this section is to provide an understanding of what is defeating evidence. Otherwise known as defeaters, there are different ways epistemologists use the term 'defeater.' I will examine some of the ways the term defeater is used. This will enable us to better understand Greco's view of defeaters. After Greco's account of defeaters is given, I will carve out a place for it within the dialectic concerning the skeptical argument from epistemic superiors who disagree with us. More specifically, because Greco's account of knowledge requires a subjective justification condition, we'll see how disagreement from epistemic superiors sometimes can and sometimes cannot provide one with a defeater against knowledge.

1.1 What Is A Defeater?

The term 'defeater' admits of many different meanings amongst epistemologists. It was first put forth by John Pollock whose seminal work on defeaters distinguished two kinds: undercutting defeaters and rebutting defeaters.[293] In general, a defeater is a reason to no longer hold a belief. More exactly, a defeater is a doubt, belief, or experience that indicates one's belief is either false, unreliably formed, or irresponsibly retained. Defeaters, then, can defeat a belief's justification in different ways.

Rebutting defeaters give you a reason or a new belief or experience for thinking your belief is false. For example, you believe there is no more syrup left in the pantry. You look in the pantry one last time before going to the store and see a new bottle of syrup out of place; it's behind a bag of flour. This perceptual experience defeats the belief that you lack syrup. Undercutting defeaters, however, defeat beliefs in a different way: namely, by indicating a belief's source is unreliable thereby undercutting the belief's justification. For example, you believe that the National Football League (NFL) record for consecutive games in which a player throws a touchdown pass is 47 games achieved by none other than quarterback, Johnny Unitas. You form this belief on the basis of a particular football almanac. Later, however, you find out that your football almanac has been discredited by NFL statisticians and analysts who discovered over fifteen-hundred errors out of the book's two-thousand statistics. Since this book is the only source from which your belief derived, you have strong evidence your belief is no longer justified and most likely false. Thus, this belief should no longer be retained.[294]

---

[293] See John Pollock, *Contemporary Theories of Knowledge* (Totowa, NJ: Rowman & Littlefield, 1986), p. 37-39.

[294] Undercutting defeaters provide good reasons for thinking the source of your belief is unreliable, but they—unlike rebutting defeaters—don't necessarily undercut your belief's justification. Bergmann provides a case involving a child telling you that the door is locked to flesh out how undercutting and rebutting defeaters defeat your belief. So consider a case where you form the testimonial belief 'The door is locked' from the

It should be pointed out that not all believed defeaters are actual defeaters, despite their ability to do defeating work.[295] An actual defeater is a true belief that defeats a belief and also enjoys the status of knowledge; whereas a believed defeater efficaciously defeats the belief (i.e., the *defeatee*), but fails to be an instance of knowledge. That is, it may, in fact, be a false belief, yet still function as a defeater. In order for a belief to function as a defeater it need not be true nor a belief enjoying the status of knowledge; it only needs to be a belief you think defeats your belief.

For example, suppose you irresponsibly believe—as a result of consulting a Ouija board—that your neighbor hates you and is bent on vandalizing your property. You wake up one day and see your car window is shattered. From this you form the belief that <One of the neighbors' kids must have broken the window on accident with the ball I saw them playing with yesterday>. This belief is, in fact, true. But you don't know that it's true because you also believe <Breaking my car window is just the kind of thing my menacing neighbor wants to do to me>. After all, you're anticipating such mischievous acts from your neighbor given the belief you formed by consulting the Ouija board. This false belief, then, defeats your true belief that the neighbor's kids broke your window by accident. So not all believed defeaters are actual defeaters because they may be false beliefs or epistemically poor reasons. An actual defeater *should* defeat one's belief. But it may not do

---

child's testimony. But then the child's parent says the child doesn't know how to tell the difference between a locked and unlocked door. Now you have an undercutting defeater for your belief's justification—you're given strong evidence that the source of your belief is unreliable. You don't gain knowledge that the door is locked or unlocked; you only learn that the belief's source is suspect. But if you turn the door's handle and discover it's unlocked—prior to hearing the child's parent say the boy can't tell the difference between a locked and unlocked door—you would have a rebutting defeater for your testimonial belief. Why? Because you have a new experience or belief that your previous belief is unjustified. See Bergmann, *Justification Without Awareness* (Oxford: Oxford University Press, 2006), p. 159.

[295] I borrow this distinction from Michael Bergmann. See Michael Bergmann, *Justification Without Awareness* (Oxford: Oxford University Press, 2006), p. 160.

so as a result of one failing to believe what one should believe.[296] This kind of defeater has been deemed a normative defeater. A normative defeater is a doubt or belief that S *ought* to have and that indicates that S's belief that p is either false or unreliably formed or sustained.[297]

1.2 Greco's Account of Defeaters

1.2.1 Staying Consistent with Externalist Principles

Defeaters prevent beliefs from being instances of knowledge by rebutting the belief's content or undermining the belief's justificatory support. Defeaters also satisfy a subjective or rationality condition for knowledge that is internalist. This internalist no-defeater condition amounts to this: S knows that p or S's belief that p has warrant or S's belief that p is justified only if S doesn't believe that her belief is defeated by other things S believes. The "other things" S believes functions, then, as a believed defeater that can either rebut or undercut a belief's warrant/justification.[298]

What is unique about Greco's account of defeaters is that he parts company with reliabilists who go internal. That is, Greco rejects the move virtually all reliabilists make: namely, take an internalist approach to explain defeating evidence. In other words,

---

[296] Alvin Plantinga calls actual defeaters 'purely alethic rationality defeaters' because he wants to explain how a defeater does its defeating work within a cognitive system that is immune from error. There are no false beliefs or non-alethically aimed cognitive processes at work within such a cognitive system. See Alvin Plantinga, "Reply To Beilby's Cohorts" in James Beilby, ed., *Naturalism Defeated?: Essays on Plantinga's Evolutionary Argument Against Naturalism* (London: Cornell University Press, 2002), p. 209.

[297] See Jennifer Lackey, *Learning From Words* (Oxford: Oxford University Press, 2008), p. 45. Lackey also takes there to be what she calls a psychological defeater: A psychological defeater is a doubt or belief that is had by S and that indicates that S's belief that p is either false or unreliably formed or sustained. She also accepts that psychological defeaters can be either rebutting defeaters or undercutting defeaters. Moreover, she agrees that psychological defeaters remove a belief's justification because they are recognized as doing so by subjects, regardless of whether or not the defeater is true, reliable, or has positive epistemic status.

[298] In addition to Bergmann, Lackey, and Plantinga, I also could have included many other epistemologists like Alvin Goldman and Robert Nozick who also have an internalist, no-defeater condition for knowledge. See Alvin Goldman, "What is Justified True Belief?" in George Pappas, ed., *Justification and Knowledge* (Dordrecht: D. Reidal, 1979) and Robert Nozick, *Philosophical Explanations* (Cambridge, MA: Harvard University Press, 1981).

reliabilists recognize that *de facto* reliable processes alone that produce a true belief is too weak for knowledge. Many counter examples to reliabilism elicit this intuition. So, in order to account for these counter examples, reliabilists include a no-defeater condition which gets fleshed out within an internalist framework. So, even though reliabilism is a version of externalism, many reliabilists become a hybrid of sorts once they adopt this internalist condition—specifically, a no-defeater condition for knowledge. Bergmann, for example, defines the no-defeater condition (NDC) as follows:

NDC is satisfied by *S*'s belief that *p* if and only if *S* does not believe (and would not upon reflection) that her belief that *p* is defeated.[299]

Greco thinks reliabilists who go internal relative to defeating evidence ignore one of the most important insights reliabilism offers: namely, that all evidence ought to be understood in terms of "contingent belief-forming dispositions of cognitive agents."[300] Internalists, however, characterize defeating evidence as that which supervenes on one's mental states. This defeating evidence is accessed via reflection alone which is what makes it internal. Indeed, this strategy helps intuitively explain why it seems like S's belief loses its positive epistemic status once S encounters strong—albeit misleading—counterevidence, regardless of whether or not S's belief was reliably formed. E.g., consider one of Bonjour's cases that has been a burr in the reliabilist's saddle; for this kind of case exploits how a reliably formed true belief still seems to lose its positive epistemic status, even though the counterevidence is misleading. Herein lies the rub because what matters for a reliabilist's/externalist's account of knowledge is that S's belief was the product of a reliable process.

---

[299] Michael Bergman, "Internalism, Externalism, and the No-Defeater Condition," *Synthese* 110, 3 (1997): 407. Quoted from John Greco, *Achieving Knowledge* (Cambridge: Cambridge University Press, 2010), p. 158.

[300] John Greco, *Achieving Knowledge*, p. 162.

> Case I. Samantha believes herself to have the power of clairvoyance, though she has no reasons for or against this belief. One day she comes to believe, for no apparent reason, that the President is in New York City. She maintains this belief, appealing to her alleged clairvoyant power, even though she is at the same time aware of a massive amount of apparently cogent evidence, consisting of news reports, press releases, allegedly live television pictures, etc., indicating that the President is at that time in Washington, D.C. Now the President is in fact in New York City, the evidence to the contrary being part of a massive official hoax mounted in the face of an assassination threat. Moreover, Samantha does in fact have completely reliable clairvoyant power, under the conditions that were then satisfied, and her belief about the President did result from the operation of that power.[301]

As a result of such counterexamples, many externalists add a no-defeater condition that gets cashed out in internalist's terms. E.g., Greco says something like the following principle is employed by internalists and externalists who require a NDC.

Internalist Defeat Principle (IDP). $(p)(q)[\{P(p/q)<.5\} \& \sim\exists(r)\{P(p/q\&r) > .5\} \rightarrow qDp]$

Here is how Greco interprets IDP, which seems to be the driving intuition for why Samantha doesn't know in Case 1 above: "Let p and q range over propositional contents of S's beliefs and let r range over the set of S's relevant evidence. If the probability of p on q is low, and if there is no further relevant evidence r such that adding that evidence to q makes the probability of p high, then q defeats any [positive epistemic status] attaching to S's belief that p."[302]

IDP, according to Greco, explains "the intuitive appeal of Bonjour's […] counterexamples, and why so many externalists agree that Samantha [doesn't] know …."[303] Because the belief p—<The president is in NYC>— is based on Samantha's belief that she's clairvoyant. But she has counterevidence q that suggests the president is not in NYC.

---

[301] Laurence Bonjour, "Externalist Theories of Empirical Knowledge," *Midwest Studies in Philosophy* 5 (1980): 53-73.

[302] John Greco, *Achieving Knowledge*, p. 160-161.

[303] Ibid.

And there is no other evidence r that functions as a *defeater-defeater* towards q. So, q defeats p's positive epistemic status. This is why externalists employ an internalist condition in form of principle IDP so as to explain why someone like Samantha doesn't know despite her belief being the product of a reliable source. That is to say, if S believes that p and believing p is both true and the result of a reliable process, but S also has a lot of strong—*albeit* misleading—counterevidence that entails not-p, then any positive epistemic status S has for her belief that p is defeated.

The internal reasoning showcased in principle IDP, then, explains how Samantha's defeating evidence supervenes on her mental states that she can reflectively access. But Greco rejects this strategy because it faces a dilemma:

> (1) Reliabilist must understand defeaters in either an external or internal way.
> (2) If the reliabilist understands defeaters in an external way, then they are still stuck with Bonjour-like counter-examples.
> (3) If reliabilists understand defeaters in an internal way, then they are either (a) making their theory sound *ad hoc* (because reliabilists care about evidence that reliably supports a belief) or (b) making their theory incoherent.
> Therefore,
> (4) A no-defeater condition for reliabilism results in being either ad hoc or theoretically incoherent.[304]

Premise (1) states that defeaters must be understood as a relation one has internal access to or not. More specifically, either all defeat relations supervene on the content of propositions internally accessed or they don't. So, either IDP is true or false relative to defeating evidence. Premise (2) provides the driving intuition for why some externalists adopt an internalist no-defeater condition (NDC) for their account of knowledge. For Bonjour-like counter examples intuitively show externalist's accounts of defeating evidence to be inadequate. Premise (3) offers the consequences of choosing a particular horn (i.e., either premise (2) or (3)): namely, (a) this horn exposes how externalists that go internal

[304] Ibid., p. 160.

abandon the very principle that makes them an externalist: namely, all evidence, including defeating evidence, is to be understood as contingent belief-forming dispositions by cognitive agents. But going internal utilizes something like IDP which is a different account of defeating evidence than what governs externalism. So, it looks *ad hoc* to include a NDC just to parry such counterexamples. Horn (b) charges externalists who go internal as providing an account of defeating evidence that is incoherent. The reason why is this: on the one hand, reliabilism offers an account of evidence *in favor of* a belief by appealing to reliable sources or processes. Yet, on the other hand, when it comes to evidence *against* a belief, they abandon appeals to reliable sources and processes and go internal which relies upon a principle like IDP. So, it undermines the internal consistency of a reliabilist theory of knowledge and evidence.

In light of this, Greco rejects premise (2). That is, Bonjour-like counterexamples are not the death knell for reliabilists. For reliabilists can maintain an externalist account of defeating evidence and still plausibly respond to Bonjour-like counterexamples. What's more, one need not abandon the great insight of reliabilism: namely, that all evidence ought to be understood in terms of contingent belief-forming dispositions comprising an agent's cognitive character.

1.2.2 How Greco Stays True to an Externalist Account of Defeating Evidence

So how does Greco stay true to the reliabilist maxim that all evidence ought to be understood in terms of contingent belief-forming dispositions of cognitive agents yet account for our pre-theoretical intuitions towards cases involving a reliably formed true belief vis-a-vis counterevidence that comes from *epistemic superiors*? As we learned in Chapter Five, he adds a subjective justification condition. Specifically, he provides an account of epistemic responsibility that stays true to the externalist's maxim that all

evidence must be understood in terms of contingent belief-forming dispositions of a cognitive agent. He defines his subjective justification condition for knowledge as follows:

Resp. S's belief that p is epistemically responsible if and only if S's believing that p is properly motivated; that is, if and only if S's believing that p results from intellectual dispositions that S manifests when S is motivated to believe the truth.[305]

We are now in a position to see how Greco defines a defeater for subjective justification without going internal.

Greco's Subjective Justification Defeater (GSJD): *S*'s belief that *q* is a *defeater of S's subjective justification* for believing that *p* if and only if (i) *S* believes that q, (ii) prior to believing that *q*, *S* was subjectively justified in believing that *p*, and (iii) when *S* is properly motivated in present conditions, *S* has a disposition to give up believing that *p* upon believing that *q*.[306]

Because evidence for knowledge and counterevidence for defeaters must be grounded in an agent's cognitive character, the evidence that is either for or against a belief gets cashed out in terms of an agent's *de facto* cognitive dispositions. When *S* is properly motivated, a belief that q will defeat *S*'s justification for believing p when *S* is disposed to give up p as a result of believing q. So, in Case I above, if the clairvoyant's defeater system is like ours, then the clairvoyant will "as a matter of fact [be] disposed to give up beliefs of the relevant sort in the face of reasons of the relevant sort…."[307] Therefore, as long as one is properly motivated, when confronted with some kind of evidence, one will either be *disposed* to give

---

[305] Ibid., p. 43.

[306] Ibid., p. 167

[307] John Greco, *Achieving Knowledge*, p. 168. I say as long as the clairvoyant's defeater system is like ours because Greco is sympathetic to saying that were the clairvoyant disposed to retaining the belief vis-à-vis such counterevidence, as long as the clairvoyant faculty is *stable*—that is, the frequency of outputs are not few and far between—and *well integrated*, cooperative, and communicates well with other modules and *dispositions* within the cognitive system, then the clairvoyant has knowledge. The problem of strange and fleeting processes leveled against reliabilism (*a la* Bonjour's clairvoyants or a la Plantinga's serendipitous brain lesion case) can, according to Greco, be dissolved once the cases are better described. As long as the agent's character is an important cause as to why a true belief is obtained, then the agent deserved credit for getting the truth. But if the true belief derives from some faculty or process that is not well integrated with the agent's cognitive character or has infrequent outputs, then the agent doesn't deserve the kind of credit needed for knowledge attributions. See John Greco, *Achieving Knowledge*, chpt. 9.

up a belief or be *disposed* to retain a belief, according to one's *de facto* cognitive dispositions constitutive of one's cognitive character.

Greco thinks this characterization of defeaters is consistent with and builds upon an important point made by Gilbert Harman, and I quote at length:

> [L]ogical relations [like the one in IDP] cannot by themselves serve as evidential norms. Whenever we have two logically incompatible propositions, there must be some additional factor that determines which we ought to accept and which we ought to reject. The same point holds with respect to other semantic relations: where *q* makes *p* improbable, this fact cannot by itself determine whether we ought to give up *q*, give up *p*, or continue to hold both. The additional factor, according to a virtue theoretic account of epistemic normativity, must regard the contingent dispositions that make up virtuous cognitive character.[308]

To better appreciate how this account is (i) psychologically plausible, (ii) makes sense of our pre-theoretical intuitions towards defeating evidence, and (iii) accommodates Harman's epistemic point, let's look at some cases. The first set of cases illustrates how S can have knowledge despite potentially strong counterevidence; the second case illustrates how S's knowledge is defeated.

2. HOW GRECO'S AGENT RELIABILISM CAN ACCOUNT FOR DEFEATING EVIDENCE FROM EXPERT DISAGREEMENT

In Chapter Three we looked at how reliabilism suffers from an acute problem of accounting for counterevidence in a way that preserves our pre-theoretical intuitions about defeating evidence. Indeed, let's revisit Feldman's challenge where he calls out reliabilists to provide an account of defeating evidence, esp., one that takes into account counterevidence in the form of disagreement.

---

[308] Ibid., p. 168.

> [E]xternalists […] must face the fact that there are puzzles about how to deal with cases in which one is faced with the prospect of respecting the evidence [of disagreement with an epistemic peer.][…][I]t is one thing to say that we can have knowledge or justified belief without having information about the sources of our beliefs [so says the externalist]. It is quite another to say that our knowledge survives the acquisition of evidence that our reasons are not so good or that our processes are not so reliable [once we face disagreement with an epistemic peer].

With what we have learned thus far regarding Greco's account of defeating evidence, we are now in a position to provide an externalist response to Feldman. More specifically, I have laid out the resources for providing a Grecoian response to how an externalist can account for defeating evidence from disagreeing epistemic peers or epistemic superiors.

2.1 A Proviso To A Grecoian Account of Defeating Evidence Via Disagreement

It's important to remember that I am working within a particular project in epistemology: namely, a project of explanation. Recall back to Chapter Five where I distinguished the difference between the project of explanation (PE) versus the project of vindication (PV).

