VIRTUE EPISTEMOLOGY AND THE ANALYSIS OF KNOWLEDGE

Ian M. Church

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews

2012

UNIVERSITY OF ST ANDREWS

THE ST ANDREWS-STIRLING JOINT PROGRAMME IN PHILOSOPHY

&

VIRTUE EPISTEMOLOGY AND THE ANALYSIS OF KNOWLEDGE

A THESIS SUBMITTED BY:

IAN M. CHURCH

TO THE FACULTY OF PHILOSOPHY
IN CANDIDACY FOR THE DEGREE OF:

DOCTOR OF PHILOSOPHY

ST ANDREWS, SCOTLAND
MARCH 2012

DECLARATIONS

I, Ian M. Church, hereby certify that this thesis, which is approximately 75,000 words in length, has been written by me, that it is the record of work carried out by me and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in September, 2008 and as a candidate for the degree of Doctor of Philosophy in August, 2009; the higher study for which this is a record was carried out in the University of St Andrews between 2008 and 2012.

Date: _____ Signature of Candidate: _____

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of Doctor of Philosophy in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date: _____ Signature of Supervisor: _____

COPYRIGHT

A man will be imprisoned in a room with a door that's unlocked and opens inwards; as long as it does not occur to him to pull rather than push it.

– *Culture and Value* (1980), Ludwig Wittgenstein, 42e

ABSTRACT

This thesis centers on two trends in epistemology: (i) the dissatisfaction with the reductive analysis of knowledge, the project of explicating knowledge in terms of necessary and jointly sufficient conditions, and (ii) the popularity of virtue-theoretic epistemologies. The goal of this thesis is to endorse *non-reductive virtue epistemology*. Given that prominent renditions of virtue epistemology assume the reductive model, however, such a move is not straightforward—work needs to be done to elucidate what is wrong with the reductive model, in general, and why *reductive* accounts of virtue epistemology, specifically, are lacking.

The *first* part of this thesis involves diagnosing what is wrong with the reductive model and defending that diagnosis against objections. *The* problem with the reductive project is the Gettier Problem. In Chapter 1, I lend credence to Linda Zagzebski's grim 1994 diagnosis of Gettier problems (and the abandonment of the reductive model) by examining the nature of luck, the key component of Gettier problems. In Chapter 2, I vindicate this diagnosis against a range of critiques from the contemporary literature.

The *second* part involves applying this diagnosis to prominent versions of (reductive) virtue epistemology. In Chapter 3, we consider the virtue epistemology of Alvin Plantinga. In Chapter 4, we consider the virtue epistemology of Ernest Sosa. Both are seminal and iconic; nevertheless, I argue that, in accord with our diagnosis, neither is able to viably surmount the Gettier Problem.

Having diagnosed what is wrong with the reductive project and applied this diagnosis to prominent versions of (reductive) virtue epistemology, the *final* part of this thesis explores the possibility of *non-reductive* virtue epistemology. In Chapter 5, I argue that there are three strategies that can be used to develop non-reductive virtue epistemologies, strategies that are compatible with seminal non-reductive accounts of knowledge and preserve our favorite virtue-theoretic concepts.

ACKNOWLEDGMENTS

When my wife and I moved to St Andrews in September 2007, neither one of us had lived more than 45 minutes away from where we were born in the state of Indiana. That first year, while I pursued my MLitt in the St Andrews-Stirling Joint Programme, was marked with anxiety and uncertainty. By the time I started on my PhD, however, we were flourishing in our new environment. We had grown to love not only St Andrews, Scotland, but the United Kingdom, Europe, and indeed the (now much bigger) world. My tenure at the University of St Andrews (and the Joint Programme) has been defined by the highest standards in philosophical teaching, research, supervision, and fellowship, no doubt; but what is more, it has been defined by a new and expanding global perspective. I am immensely grateful for the faculty, staff, and students of the St Andrews-Stirling Joint Programme in Philosophy and the broader St Andrews community.

The greatest thanks for the successful completion of this doctoral thesis go to my thesis committee. First of all and preeminently, my primary supervisor, Patrick Greenough, deserves my special gratitude. Patrick tirelessly and cheerfully saw my research through turbulent and transitional times. His wisdom, guidance, encouragement, and friendship were absolutely invaluable to my research and my development as a scholar. Secondly, Philip Ebert, my secondary supervisor, was a tremendous help—offering psychological, prudential, and academic support and assistance far beyond the call of duty. Thirdly, I would like to thank Stewart Cohen, my second secondary supervisor. Stew provided excellent feedback and critical commentary during the writing of this thesis, and his guidance was indispensable. Finally, I am incredibly grateful to my thesis examiners, John Greco and Katherine Hawley. Their challenging critiques of my work and perceptive feedback during the examination were extremely helpful.

There are a host of other faculty and staff within the academic communities of the University of St Andrews and the University of Stirling who were also very valuable to me throughout my tenure as a graduate student. In particular, I am grateful to Katie Allen, Mike Arrowsmith, Yuri Cath, Shaun Darby, Dylan Dodd, John Haldane, Colin Johnston, Lisa Jones, Nikki Leavitt, Grant Macaskill, Peter Milne, Stephen Read, Sonia Roca-Royes, Brian Weatherson, and Lucie Wight. What is more, I owe a special debt to Alan Millar, Tim Mulgan, and Alan Torrance. First, Alan Millar, who was my supervisor for the first two years of my doctoral studies, offered extensive and valuable feedback and guidance during that time. Secondly, Tim Mulgan's kindness, care, and friendship have been a tremendous support and inspiration throughout my tenure in St Andrews. Thirdly and finally, Alan Torrance, Professor of Systematic Theology, has been not only an amazing help regarding

my research but also an astonishing example of character, strength, and virtue; his friendship and guidance have helped me learn more of what it means to be a Christian.

My friends and colleagues (within the Joint Programme and abroad) were also a constant source of support, encouragement, and academic dialogue. In particular, I would like to thank Shawn Bawulski, Tony Bolos, Paul Dimmock, Inés de Asis, Brannon Ellis, Andrew Hay, Spencer Johnston, Sheree Lear, Blake Roeber, and Luke Tallon. I owe a special debt of gratitude to Felix Pinkert, who has been a tremendous source of both camaraderie as well as philosophical and spiritual edification.

During our tenure in St Andrews, my wife and I attended the St Andrews Free Church of Scotland. We are both immensely thankful for the care and support over the years of the minister, Alasdair I. MacLeod, and that congregation.

Substantial portions of the research in this thesis has been carried out and presented outside the St Andrews-Stirling Joint Programme. I am exceedingly grateful for the helpful feedback I have received from the faculty and students at Ball State University, the City University of New York, Rutgers University, Taylor University, the University of Edinburgh, and the University of Notre Dame. In particular, I would like to thank Michael Levin, Alvin Plantinga, Duncan Pritchard, Ernest Sosa, James Spiegel, and Chris Tucker.

I am also extremely grateful to Brian Kinney, who copyedited the final version of this thesis.

Finally, I would like to thank my wife, Corrie, to whom this thesis is dedicated. During our time in St Andrews, she has worked tirelessly as a graphic designer, an artist, a homemaker, and the mother of our two boys, Aed and Asher—all in support of my graduate studies. She has been a tremendous friend and companion, and she is truly the heroine of this thesis.

DEDICATION

I dedicate this thesis to my wife, Corrie. Without her love, encouragement, support, inspiration, and friendship this thesis would not have been possible. Her virtues are beyond analysis.

TABLE OF CONTENTS

INTRODUCTION

Epistemology is on the move. On the one hand, there is a growing dissatisfaction with the project of explicating knowledge in terms of a reductive analysis, in terms of necessary and jointly sufficient conditions. Ever since 1963 when Edmund Gettier challenged the sufficiency of the standard analysis of knowledge with a series of counterexamples, all attempts to defend it have been shown either to lead to further Gettier-style counterexamples or to produce an analysis of knowledge that is unfeasible.[1] On the other, there is the ascendancy and popularity of virtue-theoretic epistemology. Tracking our intuitions across a range of cases, virtue epistemology offers new and exciting solutions to a variety of problems afflicting other accounts of knowledge. *The goal of this thesis is to endorse both of these trends, to endorse the move away from the reductive analysis model while affirming virtue-theoretic epistemology.*

Unfortunately, however, within the contemporary literature, these trends are at odds. Prominent renditions of virtue epistemology assume the reductive analysis model, and as such they stand in the way of any movement toward non-reductive virtue-theoretic accounts.[2] In this thesis, I do three things. First, I elucidate and defend a diagnosis of the

[1] See "Is Justified True Belief Knowledge?" (1963).

[2] Perceptively, John Greco's recent work, *Achieving Knowledge* (2010)—bearing every indication of a new seminal work within virtue epistemology—shows signs of moving away from this trend. To quote Greco: "Knowledge is a kind of success from ability. . . . Is [this] intended as an 'analysis' of knowledge? Not in any traditional sense. . . . As Timothy Williamson has pointed out, these sorts of analysis make little sense outside the context of earlier philosophical projects. Do I intend, at least, to give necessary, sufficient and informative conditions? No, because to say that knowledge is a kind of success from ability is not to give *sufficient* conditions. On the contrary, it is to give a species without giving a difference. Nevertheless, the account is informative in a straightforward sense: it provides insight into what knowledge is by identifying it as an instance of a more general familiar kind" (2010, 3–4). That said, however, Greco does not seem to endorse the whole-hog abandonment of the reductive analysis project in favor of alternative models that I will be

dominant reason for abandoning the reductive model, namely, the Gettier Problem. Second, I apply this diagnosis to two of the major players in contemporary virtue epistemology, Alvin Plantinga and Ernest Sosa—arguing that virtue epistemology simply offers no new cures. Third and finally, I conclude by explicating by way of prolegomena some ways in which we might pursue non-reductive virtue epistemology.

Before getting started, we need to do some brief stage setting. In this introduction, I first elucidate and outline in very broad strokes the two trends in contemporary epistemology that motivate my thesis: the growing pessimism regarding the analysis of knowledge and the ascendancy of virtue epistemology. In Section 1, I chart the growing dissatisfaction with the reductive analysis model of knowledge. In Section 2, I explicate the contours of the virtue-theoretic epistemology that dominates the current epistemic landscape. In Section 3, I conclude by summarizing in detail the rest of the thesis, chapter by chapter.

Section 1: Current Trends in Epistemology, Part 1: A Growing Dissatisfaction

Before we consider the post-Gettier fallout that has lead so many epistemologists to be pessimistic regarding the analysis of knowledge, we should be clear regarding just what it means to pursue a reductive analysis of knowledge. Consider Matthias Steup's explanation:

> The objective of the analysis of knowledge is to state the conditions that are individually necessary and jointly sufficient for propositional knowledge: knowledge that such-and-such is the case. . . . The concept to be analyzed—the analysandum—is commonly expressed using the schema "*S* knows that *p*", where "*S*" refers to the knowing subject, and "*p*" to the proposition that is known. A proposed analysis consists of a statement of the following form: *S* knows that *p* if and only if __. The blank is to be replaced by the analysans: a list of conditions that are individually necessary and jointly sufficient. To test whether a proposed analysis is correct, we must ask (a) whether every possible case in which the conditions listed in the analysans are met is a case in which *S* knows that *p*, and (b) whether every possible case in which *S* knows that *p* is a case in which each of these conditions is met. When we ask (a), we wish to find out whether the proposed analysans is sufficient for *S*'s knowing that *p*; when we ask (b), we wish to determine whether each of the conditions listed in the analysans is necessary. (Steup 2006, ¶ 1)

---

advocating in this thesis. While Greco explicitly declines to provide a thoroughgoing reductive analysis, he does eventually seem to take himself as providing *enough* of an analysis (conditions on knowledge that are necessary and sufficient *enough*) to achieve viable, lasting progress against the Gettier Problem, a point that this thesis takes issue with (see 2010, 12–13, 73–80).

The reductive analysis of knowledge, in sum, is the project of explicating knowledge in terms of a list of necessary and jointly sufficient conditions. While this is certainly right, more needs to be said. After all, *being Tully* is both necessary and sufficient for *being Cicero*, but that does not mean that *being Tully* is a reductive analysis of *being Cicero*. The reductive analysis of knowledge not only entails the pursuit of necessary and jointly sufficient conditions, it entails the pursuit of necessary and jointly sufficient conditions that are purported to be conceptually prior to knowledge. The reductive analysis of knowledge project is the project of building up to knowledge, so to speak, from various necessary conditions that are taken to be conceptually prior to knowledge—conditions such as truth, belief, etc. Just as terms like "bachelor" and "vixen" yield reductive analyses in the more primitive terms of unmarried men and female foxes respectively, "knowledge," it is assumed, will yield a reductive analysis in the purportedly more primitive terms of truth, belief, and *whatever else makes true belief knowledge* (call it warrant).

Epistemologists have always wanted to know what knowledge was, and in at least the past 40 years this inquiry has taken the shape of a reductive analysis of its necessary and jointly-sufficient conditions. We may at the outset wonder whether such an approach is fundamentally misguided; after all, why should we think knowledge yields a reductive analysis in the first place?[3] Why would we think "knowledge" is like "bachelor" or "vixen"? In any case, while I will in this thesis eventually suggest that such an approach to propositional knowledge *is* fundamentally misguided, all we need to appreciate now is that this approach is precisely what has dominated contemporary epistemology.

Enter Gettier. It's a story that we all know well—a story that is retold in almost every introduction to epistemology module. Since the time of Plato, so the story goes, we were all raised on our mothers' knees learning that knowledge consisted of justification, truth, and belief; since time immemorial, *justified true belief* has been taken to be the necessary and jointly-sufficient conditions for knowledge.[4] But then, with the publication of Edmund Gettier's 1963 paper "Is Justified True Belief Knowledge?" everything changed. Consider the following classic counterexample:

---

[3] After all, to quote William G Lycan, "It is well to remind ourselves that no effort of analytic philosophy to provide strictly necessary and sufficient conditions for a philosophically interesting concept has ever succeeded" (2006, 150).

[4] To be sure, it is not at all clear that Plato had anything like a reductive analysis in mind when he explicated knowledge in such terms in the *Meno*. See Plato 2002.

> **Classic Case**: Smith and Jones are applying for the same job. Smith has very strong evidence for thinking that Jones will get the job (e.g., the employer tells Smith that he will hire Jones, etc.), and for thinking that Jones has 10 coins in his pocket (e.g., Jones emptied his pockets in front of Smith and then clearly, slowly, in good lighting, and perhaps even counting out loud, placed 10 coins in his pocket). As such, Smith forms a belief in the general proposition that "the man who gets the job has 10 coins in his pocket." As it turns out, however, Smith gets the job and he happens to also have 10 coins in his pocket. (paraphrased from Gettier 1963, 122)

In this case, Smith seemingly has a justified true belief that "the man who gets the job has 10 coins in his pocket," but, as almost everyone agrees, surely Smith's belief is not knowledge. As such, *justified true belief* is simply not sufficient for knowledge.[5]

As such, epistemologists quickly tried to find ways to save or repair the *justified true belief* analysis—typically by trying either to strengthen the justification condition or to add more conditions (e.g., *justified true belief* plus some fourth condition). What ensued was a confusing and cacophonous array of justification (e.g., justification as evidence, justification as reliability, internalist justification, externalist justification, etc.) and additional conditions on knowledge (e.g., safety conditions, sensitivity conditions, defeasibility conditions, etc.). All of this sought to understand knowledge via reductive analysis, all of which assumed a *warranted true belief* analysis of knowledge (where "warrant" stands for "whatever turns true belief into knowledge"). Unfortunately, none of these proposals (in any combination) achieved any lasting success against Gettier counterexamples—either falling into further Gettier-style counterexamples or leading to unpalatable conclusions (such as radical skepticism).[6] Providing a feasible reductive analysis of knowledge, it seems, is a Sisyphean endeavor.