The latter project traces back to challenges leveled by Pyrrhonian skepticism. Such forms of skepticism typically want (a) a fully general account of knowledge that employs a non-question begging argument for the conclusion that we do indeed have knowledge. Thus, it demands a *vindication* of knowledge—that is, it demands an account of knowledge that must be fully general insofar as accounting for all the knowledge we claim we have. And (b) it must be non-circular insofar as providing a non-question begging argument that vindicates our knowledge claims. In other words, an account that vindicates knowledge must do so in a manner that does not presuppose any knowledge to begin with. The former project traces back to Plato's *Theatetus* wherein an answer to the question "What is knowledge?" is pursued.

Greco sees PV making demands that are ill conceived (at best) and incoherent (at worst). For he says, PV is, in the language of proof, "asking that we give a proof but without employing any premises. […] [And] why should anyone think [PV] is possible or desirable?"[309] For example, why think we should concede to the skeptic's demands of vindicating that we have perceptual knowledge without relying on our perceptual faculties to do so? Or why concede to the demands of the skeptic who wants us to vindicate that we have knowledge of the external world without relying upon any knowledge of the external world?[310] This is why Greco says twentieth century epistemologists "clearly privileged the project of explanation over the project of vindication […]."[311] Furthermore, Greco adds, "there is an emerging consensus that [PV] is somehow flawed or misguided, and that the proper task of epistemology is [PE]."[312]

Skeptics typically demand some condition C to be satisfied in order for one to acquire knowledge. The complaint against such demands is that humans rarely ever satisfy it; yet our evaluations of cases wherein one fails to satisfy condition C are also cases where most peoples' pre-theoretical judgments are that knowledge *is* acquired. These judgments are taken to be intuitively correct. For example, consider the kind of demands made when requiring one to vindicate the reliability of one's faculties in a non-circular way. Such

---

[309] John Greco, "Religious Knowledge in the Context of Conflicting Testimony" in *Proceedings of The American Catholic Philosophical Association* 83, (2009): 65.

[310] I borrow these examples from Greco, "Religious Knowledge in the Context of Conflicting Testimony," p.66. Similarly, Paul K. Moser says such demands suffer from a kind of incoherence. He likens the skeptic's demand to requiring one to stand somewhere while at the same time not being allowed to stand anywhere. For, necessarily, if all sources of knowledge are under question—that is, the fully general requirement—and in need of verification, then "none will be non-question-begging [which is] analytically true, and importantly analytically true." See Paul K. Moser, "Skepticism Undone?" in John Greco, ed., *Ernest Sosa and His Critics* (Oxford: Blackwell, 2004), p. 141.

[311] John Greco, "Epistemology" in Constantin V. Bounds, ed., *Columbia Companion to Twentieth-Century Philosophies* (New York: Columbia University Press, 2007), p. 172.

[312] John Greco, *Achieving Knowledge*, p. 5.

demands are unreasonable given one's faculties are the only instruments that can be used to complete the task. The skeptic's demands prevent one from having enough ground for any rational inquiry to stand upon. This is why it's unreasonable to demand a vindication for the reliability of one's faculties without allowing any of the faculty's deliverance to serve as evidence. Greco denies that this kind of circularity is a problem because using one's faculties to support the reliability of one's faculty is a logical consequence of a fully general epistemology.[313] To require some other means would be impossible since we can only form beliefs by using the very faculties from which beliefs are formed. Indeed, even if we assume there is something like an omniscient being—say, a paragon of a knower—this person, too, would not be able to satisfy the skeptic's demands. Such a consequence seems to be a *reductio* to the vindicatory demands and why it's unreasonable.[314] Given the logical entailment of how any cognizer that thinks and forms beliefs in virtue of its cognitive faculties, it seems the demand for a vindication of one's cognitive faculties, without bottoming out in circularity or a regress, is assuming something about the nature of knowledge that is false.

This is one of the main reasons Greco thinks the project of epistemology should be PE and not PV. The task should be to explain how knowledge is possible. What's more, PE should provide understanding of "what knowledge is, and [provide] an account of

---

[313] John Greco, "Epistemic Circularity: Vicious, Virtuous, and Benign," in *International Journal for the Study of Skepticism* 1 (2011), p. 109.

[314] Keith DeRose, in a paper presented at Fordham University, made the point that our inability to satisfy this demand is a problem—assuming it is one—for not only humans but for every cognizer, including God. Because cognizers can only think in virtue of the way they think; it's a logical consequence of how cognition works. This reference to DeRose's point is taken from John Greco's "Putting Skeptics In Their Place," p. 186, 29n. William Alston also credits Alvin Plantinga with making the same point: namely, the inability to provide a non-question begging response for the reliability of our faculties is not something unique to the human cognitive situation. Indeed, Alston says even "God couldn't *establish* the reliability of some belief-forming practice without using some doxastic practice to do so. And then the [regress problem] would apply." See William P. Alston, *Perceiving God*, p. 150, fn., 7.

epistemic normativity…."[315] Such an approach removes any worries of begging questions in a dialectically inappropriate way. And if one demands an answer to how we know knowledge is possible we concede such an answer has not been provided. But this should not worry us given it only admits that "the project of explanation is not the project of vindication."[316] But Greco clearly wants to give the skeptic her due. For he thinks skepticism is a heuristic device that sheds light on the nature of knowledge and evidence. PE, then, should explain where skeptical arguments go wrong even if skepticism is assumed at the outset to be false. Indeed, Greco indicts those who don't take skepticism seriously and exposes the fallacious reasoning behind dismissive responses towards it.[317]

Thus, the purpose of this chapter is to explain how an externalist can account for defeating evidence—specifically, defeating evidence sourced in disagreement. Thus, it's important to bear in mind that I don't intend to provide an account of how an externalist can *vindicate* knowledge in the face of disagreement. For such projects, I'm assuming, make demands that are ill conceived (at best) and incoherent (at worst).[318] Nevertheless, externalists (in general) and reliabilist (in particular) still need to account for how disagreement—as a kind of counterevidence—can be defeated and knowledge retained. Moreover, it must also explain how disagreement can defeat knowledge. A response to such problems is where we now turn.

---

[315] John Greco, *Achieving Knowledge*, p. 5. Greco thinks an account of knowledge should do even more. In addition to explaining the nature of knowledge, he thinks the project of explanation should also explain our concept of knowledge or term "knowledge" and its cognates, as well as how the language of epistemic evaluations (e.g., knowledge attributions) function both semantically and pragmatically.

[316] Ibid., p. 5.

[317] See John Greco, *Putting Skeptics in Their Place: The Nature of Skeptical Arguments And Their Role In Philosophical Inquiry* (Cambridge: Cambridge University Press, 2000).

[318] See Chapter 2.

2.2 Counterevidence Via Disagreement as a Full-Defeater

This case is designed to show how S can have a reliably formed true belief that gets defeated by misleading testimonial counterevidence coming from an epistemic superior; more specifically, the misleading counterevidence defeats S's belief's positive epistemic status; if S retained her belief, she would be epistemically irresponsible.

> TRICK SOCKS: You see a white, fluffy sock from a couple feet away in a well lit room. This belief—<That object is a white sock>—is both true and the product of a reliable belief forming process. All is well with your knowledge of what color the sock is until your friend, Julie, tells you the sock is one of her trick socks she constructed in the lab: although it appears white, the color of the socks are really blue. Julie is your epistemic superior when it comes to matters related to nanoengineering. You know her research at MIT involves the construction of weird objects such as these. Indeed, she tells you that the surface of these socks have rotating micro-discs on the surface, which then refracts light in such a way that it can make any colored surface appear white. Indeed, Julie's colleagues came into the room corroborating her testimony when they tell her to return the trick socks back to the lab because they need to run some more tests on them. The wrinkle in this case is that one of Julie's colleagues, Ryan, is playing a joke on everyone. Before anyone arrived at the lab, he switched the trick socks with an ordinary pair of white socks.[319]

Intuitively, you would be epistemically irresponsible for retaining your belief that the socks are white in the face of all this counterevidence from epistemic superiors. Despite your belief being true and reliably produced, such misleading counterevidence (which is similar to the kind of misleading evidence the clairvoyant is exposed to) will elicit a defeater to anyone who is properly motivated to get the truth. For our pre-theoretical intuitions about counterevidence suggests that you wouldn't know despite your belief being reliably formed and true.

[319] This kind of case was inspired and adopted from Bryan Frances's case taken from his article "When a Skeptical Hypothesis Becomes Live," *Noûs* 39, 4 (2005): 559-95.

2.3 A Potential Defeater Via Counterevidence From Disagreement

The following cases illustrate how one can retain knowledge vis-à-vis counterevidence from disagreeing epistemic superiors. Counterevidence from epistemic superior disagreement can be a potential defeater that gets defeated by one's dispositions functioning as a defeater-defeater.

> Case A. Jamie is an undergraduate student taking a class from a recognized expert in the field of philosophy. The professor provides an argument in class that Jamie can find nothing wrong with. It's a reworked argument associated with Zeno's Paradox involving motion. The conclusion of the argument is that motion is impossible because one can't traverse an infinite number of parts comprising a finite distance. Jamie believes that each of the premises is true and that the form of the argument is valid. Thus, Jamie can't find anything wrong with the argument. Nevertheless, Jamie knows the car is moving while driving home from school.
>
> Case B. Riley is impressed by the professor's argument for idealism. According to Riley, the professor has provided a sound version of Berkeley's argument for idealism, according to which, objects exist but only are immaterial. Such immaterial objects are not hard, heavy objects in the world but, rather, bundles of ideas in the mind, and thus, dependent upon the mind. Riley believes both that her professor is her epistemic superior and that her professor has provided an argument that concludes that nothing exists when it is not perceived. Furthermore, Riley believes there is nothing wrong with the argument; she believes the premises are true and the form is valid. Nevertheless, Riley knows that she has two hands, even when she does not perceive them; that is, she knows her hands are external objects and not immaterial, mind-dependent ones.
>
> Case C. Peyton believes the professor is an epistemic superior. Today Peyton is persuaded that the professor's argument for determinism is sound and that if determinism is true, then humans lack free will. Peyton believes the premises are true and the argument's form valid. What's more, Peyton can't find anything wrong with the argument. Nevertheless, Peyton knows that humans are sometimes free. Indeed, Peyton knows she could have stayed in the classroom longer than she did. [320]

[320] These case were used by Greco and I have made some minor amendments to them. See John Greco, *Achieving Knowledge*, p. 165.

> Case D. Suppose Julie believes Alex Byrne's statement that we must have a mechanism for detecting our own mental states. (And suppose there really is such a mechanism operative when one introspects to discover one's occurrent mental state.) Jackie is a world renowned neurologist and recognized by Julie as her epistemic superior relative to neurology. One day Julie visits Jackie's office to participate in her neurological experiments involving the introspective-mechanism located in the brain. Jackie performs surgery on Julie's brain so that she can now read the outputs of her introspective-mechanism on a monitor. Jackie can learn from the monitor when Julie is experiencing a tickle, itch, and other phenomenal mental states. But there is a moment during the experiment when Julie reports to Jackie that she is experiencing excruciating pain in her tooth. However, Jackie consults the monitor and asks Julie to introspect again because the monitor is not reporting any such state. Julie does and reports back the same excruciating pain emanating from her tooth. Jackie adjusts some connections, including the one connected to Julie's introspective-mechanism. Jackie then consults the monitor again and tells Julie that she really isn't in pain. Indeed, Jackie insists that Julie is either lying or playing a trick on her because otherwise she'd be able to detect it. Be that as it may, Julie knows that she is pain—indeed, she feels how hurtful the pain is![321]

Despite Jamie, Riley, and Peyton all lacking the ability to see where the arguments from their epistemic superiors go wrong, we intuitively take them all to still know despite the counterevidence suggesting otherwise. Indeed, even though they believe the premises are true and the argument's form is valid, it still seems like Jamie knows the car is moving. Similarly with Riley and Peyton: we intuitively take Riley to know she has hands and Peyton to know that she has free will. Jamie, Riley, and Peyton have retained their knowledge given they're—*as a matter of fact*—disposed to retain their beliefs. And—*as a matter of fact*—they are not disposed to give them up in the face of theoretical arguments that conflict with their perceptual beliefs. And shifting from counterevidence sourced in arguments to counterevidence from testimony by an epistemic superior, we still intuitively take Julie to know that she is in pain.

---

[321] This case was inspired by Tomas Bogardus. See Tomas Bogardus, "A Vindication of the Equal-Weight View," *Episteme* (2009): 327. The reference to Alex Byrne's claim comes from Alex Byrne, "Introspection," *Philosophical Topics* 33, 1 (2005): 79-104.

2.4 How Counterevidence From Disagreement Can Both Defeat and Not Defeat A Belief

So now the questions are these: In light of Greco's account of defeaters being grounded in an agent's *cognitive character*, how does he explain how people in Cases A-D lack a defeater from disagreement with cognitive superiors thereby retaining knowledge? And how is it that in the TRICK SOCKS case one does have a defeater thereby preventing one from knowing?

Answer: Greco would say the salient feature that distinguishes the two cases is this: "*as a matter of fact*, we are disposed to be moved by considerations of the sort presented [by the testimony of epistemic superiors], and *as a matter of fact* we are disposed not to be moved by the theoretical considerations [put forth by epistemic superiors] that contradict common sense."[322] In other words, because a properly motivated agent would be disposed to no longer believe the socks are white, such an agent has a defeater—manifested as a contingent belief-forming disposition. Jamie, Riley, Peyton, and Julie—when properly motivated to get the truth—will, however, lack a disposition to give up their respective beliefs.[323] One may wonder what explains the difference between a defeater system that yields (a) a disposition to retain a belief and (b) a disposition to give up a belief? I take up this question below in § 2.5.

---

[322] John Greco, *Achieving Knowledge*, p. 171 (emphasis Greco's).

[323] Greco's account of defeaters being contingent belief-forming dispositions is consistent with the idea behind Plantinga's notion of a *partial defeater*. For sometimes counterevidence will result in one forming a disposition that results in having less confidence in the belief's truth status or one may have a weaker conviction that the belief is true. So we could describe a variant of TRICK SOCKS such that Ryan—who is an unreliable testifier given his penchant for playing jokes on people—tells you that he switched the trick-sock with a normal sock. You aren't sure whether Ryan is telling the truth or playing another one of his games. Be that as it may, it still seems like Ryan's testimony can provide a *partial defeater*, which results in you having a *weaker disposition* to believe that the sock is not white.

In the next chapter more will be said about how knowledge is possible for Jamie, Riley, Peyton, and Julie vis-à-vis epistemic superior disagreement. But since we are discussing Greco's view of defeaters, it's important to explain how defeater systems operate. It strengthens Greco's account of defeating evidence brought about by epistemic superior disagreement if it can explain why one doesn't have knowledge in cases like TRICK SOCKS, but in cases like A-D one does retain knowledge. In other words, how can Jamie, Riley, Peyton, and Julie retain knowledge despite believing their cognitive superiors believe otherwise? Moreover, they believe there are sound arguments that entail the denial of their beliefs; so how can they still know? More specifically, how can they still meet the responsibility (or subjective justification) condition for knowledge?

Part of Greco's answer includes an explanation of how our defeater systems might operate. For example he says:

> Suppose that our cognitive systems are modular in the following sense: beliefs produced by some parts of the system tend to be insensitive to beliefs produced by other parts of the system. For example, suppose (as seems to be the case) that our perceptual beliefs tend to be insensitive to beliefs produced by highly theoretical reasoning, for example philosophical reasoning. Suppose also that our perceptual faculties are highly reliable, and that our theoretical reasoning faculties are much less reliable, so that the sort of insensitivity described actually contributes to the overall reliability of the cognitive system.

If something like this is close to being true, then Jamie, Riley, Peyton, and Julie will be disposed to continue believing their respective beliefs and not be disposed to believe the testimony or conclusion of an argument. For their defeater system is such that beliefs produced by some modules (e.g., theoretical reasoning) will yield to beliefs produced by other modules (e.g., perception) when they are properly motivated. This should come as no surprise. One need only consider how much science has advanced over the years as a result of enhancing our ability to perceive the world. Many beliefs sourced in theories were

defeated as a result of newly formed beliefs sourced in perception defeated them. Our defeater system yielded to perceptual modules over theoretical modules.

Indeed, this is why something like principle IDP above should not govern Jamie's, Riley's, Peyton's, and Julie's response to the counterevidence leveled against their beliefs. Why? Because if our defeater system works this way, then the probability relations in principle IDP are not going to capture the kind of sub-personal processes at work within one's defeater system. In other words, the kind of work a defeater system does to monitor for defeaters eludes one's reflective access. So, the kind of defeating evidence in play will go beyond the scope of evidence captured by the mental states one has reflective access to. The operative defeat relations, then, will not *always* supervene upon one's reflective mental states, but rather, derive from the defeater system, manifesting as a disposition.[324] And the provenance of such a defeater is not accessible via reflection alone like IDP requires.

In TRICK SOCKS, however, the defeater system *does* yield to the theoretical reasoning module over the perception module. Here our own best thinking disposes us to give up a reliably formed true belief. But had we not given up our belief, intuitively we would be epistemically irresponsible and lack knowledge, despite the belief being true and the product of a reliable process. Because as long as we are sufficiently motivated to get the truth, our defeater system will dispose us to no longer believe the socks are blues as a result of the counterevidence from epistemic superiors and other relevant background information at work within the defeater system. One may wonder why the theoretical module sometimes should trump other modules, and sometimes should yield to them? I

[324] Greco qualifies this point about certain belief-producing modules having priority to other modules operative in the defeater system by saying this: "The claim is not that our perceptual faculties are (or ought to be) completely insensitive to contrary beliefs from our less reliable theoretical reasoning faculties. The claim is that they are (and ought to be) insensitive to some of those beliefs, some of the time." Greco goes on to say that if this is right, then principle D is false because the defeat relations depicted in it are not genuine defeaters, esp., since there is no corresponding disposition when one is sufficiently motivated to get the truth. See John Greco, *Achieving Knowledge*, p. 166.

offer a response to this question below in § 2.5. But I will need to take a circuitous route first to get there in order to the set the stage aright.

2.5 What Explains the Difference Between Epistemically Correct and Incorrect Dispositions?

2.5.1 Setting Up The Stage for a Plausible Answer…

Let's take stock: I have explained how counterevidence from disagreement can sometimes level a defeater and sometimes not. A defeater system will sometimes produce a disposition to retain, give up, or reduce confidence in a belief in the face of disagreement. An explanation for this invokes the idea that one's defeater system will process the outputs of competing modules and produce a defeater, manifested in a belief-forming disposition. The result of this process is a belief forming disposition to retain, give up, or reduce confidence in the targeted belief.

Now I need to explain the difference between an epistemically correct disposition versus an epistemically incorrect one. The question I aim to answer, then, in this section is this: What explains the difference between epistemically correct dispositions and incorrect ones?[325]

Although Greco never directly answers this question, I think he has provided enough theoretical resources to provide a plausible explanation. For he has already provided a proposed solution to the problem of strange and fleeting processes, which asks a similar question to the one we are asking. The question behind the problem of strange and fleeting processes is this: What explains our intuition that S doesn't know despite the targeted belief deriving from a perfectly reliable process? In other words, what explains the difference between reliable processes that produce true beliefs we intuitively attribute as an

[325] I'm grateful to Martin Montminy for pushing me to answer this very question and providing such helpful comments.

instances of knowledge versus others that don't? Once an answer to this question is provided, we will be in a position to offer an answer to the question concerning correct and incorrect epistemic dispositions.