---

[5] We need not concern ourselves with mechanics or scope of Gettier counterexamples at this point. A full diagnosis of the Gettier Problem will be provided in Chapter 2 of this thesis.

[6] Consider the following diagnosis by Linda Zagzebski: "The moral drawn in the thirty years since Gettier published his famous paper is that either justified true belief (JTB) is not sufficient for knowledge, in which case knowledge must have an 'extra' component in addition to JTB, or else justification must be reconceived to make it sufficient for knowledge. I shall argue that given the common and reasonable assumption that the relation between justification and truth is close but not inviolable, it is not possible for either move to avoid Gettier counter-examples. What is more, it makes no difference if the component of knowledge in addition to true belief is identified as something other than justification, e.g. warrant or well-foundedness. I conclude that Gettier problems are inescapable for virtually every analysis of knowledge which at least maintains that knowledge is true belief plus something else" (1994, 65).

INTRODUCTION

Some commentators have put the state of contemporary post-Gettier epistemology rather mildly:

> Gettier sparked a period of pronounced epistemological energy and innovation—all with a single two-and-a-half page article. There is no consensus, however, that any one of the attempts to solve the Gettier challenge has succeeded in fully defining what it is to have knowledge of a truth or fact. So, the force of that challenge continues to be felt in various ways, and to various extents, within epistemology. (Hetherington 2005, ¶ 1)

Others, however, have been perhaps a bit more pessimistic (and, by my lights, perhaps a bit more accurate):

> Since Gettier refuted the traditional analysis of knows as has a justified true belief in 1963, a succession of increasingly complex analyses have been overturned by increasingly complex counterexamples… The pursuit of analyses is a degenerating research programme. (Williamson 2000, 30–31)

> Since [the Gettier Problem] is demonstrably unsolvable, it follows not only that the tripartite [*warranted true belief*] account is logically inadequate as it is, but also that it is irretrievably so in principle. [The Gettier Problem] is not a mere anomaly, requiring the rectification of an otherwise stable and acceptable account of propositional knowledge. It is proof that the core of the approach needs to be abandoned. (Floridi 2004, 76)

In any case, the situation seems dim. Knowledge has not yet yielded a viable reductive analysis, and there is no indication that it will any time soon. As such, many philosophers have grown pessimistic regarding such a project. Indeed, many philosophers have been tempted either to give up on knowledge altogether or to start looking for an alternative model from which to do epistemology.

Of course, it is the contention of this thesis that this latter option, looking for an alternative model from which to do epistemology, is generally the right move. Dogmatically giving up on knowledge altogether seems a bit harsh, especially if other, viable options are available. And, given the growing pessimism regarding the reductive analysis project, there seems to be sufficient ground for exploring such alternatives in earnest. For the lovers of virtue-theoretic epistemology, however, there are some major hurdles. The prominent iterations of virtue epistemology in the contemporary literature assume the reductive model, and, what is more, some of them even openly purport to have solved the Gettier Problem, to have provided viable reductive analyses of knowledge that do not yield Gettier counterexamples. As such, anyone who is duly pessimistic regarding the reductive analysis project, yet an advocate of virtue-theoretic epistemology, simply must address the

prominent reductive iterations of virtue epistemology if a move toward non-reductive virtue epistemology is to be sufficiently warranted.

Section 2: Current Trends in Epistemology, Part 2: Enter Virtue Epistemology

Despite the growing trend to abandon the reductive analysis of knowledge in favor of alternative epistemic models, reductive virtue-theoretic epistemologies have for nearly 30 years become increasingly popular and commonplace.[7] The move away from the reductive model has not yet come to bear on virtue epistemology. In this section, I briefly explicate the basic tenets of virtue epistemology, particularly of the sort that most permeates the contemporary literature—further explicating the state of affairs that motivates this thesis.[8]

What is virtue epistemology? To be sure, as it has surged in popularity over the past 30 years, virtue epistemology has developed into a multitude of positions; nevertheless, every variant of virtue epistemology holds to two basic resolutions: (1) that epistemology is a normative discipline and (2) that "intellectual agents and communities are the primary source of epistemic value and the primary focus of epistemic evaluation" (Greco and Turri 2011, para. 1). The former amounts to (a) a rejection of Quine's proposal in "Epistemology Naturalized" (1969) that epistemologists should give up on attempts to discern what is reasonable to believe in favor of projects within cognitive psychology and (b) a call for epistemologists to "focus their efforts on understanding epistemic norms, value and evaluation" (Greco and Turri 2011, sec. 1).[9] And to better understand the second resolution think of virtue *ethics'* niche within moral philosophy. For the two titans of moral philosophy, Kantian deontology and utilitarianism, the starting place for moral evaluation is *action*. For Kantians and for utilitarians, the question to ask when doing ethics is "What

[7] Of course, virtue epistemology's philosophical roots go much deeper than the past 30 years. To quote John Greco and John Turri, "Practitioners [of virtue epistemology] draw inspiration from many important historical philosophers, including Plato (Zagzebski 1996, 139), Aristotle (Greco 2002, 311; Sosa 2009a, 2: 187; Zagzebski 1996, passim), Aquinas (Roberts and Wood 2007, 69–70; Zagzebski 1996, passim), Descartes (Sosa 2007, ch. 6), Kierkegaard (Roberts and Wood 2007, 29–30) and Peirce (Hookway 2000). Hints of [virtue epistemology] can also be found in Hume (1748), Reid (1785), Russell (1948) and Sellars (1956)" (2011, sec. 2).

[8] In accord with the surging popularity of virtue epistemology, a number of excellent edited collections have been produced. For example, see Axtell 2000; Brady and Pritchard 2003; DePaul and Zagzebski 2003; Fairweather and Zagzebski 2001; Greco 2004; Kvanvig 1996; Steup 2001.

[9] See Quine 1969. Also see McDowell 1994, 133; Sosa 1991, 100-105; Zagzebski 1996, 334-338.

should I do?" (Kantians roughly answering "act in accord with what you can will to be a universal maxim"; utilitarians roughly answering "act in accord with what brings about the greatest happiness for the greatest number.") For virtue ethicists, however, the starting place for moral evaluation, where the rubber meets the road, so to speak, is the *agent*—his or her character—and subsequently the virtue ethicist asks a different question, "How should I live?"[10] To put it roughly, then, instead of focusing on the beliefs of agents (whether or not they are justified, safe, etc.), virtue epistemologists predominantly focus on the agent himself or herself—on whether he or she has the right sort of epistemic character, the right sort of cognitive faculties, whether he or she is epistemically virtuous or not. To be sure, other theories of knowledge will give some account of epistemic virtues—good memory, intellectual courage, etc. —but usually in terms of knowledge; the radical claim that virtue epistemology makes, however, is that knowledge is defined in terms of virtue.[11]

Virtue epistemology, so defined, has developed by and large into two distinct schools: agent-reliabilism and responsibilism or neo-Aristotelianism.[12] The primary difference between the schools is their application of "virtue" terminology. Agent-reliabilism, being modeled along reliabilist lines, applies virtue terminology mechanically—in the same way we might talk about a virtuous knife. In other words, just as we might call a knife virtuous if it does what it is supposed to do (cut things, be sharp, etc.), agent-reliabilism calls various cognitive faculties such as memory, perception, etc., virtuous insofar as they are reliably functioning the way they are supposed to. That is, agent-reliabilism focuses on the reliable functioning (virtuous functioning) of a given agent's cognitive faculties. Neo-Aristotelianism, on the other hand, applies virtue terminology in a way we are perhaps more familiar with—in terms of specific character traits such as open-mindedness, intellectual courage, intellectual perseverance, etc.

[10] For more on virtue ethics and its distinctiveness, see Anscombe 1958.

[11] See Pritchard 2005, 186; Greco 2010, 17–46.

[12] Chief proponents of agent-reliabilism include Alvin Plantinga (1993b), Ernest Sosa (2007), and John Greco (2010). Chief proponents of responsibilism or neo-Aristotelianism include Lorraine Code (1987), James Montmarquet (1993), and Linda Zagzebski (1996).

By far, however, agent-reliabilism virtue epistemology is the most popular version of virtue epistemology, and as such it will get the majority of our attention in this thesis.[13] (In so doing, however, we will indeed see the shortcomings of Neo-Aristotelian virtue epistemology as well, if only in passing.) Agent-reliabilism virtue epistemologies, like the ones we will be considering, developed out of a dissatisfaction with process-reliabilism—the view that, to put it roughly, S knows a true proposition *p* if and only if *p* was formed by a reliable process. Naturally enough, there are some serious concerns for such a view (i.e., the "generality problem"); however, the agent-reliabilists' primary concern was that knowledge ascriptions based on reliable processes do not always appropriately involve a given agent—that process-reliabilism seems to allow agents to "possess knowledge even though the reliability in question in no way reflects a cognitive achievement on their part" (Pritchard 2005, 187). One way this occurs is when a given reliable process does not appropriately relate to facts. For example, consider the following case by John Greco:

> **René and the Gambler's Fallacy:** René thinks he can beat the roulette tables with a system he has devised. Reasoning according to the Gambler's Fallacy, he believes that numbers which have not come up for long strings are more likely to come up next. However, unlike Descartes' demon victim, our René has a demon helper. Acting as a kind of epistemic guardian, the demon arranges reality so as to make the belief come out as true. Given the ever present interventions of the helpful demon, René's belief forming process is highly reliable. But this is because the world is made to conform to René's beliefs, rather than because René's beliefs conform to the world. (Greco 1999, 286)[14]

Though René's beliefs happen to be based on a reliable process, it is completely accidental—it is certainly no thanks to any effort of René's. Even though René's beliefs are formed by a reliable process—the helper demon—intuitively he does not have knowledge.

Another way that knowledge ascriptions based on reliable processes do not always appropriately involve a given agent is in the case of reliable cognitive malfunctions. Consider a case originally developed by Alvin Plantinga in which our protagonist has a brain lesion that causes him to believe he has a brain lesion:

> **Brain Lesion:** Suppose . . . that S suffers from this sort of disorder and accordingly believes that he suffers from a brain lesion. Add that he has no evidence at all for this belief: no symptoms of which he is aware, no testimony on the part of physicians or other expert witnesses, nothing. (Add, if you like, that he

[13] Regarding the popularity of agent-reliabilism virtue epistemology over Neo-Aristotelian virtue epistemology, see John Greco's "Two Kinds of Intellectual Virtue" (2000).

[14] Also quoted in Pritchard 2005, 187.

> has much evidence against it, but then add also that the malfunction induced by the lesion makes it impossible for him to take appropriate account of this evidence.) Then the relevant [process] will certainly be reliable but the resulting belief—that he has a brain lesion—will have little by way of warrant for S. (Plantinga 1993a, 199)[15]

Again, though S's belief that he has a brain lesion is formed via a reliable process we would not ascribe knowledge to it since it is only formed out of a glitch in S's cognitive equipment. It is simply accidental that the brain lesion causes S to form the said belief, and as such S had nothing to do with its formation. Though there may be some worries as to what constitutes a given person's cognitive equipment (Why is S's brain lesion not a part of his cognitive equipment? What if he had the lesion since birth?), we nevertheless have strong intuitions that beliefs formed as a direct result of a cognitive malfunction cannot be knowledge.[16] S's belief that he has a brain lesion cannot be knowledge simply because, though formed via reliable process, the agent, S, was not appropriately involved in its formation.

It is the agent-reliabilists' focus on the epistemically virtuous agent (i.e., the agent with mechanically sound cognitive equipment) that is the cornerstone to their epistemology. It is their particular focus on the properly functioning human knower (and his/her cognitive competencies) that is meant to distinguish their accounts from all others. The agent-reliabilists' notions of proper function, cognitive competency, and intellectual achievement are, in contrast with other theories of warrant, meant to rightly connect a given agent to the facts (cf. René and the Gambler's Fallacy) and, of course, rightly preclude knowledge from cases of malfunction (cf. Brain Lesion). According to the agent-reliabilists, other theories of warrant fail due to their inability to track warrant ascription in accord with the proper functioning of the relevant cognitive faculties behind a given belief's genesis. It is this special focus on the epistemic agent and his or her properly functioning cognitive competencies that agent-reliabilism virtue epistemologists see as "the rock on which [competing] accounts of warrant founder" (Plantinga 1993b, 4).

Having briefly explicated the basic tenets of virtue-theoretic epistemology, it is worth noting that none of them demand a reductive analysis of knowledge. There is nothing about virtue epistemology *qua* virtue epistemology that requires an analysis of knowledge in

[15] Also quoted in Pritchard 2005, 188.

[16] See Pritchard 2005, 188; Greco 2003a, 356-357.

terms of necessary and jointly sufficient (and conceptually primitive) conditions. As such, it seems as though the *only* reason virtue epistemology developed along reductive lines is that such were the suppositions of its time. Be that as it may, given the sheer prominence of reductive virtue epistemologies, anyone interested in exploring the possibility of *non-reductive* virtue epistemology, in keeping with the trends explicated in Section 1, simply must address the reductive counterparts and motivate their abandonment. And that is the goal of this thesis: to explicate and defend a diagnosis of the main reason for abandoning *any* reductive account of knowledge, the Gettier Problem; to apply this diagnosis to prominent versions of reductive virtue epistemology, elucidating their shortcomings in light of their commitments to the reductive analysis project; to finally explore by way of prolegomena the possibility of non-reductive virtue epistemology.

Section 3: Thesis Outline

Having elucidated the two current trends in contemporary epistemology that motivate my thesis, the goal of making room for non-reductive virtue epistemology, we are ready to consider in more detail how this project will be undertaken. As I have noted, this thesis can be divided into three sections, movements, or parts: (i) the diagnosis of the problem afflicting reductive analyses, (ii) the application of the diagnosis to prominent renditions of reductive virtue epistemology, and (iii) the exploration of viable alternatives to the reductive model. These parts will be subdivided into five chapters. Part I consists of Chapter 1 and Chapter 2, Part II consists of Chapter 3 and Chapter 4, and Part III consists of Chapter 5. My goal in this final section of the Introduction is to briefly explicate and summarize these five chapters, giving the reader a broad idea of what to expect throughout this thesis.

Chapter 1: Luck and Gettier Problems. The main reason for the growing dissatisfaction with the reductive analysis project is no doubt the Gettier Problem. The inability to define knowledge in terms of necessary and jointly sufficient conditions that are not susceptible to Gettier counterexamples is perhaps the greatest difficulty facing reductive accounts. So, to diagnose what is wrong with the reductive analysis project, we need to diagnose the Gettier Problem; we need to understand when and why Gettier counterexamples occur. This is precisely what we will be doing in this chapter.