The problem of strange and fleeting processes is a particular problem that plagues reliabilism. For example, consider Plantinga's case involving a subject who has a serendipitous brain lesion.[326] The subject has no evidence of having a brain lesion—no neurologist has ever given the subject a diagnosis of one. This undetected brain lesion happens to cause the subject to form the belief that she has a brain lesion. The belief turns out to be true and the brain lesion is a perfectly reliable belief producing process.

But when we stop to consider whether or not this person knows she has a brain lesion, we intuitively take such a true belief to fall short of knowledge. Without having any evidence that one has such a reliable process and believing that one has a brain lesion willy-nilly—without any evidence for or against it—is not an intuitive instance of knowledge. Similarly, there is a case offered by Keith Lehrer involving Mr. Truetemp.[327] Here is a case where one unknowingly has been equipped with a device implanted in the brain that reliably produces true beliefs about the ambient temperature outside one's skin. Whenever Mr. Truetemp thinks about what the present temperature is, he forms a true belief of the ambient temperature with 100% accuracy. But like the subject with a brain lesion, Mr. Truetemp has no evidence of this device and has no track record to verify how reliable he is at getting the truth of the temperature.

Most epistemologists intuitively take Mr. Truetemp's true belief to fall short of knowledge, as well. Indeed, it's a case aimed to show that reliable processes alone are

---

[326] See Alvin Plantinga, *Warrant: The Current Debate* (Oxford: Oxford University Press, 1993), p. 42.

[327] See Keith Lehrer, *Theory of Knowledge* (New York: Routledge, 1990), p. 163.

insufficient for knowledge. Both the brain lesion and thermometer device are considered strange and fleeting processes, which challenge reliabilists' accounts of knowledge. Fleeting processes are reliable processes the subject has no awareness of. They are like a math student that employs a certain algorithm willy-nilly to perform a math problem that turns out to be the correct algorithm for answering the targeted problem. The student thereby forms a true belief. But we intuitively take the student as not knowing the right algorithm was employed. Therefore, the true belief is not an instance of knowledge because it is the result of luck or accident.[328] Were someone to form a belief this way, we would not attribute one with knowledge.

But why do these true beliefs sourced in a reliable process fall short of knowledge? Greco's answer to this question falls out of this virtue theoretic account of knowledge. A main component of which is that *knowledge entails some kind of intellectual credit*.[329] Attributing one's success as credit worthy requires the success to be the result of one' ability; for we don't attribute credit to an action that a person didn't appropriately bring about. But if the act *occurs* because of one's ability, then it's an appropriate candidate, potentially deserving of credit. Thus, the reason we don't attribute the kind of intellectual credit worthy of knowledge to (a) Mr. Truetemp, (b) the subject with a brain lesion, or (c) the math student is that their beliefs are brought about by belief forming processes that are too strange to be considered an ability that is a part of their *cognitive character*.

Strange processes, then, don't seem constitutive of one's character. Why? Answer: because abilities constitutive of character are *stable*. That is, abilities, on Greco's account, have a modal status whereby they will be stable enough to genuinely be considered a part

[328] See John Greco, *Achieving Knowledge* (Cambridge: Cambridge University Press, 2008), p. 149.

[329] See chapter 5, § 2.

of one's character, and thus, extant across close possible worlds.[330] This is important because we don't want to attribute credit to others' successes if the targeted achievement is the result of luck or accidentality. For example, it's possible that a very rare string of good luck can result in a neophyte-archer hitting the bull's eye ten times in a row from, say, forty yards away. Indeed, we can even consider miraculous cases involving uncanny strings of good luck involving the neophyte-archer hitting the bull's eyes but this time blindfolded. The point is this: in order to remove luck's tentacles so that a legitimate credit-worthy attribution can be made, the ability in question must have a high rate of success across nearby worlds. And because Greco adopts a modal account of an ability, the ability that is a salient part-cause to the achievement must also be successful across a range of relevantly close worlds.[331] Such candidates whose abilities are stable—thereby neither strange nor fleeting—can then be the appropriate referent of credit attributions. This account explains why the neophyte-archer doesn't deserve credit for hitting the bull's eyes because the achievement wasn't brought about by an ability that will be successful in other instances let alone in a range of close by possible worlds.

Notice, then, how this explains why the math student lacks the ability to successfully pick the right algorithm. As Greco points out, the problem is not that the algorithm is unreliable. Rather, the process used by the student is *unstable* because he chose the algorithm willy-nilly.[332] There are many nearby worlds wherein he chooses the wrong one—just as there are many nearby worlds where the neophyte-archer miserably misses the target. Therefore, we say the student lacks the ability to choose the correct algorithm for

---

[330] See Chapter 5, §2.3.

[331] Greco's view of abilities was given in chapter 5. , §2.3 Stated more formally: S has an ability to achieve result R in conditions C if and only if across the range of close possible worlds where S is in C, S achieves R in C with a high rate of success. See Greco, *Putting Skeptics in Their Place*, p. 212.

[332] John Greco, *Achieving Knowledge*, p. 150.

the relevant math problem. Why? Answer: the student's choice was the result of an unstable process contrary to that which flows out of an ability. Therefore, the math student is not deserving of credit for picking the right algorithm because it was the result of an unstable albeit lucky process.

This view of attributing one with credit only if the success is due to one's abilities—which constitute one's character—applies to moral, athletic, and other actions, as well. For we neither attribute moral credit nor athletic credit to agents' actions unless their actions are *appropriately theirs*. Such actions must be the result of abilities which are constitutive of character, and therefore, the kind of things which can be deserving of credit. This is why simple reliable processes are too weak to attribute one with the kind of credit needed for knowledge. For the reliable process needs to be, according to Greco, a part of one's normal and psychologically healthy character conceived as either a virtue or ability.[333] And such virtues or abilities will have a modal status thereby contributing to stable processes—not fleeting ones—across close possible worlds unlike the strange processes involved in Mr. Truetemp and brain lesions.

Therefore, a Grecoian account of credit attributions requires the referent of credit to be an *ability* that is stable as a result of being constitutive of one's character. And for anything to be a part of one's character it's going to entail a degree of stability that exists across a range of close possible worlds given abilities have a modal status. This is why Mr. Truetemp lacks the kind of credit needed for knowledge because his ability is not a part of his cognitive character. Indeed, in many close possible worlds he lacks this implanted device, which is contrary to an ability. The same reasoning applies to the case involving a brain lesion: the brain lesion is presumably the result of some random process that would

[333] John Greco, *Achieving Knowledge*, p. 151,

not occur in nearby possible worlds. It's important to point out, however, that many of these cases are *underdescribed*. And as we'll see below in § 2.5.2, once we describe the cases more thoroughly, our intuitions may indeed shift.

So, one may object that the brain lesion is not a strange process after all because it may perhaps be part of one's DNA and thus meet the stability condition. Here Greco will say that it's not enough for an ability to be stable; it must also be *well integrated*: "[T]he cognitive process associated with the brain lesion are not sufficiently integrated with […] the person's [other] cognitive dispositions so as to count as being part of cognitive character. […]The present idea is that a disposition is part of cognitive character only if it is both (a) stable in the relevant sense, and (b) well integrated with […] the person's [other] cognitive dispositions."[334] Thus, for an ability to be well integrated, it must have frequent outputs that are sensitive to and cooperative with other modules in one's overall cognitive system. Abilities or processes that are infrequent, insensitive, and uncooperative with other dispositions governing the formation and evaluation of a belief are neither stable nor well integrated and thus not a part of one's cognitive character.[335]

What's interesting about the requirement of an ability having to be both stable and well integrated to count as part of one's cognitive character is this. Depending on how we describe the case involving subjects with, say, the power of clairvoyance, they may very well be attributed with knowing after all. Greco thinks people who have conflicting intuitions about whether or not the clairvoyant knows <The president is in New York City> is a result of the case being *underdescribed*. Were we to describe the case in such a way that the clairvoyant module has a sufficient number of outputs so as to be well integrated with the

[334] Ibid., p. 152.

[335] Ibid.

overall cognitive system—that is, the clairvoyant module is sensitive and cooperative with other modules, counterevidence, evidence, and dispositions governing the formation of beliefs—then it seems very plausible to attribute the clairvoyant as having knowledge.[336]

At this point, one may wonder how this story of integration goes. Fortunately, one will fall out straightforwardly in my answer to question we set out to answer: namely, "What explains the difference between epistemically correct and incorrect dispositions?" We turn now to the answer of this question.

2.5.2 An Attempt at Explaining the Difference Between Epistemically Correct and Incorrect Dispositions

From what we have learned above: namely, why reliable processes alone are too weak for knowledge as evidenced by cases involving strange and fleeting processes, I think a plausible explanation for the difference between epistemically correct and incorrect dispositions falls out straightforwardly. For the question at hand is 'What is the difference between a defeater system that produces epistemically correct dispositions versus epistemically incorrect dispositions?'

This question arises from cases A-D above. For example, what makes it the case that Riley's disposition to retain her belief that she has hands is an epistemically correct disposition despite the fact that she is convinced that her epistemic superior's argument for idealism is sound? The first way of explaining the difference between an epistemically correct disposition to retain a belief versus an incorrect one is by asking whether or not the defeater system is well-integrated with the overall cognitive system. One way to think of a defeater system's integration is in terms of its frequency of outputs. A module that doesn't produce a sufficient number of outputs will be a fleeting one thereby preventing it from

[336] Ibid., p. 154.

being well integrated with the other modules within the overall cognitive system. Because in order for the defeater system to be reliable, each module must establish some kind of track record so that the defeater system can update itself and "know" which modules, given the context and content of the belief, should take precedence over the outputs of other competing modules. Over time the defeater system becomes more integrated with other standing beliefs within one's overall cognitive system and modules that are producing new beliefs. A module that produces beliefs at a very infrequent rate, then, will not allow the kind of integration needed for the defeater system to adjudicate between competing outputs so as to establish which module is more reliable given the extant circumstances.

Integration$_1$ = A disposition to retain a belief, give up a belief, or reduce confidence in a belief is well integrated only if the defeater system from whence it came produces enough outputs so as to establish a relation with and sensitivity to other occurrent and standing beliefs within the cognitive system.[337]

An epistemically incorrect disposition, then, will lack integration with other beliefs operative within the cognitive system. So, if Riley's disposition to retain the belief that she has hands vis-à-vis her belief that her epistemic superior's argument for idealism is sound, then it can certainly be an epistemically incorrect one. Because unless Riley's defeater system is well integrated as result of producing a sufficient number of outputs to be sensitive to her other beliefs, then her disposition is an epistemically incorrect one. Notice that Integration$_1$ is only necessary and not sufficient for the disposition to be deemed appropriately Riley's and a part of her cognitive character. For in addition to a defeater system's output having to be well integrated with other beliefs, it must also be well

[337] I take occurrent beliefs to be the one's we're currently accessing through reflection alone. Standing beliefs extant in the cognitive system but are not currently accessed via reflection. They are beliefs we formed in the past, not currently before our conscious reflection, and not currently invoked to explain anything. Nevertheless, standing beliefs are operative within the cognitive system as evidence when they light up upon being asked a question about, say, whether you belief such-and-such is a trustworthy person.

integrated with evidence, counterevidence, cognitive modules, monitoring devices, other sub-systems within the overall cognitive system. Let's call this Integration$_2$.

Integration$_2$ = A disposition to retain a belief, give up a belief, or reduce confidence in a belief is well integrated only if the defeater system from whence the disposition came is sensitive to evidence for or against the targeted belief.

An epistemically incorrect disposition, then, will lack integration with other modules and monitoring devices operative within the overall cognitive system. Were Riley's disposition to retain her belief the result of a defeater system that is not sensitive to counterevidence, then it will be an epistemically incorrect disposition. Ultimately, the difference between an epistemically correct disposition and an epistemically incorrect one will involve facts that are conducive to reliability.

Thus far, I have described the kind of integration required for a defeater system to produce an epistemically correct disposition. Now I want to move from a level of abstraction to a more concrete level. So let's consider the difference between an epistemically correct and incorrect disposition in more detail. I will attempt to describe the differences by appealing to facts that are conducive to reliability. But we must keep in mind that this is a reliabilist account.

2.5.3 Riley's Perceptual Belief Versus Her Theoretical Belief

In this section and the following section (§ 2.5.4) I will be slipping into some internalist's perspectives and language to capture an intuitive sense of what it means to use one's best thinking. But one should recognize that I am assuming there is much more cognitive work, processing, updating, and correcting going on at the sub-conscious level when we are using our own best thinking. This is important because I don't want the external facts that are conducive to making a defeater system more reliable to get lost. For such facts elude what the subject can be aware of through reflection alone.

There seems to be some minimum conditions that need to be satisfied in order for Riley's disposition to be an epistemically correct one so as to avoid being epistemically naïve. First, Virtuous-Riley's disposition to retain her belief will be the result of a defeater system's output that is well-integrated and stable, as explained in Integration $_{1+2}$. So the defeater system is well integrated with other beliefs, modules, monitoring devices, systems, sub-systems, etc., within Virtuous-Riley's overall cognitive system. Thus, the defeater system is neither like Mr. Truetemp's device nor the brain lesion. Second, Virtuous-Riley is aware of no previous instances of when a mere argument provided conclusive evidence that caused her to give up her perceptual belief—especially, a perceptual belief of a middle-sized object in epistemically friendly environment. That is, Virtuous-Riley is aware of no previously established track record of beliefs sourced in arguments providing conclusive evidence that undercut or rebut her perceptual belief of middle-sized objects perceived in epistemically friendly environments. This kind of integration and track-record of competing sources from which beliefs derive accounts for why Virtuous-Riley's cognitive system (in general) and defeater system (in particular) becomes more reliable over time. From the start of Virtuous-Riley's cognitive life, her defeater system has undergone countless feedback loops updating its ability to process the outputs from different modules, adjudicate them, and produce the most reliable belief-forming disposition available.

These two points help explain why Riley's disposition to retain her belief in the face of counterevidence is an epistemically correct one. But we can easily envisage an additional explanation that can quickly get very detailed and complicated. For consider a variant of Case B: Virtuous-Riley believes the professor's argument for idealism is sound. But her defeater system still produces a reliable belief-forming disposition to retain her belief that

she has hands. What makes this disposition an epistemically correct one? Here is one way at providing an answer. We can say Virtuous-Riley acted in accordance with a professional Neo-Moorean epistemologist. This is to say, she knows she has hands because she perceives them. And she, therefore, infers that the argument for Idealism must be unsound. Of course, she may have never heard of G.E. Moore, much less known she performed a Moorean-Shift or that knowledge is enclosed under entailment. So she didn't act *from* her previous knowledge of a particular tradition in epistemology. Nor did she falsify one of the premises in the argument. But she did act in *accordance with* a Moorean Shift—even if she didn't act from her understanding of Moore's epistemology. Be that as it may: her defeater system can still produce an epistemically correct disposition given the kind of updating and track record that's been established over the years between beliefs sourced in perception versus beliefs sourced in arguments. Because this perceptual belief takes place in an epistemically friendly environment, Virtuous-Riley knows there is nothing funny going on in classroom wherein she formed her belief. She knows she is not watching some optical illusion on the computer or watching an illusionist perform an illusion on stage. All this background information is operative within her cognitive system—whether Riley is consciously aware of it or not—and informing the defeater system of which module should take precedence over the other, given the context and content of the contrary beliefs. As a result, Virtuous-Riley's defeater system manifests a reliable belief-forming disposition to retain the belief that she has hands. And even though she may not be able to see what is wrong with the argument for idealism, Virtuous-Riley's defeater system has had enough time to establish the kind of reliability needed to produce a reliable disposition that privileges the perceptual module's output over the theoretical reasoning module's output. The defeater system's output is reliable, and thus, epistemically correct.

Let's contrast Virtuous-Riley's disposition with Vicious-Riley's disposition to retain belief in the face of counterevidence to better understand the difference between epistemically correct and incorrect dispositions. So what makes Vicious-Riley's defeater system produce an epistemically incorrect disposition? First, she would lack Integration $_{1+2}$. A defeater system operating within a fragmented cognitive system is going to be too strange and fleeting. Thus, the defeater system can't produce an epistemically correct disposition because it's a disposition that is not appropriately hers. It's a weird output like the one clairvoyant, Mr. Truetemp, or the subject with the brain lesion receives thereby preventing it from being appropriately hers. Second, Vicious-Riley lacks any awareness of which sources of belief are more reliable than others; moreover, she lacks any awareness of what makes an environment an epistemically friendly one in comparison to an epistemically hostile one involving, say, trick lights, mirrors, diversions, strange surfaces refracting light differently, etc. She is aware, of course, that her belief that she has no hands is sourced in an argument. And she is aware that her belief that she has hands is sourced in perception. But Vicious-Riley suffers from a defeater system that, perhaps, gives precedence to beliefs sourced in arguments over other sources whenever her blood sugar level is at an even number. Or, perhaps, her defeater system randomly gives priority to one module over another. The upshot is that Vicious-Riley's defeater system will not be able to undergo the kind of feedback, processing, monitoring, and adjudication needed to update the cognitive system so as to increase the defeater system's reliability. Were Vicious-Riley's defeater system to produce a disposition to believe she has hands, it would not be an epistemically correct one. And it would clearly be an epistemically worse disposition than Virtuous-Riley's disposition.

2.5.4 Riley's Perceptual Belief Versus Her Testimonial Belief

Let's consider, again, Case D involving Riley. But this time, instead of following her epistemic superior's argument and believing it to be sound, she hears a piece of testimony. The professor tells her that external objects, including her hands, do not exist. Suppose Riley's defeater system produces a belief forming disposition to retain her belief that she has hands. What makes it an epistemically correct one?

In order for Riley to become Virtuous-Riley she will, again, need to satisfy some conditions. First, her defeater system will be well-integrated thus satisfying Integration $_{1+2}$. Second, Riley must be aware of some facts and perform in ways that are conducive towards making testimony a reliable source for knowledge. So, for example, Riley will satisfy the following conditions for being a *virtuous receiver of testimony*. That is, she will be able to (a) independently assess whether or not such external objects (like her hands) exist; (b) independently assess whether the speaker is a reliable source of information (in general) and relative to a domain and context (in particular); (c) assess whether the testimonial report is coherent and consistent with the rest of the things she believes; (d) employ a non-reasoning faculty that discriminates—(at some level)—reliable testimony from unreliable testimony by perceiving facial expressions, body language, and/or speech patterns that indicate sincerity and competency.[338] All these conditions play a role in equipping Riley's

[338] I'm indebted to John Greco for pointing out the fruitful implications a virtue-theoretic approach contributes to the epistemology of testimony by shifting the focus of the reliability of the testifier (reductionism) and testimony (anti-reductionism) to the receiver of testimony (virtue theoretic account of testimonial knowledge). See John Greco, "Recent Work in Testimonial Knowledge" *American Philosophical Quarterly* (2010) and "Testimonial Knowledge and the Flow of Information" *Epistemic Evaluation: Point and Purpose in Epistemology*, John Greco and Dave Henderson, eds., (Oxford: Oxford University Press, *forthcoming*). Jennifer Lackey makes similar points regarding what she calls autonomous reception of testimony in contrast to non-autonomous reception of testimony, which is akin to a parrot repeating back the words of its trainer. See Jennifer Lackey, "Disagreement and Belief Dependence," Jennifer Lackey and David Christensen, eds., *The Epistemology of Disagreement: New Essays* (Oxford: Oxford University Press, 2013).

overall cognitive system with an "epistemic filter" that allows that which is epistemically good to pass through and catch that which is epistemically bad.[339]

What's more, conditions (a)-(d) contribute to how it is Riley's defeater system can establish a reliable track-record for testimony—specifically, a track-record for certain kinds of reports and the kinds of contexts they are given in to establish reliability. Having this kind of track-record, then, will provide the defeater system with information for how to process this instance of testimony so that it can deliver the most reliable output. For without (a)-(d), the defeater system cannot undergo the kind of countless feedback loops needed to update its ability to process instances of testimony, adjudicate them amongst conflicting outputs from other modules, and produce the most reliable belief-forming disposition available.

So in case B, Virtuous-Riley is aware that the counterevidence for her belief that she has hands is sourced in testimony. She is aware that her belief that she has hands is sourced in perception. She satisfies (a) because she can assess—apart from the testimony—whether or not she has hands. Regarding (b), she is aware the speaker is an epistemic superior regarding the discipline of philosophy. Furthermore, she's aware of the context that the speaker and she are in, namely, a classroom and not a magician's set. But, regarding (c), she is aware that the piece of testimony uttered by the professor is controversial and that many other epistemic superiors in philosophy believe otherwise. Moreover, she's aware that this piece of testimony is inconsistent with her perceptual belief that she has hands and background beliefs for why idealism is false. Regarding (d), she takes her professor's testimony as a legitimate potential defeater as a result of believing the professor to be generally reliable and sincere. That is, Virtuous-Riley believes her professor

[339] Jennifer Lackey, "Disagreement and Belief Dependence," p. 251.

to be presently alert, articulate, competent, sincere, and clearly not suffering from any cognitive impairing drugs. Be that as it may, despite (d)'s conduciveness to the reliability of this token piece of testimony, (a)-(c) is sufficient to undermine the epistemic significance of this counterevidence from her epistemic superior's testimony.

The reason why this counterevidence from testimony is undermined is because Virtuous-Riley's defeater system produced a reliable belief-forming disposition to retain her belief. And an explanation for why Virtuous-Riley's disposition is an epistemically correct one is that she not only satisfies conditions (a)-(d) relative to testimony; but her defeater system can be a reliable one as a result of establishing a track-record for specific *kinds* of testimony. So, for example, speakers that *aren't* lucid, clear, confident and articulate will inform the cognitive system that these are indications of unreliability. Moreover, certain kinds of testimony will get categorized from non-controversial to controversial with categories in between. At the controversial end, will be, say, skeptical claims, reports on whether or not one has ever been convicted of a crime, and politicians' attacks on political opponents' characters before an election. Such controversial reports have a track-record of being unreliable, and thus, indicate unreliable testimony. At the non-controversial end of the spectrum are reports on, say, the time of day, where the nearest food market is located, and what day the trash gets picked up. Such uncontroversial reports have a reliable track-record, and thus, indicate reliable kinds of testimony.

Given enough time, Virtuous-Riley's defeater system can achieve Integration $_{1+2}$ and undergo the kind of feedback loops needed to update the defeater system's ability to produce reliable belief-forming dispositions. Thus, once Virtuous-Riley is confronted with counterevidence from her professor's testimony, she is able to virtuously receive it in such a way that her defeater system produces a reliable disposition to retain her belief. This

disposition is intuitively an epistemically correct disposition. All of this tracks with the good reasoning Virtuous-Riley manifests. But it does not exhaust the kind of cognitive processing going on at sub-conscious level and history of updates that are also contributing to the defeater system's output. But it does provide the kind of details needed to explain why Virtuous-Riley's disposition is an epistemically correct one in comparison to Vicious-Riley's epistemically incorrect disposition as we'll see below.

In light of this, it's easy to see how Vicious-Riley can have an epistemically incorrect disposition and one that is worse than Virtuous-Riley's disposition. For if Vicious-Riley fails at satisfying conditions (a)-(d), she would not be a *virtuous receiver of testimony*. For not being able to neither assess a piece of testimony nor perceive signs that indicate a speaker's reliability and sincerity will result in Riley being epistemically gullible. And virtually all dispositions produced by her defeater system will be epistemically incorrect ones. Vicious-Riley's inability to assess the professor's claim for coherency and consistency with her background knowledge and other beliefs prevents her defeater system from having Integration $_{1+2}$. This results in an epistemically incorrect disposition, as well. Therefore, if Riley can't discriminate a reliable, sincere speaker from an unreliable, insincere one, then her disposition will be an epistemically incorrect one. And if Vicious-Riley isn't aware of what kinds of reports, contexts, and domains of inquiry are more reliable than others, then her disposition will be epistemically worse than Virtuous-Riley's disposition. Furthermore, if Vicious-Riley's defeater system lacks Integration $_{1+2,}$ then her cognitive system can't establish a track-record for testimony. This will prevent the defeater system from producing the kind of feedback loops needed to update the defeater system's ability to produce reliable belief-forming dispositions. For when there are two or more competing modules producing conflicting beliefs, Vicious-Riley's defeater system can't produce a

belief-forming disposition that privileges one over the other in an effort to provide the most reliable belief-forming disposition available.

2.5.5 Your Testimonial Belief in TRICK SOCKS

At this point, it's helpful to consider why Riley's disposition to retain her belief in the face of conflicting testimony is epistemically correct but your disposition to retain belief in the face of conflicting testimony in TRICK SOCKS would not be epistemically correct. One way to explain the difference between you and Riley is showing how you fail at being a virtuous receiver of testimony if you were disposed to retain your belief. So let's revisit the conditions for being a virtuous receiver of testimony to see which ones you would violate.

Condition (a) is the ability, in principle, to independently assess a speaker's claim. This involves being able to consult different standards, methods, and authorities to verify a piece of testimony even if one doesn't actually engage in the process. In other words, in principle, one can reduce the piece of testimony to a species of inductive knowledge if one wanted to do the work. One need not embark upon such a project in order to get testimonial knowledge.[340] But it could be done.

Condition (b) requires a virtuous receiver of testimony to be able to independently assess whether the speaker is a reliable source of information (in general) and a reliable source of information relative to a domain and context (in particular). Here one is consciously aware of whether the speaker is an expert or not. That is, one knows whether or not the speaker is reliable at disseminating accurate information. The virtuous receiver of testimony can assess, for example, whether the speaker is an authority, expert, or cognitive superior relative to a specific domain. She knows, for example, that reports given

[340] See Chpt. 5, § 2.4.2 for the difference between reductionism versus non-reductionism relative to the epistemology of testimony. In short, reductionism takes testimonial knowledge to be a species of inductive knowledge whereas non-reductionism takes testimony to be a basic source for knowledge.

in the *National Inquirer* are not as reliable as reports given in *National Geographic*.[341] And she knows, for example, that small children are reliable at reporting their names but not at diagnosing what's wrong with a vehicle's catalytic converter.

A virtuous receiver of testimony satisfies condition (c) by being able to assess whether the testimonial report is coherent and consistent with the rest of the things she believes. If one is not able to recognize incompatible beliefs or contradictory statements, then the minimal standards for rationality are violated. Furthermore, she knows which reports are more reliable given what she knows about the context they are given in. For example, she knows an uncooperative witness's report given at the police station is not going to be, on average, reliable. But a mother at home telling her child where the milk is at is a reliable report.

Condition (d) overlaps with condition (b). But it's distinct because it captures operative cognitive processes going on at the sub-conscious level. In other words, the virtuous receiver of testimony is not aware via reflection alone of all these kinds of cognitive processes at work. For (d) captures a non-reasoning faculty that can discriminate—(at some level)—reliable testimony from unreliable testimony by perceiving facial expressions, body language, and/or speech patterns that indicate sincerity and competency. Again, all these conditions play a role in equipping the virtuous receiver of testimony's cognitive system with an "epistemic filter" that allows that which is epistemically good to pass through and catches that which is epistemically bad. The upshot is that one's defeater system will be better equipped at producing a reliable belief-forming disposition that will be an epistemically correct one.

---

[341] I borrow this example from Jennifer Lackey in "It Takes Two to Tango: Beyond Reductionism and Non-Reductionism in the Epistemology of Testimony," in Jennifer Lackey and Ernest Sosa, eds., *The Epistemology of Testimony* (Oxford: Oxford University Press, 2006), p. 178.

So if you were to retain your belief in TRICK SOCKS in the face of Julie's testimony, then you would be violating condition (b). This is because you know from previous experiences that Julie is not only a sincere and reliable knower of information (in general), but your epistemic superior relative to matters pertaining to nanotechnology's ability to generate optical illusions. Thus, you've established a track-record with Julie. You know from previous cases that even though you perceived X as *f*, X was really *g* because X was one of Julie's projects taken from the lab. There have been previous times Julie showed you examples of her work and explained to you how the nanotechnology was able to generate different illusions. So now you can, in principle, satisfy condition (a) if you wanted to. But because you know Julie to be a reliable and sincere testifier, you save yourself from the work of verifying it and trust her simply on the basis of her say so. Of course, were a similar situation to arise wherein some shifty stranger tried to convince you that this stone was really a cleverly concealed 2.50 carat diamond and that he will sell it to you for $50.00 then, in this context, under these conditions, you would be intellectually gullible to believe such testimony. But when your perceptual belief that X is *f* conflicts with Julie's testimony that X is *g*, you know that—in *this* context—Julie's testimony is more reliable than your perception. Put differently: your defeater system is able to assess the outputs from both modules and determine what belief-forming disposition, in this context, should be produced thereby allowing one module to take priority over the other.

Therefore, were your defeater system to produce in you a belief-forming disposition to retain your belief that the socks are colored white in the face of Julie's testimony that the socks are one of her projects and that they're really colored blue, then your disposition would be an epistemically incorrect one. For your defeater system is not reliably producing a belief-forming disposition as a result of you being a *vicious receiver of*

testimony. That is, you failed at satisfying condition (b).The key difference, then, for why your disposition to retain your perceptual belief is epistemically *in*correct and why Riley's disposition to retain her perceptual belief is epistemically correct turns on Riley's disposition being reliable as a result of satisfying Integration $_{1+2}$ and conditions (a)-(d). Furthermore, assuming you constantly violate condition (b), Riley's defeater system will be more reliable than yours. For her defeater system will have benefitted from conditions (a)-(d) contributing towards providing the kind of feedback loops needed to update the defeater system's ability to process the outputs from different modules, adjudicate them, and produce the most reliable belief-forming disposition available.

2.5.6 Are All Disposition Aimed At Getting The Truth?

There is another concern that I want to address. This problem arises from that fact that not all belief-forming processes are aimed at truth. Notice how the first way of explaining the difference between epistemically correct and incorrect dispositions appeals to processes of an alethically aimed system. More exactly, this cognitive system involved processes that derive from abilities that are reliable, stable, well-integrated, and successful across a range of close possible worlds. But it seems there can be cases where Integration$_{1+2}$ are operative and conditions (a) – (d) get satisfied at one time yet fail to be satisfied at another. How, then, can one be a virtuous receiver of testimony at one moment, yet be a vicious receiver of testimony at another moment?

For, example, suppose Riley is disposed to retain her belief that external objects exist despite also believing that an argument concluding otherwise is sound. And the person propounding this argument is believed to be her epistemic superior. One way of explaining why Riley's disposition is an epistemically incorrect one is this: Riley's disposition is the output of a psychological maintenance module that is not alethically

aimed, but rather, aimed at preserving psychological stability. That is, it's plausible to assume human cognitive systems do not *always* aim at getting truth.[342] Sometimes the aim of forming a belief is to secure some other good as when one believes the oncologists' diagnoses that one has only six months to live are false. Believing that one will live much longer than the oncologist's prognosis can secure all kinds of pragmatic goods.

Such a belief may be produced by a psychological maintenance module that produces what Plantinga calls an "optimistic overrider."[343] This optimistic overrider is a belief forming disposition that is not alethically aimed but, rather, aims to secure some other end: namely, a belief that will secure, say, some level of psychological stability so as to avoid psychological disaster. Here we see the difference between a truth-aimed motivation versus a prudential motivation. Indeed, it's in one's best interests, in this case, to be prudentially motivated to secure beliefs that have other useful properties than the property of truth. For one may not be able to commit to seeking the daunting amount of cancer treatments prescribed were one not prudentially motivated to believe that "I'm going to beat this disease!"[344] So, perhaps, Riley's disposition to retain the belief that she has hands is the product of a psychological maintenance module and not the product of an alethically aimed, reliable defeater system. Such a case, then, belies the fact that Integration_{1+2} and conditions (a)-(d) are always operative without a sufficient motivation to get the truth. For Riley must be sufficiently motivated to get the truth otherwise her defeater system will not be functioning reliably.

---

[342] This is why Greco's virtue theoretic account of knowledge has a subjective justification condition for knowledge, which goes like this: S's belief that p is epistemically responsible if and only if S's believing that p results from intellectual dispositions that S manifests when S is motivated to belief the truth. See Chapter 5, § 2.2.

[343] See Alvin Plantinga, "Reply to Beilby's Cohorts," in *Naturalism Defeated?: Essays On Plantinga's Evolutionary Argument Against Naturalism* (London: Cornell University Press, 2002), p. 207.

[344] I'm grateful to Martin Montminy for pressing me on this point.

Thus, we see how evaluations surrounding instances of knowledge require a subjective justification component, in additional to a reliability component. And if Riley is not sufficiently motivated to get the truth, then her disposition to retain her belief can very well be the product of a psychological maintenance module that's not truth-aimed but, rather, aimed at some other end. Indeed, lacking a sufficient level of motivation to get the truth can result in the formation of beliefs sourced in processes aimed at securing prudential or moral ends, and not epistemic ones. Therefore, a reliable defeater system can be overridden by non-alethically aimed processes if Riley lacks a sufficient motivation to get the truth.[345]

A principled answer, then, to the question: What is the difference between an epistemically correct and epistemically incorrect disposition sourced in a defeater system is this.

> **Principle of Epistemically Correct Dispositions (PECD)**: A belief forming disposition that is a product of a defeater system is an epistemically correct one only if the defeater system is well integrated with other beliefs, sub-systems, modules, awareness of facts conducive to reliability, and well integrated with one's overall cognitive system and is functioning reliably because one is sufficiently motivated to get the truth.[346]

---

[345] It's worth pointing out that Greco takes a motivation to get the truth to be a default position in human cognitive systems. So, he doesn't require an intentional act of the will thereby resulting in one wanting to get the truth in all cases of knowledge. For example, a person can know that a truck is bearing down on her through her perceptual abilities without having to muster up a motivation to get the truth because there is a default level of motivation for the truth already operative at the subconscious level. See John Greco, *Putting Skeptics in Their Place*, p. 191.

[346] One may wonder how we can tell whether a system is motivated to get the truth rather than being, say, prudentially motivated. It's important to point out that one need not have a perspective on the kinds of conditions that need to be satisfied in order to know. In other words, a justified belief needs to be grounded in one's cognitive character as pointed out above in our discussion regarding strange and fleeting processes. So, one need not have a perspective or beliefs *about* the reliability of one's cognitive character. Consider athletes or people with disabilities (e.g., autism) that have the ability to reliably achieve medals and true beliefs, respectively. They don't, according to a virtue theoretic account of knowledge, need to explain how it is one can achieve the targeted athletic feat or true belief. I say more about this in the next chapter where I think such a demand confuses the difference between the conditions that need to be satisfied for knowledge versus the conditions that need to be satisfied for understanding. But since this question concerns a subjective justification for knowledge, I think the way to answer it is in terms of whether or not the agent manifests intellectual virtues versus intellectual vices. If an agent is motivated to seek psychological stability, comfort, attention, or is simply motivated by the vice of pigheadedness, then such an agent will not be

The reason why Jamie, Riley, Peyton, and Julie still know in the face of epistemic superior disagreement is that their beliefs are sourced in epistemically correct dispositions. Accordingly, they all meet both the objective and subjective justification for knowledge. Of course, were their dispositions to retain belief the result of wishful thinking or the output of a non-alethically aimed psychological maintenance module, then they wouldn't know. Such dispositions are usually the result of an insufficient motivation to get the truth.

Admittedly, like all reliabilist's accounts, there is a generality problem lurking in the background of my explanation for the difference between epistemically correct and epistemically incorrect dispositions. But once we get clear on the level of generality that's operative when describing the relevant cognitive process at work, our intuitions are clearer as to whether or not Jamie, Riley, Peyton, and Julie have epistemically correct dispositions. And Greco thinks such an account rules correctly whether the targeted disposition is sourced in a reliable defeater system or some unreliable source like a psychological maintenance module aimed at securing non-alethic, prudential ends. For Greco says were the dispositions to retain belief in cases A-D the result of mere prudential reasoning, then "at a higher level of generality, [Jamie, Riley, Peyton, and Juie are violating their] own best thinking and so [are] unjustified. At a very narrow level of generality, specific to this odd case requiring psychological stability over truth, [Jamie, Riley, Peyton, and Julie come] out unjustified because the process described at this level does not include proper motivation to believe the truth."[347]

---

reliable at monitoring for defeaters. An agent that manifests intellectual vices will therefore not be intellectually responsible. Thus, the agent will not satisfy a subjective justification condition for knowledge and not be attributed with the kind of credit needed for knowledge.

[347] This response was obtained from my correspondence with John Greco.

But perhaps Jamie, Riley, Peyton, and Julie have multiple motivations. Were a case like this described wherein a motivation to get the truth is accompanied by a prudential motivation to secure some non-alethic end, then our intuitions will be fuzzy. Such a mix of intuitions, according to Greco, will generate noise thereby making it unclear as to whether or not the subjects in cases A-D have epistemically correct dispositions.[348]

2.5.7 A Potential Objection

I acknowledge that such an explanation for the difference between an epistemically correct disposition and epistemically incorrect one sourced in a defeater system will not satisfy everyone. For there are still philosophers who want an explanation that vindicates that a belief forming disposition sourced in a defeater system is reliable. But not only does this demand confuse the difference between (a) the external conditions S needs to satisfy in order to know and (b) S having a perspective via reflection alone on whether or not S satisfies such conditions; this demand is a variant of a general requirement for all sources of knowledge that ultimately leads to skepticism. Indeed, such a requirement ultimately renders knowledge impossible. This is why, as pointed out above, I'm working within an explanatory approach to epistemology, which traces back to Plato, as opposed to a vindication approach, which traces back to Pyrrhonian skepticism.[349]

2.6 Summary

It's important to recall that this account of a defeater system builds upon Gilbert Harman's point that logical relations alone can't function as evidential norms. Recall Greco's point that "[w]henever we have two logically incompatible propositions, there must be some

---

[348] Ibid.