According to Linda Zagzebski's 1994 diagnosis of Gettier problems, if whatever we take to bridge the gap between true belief and knowledge (call it warrant) bears some

violable relationship to truth, then it will be possible for that belief to be so warranted and true for reasons unrelated to the warrant (Zagzebski 1994, 65). That is to say, so long as we reasonably assume that warrant is neither divorced from truth nor inseparable from it (such options being epistemically unpalatable), Gettier cases are "inescapable" (Zagzebski 1994, 65).

As eloquent and intuitive as Zagzebski's diagnosis of Gettier problems may be, the prominent advocates of virtue epistemology have apparently not taken it altogether seriously. If Zagzebski is right, this has very dim consequences for any palatable reductive analysis of knowledge; however, virtue epistemologists continue to work from such a model. Though there is a common trend in contemporary epistemology to abandon the standard analysis of knowledge in favor of alternative models, this trend has not yet extended to those working within virtue epistemology.

In Chapter 1, I add credence to Zagzebski's diagnosis (and the movement away from the standard analysis) by analyzing the nature of luck. It is widely accepted that the lesson to be learned from Gettier problems is that knowledge is incompatible with luck or at least a certain species thereof.[17] As such, understanding the nature of luck is central to understanding the Gettier Problem.[18] Thanks by and large to Duncan Pritchard's seminal work, *Epistemic Luck* (2005), a great deal of literature has been developed recently concerning the nature of luck and anti-luck epistemology. The literature, however, has yet to explore the very intuitive idea that luck comes in degrees.[19] I propose that once luck is recognized to admit degrees, even the slightest non-zero degree of luck (of the relevant sort) precludes knowledge. Connecting this to Zagzebski's thesis, I propose that a given theory of warrant must guarantee truth in order to avoid Gettier counterexamples (or subsequently deny that warrant bears any relationship whatsoever to the truth), simply because a sufficient reductive analysis of knowledge (whether a virtue epistemology or not) cannot allow for knowledge that is even marginally lucky.

---

[17] See Dancy 1985, 134; Zagzebski 1999, 99-101; Pritchard 2005, 4-5; Steup 2006, §2.

[18] While there are well over 100 different Gettier cases, they are united by common theme: a warranted belief that is only luckily true. As L. Floridi notes, though there are all of these cases, nobody would expect them to be answered in a case-by-case basis. (2004, 64).

[19] Someone may occasionally elucidate safety as coming in degrees (see Smith 2009), but this is the closest, to my knowledge, philosophers have come to describing luck in terms of degrees.

I justify this proposal by treating Pritchard's work as archetypal. Motivated by the intuition that luck suits degrees, I do two things in Chapter 1: (i) I argue for a consistent extension of Pritchard's modal account of luck (hereafter MAL—perhaps the most complete theory of luck to date) that suits degrees of luck; and (ii) I argue that such an extension highlights the inability of Pritchard's safety theory (the linchpin of his anti-luck epistemology) to feasibly avoid Gettier counterexamples. Though we will only focus on Pritchard's account of luck and corresponding anti-luck epistemology, *prima facie* parallel arguments can be made for divergent accounts as well. And this is precisely what we will see when we consider the prominent virtue epistemologies of Chapters 4 and 5. If the only way to avoid Gettier problems is to prohibit beliefs that are even marginally lucky from being knowledge, as I hope to demonstrate using Pritchard's account, then the diagnosis for the reductive analysis project is truly bleak indeed.

Chapter 2: Objections. Now, as I see it, there are four objections someone might level against such a diagnosis of the Gettier Problem. First, someone may level the objection that assuming an inviolable relationship between warrant and truth does not lead to radical skepticism. Likewise, someone may level the objection that assuming a close (though not inviolable) relationship between warrant and truth does not always lead to further Gettier counterexamples. Third, someone might level the objection that assuming that warrant bears no relationship to truth is not, in fact, counterintuitive. Fourth and finally, someone may level a meta-objection that calls into question the epistemic import of Gettier-style counterexamples altogether. While in the latter half of this dissertation (i.e., in Part II) we will see the proposed connection between radical skepticism and the inviolable warrant/truth relation vindicated at all the relevant instances,[20] it is nevertheless important that we guard against the latter three types of objections, especially insofar as these sorts of objections are manifest in the contemporary literature.

And that is precisely what I do in Chapter 2. First, we consider Daniel Howard-Snyder, Frances Howard-Snyder, and Neil Feit's paper, "Infallibilism and Gettier's Legacy" (2003), which directly objects to our conclusion in Chapter 1 that a close but not inviolable relationship between warrant and truth will always lead to Gettier counterexamples; I argue that their proposed solution to the Gettier Problem is, ironically, only successful insofar as it conforms to the proposed diagnosis. Second, we consider recent work by Stephen

[20] Pursuing an infallibilist response to the Gettier Problem is, after all, very common. See Daniel Howard-Snyder, Frances Howard-Snyder, and Feit 2003 and Zagzebski 1994.

Hetherington (particularly from his book *Good Knowledge, Bad Knowledge*, 2001), which proposes an analysis of knowledge that assumes a divorced relationship between warrant and truth; I argue that Hetherington's epistemology is hopelessly counterintuitive and that it need not be advocated in light of viable alternatives. Third and finally, we consider landmark papers by Brian Weatherson (2003) and Jonathan M. Weinberg, Shaun Nichols, and Stephen Stich (2003), which call into question the epistemic import of Gettier counterexamples; I argue that such criticisms either lead to intractable forms of skepticism or are ultimately incoherent.

Chapter 3: Plantinga's Virtue Epistemology. Chapter 1 advocates a diagnosis of the Gettier Problem which predicts that Gettier counterexamples cannot be feasibly avoided within the reductive analysis project. In that chapter I propose that the only way to avoid Gettier counterexamples, given the nature of luck, is to assume that warrant is either completely divorced from truth or, conversely, necessarily wed to it—neither option seeming viable. In other words, in accord with Zagzebski's 1994 diagnosis, I have argued that as long as we reasonably assume that warrant bears a close but not inviolable relationship to truth, the Gettier Problem is unavoidable (Zagzebski 1994, 65).

With this diagnosis in hand, we consider in Chapter 3 what is perhaps the most iconic virtue epistemology of the twentieth century, Alvin Plantinga's analysis of knowledge in terms of properly functioning cognitive faculties—a view that developed throughout his monumental warrant trilogy: *Warrant: The Current Debate* (1993), *Warrant and Proper Function* (1993), and *Warranted Christian Belief* (2000). Plantinga's epistemology stands in stark contrast with our proposed diagnosis and the growing trend in contemporary epistemology to abandon the analysis of knowledge—offering what is meant to be a viable reductive account that is immune to Gettier counterexamples. Our goal in Chapter 3 is to apply the diagnosis of the Gettier Problem developed in Chapter 1 to each iteration of Plantinga's epistemology and argue that, in accord with Zagzebski's diagnosis, Plantinga's reductive analysis of knowledge repeatedly becomes either unpalatable or unable to surmount the Gettier Problem. First, we will elucidate and critique Plantinga's analysis of knowledge as it is found in *Warrant and Proper Function* (1993). Second, we will elucidate and critique the proposed modifications to Plantinga's original account found, first, in "Respondeo" (1996) and "Warrant and Accidentally True Belief" (1997) and, then, in *Warranted Christian Belief* (2000). In every instance, we will find our proposed diagnosis of Gettier problems

vindicated, with each iteration and proposal failing precisely along the lines our diagnosis predicted.

Chapter 4: Sosa's Virtue Epistemology. While Alvin Plantinga's reductive analysis of knowledge is perhaps the iconic virtue epistemology of the twentieth century, Ernest Sosa's recent work in *A Virtue Epistemology* (2007) and *Reflective Knowledge* (2009) is arguably the seminal virtue epistemology of the last decade. Where Plantinga tries to satisfy the high demands of reductive analysis through an increasingly complex and intricate account of warrant, Sosa's virtue epistemology, in contrast, is simple and elegant. Knowledge, for Sosa, is *apt belief*, belief that is "true because competent" (2007, 1:23). And what is more, this straightforward analysis of knowledge is iconic in its own right—being representative (at least in form) of an approach taken by other prominent virtue epistemologists, not least Linda Zagzebski in *Virtues of the Mind* (1996) and "What is Knowledge?" (1999); and John Greco in "The Nature of Ability and the Purpose of Knowledge" (2007), "Knowledge and Success from Ability" (2009), and *Achieving Knowledge* (2010).

In Chapter 4, we will elucidate Sosa's analysis of knowledge and consider it in light of the diagnosis of the Gettier Problem proposed in Chapter 1. And insofar as Sosa's virtue epistemology has not succumbed to a series of modifications and iterations, our work in this chapter can be relatively straightforward. First, we will elucidate Sosa's theory of knowledge—unpacking its virtues and its various dimensions. Second, we will consider it in light of our proposed diagnosis, arguing that the simplicity of Sosa's analysis of knowledge is deceptive—hiding dubious ambiguity. I argue that as soon as we try to elucidate *knowledge as apt belief*, we find that under any plausible interpretation, Sosa's analysis, like Plantinga's, either succumbs to Gettier counterexamples or yields a radical and untenable form of skepticism. This, again, is in keeping with our proposed diagnosis of Gettier counterexamples.

Chapter 5: Prolegomena to Non-Reductive Virtue Epistemology. Despite the grim and seemingly hopeless history of the reductive model, the majority of work done in contemporary epistemology toward defining knowledge is nevertheless still pursued via reductive analysis.[21] The reductive analysis project is still highly prevalent; its ghost has by no means

[21] I assume that there is a difference between "defining" knowledge and providing an "analysis" of knowledge, that abandoning the reductive analysis project does not mean the abandonment of attempts to define knowledge. Just as we might still define the concept "red" without defining it in terms of a reductive analysis, I take it that we can define the concept "know" without defining it in terms of a reductive analysis.

yet been exorcised. While there is, no doubt, a growing trend to abandon the reductive analysis project, advocates of such a move are plainly in the minority. One of the primary and simple goals of this thesis is to lend credence to this trend by criticizing some of the chief proponents of what is perhaps the most popular species of reductive analysis to date, virtue epistemology.

While I advocate the trend to abandon the reductive analysis of knowledge in favor of alternative epistemic models by criticizing dominant forms of virtue epistemology, it is nevertheless my conviction that accounting for virtue within epistemology is ultimately meritorious—that there are significant advantages to pursing virtue-theoretic epistemology even if we must renounce the analyzability of knowledge. Having (i) championed contemporary virtue epistemology as one of the most sophisticated and popular species of reductive analysis to date and (ii) argued that, given the proposed diagnosis of the Gettier Problem, we have very good reason to doubt that any reductive analysis (virtue-theoretic or otherwise) can feasibly surmount Gettier counterexamples (as demonstrated in the work of Plantinga and Sosa), I nevertheless want to explore in this chapter the possibility of adapting virtue-theoretic epistemology to suit an alternative epistemic model. In Chapter 5, I first offer motivation for such a project by explicating some of what is so attractive about virtue-theoretic epistemology, giving us reason to preserve it via an alternative model. Second, I elucidate and outline the seminal alternative epistemic model to date, namely, the non-reductive model of knowledge developed by Timothy Williamson in *Knowledge and Its Limits* (2000). Third, I explore how virtue might be incorporated within a non-reductive model, generally. Finally, I reconsider the possibility of non-reductive virtue epistemology in light of Williamson's specific non-reductive epistemology and the respective virtue-theoretic concepts of Plantinga and Sosa.

PART I:
DIAGNOSING THE PROBLEM

1 LUCK AND GETTIER PROBLEMS

The dialectic has been set. In the Introduction, we elucidated two trends in contemporary epistemology: a growing dissatisfaction with the reductive analysis of knowledge, on the one hand, and the widespread popularity of virtue epistemology, on the other. And given that prominent virtue epistemologists insist on defining knowledge via reductive analysis, these trends are at odds. While my goal is ultimately, in the final chapter, to endorse *both* the move away from the reductive analysis and the move toward virtue-theoretic epistemology (via an alternative non-reductive model), I first want to explore (i) the problem with the reductive analysis project and (ii) how this applies to contemporary virtue epistemologies. In this chapter and the next, I elucidate and defend a diagnosis of what is wrong with the reductive analysis of knowledge, the reason for the growing dissatisfaction. Then, in Chapters 3 and 4, I use this diagnosis to elucidate and explain the corresponding failures of two eminent virtue epistemologies, championing them as among the most robust, seminal, and popular reductive analyses to date.

No doubt, the main reason for the growing dissatisfaction with the reductive analysis project is the Gettier Problem. The inability to define knowledge in terms of necessary and jointly sufficient conditions that are not susceptible to Gettier counterexamples is perhaps the greatest problem with reductive accounts. So, to diagnose what is wrong with the reductive analysis project, we need to diagnose the Gettier Problem; we need to understand when and why Gettier counterexamples occur. This is precisely what we will be doing in this chapter.

According to Linda Zagzebski's 1994 diagnosis of Gettier problems, if whatever we take to bridge the gap between true belief and knowledge (call it warrant) bears some violable relationship to truth, then it will be possible for that belief to be so warranted and

true for reasons unrelated to the warrant (Zagzebski 1994, 65).[22] Consider it this way. There are three relationships warrant can bear to truth (hereafter WT-relations) within any given reductive analysis:

> **Inviolable WT-relation**: Warrant necessarily entails truth. (Zagzebski 1994, 69–70)[23]
>
> **Close (but not Inviolable) WT-relation**: Warrant generally tracks truth, but not necessarily so.[24]
>
> **Divorced WT-relation**: Warrant bears no relation to truth whatsoever.[25]

Zagzebski's diagnosis, then, is simply that so long as a given reductive analysis assumes a Close WT-relation, Gettier counterexamples will be unavoidable.[26] That is to say, so long as

---

[22] While Zagzebski does uses the term "justification" instead of "warrant," I think this is the most charitable way to understand her diagnosis. Given that she wishes to apply her diagnosis to all versions of the reductive analysis and not just to the classic "justified true belief" accounts, she must surely be taking "justification" here to mean something quite broad—something like "that which with true belief makes knowledge," what we might otherwise identified as "warrant." (For other instances of "warrant" being used in this way see Plantinga 1993a, 3; D Howard-Snyder, F Howard-Snyder, and Feit 2003, 1; Pritchard 2005, 5). Say we add a strong defeasibility condition to our analysis of knowledge, one that prohibits S from knowing *p* unless there is no evidence against *p* strong enough to usurp S's belief that *p*, were S to be aware of it. Such a strong condition, epistemically infeasible as it may be (it seems like such a condition would uncomfortably reduce the amount of things that we would like to think we know), we can assume for the time being blocks all Gettier cases. As such, if we were to give Zagzebski's use of justification a more narrow reading, one that signified only, say, an internalistic condition on knowledge, then her thesis is simply false: we could assume that justification, so construed, is closely related to the truth and still avoid Gettier cases by additional conditions such as a strong defeasibility condition. Additionally, given that warrant is something that comes in degrees, I take it that any *warranted true belief* analysis of knowledge assumes a sufficient degree of warrant.

[23] See Zagzebski 1994, 69-70.

[24] See Zagzebski 1994, 73.

[25] See Zagzebski 1994, 72. While this is not a terribly popular view, it does have its advocates. In his book, *Good Knowledge, Bad Knowledge* (2001), Stephen Hetherington argues explicitly for an analysis of knowledge that assumes a divorced relationship between truth and warrant. We will consider Hetherington's account in Chapter 2.