[349] Because my approach is a Grecoian one, additional reasons for this approach are given in chapter 5, §2.1.2 and subsequent reasons are provided in subsequent chapters 7, §1.4 and chapter 8.

additional factor that determines which we ought to accept and which we ought to reject. The same point holds with respect to other semantic relations: where *q* makes *p* improbable, this fact cannot by itself determine whether we ought to give up *q*, give up *p*, or continue to hold both."[350]

The additional factor that determines what we ought to do is grounded in a belief forming disposition that traces back to a defeater system that is a part of one's cognitive character. Sometimes we will have internal access to *some* of the processes involved in our belief forming disposition, sometimes not. Inferential belief forming processes are clearly better candidates for accessing the kind of evidence that contributes to the formation of a belief rather than non-inferential belief forming processes. And this seems psychologically plausible, too, once we put ourselves in the shoes of Jamie, Riley, Peyton, and Julie. The additional factor that determines what we ought to do when we believe that, say, (a) the skeptical argument for the denial of motion is sound but we also believe that (b) we are moving in our ride home from school is going to be our contingent belief forming disposition. This assumes, of course, that we are sufficiently motivated to get the truth. This is what it means to say a defeater system will privilege one module's output over another one as evidenced by our disposition to, in this case, retain our belief that motion exists. But it's also important to keep in mind that this account is working within the epistemic project of explanation. So, I don't claim to have internal access of my defeater system, cognitive modules, and processes involved in producing my belief-forming dispositions. For having such a perspective on the reliability of these processes is not required for knowledge even if they are required for some level of understanding.

[350] John Greco, *Achieving Knowledge*, p. 168.

Therefore, I acknowledge that my account does not vindicate the reliability of a defeater before allowing it to be a source that can ground knowledge.

Let's now take stock: in this chapter I have done five things. First, I provided Greco's account of defeating evidence. This reliabilist account stayed true to the insight that all evidence should be understood in terms of contingent belief-forming dispositions manifested by a cognitive agent's character. Thus, it's an account that avoids having to go internal relative to defeating evidence, which avoids the criticisms of being either *ad hoc* or inconsistent with reliabilism (or both).

Second, I showed how Greco's account of defeaters intuitively accounts for how knowledge can survive the acquisition of counterevidence—specifically, cognitive superior disagreement. I did this by providing cases that elicited the intuition.

Third, I also showed how this account of defeaters can intuitively account for how disagreement can defeat knowledge. I did this by providing a case that also elicited the intuition.

Fourth, I offered an explanation for the difference between epistemically correct dispositions and epistemically incorrect ones. This explanation stayed consistent with reliabilist's theories of knowledge and evidence. And it appealed to theoretical resources that assist reliabilism in also responding to the problem of strange and fleeting processes.

Last, I considered an objection that many philosophers will level at this account of defeating evidence and its relation to one's defeater system: namely, an account that vindicates the reliability of one's defeater system's outputs. And although I take this objection to be misguided, I understand the driving intuition behind it. I trust that what I have to say in Chapter Seven will bear this remark out.

I have thus responded to Feldman's demand for an externalist account of defeating evidence. But now I do well to explain how knowledge is possible vis-à-vis disagreement with epistemic superiors—specifically, where the argument from SAFES goes wrong. We thus move on to Chapter Seven where I undertake this task.

7

A Grecoian Response to The Skeptical Argument From Epistemic Superiors (SAFES)

1. A GRECOIAN RESPONSE THAT EXPLAINS HOW KNOWLEDGE IS POSSIBLE VIS-À-VIS COGNITIVE SUPERIOR DISAGREEMENT

Any plausible response to skepticism should explain what knowledge is and explain why the skeptic's assumptions about what's required for knowledge are false. In Chapter Five, I provided Greco's account of knowledge. It will be helpful keeping in mind that for Greco knowledge is success through ability, which prevents one from acquiring knowledge luckily or by accident. In § 1.1, I familiarize the reader with the argument from SAFES. In § 1.2, I will argue for why the argument from SAFES is unsound. In § 1.3 I offer an error theory that explains epistemic superiority in a way that doesn't entail the falsity of an inferior disputant's belief. In § 2, I pose a challenge to my Grecoian response to SAFES and argue why the challenge is ultimately unsuccessful.

1.1 Where Does The Argument From SAFES Go Wrong?

Let's revisit Greco's account of knowledge, which I laid out in Chapter Five. There he said knowledge is a success from ability and not the result of luck or accidentality.[351] In cases of knowledge, S knows because there is a causal explanation for how S acquired a true belief in a credit deserving way. S's abilities explain why S acquired a true belief in a way that avoids the undermining effects of luck or accident. More formally:

[351] What is known as the *incompatibility thesis* dates back to Plato and basically says that epistemic luck is incompatible with knowledge. Because true beliefs can be acquired through guessing or wish fulfillment or reading tea leaves, epistemologists have reached a virtual consensus that knowledge requires more than mere true belief. Justification was believed to be the difference between true belief and knowledge until Edmund Gettier's 1963 paper showed otherwise as explained further below in note 52.

S knows that p *if and only if*

1. *p* is true;
2. S believes that *p*; and
3. S believes the truth because S's belief is produced by intellectual ability.[352]

So when we consider, say, Case B, we intuitively take Riley to know she hands despite believing the argument for idealism is sound. We are therefore assuming premise (1) is true: namely, Riley's hands are real, external objects in the world. Riley believes this and thus satisfies premise (2). When we consider premise (3), however, we can say she satisfies it in the following two ways.

First, she believes she has hands, which are real, external objects in the world as a result of an intellectual ability: namely, perception. Her reliable perceptual ability explains why she believes truly. And because her perceptual ability is constitutive of her cognitive character, she is deserving of the kind of intellectual credit required for knowledge. This prevents the causal explanation from appealing to that which can undermine knowledge—specifically, knowledge-undermining luck or accidentality.

But Riley faces a potential defeater once she believes the argument for idealism is sound. The causal explanation for Riley's true belief shifts from perception to her defeater system. Recall that in Chapter Six we learned that one's defeater system can constitute reliable belief-forming dispositions. Riley's defeater system is reliable and also underwrites an intellectual ability constitutive of her cognitive character. In Case B, it produced a reliable belief-forming disposition to retain her belief that she has hands in the face of a potential defeater. This intellectual ability explains why she believes truly. And because it's constitutive of her cognitive character, it's appropriately hers. So, Riley believing truly—in this instance—constitutes an achievement. As Chapter Six, § 2.5.1 explained, we don't attribute credit to agents for actions that are the result of processes that aren't appropriately

[352] John Greco, *Achieving Knowledge*, p. 12.

theirs. Riley thus satisfies the conditions for knowledge. But if Riley knows she has hands despite being aware of a potential defeater, then what is wrong with SAFES?

Let's reconsider, then, the argument SAFES so we can better understand where it goes wrong:

SAFES

(1) If …
   a. S is aware that X is an epistemic superior, and
   b. S is aware that X believes *t*, and
   c. S understands that (if *t*, then not-p), and
   d. S can't rule out *t*,

   …then S doesn't know p.

(2) (a) S is aware that X is an epistemic superior, (b) S is aware that X believes *t*, (c) S understands that (if *t*, then not-p), and (d) S can't rule out *t*.

Therefore,

(3) S doesn't know that p.

1.2 Why Clause (d) In Premise 2 Is False

Before I argue that clause (d) in premise 2 is false, we do well to first understand the difference between inferential and non-inferential belief forming processes. This distinction is important because many instances of knowledge are not the result of inferential belief-forming processes. For when it comes to knowledge, a belief's *etiology* matters. There will be instances of knowledge wherein one's disposition to believe is the result of making valid inferences from reliable evidence. And there will be other instances of knowledge wherein one's disposition to believe is the result of non-inferential processes involving, say, perception, introspection, memory, logical intuition, (*some*) testimony, etc.

So, for example, when one forms a perceptual belief, one doesn't infer it from other reasons or premises in an argument. Rather, it's a direct, non-inferential process.

Reliabilists, like Greco, likens our cognitive abilities to a model of information processing: "we have powers of information uptake (perception), information retrieval (memory), and information processing (reasoning)."[353] But perception does not uptake propositions; rather, it uptakes "properties of appearances, such as intensities and contrasts, and operates on them so as to produce object-level representations." [354] Therefore, perception doesn't resemble an executive function of inferential processing of propositions. What's more, much of the perceptual processing going on happens at the sub-personal level. So, the information getting processed during the uptake of information is not of beliefs or propositions like it is with reasoning processes. Indeed, the uptake of perceptual information is not of beliefs about, say, what wavelengths are impinging upon the eye's retina.

This is important because it's false to assume, like many skeptics do, that all knowledge-producing processes fit the pattern of reasoning processes. Instances of knowing can be the result of immediate, direct processes that don't involve reasoning patterns manifested in inductive, deductive, or abductive arguments. Indeed, many instances of knowing through perceptual processes are without any reasoning processes at all because reasoning involves making inferences from prior beliefs that function as premises in an argument. This is why knowledge-producing processes involving perception, introspection, memory, logical intuition, (*some*) testimony, etc. are not a kind of reasoning.

As important as this distinction is, however, it's not going to be enough to get someone like Riley in Case B off the hook. For SAFES doesn't assume *all* instances of

---

[353] John Greco, "Skepticism About the External World," John Greco, ed., *The Oxford Handbook of Skepticism* (Oxford: Oxford University Press, 2008), p. 122.

[354] Ibid., p. 123.

knowledge must fit the pattern of reasoning processes. It doesn't assume that someone like Riley doesn't know that she has hands prior to encountering the counterevidence from the argument for idealism.[355] What SAFES says is that when Riley satisfies conditions (a)-(d), she doesn't know she has hands because her belief has been defeated by this *live* counterevidence. The argument from an epistemic superior or the testimony from an epistemic superior has provided the kind of counterevidence that levels a defeater against her belief.

Notice how SAFES seems to be the kind of argument that explains the intuition for why we intuitively take the clairvoyant as not knowing the president's whereabouts when faced with such strong counterevidence. Recall that the clairvoyant encounters strong counterevidence from multiple news anchors functioning as epistemic superiors in that they are providing *live* footage of the president walking around the White House. This conflicts with the clairvoyant's true belief (that derives from a reliable clairvoyant faculty) that the president is in New York City. And because the clairvoyant can't rule out such strong counterevidence, we intuitively take the clairvoyant's reliably formed true belief to fall short of knowledge. Similarly, Riley—the fledging philosopher that she is—can't rule out the argument for idealism. Indeed, the argument appears to her sound! That is to say, the argument *appears* to be sound similar to how holograms appear as real objects. Both the argument and optical illusion are convincing, but one still knows the conclusion is false and the hologram isn't realm, respectively. But she doesn't even know enough about the theory of idealism to begin dismantling it brick by brick. So, it seems Riley doesn't know she has

[355] It doesn't matter if the counterevidence from an epistemic superior comes by way of testimony or an argument put forth by an epistemic superior. In Chapter Six, I interacted with both kinds of counterevidence. And although Case B includes an argument, one can, if one prefers, think of Case B as involving testimony or a combination of both testimony and argument since the professor tells Riley idealism is true and provides an argument for idealism.

hands once she (a) is aware that her professor is an epistemic superior; (b) is aware that her professor believes idealism is true on the basis of an argument she believes to be sound; (c) understands that (if idealism is true, then her hands are not external objects in the world); and (d) can't rule out the argument for idealism.

The trouble is that most of us share the intuition that Riley *does* know despite the counterevidence. *So where does SAFES go wrong?* After the work that's been done in Chapter Six, it's clear that one of the culprits in SAFES is clause (d) in premise (2). Admittedly, the locution "can't rule out" is a bit ambiguous. So let's disambiguate it. In one sense, it's true that Riley can't rule out the counterevidence. This is because she believes the argument for idealism is sound—that is, she believes the premises to be true and believes the form of the argument is valid. But it's false to assume that in order for Riley to know she has hands, she needs to prove which premise in the argument is false or prove the argument's form is invalid. Recall that in Chapter Six, I argued for a reliabilist's notion of defeating evidence understood in terms of an agent's cognitive character manifesting a contingent belief-forming disposition. This was to further build upon Gilbert Harman's point that logical relations alone can't function as evidential norms. Recall Greco's point that "[w]henever we have two logically incompatible propositions, there must be some *additional* factor that determines which we ought to accept and which we ought to reject. The same point holds with respect to other semantic relations: where *q* makes *p* improbable, this fact cannot by itself determine whether we ought to give up *q*, give up *p*, or continue to hold both."[356] The additional factor that determines what we ought to do is grounded in a belief-forming disposition that traces back to a defeater system that is constitutive of one's cognitive

[356] John Greco, *Achieving Knowledge*, p. 168 (emphasis mine).

character. This is where Riley and others get their defeater-defeater against the potential defeater leveled by SAFES.

So, with this in mind, there is *another* sense in which Riley *can* rule out idealism: namely, she knows she has hands because her belief-forming disposition to retain that belief is produced by a reliable defeater system! And since she knows she has hands, she knows there must be *something* wrong with the argument even if she can't put her finger on it.[357] And failing to know where the argument for idealism goes wrong doesn't prevent Riley from both knowing she has hands and that there is something wrong in the argument. For this is another false assumption pointed out long ago by the late William P. Alston. Alston argued that whenever one confuses the difference between (a) an agent satisfying the conditions for knowledge with (b) an agent knowing how it is those conditions guarantee knowledge, one is guilty of "levels confusion."[358] Therefore, premise (2) in SAFES is false because Riley *can* rule out idealism in virtue of knowing that she has hands.[359] At first blush, it may seem that Riley is merely digging her heels in. But her response sits in company with a rich tradition rooted in Thomas Reid and G.E. Moore. Otherwise known as the Moorean Shift, Riley knows that idealism is false because she

---

[357] I'm grateful for Martin Montminy for recommending that I stick to this solution rather than offering a number of others that would have made this chapter more cumbersome to read.

[358] Again, to be guilty of levels confusion is when one assumes that in order to know something they need to know how it is that they know it. That is, one assumes one needs evidence, arguments, and/or reasons in order to know how it is one knows. This confusion was diagnosed by William Alston and called just that--'levels confusion.' See William Alston, "Level Confusion in Epistemology," in *Midwest Studies in Philosophy* 5 (1980):135-150.

[359] One might think that we can rule out clause (d) in premise 1. But this move would require one to deny that knowledge is closed under entailment which is, for most people, too high of a price to pay. I am not willing to pay that price for reasons stated back in Chapter 4 when I first laid out the argument for SAFES. I'm grateful to Martin Montminy for raising this point to me.

knows she has hands. And she knows she has hands because her reliable defeater system produced a belief-forming disposition to retain her reliably formed perceptual belief.[360]

1.3 Epistemic Superiors Have (at least) *Weak* Understanding That Surpasses Their Epistemic Inferiors

At this point I want to provide an error theory that not only preserves the concept of an epistemic superior but explains how epistemic inferiors like Riley can know that p while at the same time an epistemic superior believes not-p. In Chapter Four, § 2.2, I listed some characteristics that explain why, say, Linda Zagzebski, is my epistemic superior. Some of those characteristics are as follows: (a) Zagzebski is more informed than I am on topics involving p; (b) she has more raw intelligence than me; (c) she has thought about and investigated whether p is true longer and more carefully than me; (d) she has thought about the topics surrounding p longer and in more depth than me; (e) she is more intellectually careful than me; (f) she is less biased than me; (g) she has a significantly greater publication and citation record—that is, she is recognized in the field as an expert given the number of peer reviewed publications she has written and the number of times her work has been cited. But this kind of asymmetry between an epistemic superior and an epistemic inferior may, at times, only be capturing how the former has a superior kind of *understanding* than me. It doesn't guarantee such superiors are more reliable at getting the truth, esp., in domains wherein disagreement abounds. For we lack a noncircular proof to demonstrate Zagzebski's philosophical methods are more reliable than some other experts in the field. So, to say that A is an epistemic superior to B is still consistent with saying the epistemic

[360] Recall back in the Introduction of this dissertation that I'm assuming a certain methodology. This methodology assumed the following three principles which trace back to Reid and Moore and are endorsed by Greco. M1: One should not try to prove what is not known by proof. M2: Rather than trying to prove that external things exist, or that we know that external things exist, one should take a close look at the skeptic's reasons for saying that we do not know. M3: Common sense has defeasible authority over philosophical theory.

inferior can still *know* that p, despite lacking the epistemic superior's understanding. In other words, Zagzebski may understand *why* or *how* not-p is the case in a way that doesn't entail the falsity of p; whereas I may only know *that* p is the case.

But if Zagzebski and I disagree about the truth status of some targeted proposition, p, then how can I still know that p if Zagzebski understands why not-p is the case? In other words, doesn't understanding *why* or *how* p is not the case entail knowing that p is false? In what follows I apply Greco's account of understanding in order to explain how there is logical space for cases where an epistemic superior can understand how p is false yet a disputant can still know that p, despite being aware of the extant disagreement with an epistemic superior. Let's revisit, then, these different kinds of understanding.

(a) ***Real Objects***: S understands the "real" dependency relations [and their *relata*] in the world comprising a system that X is located within (e.g., ecosystems, economies, a machine, or historical event).
(b) ***Representations***: S understands a representation [replete with its own dependency relations and *relata*] of a real system in the world (e.g., a theory, narrative, model, or a set of equations).
(c) ***Relation Between Objects and Representation***: S understands the relations between a real system and a representation (e.g., relations between a diagram and a machine that it represents or the relations between a narrative of a historical event and the actual event it represents).

Given these different kinds of understanding, the explanation falls out of it straightforwardly.[361] For we see how one can understand something in three different ways. This all depends on what objects of knowledge function as the *relata* in one's systematic knowledge of dependency relationships. For one can know the real objects within different dependency relationships, or one can know the representations of such objects, or both.

[361] N.B. this response doesn't commit one to the controversial claim that knowledge and understanding are irreducibly distinct epistemic goods. Although Greco takes the two to be distinct, this response doesn't assume one has to hold this view. So although I use the language of knowledge and understanding, one can interpret these as synonyms that are being used interchangeably.

Let's see how this works by reconsider Julie and Jackie from Case D in Chapter Six. Jackie, recall, is an epistemic superior with respect to the discipline of philosophy of mind and cognitive science. She satisfies the criteria of an epistemic superior laid out in Chapter Four. For example, Jackie has an impressive publication record and her work is often cited in many peer reviewed journals. She also has a facile understanding of the history and development of philosophy of mind and cognitive science. Jackie can explain the nuts and bolts of all the main competing theories and the subtle differences between their representatives. Furthermore, she understands a lot about mental states—specifically, the different ontological theories of mental states (in general) and phenomenal mental states (in particular). In short, Jackie has spent more time and investigated the topics surrounding the case at hand than Julie.