[26] Zagzebski's analysis, I think, rightly identifies the target of Gettier cases as conceptual necessity instead of metaphysical necessity as Timothy Williamson has argued in *The Philosophy of Philosophy* (2007). As Jonathan Schaffer pointed out in his presentation, "Modalities and Methodologies," viewing the target of Gettier cases as metaphysical is "too weak" such that "[even] if Gettier cases were metaphysically impossible, [Gettier's]

we reasonably assume that warrant is neither divorced from truth nor inseparable from it (such options being epistemically unpalatable), Gettier cases are "inescapable" (Zagzebski 1994, 65).

As eloquent and intuitive as Zagzebski's diagnosis of Gettier problems may be, the prominent advocates of virtue epistemology have apparently not taken it altogether seriously. If Zagzebski is right, this has very dim consequences for any palatable reductive analysis of knowledge; however, virtue epistemologists continue to work from such a model. Though there is a common trend in contemporary epistemology to abandon the standard analysis of knowledge in favor of alternative models, this trend has not yet extended to those working within virtue epistemology.

In this Chapter, I add credence to Zagzebski's diagnosis (and the movement away from the standard analysis) by analyzing the nature of luck. It is widely accepted that the lesson to be learned from Gettier problems is that knowledge is incompatible with luck or at least a certain species thereof.[27] As such, understanding the nature of luck is central to understanding the Gettier Problem.[28] Thanks by and large to Duncan Pritchard's seminal work, *Epistemic Luck* (2005), a great deal of literature has been developed recently concerning the nature of luck and anti-luck epistemology. The literature, however, has yet to explore the very intuitive idea that luck comes in degrees.[29] I propose that once luck is recognized to admit degrees, even the slightest non-zero degree of luck (of the relevant sort) precludes knowledge. Connecting this to Zagzebski's thesis, I propose that a given theory of warrant must guarantee truth in order to avoid Gettier counterexamples (or subsequently deny that warrant bears any relationship whatsoever to the truth), simply because a sufficient reductive analysis of knowledge (whether a virtue epistemology or not) cannot allow for knowledge that is even marginally lucky.

---

counterexample to the proposed definition of knowledge would still stand" (Schaffer 2009). See Jackson 2009; Williamson 2009a.

[27] See Dancy 1985, 134; Zagzebski 1999, 99-101; Pritchard 2005, 4-5; Steup 2006, §2. In Riggs 2007, Wayne Riggs explores some of the reasons it is so important that knowledge preclude (the relevant sort of) luck.

[28] While there are well over 100 different Gettier cases, they are united by common theme: a warranted belief that is only luckily true. As L. Floridi notes, though there are many cases, nobody would expect them to be answered in a case-by-case basis (2004, 64).

[29] Someone may occasionally elucidate safety as coming in degrees (see Smith 2009), but this is the closest, to my knowledge, philosophers have come to describing luck in terms of degrees.

I justify this proposal by treating Pritchard's work as archetypal. Motivated by the intuition that luck suits degrees, I do two things in this chapter: 1) argue for a consistent extension of Pritchard's modal account of luck (hereafter MAL—perhaps the most complete theory of luck to date) that suits degrees of luck; 2) argue that such an extension highlights the inability of Pritchard's safety theory (the linchpin of his anti-luck epistemology) to feasibly avoid Gettier counterexamples. Though we will only focus on Pritchard's account of luck and corresponding anti-luck epistemology, *prima facie* parallel arguments can be made for divergent accounts as well. And this is precisely what we will see when we consider the prominent virtue epistemologies of Chapters 3 and 4. If the only way to avoid Gettier problems is to prohibit beliefs that are even marginally lucky from being considered knowledge, as I hope to demonstrate using Pritchard's account, then the diagnosis for the reductive analysis project is truly bleak indeed.

Section 1: Degrees of Luck

In *Epistemic Luck*, Duncan Pritchard notes that in the contemporary literature luck is often conflated with accidents, chance, or a lack of control; however, whatever luck is, Pritchard notes that such conceptions of luck do not sufficiently characterize it. For example, if S is purposely playing the lottery, it may be lucky if S wins, but it would be strange to call such a lottery win an *accident* (Pritchard 2005, 126). The problem with conflating luck with *chance*, according to Pritchard, is that most of us would say that in order for a given event to be lucky it must affect an agent; with chance, however, no agent needs to be involved. It may be a matter of *chance* that a landslide did or did not occur on such and such a mountain, but most of us would not call such an event a matter of luck if no one was affected (Pritchard 2005, 126). Pritchard notes that conflating luck with an absence of control is perhaps the most common characterization of luck. This is due to an influential paper by Thomas Nagel in which he so defines a particular species of moral luck (Pritchard 2005, 127).[30] Such a characterization, however, has been recognized to be at best a necessary condition on (moral) luck; it is, after all, out of our *control* that the sun rose this morning, but few of us would describe its rising as lucky (Pritchard 2005, 127).[31] What is more, typifying luck as a lack of control is doubly problematic when it comes to *epistemic* luck, given that many

[30] See Nagel 1979, 25. Similar accounts can be found in Zimmerman 1993; Greco 1995.

[31] Also see Statman 1991, 146; Latus 2000, 167.

epistemologists would adhere to doxastic involuntarism for a whole host of perceptual beliefs (Pritchard 2005, 127). It seems reasonable to think that I have no control over whether or not I believe that "I am now being appeared to redly" when I am so appeared to, but few would identify such a belief as lucky. In light of this, Pritchard notes that though it may seem to be intuitive to characterize luck in terms of accidents, chance, or a lack of control, "there is no straightforward way available of accounting for luck in these terms"(Pritchard 2005, 127).

Pritchard puts forward MAL as a theory of luck that appropriately tracks our intuitions across the relevant cases. It consists of two conditions:

> **L1:** If an event is lucky, then it is an event that occurs in the actual world but which does not occur in a wide class of the nearest possible worlds where the relevant initial conditions for that event are the same as in the actual world. (Pritchard 2005, 128)

> **L2:** If an event is lucky, then it is an event that is significant to the agent concerned (or would be significant, were the agent to be availed of the relevant facts). (Pritchard 2005, 132)

MAL seems to accurately identify paradigm cases such as lottery wins as lucky (contra an *accident* conception of luck), and it also seems to prevent sunrises, being appeared to redly, and isolated (insignificant to an appropriate agent) landslides, from being deemed lucky as well (contra *chance* and *lack of control* conceptions of luck). To be sure, Pritchard notes that there is an inherent vagueness to MAL, though we do seem to have a good intuitive grasp as to how L1 and L2 are meant to function. Pritchard is willing to let us be guided by our intuitions, and, given our current purposes, so am I.

* * *

Aside: In "Pritchard's Epistemic Luck" (2006) and "What Luck is Not" (2008), Jennifer Lackey raises an objection to MAL through a counterexample she calls Buried Treasure:

> **Buried Treasure:** Sophie, knowing that she had very little time left to live, wanted to bury on the island she inhabited a chest filled with all of her earthly treasures. As she walked around trying to determine the best site for proper burial, her central criteria were, first, that a suitable location must be on the northwest corner of the island, where she had spent many of her fondest moments in life, and secondly, that it had to be a spot where rose bushes could flourish, since these were her favorite flowers. As it happened, there was only one particular patch of land on the northwest corner of the island where the soil was rich enough for roses to thrive. Sophie, being excellent at detecting such soil, immediately located this patch of land and buried her treasure, along with seeds for future roses to bloom, in the one and only spot that fulfilled her two criteria. One month later, Vincent, a distant neighbour of Sophie's, was driving in the northwest corner of

> the island, which was also his most beloved place to visit, and was looking for a place to plant a rose bush in memory of his mother who had died ten years earlier, since these were her favourite flowers. Being excellent at detecting the proper soil for rose bushes to thrive in, he immediately located the same patch of land as Sophie had found one month earlier. As he began digging a hole for the bush, he was astonished to discover a buried treasure in the ground. (Lackey 2006, 285)[32]

The fact that Vincent found the buried treasure intuitively seems lucky even though it does not appear to meet L1. According to Lackey, Vincent finds the treasure both in the actual world and in a wide range of nearby possible worlds. If one is not convinced of this being the case, Lackey notes that it can be modified to make this more apparent without doing damage to our initial intuition (e.g., the topography of the island is invariant; the only flower that Sophie and Vincent's mother have ever liked is roses; Sophie has always had this specific detailed plan to bury her possessions once she was informed of her illness; etc.) (Lackey 2006, 286). To be sure, Buried Treasure does not seem to be an isolated case. According to Lackey, to make additional cases all someone has to do, roughly, is pick a paradigmatic instance of luck and then "construct a case involving such an event in which both its central aspects are counterfactually robust, though there is no deliberate or otherwise relevant connection between them," then modifying the case as need be so that the lucky event is bound to happen in all (or most) nearby possible worlds.

Though Lackey takes Buried Treasure to be a clear counterexample to MAL, it seems to me that it rests on a dubious conception of how possible worlds are ordered. Lackey assumes (according to my reading of her) that the closeness of possible worlds should be determined by whether, downstream of some relevant initial conditions, a given event was bound or determined to happen. But why should we think a thing like that? Pritchard does not defend a method for judging the closeness of worlds and neither does Lackey, so until we have independent reason to think the closeness of possible worlds should be judged how Lackey assumes, it is not clear that Buried Treasure even offers a counterexample to MAL let alone a reason to think that MAL is "fundamentally misguided."

To be sure, it is not just that the defender of MAL has no reason to accept Lackey's assumption concerning how possible worlds are ordered, it is that the defender of MAL has very good reason to reject it. Assuming that the closeness of a world is indeed discerned by whether, downstream of some relevant initial conditions, a given event was bound or determined to happen not only leads to hairy cases such as Buried Treasure, it

[32] Also see Lackey 2008, 261-262.

undermines some of MAL's principal goals. For example, one of Pritchard's chief objectives in proposing MAL is to better understand the luck involved in Gettier counterexamples; however, once Lackey's assumption has been made, it looks like someone could make a luck-less Gettier case. For example, if we were to expand on a classic Gettier case such that the relevant events were bound to happen (Smith was just bound to lie to Jones about owning a Ford, etc.), it would no longer exhibit luck according to Lackey's reading of MAL. This is a bizarre conclusion to say the least—after all, it is almost universally agreed that the lesson to be learned from Gettier cases is that knowledge is incompatible with (at least a species of) luck. So insofar as the defender of MAL would want to say that classic cases exhibit luck even if the relevant events were predetermined somehow, I think we have not only a blatant reason for thinking that the defender of MAL should reject Lackey's assumption but, what is more, that Lackey's assumption simply was never a part of MAL even at its conception.

The reason we think Vincent is lucky to find the treasure is not based on the event's causal indeterminacy but simply on the fact that planting roses does not usually yield treasure—planting roses is not apt for finding treasure. Instead of determining modality in MAL in terms of causality we should, I think, determine modality in terms of something like aptness. Consider Buried Treasure once again. We could indeed evaluate modality in terms of the causal-context provided in Buried Treasure, which does seem to yield the result that Vincent's specific instance of rose planting that yielded treasure would indeed occur in most close possible worlds. As such, given that we take Vincent's finding the buried treasure to be lucky, this seems to be contra MAL. However, we can instead evaluate modality in terms of the likelihood or aptness of any given rose-planting event yielding treasure, in which case the fact that Vincent found treasure by planting roses is an event that would occur in relatively few close possible worlds (in accord with MAL). In other words, take X_n to be the set of all rose-planting events, of which X_1, Vincent's specific rose-planting event, is one of them. Even if X_1 was predetermined to happen via some story of causal relations, we can still judge X_1 as an event that would occur in few close possible worlds in light of X_n. We take rose planting to be an inapt method for finding buried treasure (no one, barring special cases, would plant roses in hopes of making such a find), and so, insofar as this plays a large role in our luck ascriptions in Buried Treasure, we have reason to determine the modality of X_1 in terms of X_n and not how

Lackey is assuming. Of course, Lackey might reply with a case like Buried Treasure Addendum:

> **Buried Treasure Addendum:** The people in Vincent and Sophia's world have a subconscious desire to bury their treasure in locations that have such and such a smell; these locations just so happen to perfectly correspond to places with soil suitable for planting roses. After generations and generations of people with this subconscious desire, almost every place in the world suitable for planting roses has treasure buried there. Planting roses, then, is an activity that is very apt for finding treasure, though, we should add, no one is aware of this—no one has discovered the correspondence between rose soil and treasure. What is more, all of this is in addition to the circumstances outlined in Buried Treasure (e.g., Vincent and Sophia live on an island with only one place suitable for roses (which, we shall assume, was one of the very few such places without treasure already when it came time for Sophia to bury hers), both Sophia and Vincent's mother have always liked only roses, Sophia and Vincent had planned on taking their respective actions long in advance, etc.).

If we think Vincent is lucky in Buried Treasure Addendum we have a counterexample to the reading of MAL in which we judge modality in terms of aptness; however, I think we can comfortably deny that Vincent is lucky here. Finding the buried treasure may seem lucky to Vincent, which it surely would, but this does not mean it was a truly lucky event. Becoming violently ill may seem terribly unlucky to the first person who ate a Jack-O'-Lantern mushroom (also known as Omphalotus olearius; a mushroom that tastes and smells edible and indeed delicious), but that doesn't mean that it is—as a matter of fact, when you eat Jack-O'-Lantern mushrooms you get sick. Hence, quintessentially lucky events (which Lackey takes finding buried treasure to be) may not necessarily be lucky.

* * *

MAL is a theory of luck; it tells us when a given event is lucky. As such, it is simply not meant to tell us how lucky a given event is—it is not meant to tell us how luck admits degrees. Nevertheless, given that it is indeed very intuitive that luck comes in degrees, we have an interest in exploring a suitable extension of MAL that will do just this.[33] To be sure,

[33] Though the intuition that luck suits degrees can warrant an extension in and of itself, perhaps we can say a bit more in its favor. Say Jack is playing three lotteries—each of which has a jackpot prize of £1,000,000. Jack has 1 out of 1,000,000 chances of winning the jackpot of Lottery A. Jack has 1 out of 1 chances of winning the jackpot of Lottery B (Jack rigged it). Finally, Jack has 1 out of 2 chances of winning Lottery C (he bought half the tickets). As it happens, Jack wins all three lotteries. We can probably all agree that Jack's winning Lottery A is lucky. Additionally, we can probably all agree that Jack's winning Lottery B is not lucky. However, it is far from clear whether or not it was lucky that Jack won Lottery C. If we say that it was lucky, what would we say if Jack had 2 out of 3 chances of winning? Or 3 out of 4? If winning a lottery with 1 out

the most philosophically interesting part of Pritchard's *modal* account of luck is, unsurprisingly, its modal condition. As such, in developing an extension of MAL to account for luck coming in degrees, my focus will be Pritchard's modal condition, L1—generally leaving Pritchard's significance condition, L2, as an assumed necessary condition for luck.

But how should degrees of luck be modeled? At first blush, one may think degrees should be modeled according to the number of nearest possible worlds in which the event in question occurs, perhaps extending L1 like this:

> **L1-I**: An event will be lucky according to the degree to which it is an event that occurs in the actual world but does not occur in a wide class of nearest possible worlds where the relevant initial conditions for that event are the same as in the actual world.