The disagreement between Jackie and Julie is whether or not Julie is in pain. More specifically, Julie introspects and experiences excruciating pain. When Julie is pressed by Jackie to provide an account of how she knows, Julie replies by saying she feels pain. Julie simply introspects and knows that she is in pain because she is experiencing a sensation of hurtfulness.[362]

---

[362] We can, perhaps, describe the case in such a way so as to make Julie a little more sophisticated, but still Jackie's epistemic inferior. Perhaps Julie is a committed particularist and rejects methodism. This would bode well for Julie satisfying the subjective justification condition by providing more details to what she manifests when she is sufficiently motivated to the truth. More specifically, as a particularist, Julie believes Particularism (Chisholm, 1973) is true. This is the view that one can know p without knowing how one knows p or why the provenance of p is reliable. Particularism is an epistemological project that is consistent with Plato's approach insofar as being motivated by answering the questions 'What do we know?' and 'How do we know" in that order. Answers to such questions provide the kind of content needed to answer Plato's question "What is knowledge" since it provides information about how knowledge is possible and explains the difference between cases wherein we intuitive judge that one knows versus cases wherein one intuitively doesn't know, given the details of the case. It's an approach to epistemology that contrasts with methodism, which first aims to answer the question 'How do we know?" in order to answer the ancillary question "What do we know?" Notice how methodism is more at home with Pyrrhonian skepticism than Plato since it first requires some condition(s) to be met before knowledge can be acquired. In order to acquire knowledge one must first meet the conditions spelled out in the method or principle prescribed for getting knowledge.

But particularism takes our pre-theoretical judgments of cases to be intuitively correct insofar as, say, some subject, S, acquired knowledge or not. One's intuitive judgments, then, provide the particularist with a set of datum. The merits of the epistemological principles competing for our endorsement are

Jackie's belief in theory *t* prevents her from attributing to Julie knowledge that she's in pain. Jackie takes phenomenal mental states to be nothing over and above mere natural properties manifested as constellations of neurons firing in the brain. So if Jackie can't detect the right constellation of neurons firing in the brain, then Julie's report is somehow false. In short, Jackie is a reductive physicalist regarding *all* mental states. Indeed, Jackie understands why Julie is really not in pain because she believes she is having an illusion of pain. In other words, Julie's fist-person introspective method of determining what are her mental states lack the reliability of Jackie's third-person, publically accessible method involving the outputs of her neurological equipment. Because her equipment is attached to Julie's introspective-mechanism located in Julie's brain, she can observe all of its outputs on a monitor. And because she doesn't see the output for pain, she infers Julie is not in pain. What's more, Jackie's reductionist theory of phenomenal mental states requires her to deny that Julie's method of introspection might be a more reliable method at getting truth. So, Jackie denies Julie's testimonial reports are about any real objective neuronal properties in the brain, and thus, doesn't believe Julie is in pain.

---

evaluated on the basis of whether or not they accommodate our pre-theoretical judgments, contra methodism. So, if there is tension between some epistemic principle and our pre-theoretical judgment of a case, a particularist is going to put a primacy on the judgment over the principle. E.g., if some epistemological principle affords someone knowledge in a Gettier case, then our verdict that such a person has been Gettiered—thus, lacks knowledge—explains why the candidate principle is inadequate.

Methodists, like David Hume, start by committing themselves to some epistemological or philosophical principle about the nature of knowledge. This method of starting first with a principle believed to answer the question "How do we know?" informs one's judgment as to whether or not S has knowledge when evaluating cases. So, if there is tension between some endorsed epistemic principle and our intuitive judgment on a case being evaluated, the Methodist thinks the principle ought to trump the weight we attribute to our intuitive judgment.

There is a third alternative defended by Gilbert Harman (2003, 2004) deemed 'reflective equilibrium.' Reflective equilibrium is conservative like particularism insofar as giving priority to our intuitive judgments over principles. But it affords greater weight to principles than particularists are willing to give. Moreover, it aims to achieve a level of coherency between the weight afforded to both principles and judgments. I'm indebted to Thomas Kelly in this footnote for his work on how particularism and methodism differ and making me aware of Harman's third alternative. See Thomas Kelly, "Moorean Facts And Belief Revision, Or Can The Skeptic Win?," *Philosophical Perspectives* 19. 1 (2005):179-209.

With the different kinds of understanding on offer, it's possible that Julie knows that she's in pain, while Jackie still retains the title of being Julie's epistemic superior—despite the fact Jackie's theory is false. Furthermore, Jackie can retain a kind of superior (*albeit* weak) understanding that doesn't entail the falsity of Julie's belief. That is, Jackie's belief that Julie is not in pain can be false at the same time Jackie can still be an epistemic superior relative to Julie. Here's how.

Jackie's reductionist model of mental states and ontological commitments may enable her to understand a *representation* of a real phenomenal mental state in the world. But her belief in the representation is false in the sense that it doesn't map onto the real world. Indeed, Jackie's metaphysical view of the world is leaving stuff out. Of course, Jackie may understand the different kinds of dependency relations between the representation she believes in and an inaccurate system in the world. But because her view of the world lacks the necessary *relatum* of non-reductive, irreducible phenomenal mental states, the relations between the *real* objects and her *representation* are inaccurate. Therefore, Jackie has *weak* understanding of type (b)—Representations—that is, she understands her own representation of the real system in the world, replete with its own dependency relations. She also, perhaps, may have a *stronger* kind of understanding than the weak sense because she also understands the "alleged" relationship between *irreducible, phenomenal mental properties* in the world and truths about the relations comprising a *non-reductionist representation* of the world.

Furthermore, she may understand the kinds of dependency relations between alternative models involving irreducible mental properties and the natural properties they supervene upon. Indeed, she may understand all the different *representations* used by all the leading non-reductionists, including physicalists, property-dualists, and substance dualists

alike. That is, Jackie understands all the competing non-reductive representations and the kinds of dependency relations those representations have with different competing models of the world. After all, she's an epistemic superior when it comes to the philosophy of mind and cognitive science! But because she doesn't believe in non-reductive mental states and the kind of dependency relations that support them, she doesn't know that Julie is not in pain despite having a superior kind of understanding over Julie. But she doesn't have *strong* understanding. Why? Answer: Jackie's theory is false. For if Jackie had strong understanding, then Julie couldn't know she was in pain because no such mental states can reliably be accessed through introspection. Without Jackie believing in irreducible mental states, she lacks knowledge of a key piece of datum non-reductionists use to inform their understanding of the relation between their representation and the real mental states in the world.

Notice, however, that in order for Julie to know she's in pain, she need only experience pain through an act of introspection. And this method of knowing doesn't require Julie to have any of the weak, stronger, or strong kinds of understanding listed above. She is much like the other students in Cases A, B, and C from Chapter Six. There Jamie, Riley, and Peyton are epistemically inferior to their professors who, let's suppose, are epistemic superiors much like Jackie. They lack the kind of weak and stronger senses of understanding their professors have. But because their professors deny motion, external objects, and non-deterministic accounts of freedom, respectively, they lack strong understanding if such things exist. And despite Jamie, Riley, Peyton, and Julie lacking any deep understanding of motion, external objects, free will, and pain, they can still know about such things through non-inferential, belief-forming processes. Accordingly, they can

all meet the conditions for knowledge, despite the recognized disagreement with an epistemic superior.

They also satisfy the subjective-justification condition for knowledge. For their beliefs are epistemically responsible because their beliefs result from intellectual dispositions they manifest because they are sufficiently motivated to get the truth. And the targeted dispositions are epistemically correct ones because they are brought about by a reliable defeater system that is well-integrated and stable, and thus, constitutive of their cognitive agency. Recall back to Chapter Six where I stated this more formally:

Principle of Epistemically Correct Dispositions (PECD): A belief forming disposition that is a product of a defeater system is an epistemically correct one only if the defeater system is well integrated[363] with other beliefs and counterevidence within the cognitive system and is functioning reliably as a result of a sufficient motivation to get the truth.

So we've carved out enough logical space to make sense out of how one can retain the title of epistemic superior. The different kinds of understanding on offer allow an inferior disputant to still know the disputed belief despite lacking such understanding which doesn't entail the falsity of the inferior's belief. Because weak understanding doesn't entail knowledge, the epistemic superior can have the former without the latter; and the inferior disputant can have the latter without the former.

2. A CHALLENGE TO THE GRECOIAN EXPLANATION FOR HOW KNOWLEDGE IS POSSIBLE IN THE FACE OF EPISTEMIC SUPERIOR DISAGREEMENT

At this point in the chapter, I have argued for how one can still have knowledge in the face of epistemic superior disagreement by denying clause (d) in premise (2). There are instances

[363] Recall that in Chapter Six I offered two kinds of integration that demarcates epistemically incorrect dispositions from epistemically correct ones. Integration$_1$ = A disposition to retain a belief, give up a belief, or reduce confidence in a belief is well integrated only if the defeater system from whence it came produces enough outputs so as to establish a relation with and sensitivity to other occurrent and standing beliefs within the cognitive system. And Integration$_2$ = A disposition to retain a belief, give up a belief, or reduce confidence in a belief is well integrated only if the defeater system from whence the disposition came is sensitive to evidence for or against the targeted belief.

of knowing wherein one can rule out counterevidence as a result of a reliable defeater system producing a belief-forming disposition to retain one's belief. The skeptical threat leveled by SAFES has been shown to be unsound as a result of falsifying premise (2). Then I preserved a notion of epistemic superiority by providing an error theory. This error theory allows us to say an epistemic superior can retain such a title once we see how they have superior kinds of understanding over their epistemic inferiors, even when they lack knowledge had by their inferiors.

Be that as it may, one may still feel unsatisfied with my response to SAFES. For one may feel I've made knowledge into something unfamiliar, alien, or foreign.[364] A common *desideratum* for an account of knowledge is that our acquisition of knowledge is something we are *intimately* related to. After all, acquiring knowledge is not like something we get from the stork—it doesn't just drop onto our laps. And if a true belief did fall on our laps, we'd call it something other than knowledge. Similarly, being the recipient of a belief-forming disposition produced by a defeater system may strike one as involving a process that is alien to our cognitive agency. That it, too, seemingly just falls unto our laps.

One way to satisfy this *desideratum* was put forth by Pyrrhonian skeptics. One became intimately connected to knowledge through having a *perspective* on *how* it is one acquires truth reliably. Indeed, Greco and Sosa have traced this desire back to Sextus Empiricus. Here Sextus illustrates how counterintuitive it is to think that the acquisition of knowledge is the result of actions that are too foreign to be considered a part of our cognitive agency.

> Let us imagine that some people are looking for gold in a dark room full of treasures…[N]one of them will be persuaded that he has hit upon gold even if he has in fact hit upon it. It the same way, the crowd of philosophers has come into the world, as into a vast house, in search of

[364] I'm indebted to Wayne Riggs for helping me to see this concern more clearly.

> truth. But it is reasonable that the man who grasps the truth should doubt whether he has been successful.[365]

Notice how this process of acquiring gold instead of granite doesn't seem to be a process *we* are intimately connected to. Thus, it doesn't seem to be an achievement that is appropriately *ours*. Similarly, if we acquire knowledge the way the person in the dark acquires gold, then getting knowledge is something foreign to us; it's something that happens to us the way, say, gifts arrive in mail: our relation to them is extrinsic and not the result of our agency.[366] The upshot is that we don't think we deserve credit for such things any more than we deserve credit for receiving a gift or luckily finding gold.

It's clear how to make grasping gold something we *are* intimately related to: namely, acquiring gold in a well-lit room so we can *discriminate* gold from granite. Similarly, if we had a perspective on how we discriminate truths from falsehoods, then we can understand how knowledge *is* intimately related to us. Getting the truth wouldn't have to be like grasping gold in the dark. Indeed, it can intuitively be something we take ourselves to be responsible for getting. But as Chapter 2 took pains to explain: we cannot have that kind of perspective on how it is our true beliefs derive from reliable processes. For what we have internal access to through reflection alone are the same appearances shared by our internal

---

[365] Sextus Empiricus, Against the Mathematicians, VII, p. 259. Quoted by John Greco in *Achieving Knowledge* (Cambridge: Cambridge University Press, 2010), pp. 126-127.

[366] Goldberg has raised a very similar concern relative to disagreement when he says this, and I quote at length: "[W]hat reasons could [disputant] S1 offer in defense of the claim that [her disposition to retain belief] is not unreliable on the matter at hand? This is not a one-off case, where S1 might appeal to her long (and independently confirmable) track record on related matters. There is no independent access to the truth of the matter [i.e., there is no perspective on the reliability of one's defeater system], so S1, like [disputant] S2, has to go on whatever [disposition her defeater systems kicks out] in order to reach a judgment on the disputed matter. In light of this, consider the situation as it strikes (or should strike) S1. Given that this is a case of [disagreement with an epistemic superior] (and that S1 […] recognizes this), S1 (like S2) should conclude that chances are good that at least some party to this dispute is unreliable. In addition, S1 (like S2) should also appreciate that […] the unreliable party (whoever it is) is not in a position to discern her or his own [reliable defeater system]; otherwise, she or he would have done so, and the disagreement would have dissipated. […] [A]nd if it is in the nature of the unreliability [of one's defeater system] that is not discernible […], then this should tamper S1's confidence in her own assessment that it is not she (S1) who is unreliable." See Sandford Goldberg, "Disagreement, Defeat, and Assertion," David Christensen and Jennifer Lackey, eds., *The Epistemology of Disagreement* (Oxford: Oxford University Press, 2013), p. 179.

duplicate that is in an epistemically bad environment. Furthermore, it's inappropriate to demand a perspective on how the part-causes that contribute to getting the truth are reliable because even an omniscient being would not be able to get outside her own cognitive system to "see" how it delivers such reliably true beliefs. This is why the vindication approach is doomed to fail.

So even though having a perspective on how our cognitive faculties get the truth is a *desideratum* any *available* account of knowledge can't satisfy, I think there is a related concern that the theory elucidated in this dissertation can address, which is this: *having a true belief-forming disposition to, say, retain a belief—as a result of one's defeater system—is not something intimately related to our cognitive agency.*

In response to this concern, what I intend to offer is this. Even though having a perspective and vindication of our cognitive faculties' reliability is beyond our ken, there is still a sense in which knowledge is deeply related to our cognitive agency. Indeed, this proposal offers an explanation for how knowledge can be something we *are* intrinsically related to. Because if knowledge is a success from ability, then once we understand what ability is, we might be reassured that knowledge is not so far removed from us. Indeed, knowledge doesn't have to be akin to luckily grasping gold in a dark room. For the epistemic inferior who retains a true belief as a result of her defeater system's output can have knowledge that is intrinsically related to her cognitive agency.

2.1 What's An Ability?

In order to show how knowledge is intimately related to us, we must understand the relationship between knowledge and ability. For ability builds the bridge between cognitive agency and credit-worthy achievements.

To understand what makes an action a credit-worthy achievement, we revisit Greco's account of ability since for an achievement to be credit-worthy, it must derive from one's abilities and not as a result of luck. In Chapter Five, § 2.3, I unpacked Greco's notion of ability. Recall that Greco takes a modal view of abilities. That is, abilities are dispositional properties constitutive of an agent's *character* that can successfully achieve certain ends across a range of close possible worlds. More specifically, Greco wants to provide an account of ability that is insulated from both good and bad luck, which can undermine a credit-worthy achievement. To attribute one with a credit-worthy achievement, it must be insulated from credit-undermining luck.

For example, on the one hand, even if one is always successful at hitting baseballs, it doesn't follow one has the ability to hit baseballs *per se* given the success can be attributed to mindboggling good luck. There are cases where, miraculously, every time you happen to swing at a baseball—while closing your eyes!—you end up hitting the ball. But your achievement is the result of an inordinate string of good luck and not ability. On the other hand, one may have the ability to achieve hitting baseballs but, because of a lifetime string of bad luck, is never successful at manifesting the ability. Sicknesses, injuries, incessant episodes of severe weather, you name it—bad luck can prevent an ability from manifesting itself. So, in order to extricate the targeted ability from luck's tentacles, Greco thinks ability must entail modal success, though not necessarily actual success. The formal definition for ability, according to Greco, is this.

Ability: S has an ability to achieve result R in conditions C if and only if, across the range of close possible worlds where S is in C, S achieves R in C with a high rate of success.[367]

So S has the ability to, say, reliably believe the truth regarding subject matter M if and only if across the range of close possible worlds where S is in conditions C, S gets the truth

[367] Greco, *Putting Skeptic in Their Place*, p.212.

about M with a high rate of success. This notion of an ability helps inform us of the kind of luck Greco believes can undermine achievements: *namely, an action that is not sufficiently brought about by one's agency*. That is, one's ability to think, choose, act, etc. comprises the characteristics of one's agency. And that which takes away one's ability to think, choose, act, etc. is the kind of luck that can undermine a credit-worthy achievement. So, S doesn't deserve credit for believing the truth when S's agency isn't a salient feature for why S got the truth. Alternatively: S deserves credit for believing the truth *because* S's agency (in general) and abilities (in particular) are salient features for why S got the truth. S got the truth *because* of S's abilities.

2.2 Agency Undermining Luck

Luck can take away one's agency by affecting an agent's abilities. A manipulative scientist controlling my actions through a remote control that sends signals to an implanted chip in my brain can rob me of my agency. If I pull the trigger as a result of the manipulative scientist causing a neuronal event in my brain to cause my finger muscles to contract, I'm not responsible for firing the gun. The scientist is responsible for firing the gun and I am just a link in the intermediary causal chain. My agency has been compromised because my basic actions that are under my direct control have been commandeered by the scientist. The scientist took away my agency! So responsibility for bringing about an action is intrinsically related to matters pertaining to one's agency. This is why we endorse Kant's maxim that one *ought* to do A only if one *can* do A. Saying I ought not have fired the gun doesn't apply to me because that assumes I can prevent my finger from pulling the trigger. But I can't! My agency was commandeered by the vicious scientist.

Or consider how Jeter's agency can be unmined by an agency-impairing drug. If Jeter were inebriated to the point he could barely stand and, yet, still "achieved" hitting a

95 mph fastball once out of twenty tries, it would be the result of luck. The reason why is because his agency has been compromised and the success is because of luck and not ability. He lacks the ability to successfully hit fastballs with a high rate of success in that condition. Note that Jeter's ability to hit fastballs stems from all his habits, dispositions, and abilities that are not *all* under his direct, conscious control. He has no perspective on all the goings-on taking place in both his nervous, muscular, and skeletal systems. Indeed, Jeter may be like many athletes who are unable to even describe in detail how it is they are able to achieve their respective feat because the movement is more of an indirect response than a basic, direct action typical of actions that are under our direct, conscious control. (This is why not all athletes make good coaches and why good coaches may be terrible athletes. The ability to (a) perform well athletically, which doesn't require a conscious perspective on how it is one performs well, and the ability to (b) coach well, which does require a conscious perspective on how to perform well so it can be articulately communicated to fledgling athletes in an effort to teach them, can clearly come apart.)