That is to say, an event, E, will be lucky to degree D, where D varies directly with the proportion of nearby possible worlds where the relevant initial conditions for E are satisfied and E does not occur. L1-I rightly predicts, for example, that winning a lottery where you had 1 out of 1,000,000 chances is luckier than winning a lottery where you had 1 out of 1,000 chances simply because there will be (so it seems) proportionally fewer close possible worlds in which you win in the former case than in the latter. Changing the game, L1-I similarly predicts that winning a round of roulette by betting on black is less lucky than winning a round of roulette by betting on black 28 simply because there will be (so it seems) proportionally fewer close possible worlds in which you win in the latter than in the former. As such, L1-I appears to rightly ascribe luckiness in such paradigmatic cases; however, on further reflection, L1-I no longer seems sufficient. Consider the following case:

> **Haven:** Haven is an exceedingly safe place. Haven is so safe that it is incredibly unlikely that anyone visiting will be shot. Visitor comes to Haven and is able to leave without being shot.

For the sake of argument, let us say that Haven is so safe that when Visitor visits there are no nearby possible worlds in which he is shot; though, perhaps, there *are* distant worlds

---

of 1,000,000 odds is lucky and winning a lottery where you have 1 out of 1 is not, at what point does a lottery win stop being lucky? Whatever answer we pick, it will seem to be arbitrary. Instead of trying to decide when a lottery win becomes lucky, I think our intuition is simply that winning a lottery where you have 1 out of 2 chances is far less lucky than winning a lottery where you have 1 out of 1,000,000 chances. To account for this, perhaps we should extend L1 accordingly.

where Visitor is so harmed in Haven. Nevertheless, according to L1-I, Visitor's not being shot exhibits *zero*-degrees of luck; after all, in none of the nearest possible worlds is Visitor shot when he comes to Haven. But this does not seem correct—surely even distant possible worlds affect an event's luckiness. For example, say that Visitor came to a *perfectly* safe place—call it Heaven—a place where he remains unshot in all possible worlds both near and far. Seemingly, Visitor's not being shot in Haven is luckier than his not being shot in Heaven, yet L1-I does not account for this.[34]

As such, we should look to all possible worlds with the relevant initial conditions and not just the set of nearest ones—perhaps extending L1 like this:

> **L1-II**: An event will be lucky according to the degree to which it is an event that occurs in the actual world but does not occur in other possible worlds where the relevant initial conditions for that event are the same as in the actual world.

That is to say, an event, E, will be lucky to degree D, where D varies directly with the proportion of other possible worlds (*both near and far*) where the relevant initial conditions for E are satisfied and E does not occur. L1-II not only appropriately handles lottery/roulette cases, but it also rightly predicts that remaining unshot in Heaven is less lucky than remaining unshot in Haven. However, an additional clarification needs to be made. Though we should consider all of the relevant possible worlds in modeling degrees of luck, surely closer possible worlds should be weighted more heavily. For example, consider the following case:

> **Ill-Fated Kangaroos:** If a kangaroo loses its tail, it will presumably not be able to walk in nearly all close possible worlds—though in distant possible worlds, let's assume, kangaroos know how to use crutches and can in fact walk without tails. If a kangaroo loses an ear, presumably it will continue to be able to walk in nearly all close possible worlds—though in distant possible worlds, let's assume, ears play a more central role in a kangaroo's ability to walk. Kangaroo-Jim loses his tail and is still able to walk. Kangaroo-Bill loses an ear and is still able to walk.[35]

---

[34] In light of this objection, someone may propose that luck varies according to how different the world would have to be in order for the given event not to occur. Though modeling degrees of luck in this way provides a way to make the distinction between Visitor's visits to Haven and Heaven, it runs into other problems. For example, say I am blindly drawing marbles out of a bag where there is a 10-to-1 ratio between black marbles and red marbles. Is the possible world in which I draw a black marble any farther away than the possible world in which I draw a red one? Seemingly not.

[35] This case was inspired by the opening pages of David K. Lewis's *Counterfactuals* (2000).

Everything else being equal, an event that occurs in close possible worlds (Kangaroo-Bill being able to walk without an ear) is surely less lucky than an event that occurs only in more distant worlds (Kangaroo-Jim being able to walk without a tail)—the former occurring more readily than the latter.

As a final first approximation, the extension of L1 that seems to rightly model degrees of luck looks something like this:

> **L1-III:** An event will be lucky according to the degree to which it is an event that occurs in the actual world but does not occur in other possible worlds (where the relevant initial conditions for that event are the same as in the actual world) such that closer worlds are weighted more than distant worlds.

Consider the following diagram, where A signifies the actual world, P1-3 signify the sets of relevant possible worlds, and the dash-circle signifies the set of "the nearest possible worlds":

Figure 1.1

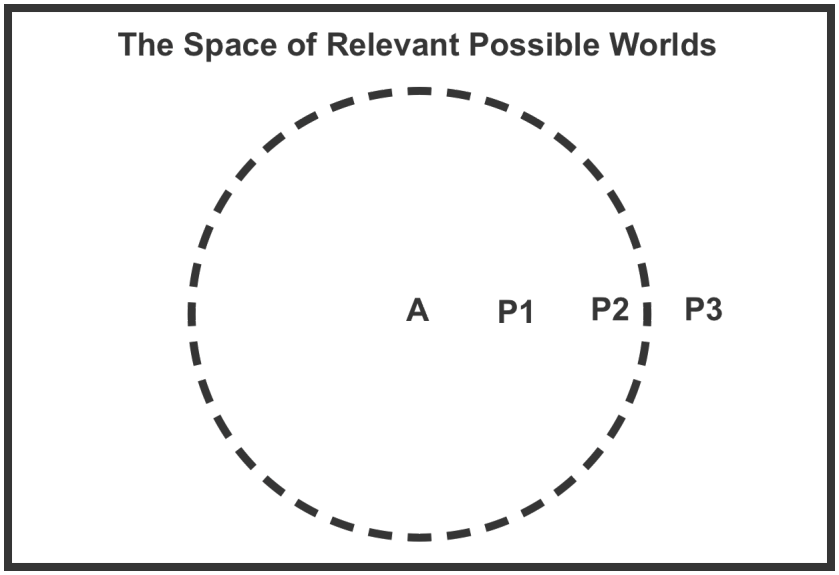

In Figure 1.1, P1, P2, and P3 are arranged according to their similarity to A, where closer proximity to A signifies a greater similarity to A. Let us say that E1, E2, E3, E4, and E5 are events that occur in A and which have, let us suppose, the same relevant initial conditions. Additionally, assume that these are conditions that P1, P2, and P3 meet. In addition to occurring in A, E1 occurs in P1, P2, and P3. In addition to occurring in A, E2 occurs in P1 and P2. In addition to occurring in A, E3 occurs in P1. In addition to occurring in A, E4 occurs in P2. Finally, E5 occurs only in A (see Table 1.1). For our consideration, let us also say that P1, P2, and P3 account for all possible worlds where the relevant initial conditions

Table 1.1

Event	Possible Worlds
E1	A, P1, P2, P3
E2	A, P1, P2
E3	A, P1
E4	A, P2
E5	A

are the same. Ordering these events from highest to lowest degree of luck (from luckiest to least lucky), they are: E5, E4, E3, E2, E1. Understandably, E5 is the luckiest event because it occurs in the actual world but in no other relevant possible world. E4 is the second luckiest event because, though it occurs in P2, it does not occur in the closest possible worlds P1—as such, it would have been easier for E4 not to have occurred than E3, which does occur in the closest possible worlds. E2 is the second to least lucky as it occurs in all but the most distant P3 worlds. E1, of course, is the least lucky as it occurs in all possible worlds.

* * *

Aside: Though unnecessary for the argumentation above, it is worth noting that one could also see luck as varying in L2, in accord with how significant the given event is to the protagonist in question. Seemingly, finding £5 on the street is lucky, but perhaps finding £100 is much luckier, even if both events occur in the same proportions of relevant possible worlds. As such, we could modify L2 accordingly:

> L2-I: An event is lucky according to the degree to which it is an event that is significant to the agent concerned (or would be significant, were the agent to be availed of the relevant facts).

Finally, given that the relationship between L1-III and L2-I is dynamic, luckiness can be graphed accordingly (where luck increases with shading):

Figure 1.2

Our intuitions may be vague on the parameters of what counts as lucky in such an analysis, but such vagueness seems to be a part of how we use the term. So the fact that this revised analysis captures this should, I think, go to its credit. What is more, it is worth noting that such an all-permeating conception of luck (after all, by this analysis most events no matter how modally probable or only remotely significant to an agent will be deemed to have traces of luck) matches Zagzebski's diagnosis of the Gettier Problem, which saw luck as irreparably saturating the human condition.[36]

* * *

Although more could easily be said, I think we can take L1-III as a more-or-less healthy extension for MAL (hereafter MAL-III) that accounts for our intuition that luck admits degrees. Not only does MAL-III tell us why some events are luckier than others in paradigmatic cases such as lottery wins, but it also helps us assess minute degrees of luckiness. However, given that luck is the central and fundamental component of Gettier cases, how we think of luck, be it in terms of MAL alone or in terms of MAL with a degrees extension, will affect how we attempt to surmount Gettier problems. In the next section, we will consider Pritchard's safety theory in light of the MAL-III extension and test how successful it is at precluding a specific Gettier case.

Section 2: Pritchard's Anti-Luck Epistemology

In developing his safety theory, Pritchard is trying to establish an analysis of knowledge that avoids Gettier counterexamples. It is almost universally agreed that the lesson to be learned from Gettier cases is that knowledge is incompatible with luck—that luck is the central and fundamental component of all Gettier cases.[37] Naturally enough, therefore, how Pritchard conceives of luck in general (MAL) will play a critical role both in how he conceives of the problem raised by Gettier cases and how it might be surmounted via safety.

Gettier cases, according to Pritchard, are caused by a specific species of luck (Pritchard 2005, 145–148). As such, the lesson to be learned from Gettier cases is not so much that knowledge is incompatible with luck *simpliciter*, but rather that knowledge is incompatible

[36] See Zagzebski 1999, 102. To be sure, whether or not luck varies in accord with significance, a full account of degrees *must* incorporate something like a modal condition.

[37] See Dancy 1985, 134; Zagzebski 1999, 99-101; Pritchard 2005, 4-5; Steup 2006, §2.

with this particular *species* of it. By these lights, any successful analysis of knowledge, therefore, must (at the very least) track knowledge ascriptions in accord with this species' absence.

Pritchard calls this species of luck that is behind Gettier cases "Veritic Epistemic Luck." According to Pritchard, a given agent's belief exhibits Veritic Epistemic Luck when the following description is met:

> **Veritic Epistemic Luck**: It is a matter of luck that the agent's belief is true. (Pritchard 2005, 146)

In other words, in light of MAL, a given agent's true belief exemplifies Veritic Epistemic Luck when: 1) in a wide class of nearest possible worlds the belief is false, and 2) the belief's truth or falsity is significant to the agent. To be sure, Veritic Epistemic Luck is *not* meant to refer to cases where it is a matter of luck that *the propositional content* of an agent's belief is true. Surely, such luck *is* compatible with knowledge; surely, for example, we can have knowledge of who won the lottery, where lightning struck, what number rolling a die produced, etc. Given that Veritic Epistemic Luck is meant to be the species of luck behind Gettier cases and subsequently incompatible with knowledge, Veritic Epistemic Luck must refer to something else. Pritchard goes on to elucidate what Veritic Epistemic Luck demands, namely, that

> the agent's belief is true in the actual world, but that in a wide class of nearby possible worlds in which the relevant initial conditions are the same as the actual world—and this will mean, in the basic case that the agent at the very least *forms the same belief in the same way* as in the actual world—the belief is false. (Pritchard 2005, 146 - emphasis mine)

So in order for something to be an instance of Veritic Epistemic Luck, not only does it have to be a matter of luck that the agent's belief is true, but it has to be lucky given the way it was formed. By Veritic Epistemic Luck, Pritchard means something like "It is a matter of luck, *given the way the belief is formed*, that the belief is true."

* * *

Aside: Although it is almost universally agreed that the lesson to be learned from Gettier cases is that knowledge is incompatible with luck—that luck is the central and fundamental component of all Gettier cases—Pritchard goes on to show that there are at least four species of epistemic luck that do not preclude knowledge. The first of these is Content Epistemic Luck:

> **Content Epistemic Luck:** It is lucky that the proposition is true. (Pritchard 2005, 134)

An event's being lucky does not impede it from being knowable. For example, Eli is walking down Placid Lane, a calm suburban road, and sees a car crash into a tree. To be sure, there are hardly ever car accidents on Placid Lane. It may be a matter of luck that a given car accident occurred; this, however, does not prevent agents from knowing that it occurred. Eli's belief that the accident occurred is warranted and only luckily true, but the luck at issue is epistemically inconsequential. Pritchard identifies the second species of benign luck as Capacity Epistemic Luck:

> **Capacity Epistemic Luck:** It is lucky that the agent is capable of knowledge. (Pritchard 2005, 134)

If it is somehow lucky that a given agent has the capacity to know a given belief, this does not prevent that agent from knowing it. Using Pritchard's example, say Jones is walking through the forest and only narrowly avoids being smacked in the face with a branch that would have blinded him (perhaps he bent down to tie his shoe right when the branch swung by); it is, therefore, lucky that Jones has the capacity to see, but this luck does not thwart Jones's future perceptual beliefs from being known (Pritchard 2005, 135). The third species of benign epistemic luck is Evidential Epistemic Luck:

> **Evidential Epistemic Luck:** It is lucky that the agent acquires the evidence that she has in favour of her belief. (Pritchard 2005, 136)

Smith just so happens to walk by his employer's door and overhear that he is going to be fired, which, let us say, is true. It is, then, a matter of luck that Smith has the evidence that he has in favor of his belief that "I am going to be fired," but such luck does not preclude Smith from knowing such a belief.[38] Finally, the fourth species of benign epistemic luck that Pritchard identifies is Doxastic Epistemic Luck:

> **Doxastic Epistemic Luck:** It is lucky that the agent believes the proposition (Pritchard 2005, 138).

Not only is it lucky that Smith overhears that he is going to be fired when he just so happens to walk by his employer's door, it is also lucky that he forms the belief that he is going to be fired. He would not have formed the belief that he is going to be fired in

[38] A version of this example was developed in Unger 1968, 159, and was referenced by Pritchard in 2005, 136.

relevant nearby possible worlds and, presumably, this realization is significant to him. As Pritchard notes, it does not look like a given event can exhibit Evidential Epistemic Luck without exhibiting Doxastic Epistemic Luck and vice versa—at least not without being contentious. For example, we might say that perception immediately forms beliefs, such that an agent who luckily perceived some evidence may be evidentially lucky but not doxastically; however, such a model only seems compatible with an externalist point of view. Conversely, we might say that a given agent had all the evidence needed for such and such a belief but did not believe it until by luck some non-evidential stimulus caused the agent to draw the appropriate conclusion; however, can evidence for X count as evidence for X if the agent fails to recognize it as such? It is not at all clear.[39] As with Evidential Epistemic Luck, Doxastic Epistemic Luck too seems epistemically benign—Smith can know he is going to be fired even if he exhibits Doxastic Epistemic Luck.