For example, if Jeter were to watch all the different movements his body makes while hitting a ball, he would be unable to explain why the mechanics of his swing generates the speed and power it does. He might not be able to explain why his hips did this or why his shoulders do that by articulating all the part-causes, which includes certain beliefs and intentions, that brought about the achievement of hitting a fastball. The reason why is because there are so many indirect movements comprising the act of "hitting a fast ball" that are not under his direct, conscious control nor are they all internally, consciously accessible in a way he could describe them individually in chronological order. Be that as it may, the ability to hit fastballs is still constitutive of and caused by his *agency*, which is comprised by different habits, dispositions, and other abilities he's responsible for because

they are appropriately his. That is, because his success is the result of his ability, Jeter deserves credit for his achievement. And this is important because Jeter's credit-worthy achievement was brought about in the right sort of way: namely, an ability that is intrinsically a part of his agency.[368]

So far we've learned how luck can directly compromise one's agency. My nervous system was compromised by the nefarious scientist and Jeter's agency was compromised by drugs. But luck can also compromise one's agency *indirectly* as a result of the conditions in the agent's environment shifting from normal ones to unlucky ones. For Greco's account of ability, recall, is intrinsically related to an environment. So, if we were transported to a world where there was no light, we would no longer have the ability to see given our ability to do so is intrinsically dependent upon an environment that has light waves. Without light waves impinging upon our retina, we will not have the ability to see. Because abilities are intrinsically dependent upon an environment, success is going to be dependent upon the environment cooperating with them. Greco, then, takes this as a general truth about ability and success: "Whether intellectual, athletic, moral or otherwise, success from ability is never the result of ability alone, but always the result of ability and enabling circumstances."[369] Jeter can't hit fastballs in the middle of a hail storm nor can we manually shift gears in a car that has an automatic transmission. Our abilities can manifest themselves only if our environment is cooperating given the intrinsic dependency relation between the two.[370]

---

[368] I'm grateful for Wayne Riggs's comments which helped sharpen this point.

[369] Greco, *Achieving Knowledge*, p. 140.

[370] The importance of understanding an ability's relation to its environment is important because it helps us better *refine* our credit-worthy attributions and their relation to credit-undermining luck. For example, take an athlete who plays the sport of golf. This athlete has the ability to always shoot in the 60's at her club's course, as long as the environmental conditions are fairly normal. That is, as long as the wind isn't consistently

So, it seems like luck can *indirectly* affect one's agency for good or for ill depending on the environmental conditions. On the one hand, good luck can enable one's agency to manifest an ability as a result of the environment cooperating and this doesn't undermine one's credit-worthy achievement. On the other hand, bad luck can compromise one's agency by not allowing an ability to manifest itself as a result of the environment not cooperating. This kind of luck doesn't undermine a credit-worthy achievement because there's no act to be evaluated as an achievement. Here it is important to note that not all luck undermines one's credit-worthy action. In other words, not all luck undermines a credit-worthy achievement that one is responsible for because the kind of luck involved—namely, a cooperating environment—is what *enables* an ability to manifest itself; it's what makes the manifestation of the ability possible.[371] It's only the kind of luck that

---

blowing 35 mph or faster and there are no destructive storms over the course and, of course, the golf course is in fair, playable conditions, this person will shoot below par. Now one might think that it's a result of bad luck that whenever this person enters into, say, the Adams Tour Golf Tournament held at her club's course she shoots in the 80's. This score prevents her from making the cut, which allows her to continue playing the last two days of the four-day tournament. (If she shot like she normally does in the 60's, she would make the cut and win the tournament every time.) So, if we constrain her ability to shoot in the 60's to a normal environment at her club's course, it may seem like her performance at the professional tournament was a result of bad luck given she has such a high rate of success of shooting in the 60's. Indeed, she has shot in the 60's every time she's played there for the past five years, totaling 1,300 rounds of consistently breaking 70! But once we broaden the environmental component as it relates to her abilities beyond weather and course conditions to include a *competitive atmosphere*, it seems this person lacks the ability to shoot in the 60's in a competitive—"game-time!"—environment. Indeed, were she to fail at shooting in the 60's in the next 10 tournaments she entered in at her golf club, it seems her ability to play golf well—that is, shooting-in-the-60's—is restricted to non-competitive environments where the atmosphere is more akin to a practice environment. Practice environments provide a low-stakes atmosphere where one can perform without worrying about winning, competing, and experiencing the risks involved in losing. There are no psychological pressures equivalent to the kinds of pressures operative in competitive environments like professional tournaments. So, it's not unlucky this golfer continues to fail at shooting in the 60's or making the cut whenever she plays in a professional tournament. Rather, she lacks the ability to shoot in the 60's in an environment that includes having to compete against other golfers in a professional competition. She simply lacks the ability to score in the 60's in competitive environments, despite having the ability to do so in practice, non-competitive environments. Not succumbing to nervousness, anxieties, and other fears of failing that inhibit one's ability to perform well is sometimes the difference between two equally performing athletes in a practice environment and the one that can perform equally well in a competitive environment. Hence, some athletes are noted for having "ice running through their veins" because they are so calm, cool, and collected while competing—not allowing the pressures of competition to affect their psychological performance. Those gifted athletes that lack *this* kind of ability is why some never become professional athletes.

[371] I'm grateful to John Greco for pointing this out to me. See John Greco, *Achieving Knowledge*, p. 139-142.

compromises one's agency—that is, *takes one's agency away or distorts it*—that is inconsistent with being responsible for a credit-worthy achievement because of ability. Therefore, some kinds of luck enable one's agency—whether it be broad or narrow—to manifest a credit-worthy achievement. Therefore, some kinds of luck don't undermine responsibility. Luck and responsibility are not always incompatible. It's only the kind of luck that distorts one's agency by affecting one's character that undermines a credit-worthy achievement.

2.3 What Makes An Ability A Part of One's Cognitive Character?

What makes an ability—specifically a cognitive ability—genuinely a part of one's *cognitive agency*? We answered this question, recall, back in Chapter Six. There I said Greco identified at least two conditions that need to be satisfied in order for a cognitive ability to be a part of one's cognitive agency (or cognitive character): (a) the ability must be *stable* in the relevant sense and (b) the ability must be *well integrated* with other cognitive abilities, modules, and dispositions. A lack of stability and integration are what makes mere *de facto* reliable processes seem unfamiliar, foreign, or alien to our agency and counterintuitive to genuine instances of knowledge. Unless the ability is grounded in an agent's *cognitive agency*, the quasi-ability is too foreign to the agent's overall *cognitive character*. Accordingly, the achievement needs to be the result of an ability that is *appropriately* a part of one's cognitive character in order to receive a credit attribution. Because, according to Greco, "to the degree that the odd cognitive power is well integrated with […] the believer's cognitive dispositions, there is a tendency to judge that the person has knowledge. To the extent that the odd power in not well integrated, there is a tendency to judge that the person does not have knowledge."[372]

[372] John Greco, *Achieving Knowledge*, p. 153.

Once we understand that in order for an ability to be a part of one's cognitive agency it must be both modally stable and well-integrated, then the case for how one's defeater system can also be appropriately a part of one's cognitive agency falls out straightforwardly. If the defeater system is neither modally stable nor well-integrated, then an agent doesn't deserve credit for being disposed to retain a true belief vis-à-vis disagreement with an epistemic superior. So the particular cases we are evaluating involving an agent's defeater system disposing one to retain a belief vis-à-vis epistemic superior disagreement must not be underdescribed. Once the details of the case include (i) how well the defeater system cooperates with other modules, evidence, counterevidence, and sub-systems thereby establishing *integration* and (ii) how frequent its outputs are thereby establishing the kind of track record needed for *stability*, we can see how one's defeater system is a part of one's cognitive agency. And the kind of achievements resulting from one's defeater system are thus genuine instances of knowledge because the defeater system meets the conditions required for being a part of the agent's agency. Therefore, because the disposition to retain a true belief is not a result of knowledge-undermining luck given the defeater system is a part of the subject's cognitive agency, knowledge is inextricably bound to agency and our initial concern is significantly allayed.

3. Summary

In this chapter I provided a Grecoian response to the skeptical argument from epistemic superior disagreement (SAFES). § 1 of this chapter offered reasons for thinking clause (a) in premise (2) is false. In § 1.2, I argued that Riley can rule out the counterevidence because she knows she has hands. How? Answer: her belief-forming disposition to retain that belief is produced by a reliable defeater system. And since she knows she has hands, she knows there must be *something* wrong with the argument even if she can't put her finger on it.

§ 1.3 of this chapter preserved a notion of 'epistemic superior' by appealing to the different kinds of understanding such a person may have. The upshot of this section is that an inferior disputant can retain knowledge in a case of disagreement with a *bona fide* epistemic superior who believes otherwise because the latter has a kind of understanding that doesn't entail knowledge. In other words, the inferior disputant can still know p despite the cognitive superior having a kind of understanding of the matter of p, yet fail to know that p is false.

§ 2 of this chapter presented a challenge to my Grecoian response to SAFES. More specifically, this section leveled a concern that my account of knowledge involves processes that are too foreign to our agency, which prevents us from receiving the credit we take instances of knowledge to entail. More specifically, a disputant who has a disposition to retain a true belief in the face of epistemic superior disagreement involves a defeater system that seems too disconnected from our agency. And any true belief-forming disposition produced by a defeater system seems to be the result of luck because it undermines agency. Such luck, then, undermines the disputant's ability to know despite having the disposition to retain a true belief. But after discussing what ability is, I argued a defeater system meets the conditions for being a part of one's cognitive agency. Knowledge, it turns out, is intimately related to agency. And the concern about defeater systems delivering us true dispositions is not so foreign to our agency after all.

BIBLIOGRAPHY

Adler, Jonathan. 1994. "Testimony, Trust, Knowing." *Journal of Philosophy* 91 (5): 264-275.

Alston, William. 1980. "Level Confusions in Epistemology." *Midwest Studies in Philosophy* 5 (1): 135-150.

_____. 1991. *Perceiving God*. Ithaca: Cornell University Press.

_____ . 1988. "The Deontological Conception of Epistemic Justification." *Philosophical Perspectives* 2: 257-299. repr. in Ram Neta and Duncan Pritchard, eds., *Arguing About Knowledge*. New York: Routledge, 2009, 324-350.

_____ . 2005. *Beyond Justification: Dimensions of Epistemic Evaluation*. Ithica, NY: Cornell University Press.

Audi, Robert. 1998. *A Contemporary Introduction to the Theory of Knowledge*. New York: Routledge.

_____ . 2001a. "Doxastic Voluntarism and the Ethics of Belief," in Matthias Steup, ed., *Truth, Duty, and Duty*. Oxford: Oxford University Press.

_____ . 2001b. "An Internalist Theory of Normative Grounds." *Philosophical Topics* 29 (1/2): 19-46.

_____ . 2006. "Testimony, Credulity, and Veracity," in Jennifer Lackey and Ernest Sosa, eds., *The Epistemology of Testimony*. Oxford: Oxford University Press.

_____ . 2011. *Rationality and Religious Disagreement*. Oxford: Clarendon Press.

Axtell, Guy. 1997. "Recent Work in Virtue Epistemology." *American Philosophical Quarterly* 34: 1-27.

Baehr, Jason. 2006. "Character, Reliability, And Virtue Epistemology." *The Philosophical Quarterly* 56 (223): 193-212.

Battaly, Heather. 2008. "Virtue Epistemology." *Philosophy Compass* 3 (4): 639-663.

Baumann, Peter. 2009. "Epistemic Closure," in Sven Bernecker and Duncan Pritchard, eds., *The Routledge Companion to Epistemology*. New York: Routledge.

Bergmann, Michael. 2009. "Rational Disagreement After Full Disclosure." *Episteme: A Journal of Social Epistemology* 6 (3): 336-353.

_____ . 2006. *Justification without Awareness*. New York: Oxford University Press.

_____ . 1997. "Internalism, Externalism, and the No-Defeater Condition." *Synthese* 110 (3): 399-417.

Bogardus, Thomas. 2009. "A Vindication of the Equal-Weight View." *Episteme* 6 (3): 324-335.

Bonjour, Laurence. 1980. "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy* 5(1): 53-73.

Burge, Tyler. 1993. "Content Preservation." *The Philosophical Review* 102 (4): 457-488.

Byrne, Alex. 2005. "Introspection," *Philosophical Topics* 33 (1): 79-104.

Chisholm, Roderick. 1966. "Freedom and Action," in Keith Lehrer, ed., *Freedom and Determinism*. New York: Random House.

_____ . 1977. *The Theory of Knowledge*, 2nd ed. Englewood Cliffs, NJ: Prentice Hall.

Christensen, David. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116 (2): 187-217.

_____ . 2009a. "Introduction: The Epistemology of Disagreement." *Episteme: A Journal of Social Epistemology* 6 (3): 231-232.

_____ . 2009b. "Disagreement as Evidence: The Epistemology of Controversy." *Philosophy Compass* 4 (5): 756-767.

Churchland, Patricia S., 1986. *Neurophilosophy: Toward A Unified Science of the Mind-Brain*. Cambridge, MA: MIT Press.

Churchland, Paul. 1989. *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*. Cambridge, MA: MIT Press.

Coady, C.A.J. 1992. *Testimony: A Philosophical Study*. Oxford: Clarendon Press.

Code, Lorraine. 1987. *Epistemic Responsibility*. Hanover, NH: University Press of New England.

Cohen, Stewart and Keith Lehrer. 1983. "Justification, Truth, and Coherence." *Synthese* 55 (2): 191-207.

_____ . 1988. "How to Be a Fallibilist," *Philosophical Perspectives* 2: 91-123.

_____ . 1998. "Two Kinds of Skeptical Argument." *Philosophy and Phenomenological Research* 58 (1): 143-159.

_____ . 2002. "Basic Knowledge and the Problem of Easy Knowledge." *Philosophy and Phenomenological Research* 65 (2): 309-329.

Conee, Earl and Richard Feldman. 2004. *Evidentialism: Essays in Epistemology*. New York: Oxford University Press.

_____ . 2010. "Rational Disagreement Defended," in Richard Feldman & Ted A. Warfield, eds., *Disagreement*. Oxford: Oxford University Press.

Craig, Edward. 1990. *Knowledge and the State of Nature*. Oxford: Oxford University Press.

Dennett, Daniel. 1978. *Brainstorms*. Cambridge, MA: MIT Press.

DePaul, Michael. 1993. Balance and Refinement" *Beyond Coherence Methods of Moral Inquiry*. New York: Routledge.

DeRose, Keith. 1995. "Solving the Skeptical Problem." *Philosophical Review* 104 (1): 1-52.

Descartes, René. 1979. *Philosophical Works*. Vol. 1. Trans. Elizabeth S. Haldane and G. R. T. Ross. Cambridge: Cambridge University Press. 1979.

David. M. and Warfield, T.A. 2008. "Knowledge–Closure and Scepticism," in Quintin Smith, ed., *Epistemology: New Essays*. Oxford: Oxford University Press.

Davidson, Donald. 1986. "A Coherence Theory of Truth and Knowledge," in Ernest LePore, ed., *Truth and Interpretation*. Oxford: Blackwell.

Doris, John. 1998. "Persons, Situations and Virtue Ethics." *Noûs*, 32 (4): 504-530.

_____ . 2002. *Lack of Character*. Cambridge: Cambridge University Press.

Dretske, Fred. 2003. "Skepticism: What Perception Teaches," in S. Luper, ed*., The Skeptics: Contemporary Essays*. Aldershot: Ashgate.

Elga, Adam. 2007. "Reflection and Disagreement." *Noûs* 41(3): 478-502.

_____ . 2010. "How To Disagree About How to Disagree," in Richard Feldman & Ted A. Warfield, eds., *Disagreement*. Oxford: Oxford University Press.

Elgin, Catherine Z. 2004. "True Enough." *Philosophical Issues* 14 (1): 113-131.

_____ . 2010. "Persistent Disagreement," in Richard Feldman & Ted A. Warfield, eds., *Disagreement*. Oxford: Oxford University Press.

Engel, Mylan Jr. 1992. "Is Epistemic Luck Compatible with Knowledge?" *Southern Journal of Philosophy* 30 (2): 59-75.

Feldman, Richard. 2005. "Respecting the Evidence." *Philosophical Perspectives* 19 (1): 95-119.

_____. 2006. "Epistemological Puzzles about Disagreement" in Stephen Heatherington, ed., *Epistemology Futures*. New York: Oxford University Press.

_____. 2007. "Reasonable Religious Disagreements," in Louis M. Antony, ed., *Philosophers Without Gods: Meditations on Atheism And The Secular Life*.

Feyerabend, Paul. 1963. "Materialism and the Mind-Body Problem." *Review of Metaphysics* 17 (September): 49-67.

Foley, Richard. 1987. *A Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.

Frances, Bryan. 2005. "When a Skeptical Hypothesis Becomes Live," *Noûs* 39 (4): 559-595.

_____ . 2008. "Live Skeptical Hypotheses," in John Greco, ed., *Oxford Handbook of Skepticism*.

_____ . 2010. "The Reflective Epistemic Renegade." *Philosophy and Phenomenological Research* 81 (2): 419-463.

Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.

Fumerton, Richard. 2010. "You Can't Trust A Philosopher," in Richard Feldman and Ted A. Warfield, eds., *Disagreement*. Oxford: Oxford University Press.

Gettier, Edmund. 1963. "Is Justified True Belied Knowledge?" *Analysis* 23 (6): 121-123.

Ginet, Carl. 1975. *Knowledge, Perception, and Memory*. New York: Springer.

Goldberg, Sanford. 2006. "Reductionism and the Distinctiveness of Testimonial Knowledge," in Jennifer Lackey and Ernest Sosa, eds., *The Epistemology of Testimony*. Oxford: Oxford University Press.

_____ . 2009. "Reliabilism in Philosophy." *Philosophical Studies* 142 (1): 105-117.

_____ . 2013a. "Defending Philosophy in the Face of Systematic Disagreement," in Diego E. Machuca, ed., *Disagreement and Skepticism*. New York: Routledge.

_____ . 2013b. "Disagreement, Defeat, and Assertion," David Christensen and Jennifer Lackey, eds., *The Epistemology of Disagreement*. Oxford: Oxford University Press.

Goldman, Alvin I. 1976. "Discrimination and Perceptual Knowledge." *Journal of Philosophy* 73 (November): 771-91

_____ . 1986. "What Is Justified Belief?" in Paul K. Moser, ed., *Empirical Knowledge*. Totowa, NJ: Rowman & Littlefield.