There are, to put it roughly, three things that typically go into belief formation: a believer, reasons for belief (mediate or immediate; internal or external), and the belief itself (along with its truth or falsity). (By "reasons," I mean to capture any relationship a given believer may have to the genesis of his or her belief. For example, if John has a brain lesion that causes him to form the belief "I am being appeared to redly," John may not have any reflective access as to why he believes he is being so appeared to, but nevertheless the brain lesion itself stands as the reason John has the belief.) Consider the following diagram:

Table 1.2

Believer	→	Reasons	→	Belief
A	B	C	D	E

On such a model, again to put it roughly, there are five places luck can reside: A) with the believer, B) with the believer's relationship with the reasons for belief, C) with the reasons, D) with the relationship between the believer's reasons and the belief itself (its truth or falsity), or E) with the belief itself. As it happens, these benign species correspond to all but one of these locations: location A with Capacity Epistemic Luck, location B with Doxastic Epistemic Luck, location C with Evidential Epistemic Luck and location E with Content Epistemic Luck; leaving us with location D, which is where, I think, Veritic Epistemic Luck is meant to reside.

* * *

[39] See Pritchard 2005, 136-141.

By these lights, to avoid Gettier cases entirely one needs to have an epistemology that does not attribute knowledge to beliefs that exhibit Veritic Epistemic Luck—luck that, given the way the belief was formed, the belief is true. So, given MAL and Pritchard's analysis of Gettier problems, he fashions his safety theory to do just this.

> **Pritchard's Safety Theory**: For all agents . . . if an agent knows a contingent proposition ψ, then, in nearly all (if not all) nearby possible worlds in which she forms her belief about ψ in the same way as she forms her belief in the actual world, that agent only believes that ψ when ψ is true. (Pritchard 2005, 163)

Once given MAL, such that something is lucky if it meets L1 and L2, and Pritchard's understanding of what species of luck is behind Gettier problems, Pritchard's Safety Theory is a natural solution. It is specially designed to prevent any belief from being deemed knowledge that is only luckily (as described in MAL) connected to the reasons behind its genesis.[40] Such a belief could be so believed by an agent in nearly all (if not all) nearby possible worlds in which it is formed in the same way (e.g., for the same reasons) as in the actual world and yet be false. When a belief exhibits Veritic Epistemic Luck, Pritchard's Safety Theory will not be met; thus, according to Pritchard's analysis, such a belief will not be knowledge. Consider the following classic case:

> **Classic Case**: Smith and Jones are applying for the same job. Smith has very strong evidence for thinking that Jones will get the job (e.g., the employer tells Smith that he will hire Jones, etc.), and for thinking that Jones has 10 coins in his pocket (e.g., Jones emptied his pockets in front of Smith and then clearly, slowly, in good lighting, and perhaps even counting out loud, placed 10 coins in his pocket). As such, Smith forms a belief in the general proposition that "the man who gets the job has 10 coins in his pocket." As it turns out, however, Smith gets the job and, unbeknownst to Smith, he happens to also have 10 coins in his pocket. (paraphrased from Gettier 1963, 122)

Smith's belief may very well be justified and true, but because it would not still be true in many nearby possible worlds in which Smith forms the belief in the same way (worlds where, for example, Smith has more or less than 10 coins in his pocket), it is not safe and therefore not an instance of knowledge. Also consider Alvin Goldman's Fake Barns case:

> **Fake Barns**: Henry is driving in the country with his son. For the boy's edification Henry identifies various objects on the landscape as they come into view. "That's a cow," says Henry, "That's a tractor," "That's a silo," "That's a barn," etc. Henry has no doubt about the identity of these objects; in particular, he has no doubt

[40] For worries regarding Pritchard's account of safety and its relationship to epistemic virtue, see Greco 2007b.

> that the last-mentioned object is a barn, which indeed it is. Each of the identified objects has features characteristic of its type. Moreover, each object is fully in view, Henry has excellent eyesight, and he has enough time to look at them reasonably carefully, since there is little traffic to distract him. . . . Suppose we are told that, unknown to Henry, the district he has just entered is full of papier-mâché facsimiles of barns. These facsimiles look from the road exactly like barns, but are really just façades, without back walls or interiors, quite incapable of being used as barns. They are so cleverly constructed that travelers invariably mistake them for barns. Having just entered this district, Henry has not encountered any facsimiles; the object he sees is a genuine barn. But if the barn on that site were a facsimile, Henry would mistake it for a barn. (Goldman 1976, 772–773)[41]

Henry's belief may very well be justified and true, but because it would not still be true in many nearby possible worlds in which Henry forms the belief in the same way (e.g., worlds where he is actually looking at a barn façade), it is not safe and therefore, so it goes, not an instance of knowledge.

According to *Epistemic Luck* (2005), in order for belief to be knowledge it must be true and safe as per Pritchard's Safety Theory. Such an analysis of knowledge, once given MAL and Pritchard's diagnostics of Gettier cases, is meant to be Gettier-proof. To be sure, we may be inclined to agree that once this much has been accepted, Pritchard has indeed established a version of the standard analysis that is immune to Gettier cases; however, as we shall soon see, adopting MAL-III calls this into question—having serious ramifications for the success of Pritchard's analysis at avoiding Gettier counterexamples.

Pritchard, relating Veritic Epistemic Luck back to his general conception of luck (MAL), says that it arises when a given "agent's belief is true in the actual world, but . . . in *a wide class of nearby possible worlds* in which the relevant initial conditions are the same as in the actual world . . . the belief is false" (Pritchard 2005, 146 - emphasis mine). However, given that luck comes in degrees, Veritic Epistemic Luck is far more insidious than that. Veritic Epistemic Luck crops up not only if the belief in question is false in "a wide class of nearby possible worlds," but rather, to at least some degree, if the belief in question is false

[41] Also see Zagzebski 1994, 66. Notably, Duncan Pritchard (2010) does not consider Fake Barn to be a Gettier case, because the protagonist (Henry in our case) does not make a "cognitive error." See Pritchard, Millar, and Haddock 2010, 35-36. According to Pritchard, in order for a case to be a Gettier case a cognitive error has to be committed by the protagonist in question. This is, I think, a rather arbitrary qualification—historically, a Gettier case has been any case in which the necessary conditions of a reductive analysis of knowledge seem to be met without being sufficient as a result of the given warrant being only luckily connected to the truth. Regardless, I will proceed, along with the majority of the contemporary literature, assuming that cases like Fake Barn are indeed Gettier cases.

in any (or nearly any) possible world whatsoever! According to MAL-III, Veritic Epistemic Luck comes in degrees. I may be very gullible and believe some dastardly fellows who, for fun, tell me there is buried treasure in my backyard, such that when my belief turns out to be true (there indeed happens to be treasure in my backyard, unbeknownst to my fellows) it exhibits a large degree of Veritic Epistemic Luck. There are any number of possible worlds, given my gullibility, where I believe whatever dastardly fellows tell me and it is false. On the other hand, I may be very cold and skeptical and (eventually) believe the testimony of legitimate experts of lost treasure, who tell me that there is treasure in my backyard based on such and such legitimate data (maps, pirate diaries, descriptions of the treasure's location, etc.) such that my belief, which happens to be true for complex reasons unrelated to the expert's data analysis, is only marginally exhibiting Veritic Epistemic Luck. As we might imagine, there are relatively very few possible worlds in which I form a calculated belief based on the honest testimony of experts in which it happens to be false.

Crucially, as we saw from the Haven/Heaven case, distant possible worlds seem to affect how lucky a given event is. The question then becomes whether or not minute degrees of Veritic Epistemic Luck are of any epistemic consequence.[42] A safety theorist like Pritchard may very well agree that luck comes in degrees, but deny that anyone fails to know a given belief when that belief is believed in the same way and yet false in *distant* possible worlds. Consider a case originally developed by Linda Zagzebski in "The Inescapability of Gettier Problems" (1994):

> **Dr. Jones and the Virus:** Smith is ill and exhibits a unique set of symptoms, S. Given these symptoms, Dr. Jones forms the belief that "Smith has Virus X," which she deduces from the true proposition that "Virus X is the only known virus to exhibit S." What is more, Dr. Jones does a blood test which verifies that Smith's body contains antibodies for Virus X, further justifying Jones's belief. Based on the evidence, it is extremely plausible that Smith has Virus X. As it happens, however, Smith's symptoms are in fact due to an unknown virus, Virus Y, which exhibits identical symptoms to Virus X; Smith exhibits antibodies for Virus X only because of an idiosyncratic feature of Smith's biochemistry which causes his immune system to maintain high levels of antibodies long past a given infection. Nevertheless, Dr. Jones's belief turns out to be true divorced from Smith's symptoms or his blood work, because Smith was infected with Virus X just before meeting with Dr. Jones—the infection being so recent that bloodwork cannot detect it and it is causing no symptoms. (paraphrased from Zagzebski 1994, 71)

[42] If I am right in thinking that even distant possible worlds can be of epistemic affect, the sufficiency of other safety theories qua safety theories may be called into question.

Dr. Jones is completely non-culpable in her belief formation—no false belief plays a causal or evidential role in her justification (Zagzebski 1994, 71). Indeed, many theorists would suppose that the evidence in question is adequate for knowledge. Seemingly, no false beliefs are in the immediate neighborhood, and Dr. Jones would believe that Smith has Virus X in a wide range of counterfactual situations (Zagzebski 1994, 71). Whether or not Dr. Jones's belief meets Pritchard's Safety Theory, and whether or not Pritchard's anti-luck epistemology is indeed Gettier-proof, will depend on whether or not in all (or nearly all) nearby possible worlds Dr. Jones would believe that "Smith has Virus X" only when it is true.

Pritchard's Safety Theory is caught between a rock and a hard place. Given the set of nearby possible worlds where Smith did not happen to *just recently* become infected with Virus X, perhaps Pritchard's safety condition is not met—in many nearby worlds in which Dr. Jones's belief is formed in the same way (on the same basis), the belief would be false. However, the case can be strengthened—we can make it such that there are fewer close possible worlds where Dr. Jones forms a false belief, as it were. Suppose we stipulated that, unbeknownst to anyone at the time, Virus X was running rampant just outside the hospital and that it is in fact incredibly unlikely that Smith would *not* have caught the virus when he did. So revised, it is less certain that Dr. Jones's belief fails Pritchard's condition—conceivably in very few (if any) nearby possible worlds would Dr. Jones's belief not be true. In other words, it looks like the luck at play in Dr. Jones and the Virus, so revised, would lead to a false belief in few (if any) nearby possible worlds.[43] Even though the Veritic Epistemic Luck involved in Dr. Jones's belief in the strengthened case is of a relatively minute degree, it is still enough to preclude knowledge. To be clear, Pritchard could strengthen his safety condition and say that the truth of Dr. Jones's belief is still (somehow) too unlikely—that the possible worlds in which Dr. Jones's belief is false though formed in the same way (worlds where Virus X is *not* running rampant outside the hospital, for example) are still too prevalent for safety to be satisfied. However, as we play this game (i.e., I keep strengthening the case, and Pritchard keeps strengthening his safety condition) we get to a point where the only possible worlds in which Dr. Jones's belief is false are very distant, and insofar as Dr. Jones's belief can still be "Gettierized" (i.e., made true for

[43] To be clear, given Pritchard's understanding of luck in general (MAL), it seems that cases like Dr. Jones and the Virus (once made suitably strong) would not officially exhibit luck, though clearly we still have a Gettier case.

reasons divorced from the way the belief in question was formed), we find ourselves in a dilemma: *either Pritchard's safety condition is eventually strengthened such that no belief counts as knowledge unless it would be true in all (or nearly all) possible worlds in which the belief is formed in the same way, drastically limiting what we know, or Pritchard's safety condition is going to fall victim to Gettier cases such as Dr. Jones and the Virus, thus making his analysis of knowledge insufficient.*

To be fair, Pritchard has revised his safety condition since *Epistemic Luck* (2005); unfortunately, however, such revisions run into the same dilemma. For example, consider the rendition of safety Pritchard endorses in "Anti-Luck Epistemology" (2007):

> **Pritchard's 2007 Safety Theory:** S's belief is safe *iff* in most near-by possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world, and in all very close near-by possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world, the belief continues to be true. (Pritchard 2007, 292)

Whatever advantages this iteration of safety may have over the one Pritchard develops in *Epistemic Luck* (2005), what is important to see for our purposes is simply that Pritchard's 2007 Safety Theory is still *only* concerned with close possible worlds. As such, it will not be able to prevent Gettier counterexamples like the strengthened Dr. Jones and the Virus case—counterexamples where the belief in question is "Gettierized" though true in all the relevant close possible worlds. Just like its predecessor, the only way Pritchard's 2007 Safety Theory can fully avoid Gettier problems is if it is so strengthened that it precludes knowledge for any belief that exhibits even the slightest degree of Veritic Epistemic Luck—a move that would radically limit what we know.[44]

Pritchard's modal account of luck is the seminal account to date, and I think it is more or less correct. What is more, I think Pritchard is exactly right to point to Veritic Epistemic Luck as the culprit behind Gettier cases. However, by not appreciating degrees of luck, the anti-luck epistemology he subsequently develops misses its mark. In this section, I argued for a consistent extension of MAL that accounted for our intuition that luck suits degrees.

---

[44] In *The Nature and Value of Knowledge* (2010), Pritchard adds a virtue condition to his analysis of knowledge; as such, he might be able to employ such a condition to surmount Gettier counterexamples like the strengthened Dr. Jones and the Virus case (see Pritchard, Millar, and Haddock 2010, chap. 3). Although we might consider such a move to be less than eloquent (safety is, after all, meant to be *the* anti-luck condition for Pritchard), we will not concern ourselves here with the details regarding whether or not such a proposal would in fact work. Nevertheless, given the pervasiveness of luck and the shortcomings of eminent virtue epistemologies in chapters 3 and 4, I think we have a prima facie reason to be pessimistic.

Once extended, however, we saw how Pritchard's analysis of knowledge seems to fall apart in a dilemma—either his safety theory will be too strong to be palatable or too weak to preclude Gettier cases like Dr. Jones and the Virus.[45]

What is more, there is no reason to think that the effects of accounting for degrees of luck is localized to Pritchard's analysis; seemingly parallel arguments could be made with divergent accounts of luck or divergent analyses of knowledge. What counterexamples such as the strengthened Dr. Jones and the Virus case show is that Gettier problems can arise out of even a minute degree of luck. As such, any anti-luck condition in a given theory of warrant that permits even the slightest degree of luck will be susceptible to Gettier counterexamples.

This is, to be sure, all in keeping with Zagzebski's diagnosis: if whatever is taken to bridge the gap between true belief and knowledge (call it warrant) does not guarantee the truth of the belief, then it will be possible to create a Gettier case. It will be possible for a belief to be so warranted and only luckily true for reasons unrelated to the warrant (Zagzebski 1994, 65). Given degrees of luck, that seems exactly right. If even the slightest degree of Veritic Epistemic Luck precludes knowledge, then the warrant of any Gettier-proof version of the reductive analysis (be it a virtue epistemology or not) must guarantee the truth of the belief in question. As Timothy Williamson noted in his book *Knowledge and Its Limits* (2000), "Since Gettier refuted the traditional analysis of *knows* as *has justified true belief* in 1963, a succession of increasingly complex analyses have been overturned by increasingly complex counterexamples," which, given that luck is indeed as permeating as I am suggesting, is just what one might expect (Williamson 2000, 30). Such a conclusion, perhaps, is simply further testament to the growing trend in the contemporary literature to abandon the standard analysis altogether and pursue alternative models from which to do epistemology.