_____ . 1999. "Internalism Exposed." *Journal of Philosophy* 96 (6): 271-293.

_____ . 2006. *Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading*. Oxford: Oxford University Press.

_____ . 2010. "Epistemic Relativism and Reasonable Disagreement," in Richard Feldman & Ted A. Warfield, eds., *Disagreement*. Oxford: Oxford University Press.

_____ . 2011. "Toward a Synthesis of Reliabilism and Evidentialism? Or: Evidentialism's Troubles, Reliabilism's Rescue Package," in Trent Dougherty, ed., *Evidentialism and Its Discontents*. Oxford: Oxford University Press.

Graham, Peter. "Liberal Fundamentalism and Its Rivals," in Jennifer Lackey and Ernest Sosa, eds., *The Epistemology of Testimony*. Oxford: Oxford University Press.

Greco, John. 1995. "Reid's Critique of Berkeley and Hume: What's the Big Idea?" *Philosophy and Phenomenological Research* 55 (2): 279-296.

_____ . 2000. *Putting Skeptics In Their Place: The Nature of Skeptical Arguments And Their Role In Philosophical Inquiry*. Cambridge: Cambridge University Press.

_____ . 2002. "How to Reid Moore." *Philosophical Quarterly* 52 (209): 544-563

_____ . 2003. "Knowledge as Credit for True Belief," in Michael DePaul and Linda Zagzebski, eds., *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford: Oxford University Press.

_____ . 2005. "Is Justification Internal?: Justification is Not Internal," in

Matthias Steup and Ernest Sosa, eds., *Contemporary Debates in Epistemology*. Oxford: Blackwell.

_____ . 2007a. "Epistemology in the Twentieth Century," in Constantin Boundas, ed., *The Columbia Companion to Twentieth-Century Philosophies*. New York: Columbia University Press.

_____ . 2007b. "The Nature of Ability and the Purpose of Knowledge," in Ernest Sosa, ed., *The Metaphysics of* Epistemology. Oxford: Blackwell.

_____ . 2008a. "What's Wrong With Contextualism?" *The Philosophical Quarterly* 58 (232) 416-436.

_____ . 2008b. "Friendly Theism," in James Kraft, ed., *Religious Tolerance Through Epistemic Humility*. Burlington, VT: Ashgate.

_____ . 2008c. "Skepticism about the External World," in John Greco, ed., *Oxford Handbook of Skepticism*. Oxford: Oxford University Press.

_____ . 2009a. "Religious Knowledge in the Context of Conflicting Testimony." *Proceedings of the American Catholic Philosophical Association,* 83: 61-76.

_____ . 2009b. "The Value Problem," in A. Haddock, A. Millar & D. H. Pritchard, eds., *Epistemic Value*. New York: Rutledge.

_____ . 2010. *Achieving Knowledge: A Virtue Theoretic Approach of Epistemic Normativity*. Cambridge: Cambridge University Press.

_____ . 2011a. "Epistemic Circularity: Vicious, Virtuous and Benign." *International Journal for the Study of Skepticism* 1 (2): 105-112.

_____ . 2011b. "Evidentialism and Knowledge," in Trent Dougherty, ed., *Evidentialism and Its Discontents*. Oxford: Oxford University Press.

_____ . 2011c. "Virtue Epistemology," in Edward N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/entries/epistemology-virtue//.

_____ . 2012. "Recent Work on Testimonial Knowledge." *American Philosophical Quarterly* 49 (1): 15-28.

_____ . 2014. "Episteme: Knowledge and Understanding" in Kevin Timpe and Craig Boyd, eds., *Virtues and Their Vices*. Oxford: Oxford University Press.

_____ . (forthcoming). "Testimonial Knowledge and the Flow of Information," in John Greco and Henderson, eds., *Epistemic Evaluation:*

Point and Purpose in Epistemology. Oxford: Oxford University Press

Kelly, Thomas. 2005a. "Moorean Facts And Belief Revision, Or Can The Skeptic Win?" *Philosophical Perspectives* 19 (1): 179-209.

_____ . 2005b. "The Epistemic Significance of Disagreement," in Tamar Gendler and John Hawthorne, eds., *Oxford Studies in Epistemology*. New York: Oxford University Press.

_____ . 2006. "Evidence." *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta, ed., http://plato.stanford.edu/entries/evidence//.

_____ . 2010. "Peer Disagreement and Higher-Order Evidence," in Richard Feldman & Ted A. Warfield, eds., *Disagreement*. Oxford: Oxford University Press.

_____ . 2013. "Evidence Can Be Permissive," in Matthias Steup, John Turri, and Ernest Sosa, eds., *Contemporary Debates in Epistemology*, 2nd ed., Oxford: Oxford University Press.

Harman, Gilbert. 1986. *Change in View: Principles of Reasoning*. Cambridge, MA: MIT Press.

_____ . 1999a. "The Nonexistence of Character Traits." *Proceedings of the Aristotelian Society* 100 (2): 223-226.

_____ . 1999b. *Reasoning, Meaning, and Mind*. Oxford: Oxford University Press.

Hawthorn, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.

Hookway, Christopher. 2003. "How To Be A Virtue Epistemologist," in Linda Zagzebski and Michael DePaul, eds., *Intellectual Virtue: Perspectives From Ethics and Epistemology*. Oxford: Oxford University Press.

Hume, David. [1738] 1979, *A Treatise of Human Nature*. Edited by L.A. Shelby-Bigge, 2nd. ed. Oxford: Clarendon Press.

_____ . [1748] 1975. *An Enquiry concerning Human Understanding and Concerning the Principles of Morals*. Edited by L. A. Shelby-Bigge, 3rd. ed. Oxford: Clarendon Press, 1975.

Kim, Jaegwon. 2006. *Philosophy of Mind*. Cambridge, MA: Westview.

King, Nathan L. 2008. "Religious Diversity and its Challenges to Religious Belief." *Philosophy Compass* 3 (4): 830-853.

_____ . 2011. "Disagreement' What's the Problem? *or* A Good Peer is Hard

to Find." *Philosophy and Phenomenological Research* 85 (2): 249-272.

_____ . 2013. "Disagreement: The Skeptical Argument from Peerhood and Symmetry," in Diego E. Mancuca, ed., *Disagreement and Skepticism*. New York: Routledge.

Kornblith, Hilary. 2010. "Belief in the Face of Controversy," in Richard Feldman & Ted A. Warfield, eds., *Disagreement*. Oxford: Oxford University Press.

_____ . 2013. "Is Philosophical Knowledge Possible?" in Diego E. Machuca, ed., *Disagreement and Skepticism*. New York: Routledge.

Kripke, Saul. 1972. *Naming and Necessity*. Cambridge, MA: Harvard University Press.

Kvanvig, Jonathan. 1992. *The Intellectual Virtues and the Life of the Mind: On the Place of Virtues in Epistemology*. Savage, MD: Rowman & Littlefield.

_____ . 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.

_____ . 2006. "Closure Principles." *Philosophy Compass*, 1 (3): 256-267.

_____ . 2011. "The Rational Significance of Reflective Ascent," in Trent Dougherty, ed., *Evidentialism and Its Discontents*. Oxford: Oxford University Press.

Lackey, Jennifer. 2006. "It Takes Two to Tango: Beyond Reductionism and Non-Reductionism in the Epistemology of Testimony," in Jennifer Lackey and Ernest Sosa, eds., *The Epistemology of Testimony*. Oxford: Oxford University Press.

_____ . 2007. "Why We Don't Deserve Credit for Everything We Know." *Synthese* 158 (3): 345-361.

_____ . 2008. *Learning from Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.

_____ . 2010a. "A Justificationist View of Disagreement's Epistemic Significance," in Adrian Haddock, Alan Millar, and Duncan Pritchard, eds., *Social Epistemology*. Oxford: Oxford University Press.

_____ . 2010b. "What Should We Do When We Disagree?" in Tamar Szabó Gendler and John Hawthorne, eds., *Oxford Studies in Epistemology* 3. Oxford: Oxford University Press.

_____ 2013. "Disagreement and Belief Dependence," in David Christensen and Jennifer Lackey, eds., *The Epistemology of Disagreement: New Essays*. Oxford: Oxford University Press.

Lehrer, Keith. 1990. *Theory of Knowledge*. New York: Routledge.

Lewicki, Paul, T. Hill & M. Czyzewska. 1992. "Nonconscious Acquisition of Information." *American Psychologist* 47 (6): 792-801.

Lewis, David. 1983. *Philosophical Papers, vol. 1*. Oxford: Oxford University Press.

_____ . 1986. *On the Plurality of Worlds*. Malden, MA: Blackwell Publishers.

_____ . 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74 (4):549-567.

Lyons, Jack. 2009. *Perception and Basic Beliefs: Zombies, Modules, and the Problems of the External World*. Oxford: Oxford University Press.

Matheson, Jonathan. 2011. "The Case For Rational Uniqueness." *Logic and Episteme* 2 (3): 359-373.

McGinn, Colin. 1993. *Problems in Philosophy: The Limits of Inquiry*. Oxford: Blackwell.

Miller, Alan. 1991. *Reasons and Experience*. Oxford: Clarendon Press.

Moser, Paul. 1989. *Knowledge and Evidence*. Cambridge: Cambridge University Press.

Montmarquet, James. 1993. *Epistemic Virtue and Doxastic Responsibility*. Lanham, MD: Rowman & Littlefield.

Moore, G.E. 1939. "Proof of an External World." *Proceedings of the British Academy* 25: 272-300.

Nagel, Thomas. 1979. *Mortal Questions*. Cambridge: Cambridge University Press.

Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge: MA: Harvard University Press.

Nelkin, Dana, K. 2008. "Moral Luck," in Edward N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/entries/moral-luck/

Orozco, Joshua. 2011. "Epistemic Luck." *Philosophy Compass* 6 (1): 11-21.

Plantinga, Alvin. 1983. "Reason and Belief in God," in Alvin Plantinga and Nicholas Wolterstorff, eds., *Faith and Rationality*. Notre Dame, IN: University of Notre Dame Press.

_____ . 1986. "Intellectual Sophistication and Basic Belief in God." *Faith and Philosophy* 3: 306-312.

_____ . 1993a. *Warrant: The Current Debate*. Oxford: Oxford University Press.

_____ . 1993b. *Warrant and Proper Function*. Oxford: Oxford University Press.

_____ . 2002. "Reply To Beilby's Cohorts," in James Beilby, ed., *Naturalism Defeated?: Essays on Plantinga's Evolutionary Argument Against Naturalism*. London: Cornell University Press.

_____ . 2011. *Where The Conflict Really Lies*. Oxford: Oxford University Press.

Plato [427-347 B.C.] 1997. "Meno," in J. Cooper, ed., *Plato: Complete Works*. Indianapolis: Hackett.

Pollock, John. 1986. *Contemporary Theories of Knowledge*. Totowa, NJ: Rowman & Littlefield.

Pritchard, Duncan. 2003. "Virtue Epistemology and Epistemic Luck." *Metaphilosophy* 34 (1/2): 106-130.

_____ . 2007a. *Epistemic Luck*. Oxford: Oxford University Press.

_____ . 2007b. 'Recent Work on Epistemic Value.' *American Philosophical Quarterly* 44: 85–110.

_____ . 2008. "The Value of Knowledge," in Edward N. Zalta, ed., *The Standford Encyclopedia of Philosophy*. http://plato.standford.edu/archives/fall2008/entries/knowledge-value/

_____ . 2009. "Knowledge," in John Shand, ed., *Central Issues of Philosophy*. Oxford: Wiley-Blackwell.

_____ . 2010. "Knowledge and Final Value," in Duncan Pritchard, Alan Millar, and Adrian Haddock, eds., *The Nature and Value of Knowledge: Three Investigations*. Oxford:
Oxford University Press.

_____ . 2011. "Evidentialism, Internalism, Disjunctivism," in Trent Dougherty, ed., *Evidentialism and its Discontents*. Oxford: Oxford University Press.

Putnam, Hilary. 1975. "The Meaning of Meaning." *Minnesota Studies in the Philosophy of Science* 7: 131-193.

Quine, Willard. V.O. 1960. *Word and Object*. Cambridge, MA: MIT Press.

_____ . 1985. "States of Mind." *Journal of Philosophy*, 82.

Reid, Thomas. [1785] 1993. "Essay on the Intellectual Powers of Man," in Sir W

Hamilton, ed., *The Works of Thomas Reid*. Charlottesville, VA: Lincoln-Rembrandt Publishing.

Riggs, Wayne. 1998. "What are the "Chances" of Being Justified?," *The Monist* 81 (3): 452-472.

_____ . 2002. "Reliability and the Value of Knowledge." *Philosophy and Phenomenological Research* 64 (1): 79-96.

_____ . 2003. "Understanding 'Virtue' and the Virtue of Understanding," in Michael DePaul and Linda Zagzebski, eds., *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford: Oxford University Press.

_____ . 2006. "The Value Turn in Epistemology," in V. Hendricks and D. Pritchard, eds., *New Waves in Epistemology*. Aldershot: Ashgate.

_____ . 2007. "Why epistemologists are so down on their luck." *Synthese* 158 (3): 329-344

_____ . 2008. "Epistemic Risk and Relativism." *Acta Analytica* 23 (1): 1-8.

_____ . 2009. "Luck, Knowledge, and Control," in Alan Millar, Adrian Haddock, and Duncan Pritchard, eds., *Epistemic Value*. Oxford: Oxford University Press.

Roberts, Robert C. and Jay Wood. 2007. *Intellectual Virtues: An Essay in Regulative Epistemology*. Oxford: Oxford University Press.

Rorty, Richard. 1970. "In Defense of Eliminative Materialism." *Review of Metaphysics* 24 (September): 112-121.

Rosen, Gideon. 2001 "Nominalism, Naturalism, Epistemic Relativism." *Noûs* 35 (s15): 69 – 91.

Scott-Kakures. 1994. "On Belief and Captivity of the Will." *Philosophy and Phenomenological Research* 54 (1):77-103.

Sextus Empiricus. [160-210 A.D.] 1997. *Against the Ethicists (Adversus Mathematicos XI),* trans. Richard Bett. Oxford: Oxford University Press.

Sharkey, Amanda J. C. & Sharkey, Noel. 2009. "Connectionism," in John Symons and Paco Calco, eds., *The Routledge Companion to Philosophy of Psychology*. New York: Routledge.

Sosa, Ernest. 1980. "The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge." *Midwest Studies in Philosophy* 5 (1): 3-26.

_____ . 1991. *Knowledge in Perspective*. Cambridge: Cambridge University Press.

_____ . 1997. "How to Resolve the Pyrrhonian Problematic: A Lesson From Descartes." *Philosophical Studies* 85 (2-3): 229-249.

_____ . 2003. "The Place of Truth in Epistemology," in Linda Zagzebski and Michael DePaul, eds., *Intellectual Virtue*. Oxford: Oxford University Press.

_____ . 2006. "Knowledge: Instrumental and Testimonial," in Jennifer Lackey and Ernest Sosa, eds., *The Epistemology of Testimony*. Oxford: Oxford University Press.

_____ . 2007a. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, vol. II. Oxford: Oxford University Press.

_____ . 2010. "The Epistemology of Disagreement," in Alan Haddock, Adrian Millar, and Duncan Pritchard, eds., *Social Epistemology*. Oxford: Oxford University Press.

Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.

Stich, Stephen. 1983. *From Folk Psychology to Cognitive Science: The Case Against Belief*. Cambridge, MA: MIT Press.

Thune, Michael. 2009. "'Partial Defeaters' And the Epistemology of Disagreement." *The Philosophical Quarterly*. 60 (239): 355-372.

_____ . 2010 "Religious Belief and the Epistemology of Disagreement." *Philosophy Compass* 5 (8): 712-724.

Unger, Peter. 1975. *Ignorance: A Case for Skepticism*. Oxford: Oxford University Press.

Van Cleve, James. 2003. "Is Knowledge Easy—Or Impossible? Externalism as the Only Alternative to Skepticism," in Steven Luper, ed., *The Skeptics*. Aldershot: Ashgate Publishing.

_____ . 2006. "Reid on the Credit of Human Testimony," in Jennifer Lackey and Ernest Sosa, eds., *The Epistemology of Testimony*. Oxford: Oxford University Press.

van Inwagen, Peter. 1983. *An Essay on Free Will*. Oxford: Clarendon Press.

_____ . 1996. "It is Wrong, Always, Everywhere, and for Anyone, to Believe Anything, Upon Insufficient Evidence," in Jeff Jordan and Daniel Howard- Snyder, eds., *Faith, Freedom, and Rationality*. Lanham, MD.: Rowman & Littlefield.

_____ . "Materialism and the psychological-continuity account of personal identity," in J. Tomberlin. ed., *Philosophical Perspectives 11: Mind, Causation and Word*. Oxford: Blackwell.

_____ . 2010. "We're Right. They're Wrong," in Richard Feldman & Ted A. Warfield, eds., *Disagreement*. Oxford: Oxford University Press.

Wedgwood, Ralph. 2010. "The Moral Evil Demons" in Richard Feldman & Ted A. Warfield, eds., *Disagreement*. Oxford: Oxford University Press.

Weiskrantz, Lawrence. 1986. *Blindsight: A Case Study and Implications*. Oxford: Clarendon Press.

Weiner, Matthew. 2003. "Accepting Testimony." *The Philosophical Quarterly* 53 (211): 256-264.

Williamson, Timothy. 2000. *Knowledge and its Limits*. New York: Oxford University Press.

_____ . 2007a. "On Being Justified in One's Head," in Mark Timmons, John Greco, and Alfred R. Mele, eds., *Rationality And The Good: Critical Essays on the Ethics and Epistemology of Robert Audi*. Oxford: Oxford University Press.

_____ . 2007b. *The Philosophy of Philosophy*. New York: Oxford University Press.

Zagzebski, Linda. 1994. "The Inescapability of Gettier Problems." *Philosophical Quarterly* 44 (174) :65-73

_____ . 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press.

_____ . 2000. "From Reliabilism to Virtue Epistemology," in Guy Axtell, ed., *Knowledge, Belief and Character: Readings In Virtue Epistemology*. Lanham, MD: Rowman & Littlefield.

_____ . 2001. "Rediscovering Understanding," in M. Steup, ed., *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue.* Oxford: Oxford University Press.

_____ . 2003. "The Search for the Source of the Epistemic Good." *Metaphilosophy* 34 (1-2): 12-28

_____ . 2007. "Ethical And Epistemic Egoism And The Ideal Of Autonomy." *Episteme: A Journal of Social Epistemology* 4 (3): 252-263.

_____. 2008. Unpublished. "Epistemic Trust and Epistemic Conflict"

_____ . 2009. *On Epistemology*. Belmont, CA: Wadsworth.

_____ . 2012. *Epistemic Authority: Trust, Authority, and Autonomy of Belief*. Cambridge: Cambridge University Press.

Michael Zimmerman. 2010. "Intrinsic vs. Extrinsic Value," in Edward N. Zalta, ed., *The Standford Encyclopedia of Philosophy*. http://plato.stanford.edu/entries/value-intrinsic-extrinsic/#WhaExtVal