There are, no doubt, a number of worries someone might level against this diagnosis of the Gettier Problem. We may worry that assuming an inviolable relationship between warrant and truth does not lead to radical skepticism. Alternatively, we may worry, with Daniel Howard-Snyder, Frances Howard-Snyder, and Neil Feit, that assuming a close (though not

[45] To be sure, Pritchard could alternatively take warrant to be completely divorced from truth—effectively denying that luck is of any epistemic consequence—but surely that is a move Pritchard would not be willing to make.

inviolable) relationship between warrant and truth does not always lead to further Gettier counterexamples.[46] We may also worry, with Stephen Hetherington, that assuming that warrant bears no relationship to truth is not, in fact, counterintuitive.[47] Fourth and finally, we may worry, with Brian Weatherson, Jonathan M. Weinberg, Shaun Nichols, and Stephen Stich, about the epistemic import of Gettier-style counterexamples altogether.[48] While in Chapters 3 and 4 of this dissertation we will see the proposed connection between radical skepticism and the inviolable warrant/truth relation vindicated at all the relevant instances, it is nevertheless important that we guard against the latter three types of objections—especially insofar as they are manifest in the contemporary literature. This will be our task in the next chapter.[49]

[46] See D Howard-Snyder, F Howard-Snyder, and Feit 2003.

[47] See Hetherington 2001.

[48] See Weatherson 2003 and Weinberg, Nichols, and Stich 2001.

[49] Significant portions of this chapter were published in Church *forthcoming*.

2 OBJECTIONS

Let us quickly take stock. Taking the standard analysis of knowledge to be *warranted true belief* (where "warrant" is whatever we take to bridge the gap between true belief and knowledge), the diagnosis advocated in the previous chapter suggested that so long as warrant is closely but not inviolably related to truth, Gettier counterexamples will be unavoidable. In other words, so long as we reasonably assume that warrant is neither divorced from truth nor inseparable from it (such options being epistemically unpalatable), Gettier cases are "inescapable" (Zagzebski 1994: 65). Consider the following figure:

Figure 2.1:

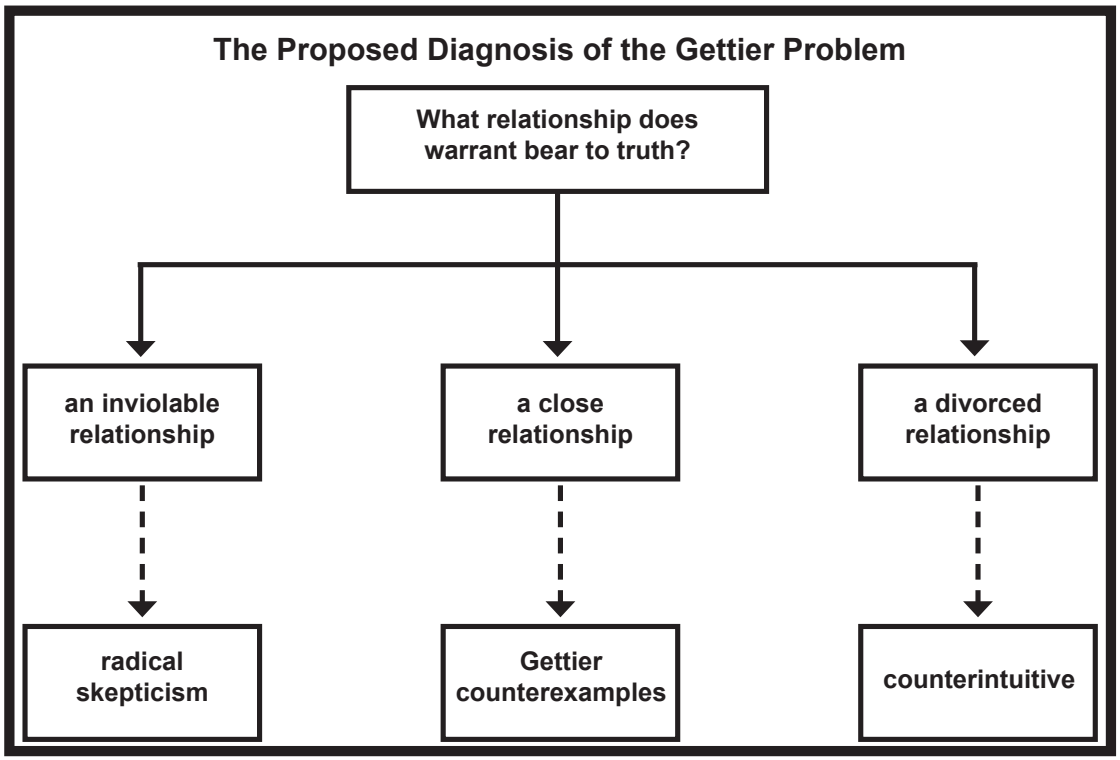

As we discussed in the previous chapter, warrant can bear one of three relationships to truth: an inviolable relationship, a close relationship, or a divorced relationship. The

problem, however, is that none of these relationships seems viable. As we will discover, presupposing an inviolable relationship between warrant and truth seems to readily lead to radical skepticism (it does at least in every analysis of warrant we will consider). Given the prevalence of luck, presupposing a close (but not inviolable) relationship between warrant and truth seems to always lead to further Gettier counterexamples. Finally, presupposing that warrant bears no relationship to truth seems completely counterintuitive.

Now, as I see it, there are four objections someone might level against such a diagnosis of the Gettier Problem. First, someone may level the objection that assuming an inviolable relationship between warrant and truth does not lead to radical skepticism. Likewise, someone may level the objection that assuming a close (though not inviolable) relationship between warrant and truth does not always lead to further Gettier counterexamples. Third, someone might level the objection that assuming that warrant bears no relationship to truth is not, in fact, counterintuitive. Fourth and finally, someone may level a meta-objection that calls into question the epistemic import of Gettier-style counterexamples altogether. While in the latter half of this dissertation we will see the proposed connection between radical skepticism and the inviolable warrant/truth relation vindicated at all the relevant instances, it is nevertheless important that we guard against the latter three types of objections, especially insofar as these sorts of objections are manifest in the contemporary literature.

And that is precisely what I do in this chapter. In the first section, we consider Daniel Howard-Snyder, Frances Howard-Snyder, and Neil Feit's paper, "Infallibilism and Gettier's Legacy" (2003), which directly objects to our conclusion in the previous chapter that a close but not inviolable relationship between warrant and truth will always lead to Gettier counterexamples; I argue that their proposed solution to the Gettier Problem is, ironically, only successful insofar as it conforms to the proposed diagnosis. In the second section, we consider recent work by Stephen Hetherington (particularly from his book *Good Knowledge, Bad Knowledge*, 2001), which proposes an analysis of knowledge that assumes a divorced relationship between warrant and truth; I argue that Hetherington's epistemology is hopelessly counterintuitive and that it need not be advocated in light of viable alternatives. Finally, in Section 3, we consider landmark papers by Brian Weatherson (2003) and Jonathan M. Weinberg, Shaun Nichols, and Stephen Stich (2003), which call into question the epistemic import of Gettier counterexamples; I argue that such criticisms either lead to intractable forms of skepticism or are ultimately incoherent.

Section 1: The Howard-Snyder/Feit Objection

In their paper, "Infallibilism and Gettier's Legacy" (2003), Daniel Howard-Snyder, Frances Howard-Snyder, and Neil Feit consider what they call three nonpartisan arguments against fallibilism (i.e., arguments that are not committed to a specific account of warrant and that are against the view that a belief can be at once warranted and false) and contend that each is lacking. Most importantly for our purposes, however, they contend that an argument for infallibilism based on the sort of diagnosis of the Gettier Problem proposed in the previous chapter is simply flawed. Contra what was said in Chapter 1, Howard-Snyder *et al.* argue that the Gettier Problem can indeed be solved fallibilistically, while assuming that warrant bears a close but not inviolable relationship to truth. In this section, we will consider this argument and Howard-Snyder *et al.*'s opposition to it, and I will argue that in light of our analysis of luck provided in the previous chapter, Howard-Snyder *et al.*'s opposition falls flat—ironically failing precisely along the lines predicted by the proposed diagnosis of Gettier problems.[50]

A typical response to the Gettier Problem is to try to make the relationship between warrant and truth so close that it is simple impossible for a warranted belief to be Gettiered (made true for reasons not predicted by the warrant)—a response that has repeatedly led epistemologists to try to endorse some form of infallibilism. The functional argument at work here, according to Howard-Snyder *et al.*, goes something like this:

1. If a belief can be at once warranted and false, then the Gettier Problem cannot be solved.
2. The Gettier Problem can be solved.[51]
3. So, a belief cannot be at once warranted and false (D. Howard-Snyder, F. Howard-Snyder, and Feit 2003, 135).

---

[50] To be sure, *fallibilism* and *infallibilism* are used here as positions within *the standard analysis of knowledge*. That said, what is important for the proposed diagnosis of the Gettier Problem is that fallibilism, within the standard analysis, cannot surmount Gettier counterexamples, which says nothing about the viability of fallibilism simpliciter. Fallibilism within a Williamsonian model, say, may very well be a viable position, but that is not of concern given the current thesis.

[51] This premise is asserting that the Gettier Problem can be solved *within the standard analysis*. One way to, in effect, "solve" the Gettier Problem would be to abandon the analysis of knowledge enterprise altogether, but that solution is not what Howard-Snyder *et al.* had in mind. Premise 2, then, is something the current thesis is pessimistic about.

The first premise, the key premise, is precisely what the proposed diagnosis of Gettier Problem would predict and this is precisely what Howard-Snyder *et al.* hope to challenge. So, if we are to defend the second branch of the trilemma portrayed in Figure 2.1, a response to Howard-Snyder *et al.* seems appropriate.

The infallibilist response to the Gettier Problem is, to be sure, an increasingly popular one; as Howard-Snyder *et al.* sweepingly point out:

> Scott Sturgeon—purporting to express "the standard post-Gettier model of propositional knowledge"—recently wrote that the Gettier Problem can be avoided only by saying that S knows that *p* if and only if S believes *p* and S's belief is fully justified, where S's belief is *fully* justified only if it is true.[52] Early on, defeasibility theorists decided that if their accounts did not have the consequence that warrant entails truth, they would be subject to Gettier cases, and so they fortified them accordingly. Causal theorists are likewise motivated by Gettier to endorse infallibilism. Recall the first words of the locus classicus of the genus: "Since Edmund L. Gettier reminded us recently of a certain important inadequacy of the traditional analysis of 'S knows that *p*,' several attempts have been made to correct that analysis. In this paper I [Alvin Goldman] shall offer still another analysis . . . , one which will avert Gettier's problem."[53] Reflecting on certain cases, Goldman concluded that one is warranted in believing *p* only if the fact that *p* is causally connected in an appropriate way with the one's believe *p*. Of course, the fact that *p* cannot be thus connected to one's believing *p* unless it is a fact that *p*, and so *p* is true; so his causal theory implies that warrant entails truth, a consequence endorsed in order to "circumvent Gettier's counterexamples."[54] As a process reliabilist, Goldman conceives of the reliability of a process in terms of how it operates in actual and relevant alternative situations. Those alternatives which are relevant are conceived in such a way that warrant entails truth. Why? To avoid Gettier cases, he says.[55] Robert Nozick added condition (iii) [his sensitivity condition] to his "truth-tracking" theory—S knows that *p* only if S wouldn't believe that *p* if *p* weren't true—precisely because he wishes to "exclude cases of the sort first described by Edmund Gettier."[56] Fred Dretske argued that one has a warranted belief that *p* only if one has a "conclusive reason" for *p*, a reason that "eliminates" the conjunction of that reason and the denial of *p* as a possible state of affairs. What motivated him to lay down this condition? "[T]he conviction (supported by Gettier-like examples) that knowledge if it embodies an evidential relation at all, must embody a strong enough one to eliminate the possibility of mistake."[57] Even Roderick Chisholm—who explicitly wished to avoid infallibilism in his response to the Gettier Problem—in the end conceded that "what is known must be evident but not defectively evident," where a proposition is not defectively evident only if the propositions on which it is based do *not* make

[52] See Sturgeon 1993, 160.

[53] See Goldman 1967, 357.

[54] See Goldman 1967, 370.

[55] See Goldman 1986, 46-55.

[56] See Nozick 1981, 173.

[57] See Dretske 1978, 42, 57.

> evident a proposition that is *false*.[58] (D. Howard-Snyder, F. Howard-Snyder, and Feit 2003, 306, 307)

What is more, it is worth noting that the prominent virtue epistemologists under consideration in this dissertation are no different; as we will see in the next couple of chapters, both Alvin Plantinga and Ernest Sosa have, in recent work, fallen into accord with an infallibilist response to the Gettier Problem. Be that as it may, the point is clear: epistemologists from across the board have, in accord with the proposed diagnosis, chosen to try to provide a feasible analysis of knowledge through inviolably connecting warrant with truth, through infallibilism.[59] It is the perceived exigency of such a move that Howard-Snyder *et al.* wish to challenge by proposing a fallibilism-friendly solution to the Gettier Problem.

Before they give us their solution, Howard-Snyder *et al.* first have us consider a classic Gettier counterexample:

> **Classic Case**: Smith and Jones are applying for the same job. Smith has very strong evidence for thinking that Jones will get the job (e.g., the employer tells Smith that he will hire Jones, etc.), and for thinking that Jones has 10 coins in his pocket (e.g., Jones emptied his pockets in front of Smith and then clearly, slowly, in good lighting, and perhaps even counting out loud, placed 10 coins in his pocket). As such, Smith forms a belief in the general proposition that "the man who gets the job has 10 coins in his pocket." As it turns out, however, Smith gets the job, and he happens to also have 10 coins in his pocket. (paraphrased from Gettier 1963, 122)

We could, of course, try to defuse such a case by saying "we cannot have warranted beliefs based on inferences from false beliefs," but as Howard-Snyder *et al.* point out, "cases that lack this feature are a dime a dozen"; further counterexamples (e.g., Fake Barns-type cases) can be easily produced (D. Howard-Snyder, F. Howard-Snyder, and Feit 2003, 307–308). That said, if we are to find a fallibilism-friendly solution, we must look elsewhere.

According to Howard-Snyder *et al.* "the distinctive feature of standard Gettier cases like these is that the reason Smith believes *p* or the processes involved in his believing *p* are not properly related to those facts that render *p* true" (D. Howard-Snyder, F. Howard-Snyder, and Feit 2003, 308). After all, being told by a reliable source that Jones is going to get the

[58] See Chisholm 1982, 45-47.

[59] This uniformity of strategy in approaching the Gettier Problem does not, of course, at all necessarily signify a real solution; it does not mean that the infallibilist strategy is a feasible one. After all, it is, as far as I can tell, the general consensus that none of the accounts Howard-Snyder *et al.* mention genuinely offer a viable solution—most of which fail precisely along lines that this thesis's proposed diagnosis would predict.

job and counting the coins in his pocket, in the end, had nothing to do with the truth of the relevant belief. "It's a matter of sheer serendipity that [Smith's] belief that *p* is true given his reasons for believing *p* or the processes involved in his believing *p*" (D. Howard-Snyder, F. Howard-Snyder, and Feit 2003, 308).

Now, where we take this "distinctive feature" of Gettier cases to be indicative of the proposed diagnosis of Chapter 1 (after all, if a given theory of warrant does not entail truth, of course it will be possible for a given belief to be "Gettiered"—to be true for reasons divorced from the warrant), Howard-Snyder *et al.* take it as directly translatable into a (fallibilism-friendly) anti-Gettier condition on knowledge. Their first proposal:

> S's belief that *p* is warranted only if S's belief that *p* is not accidentally true for S. (D. Howard-Snyder, F. Howard-Snyder, and Feit 2003, 308)

What does it mean for a belief to be "accidentally true for S"? According to Howard-Snyder *et al.*, "accidentally true for S" is shorthand for "such that what makes *p* true is not properly related to the reasons for, or processes involved in, S's believing *p*" (D. Howard-Snyder, F. Howard-Snyder, and Feit 2003, 308). What does it mean for the truth of a given belief to be "properly related" to the relevant reasons or processes involved in S's believing it? Howard-Snyder *et al.* don't say. Regardless, such a condition certainly seems to eliminate Gettier cases (especially since it seems to directly translate into "S's belief that *p* is warranted only if S's belief is not 'Gettiered'"), and it seems to do so without committing to infallibilism; according to the above condition, it will be possible for a warranted belief to be false since being false "trivially satisfies this condition" (D. Howard-Snyder, F. Howard-Snyder, and Feit 2003, 309).

As Howard-Snyder *et al.* point out, however, such a condition is arguably not adequate on at least two scores. First, we want a given anti-Gettier condition on warrant to tell us something about the nature of warrant. As they explain:

> Our first candidate fails on this score. The condition it lays down is equivalent to the following disjunctive condition: S's belief that *p* is warranted only if S's belief that *p* is false *or* nonaccidentally true for S. It divides the cases into true and false beliefs. What it says about warranted true beliefs is helpful in understanding the nature of warrant: they must be *nonaccidentally* true. But what it says about false beliefs is completely unhelpful. Any old completely unjustified and unreliably formed belief could meet this condition on warrant, so long as it was false. So this condition tells us nothing about the nature of warrant in the case of false belief. (D. Howard-Snyder, F. Howard-Snyder, and Feit 2003, 309)

So, given that Howard-Snyder *et al.* *want* a condition on warrant that tells us *something* "about the nature of warrant in the case of false belief[s]," we should look elsewhere. The second worry is that their initial proposal seems to imply that being false is an epistemic step in the right direction, since being false trivially satisfies their condition; and insofar as we should avoid avoidable implausibilities, Howard-Snyder *et al.* again acquiesce that we should look elsewhere (D. Howard-Snyder, F. Howard-Snyder, and Feit 2003, 309).

Thankfully for Howard-Snyder *et al.*'s account, such worries are easily avoided. The preferable fallibilism-friendly condition on warrant that they endorse is:

> S's belief that *p* is warranted only if S's belief that *p* would not be accidentally true for S, if it were true. (D. Howard-Snyder, F. Howard-Snyder, and Feit 2003, 309)

"[T]he distinctive feature of standard Gettier cases," according to Howard-Snyder *et al.*, remember, "is that the reason [S] believes *p* or the processes involved in his believing *p* are not properly related to those facts that render *p* true." As such, what Howard-Snyder *et al.* have done is straightforwardly convert their diagnosis of Gettier counterexamples into an anti-Gettier condition on warrant; interpreting this latter condition as simply demanding that "S's belief that *p*" will only be warranted if "the following subjunctive conditional is true: if S's belief that *p* were true, then it would also be true that what makes *p* true is properly related to the reasons for, or the processes involved in, S believing *p*" (D. Howard-Snyder, F. Howard-Snyder, and Feit 2003, 309). And (although we are still not told what it means for the truth that *p* to be "properly related" to S's believing that *p*) given (a) that this latter condition avoids the worries afflicting the former condition, and (b) that there is no reason to think this latter condition precludes the possibility of a warranted false belief, Howard-Snyder *et al.* have seemingly provided us with a fallibilism-friendly way to circumvent the Gettier Problem, without, so they would hold, sacrificing feasibility.

Now, as Howard-Snyder *et al.* point out, there are several different worries we may have for such a condition on warrant. According to Scott Sturgeon, the difficulty with finding a solution to the Gettier Problem is the difficulty of finding "the minimal link" between warrant and truth without succumbing to Gettier counterexample (Sturgeon 1993, 157). As such, we may worry that Howard-Snyder *et al.*'s condition does not provide such a *minimal* link.[60] Alternatively, we may, with Linda Zagzebski, simply worry that Howard-Snyder *et*

[60] For Howard-Snyder *et al.*'s response to this worry see Daniel Howard-Snyder, Frances Howard-Snyder, and Feit 2003, 311-312.

al.'s condition is "vague, negative, lacks practical import, . . . has little to recommend it theoretically," and "like the definition *justified true belief that is not a Gettier case* . . . ad hoc" (Zagzebski 1999, 103).[61] But so as to avoid devolving into debates regarding vagueness and just how we should understand "minimal" terminology, I want to largely put such worries aside. Instead, I want to focus on a more fundamental problem—the problem that Howard-Snyder *et al.*'s proposed condition on warrant simply doesn't work; the problem that, in accord with the diagnosis of Gettier counterexamples offered in Chapter 1, Howard-Snyder *et al.*'s proposed condition will lead either to radical skepticism or further Gettier counterexamples.

Before we do that, however, I do want to consider Zagzebski's objection that Howard-Snyder *et al.*'s condition on warrant is *ad hoc*, because regardless of whether or not it falls on the dilemma outlined in the previous chapter, Howard-Snyder *et al.*'s condition should nevertheless be rejected as adhocery. In response to Zagzebski's accusation, Howard-Snyder *et al.* say the following:

> We disagree that our condition is like "justified true belief *that avoids Gettier cases*." Our condition does not express a string of unrelated features with a proper name in it; ours leads to some understanding. More plausibly, perhaps Zagzebski's point is that *no* condition on warrant should be such that its "sole advantage is to answer counterexamples"; it must be "plausible even if no one had ever thought of Gettier cases."[62] But, is adhocery of this sort really a defect? Even if it is, it's not clear that our condition is ad hoc. Why suppose that if no one had thought of Gettier cases, no one would have thought our condition was plausible? (D. Howard-Snyder, F. Howard-Snyder, and Feit 2003, 316)

There are a couple things worth noting in response. If the Gettier Problem is, as Howard-Snyder *et al.* seem to claim, the problem of the reasons for a given belief not being "properly related" to what makes the said belief true, then Howard-Snyder *et al.*'s condition on warrant explicitly and straightforwardly amounts to "S's belief is warranted so long as it avoids Gettier cases." It is certainly no more informative. From the very beginning everyone knew that Gettier problems show that something is wrong with the fit between warrant and truth, so having a general condition that prohibits knowledge of beliefs with "a bad fit" is un-illuminating and patent adhocery. Telling us that Gettier cases involve "accidentally true beliefs" does not help in the least; we already knew *that*. As such, it is

---

[61] For Howard-Snyder *et al.*'s response to these worries see Daniel Howard-Snyder, Frances Howard-Snyder, and Feit 2003, 314-316.

[62] See Zagzebski 1996, 65.

quite clear that no one would have thought of their condition without thinking about Gettier cases, because their condition is a straightforward denial of Gettier cases. Seemingly, someone could not come up with the condition for warrant ¬*c* without also coming up with *c*.

Even if Howard-Snyder *et al.*'s proposed condition on warrant *is* somehow feasible and not *ad hoc*, I now want to argue that, based on two plausible readings of their condition, it runs into precisely the sort of dilemma predicted by our diagnosis of the Gettier Problem offered in the previous chapter—ironically lending credence to the view that infallibilism is indeed part of Gettier's legacy. I want to argue, in accord with our proposed diagnosis, that no matter how we understand Howard-Snyder *et al.*'s condition it will seemingly either run into Gettier counterexamples or lead to radical skepticism through infallibilism.

The first plausible reading of Howard-Snyder *et al.*'s proposed condition of warrant is one where it prohibits luckily true beliefs. I think it is fairly clear that Howard-Snyder *et al.* could easily be conflating "X is lucky" with "X is accidental" such that when they prohibit accidentally true beliefs (or beliefs whose reasons are not properly related to the belief's truth), they are really prohibiting something like luckily true beliefs.[63] What is more, given their extrapolation of their account (the "properly related" business), it also seems fairly clear that when they prohibit accidentally true beliefs they are really prohibiting something very much like veritically lucky true beliefs of the previous chapter. In other words, it seems like Howard-Snyder *et al.* could easily be read as prohibiting beliefs that, given the way they were formed, are only luckily true. Now, if this is right, the problems they run into are straightforward. As we noted in the previous chapter, luck comes in degrees, and as we argued, even a minute degree of luck is of significant epistemic import. And given 1) that most every belief we hold is at least minutely (veritically) lucky and 2) that Gettier counterexamples can be created out of even a minute degree of luck (consider the strengthened version of Dr. Jones and the Virus), then Howard-Snyder *et al.*'s only hope for avoiding Gettier counterexamples is to make their condition prohibit even marginally lucky beliefs; in so doing, however, they would likely be committing themselves to radical skepticism—seemingly, very few of even our most secure beliefs are completely luck-less. Given that an all-out ban on lucky beliefs would effectively make it impossible for a belief

[63] This is not, to be sure, of central importance for my subsequent argument; my worry for Howard-Snyder *et al.*'s account can run regardless of whether they are truly conflating luck with accidents. I make this point because 1) it seems like a fair one to make and 2) it offers continuity with the previous chapter.

to be warranted and false, ironically, it seems as though the only way Howard-Snyder *et al.*'s condition, so understood, can *really* avoid Gettier counterexamples is if it commits to infallibilism.

But perhaps that's not the way to read Howard-Snyder *et al.*'s condition after all, plausible though such an interpretation may be. Perhaps instead, when they prohibit accidentally true beliefs (i.e., beliefs whose reasons do not properly relate to the truth), they are prohibiting beliefs that are somehow true for the wrong reasons (whatever precisely that means)—beliefs that are true for reasons that *your* reasons, evidence, or cognitive processes would not have predicted. Surely this is precisely what is at issue in Gettier counterexamples (and seemingly Howard-Snyder *et al.* have no qualms with *directly* importing whatever they see as defining Gettier counterexamples into a condition on warrant), so perhaps *this* is the way to read their proposed condition on warrant. Sadly, however, they are going to run into similar troubles as the previous reading. Surely being "true for the right reasons" or being "true for reasons my evidence would predict" is a matter of degree. And seemingly, the vast majority of our beliefs are, to at least some minute extent, going to be true for reasons we could not have predicted. For example, my secure belief that, as I am writing this, I am a "visiting scholar" at Rutgers University is surely knowledge, but it is probably true, at least in part, for reasons my evidence does not account for—reasons like, such and such a form was filled out (which I had nothing to do with) making my tenure at Rutgers official, etc. So surely Howard-Snyder *et al.* would not want to establish an all-out prohibition on beliefs that are not *entirely* true for the right reasons, reasons predicted by my evidence, because such a prohibition would seemingly push us toward radical skepticism—denying knowledge to most of even our most epistemically secure and everyday beliefs. But unless Howard-Snyder *et al.* make such a prohibition, it looks as though Gettier counterexamples are going to be inevitable; degrees of luck, as it were, will seemingly map onto degrees of "being true for the right reasons." Consider the following example:

> **The Horticulturalist:** David is an expert horticulturalist, able to competently distinguish between the some 20,000 different species of orchid. David is presented with an orchid and asked to identify its species. Using his amazing skill he can clearly tell that this particular orchid is either going to be an X-species or a Y-species (which look quite similar), and upon even further expert analysis he comes to the conclusion that it is an X-species of orchid, which it is. However, Kevin, David's nemesis and an expert horticulturalist in his own right, decided the night before to, using his skill as a horticulturalist, make the X-species of orchid look like a Y-species of orchid. Thankfully, however, Alvin, David's other expert

> horticulturalist nemesis (who is conveniently not on speaking terms with Kevin), decided to try to trick David in the same way—arriving shortly after Kevin left, perceiving that the orchid was a Y-species, and cleverly making it look, once again, like an X-species. As such, while David's belief that the given orchid is an X-species of orchid is largely for the right reasons (he was, after all, able to narrow down the possibilities from over 20,000 to just two), he does not ultimately *know* that it is an X-species of orchid since he was effectively Gettiered by the combined efforts of Kevin and Alvin.

Howard-Snyder *et al.* could always object that the relevant belief of protagonists like David is not true enough for the right reasons, but strengthened cases can always be produced. As such, given the right-reasons reading, it once again looks as though the only way for Howard-Snyder *et al.*'s proposed condition on warrant to completely avoid Gettier counterexamples is if it prohibits any belief from being knowledge that is true for any reason not predicted by the given agent's evidence, reasons, or cognitive processes. As such, if Howard-Snyder *et al.*'s condition is to avoid Gettier counterexamples, it will, in accord with our diagnosis, seemingly lead us to radical skepticism. And insofar as it is not possible for a warranted belief to be false while satisfying "true for the *completely* right reasons" reading of their condition, it looks again as though, ironically, the only way for their condition to surmount the Gettier Problem is to acquiesce to infallibilism.

In Chapter 1, I proposed a diagnosis of the Gettier Problem, which leveled a trilemma against any version of the standard analysis of knowledge (roughly, *warranted true belief*) in accord with the three ways warrant might relate to truth. It was proposed that if we assume that warrant is close but not inviolably connected with truth (i.e., if we assume fallibilism), Gettier counterexamples will be unavoidable. Recent work by Daniel Howard-Snyder, Frances Howard-Snyder, and Neil Feit called this conclusion into question; arguably offering a fallibilist-friendly solution to the Gettier Problem. Thankfully for our proposed diagnosis, however, this solution was found lacking—being ad hoc and (based on two plausible readings) unable to surmount the Gettier Problem without requiring infallibilism (and leading to radical skepticism), which is precisely what Chapter 1's diagnosis would predict. As such, this leg of the trilemma still stands.

Section 2: The Hetherington Objection

Even if we are right that the only way to avoid Gettier counterexamples is to assume either an inviolable (i.e., infallibilistic) relationship between truth and warrant, or no relationship

at all, someone may nevertheless deny that such strategies are deeply problematic. While we will see in subsequent chapters that infallibilism is indeed infeasible (leading to radical skepticism in all of the relevant cases), our goal in this subsection is to defend the proposed diagnosis of the Gettier Problem against those who might try to defuse Gettier counterexamples by denying warrant any relationship to truth (i.e., by denying that the luck involved in Gettier cases necessarily precludes knowledge). While it is indeed almost universally accepted that warrant should bear at least *some* relation to truth, we should, for the sake of due diligence, address dissent.

For over a decade, Stephen Hetherington has been perhaps the most eminent and vocal advocate of just such a view—the view that knowledge does not require warrant to bear any relationship to truth. Starting with his landmark book *Good Knowledge, Bad Knowledge* (2001), we will, in this subsection, consider Hetherington's understanding of knowledge as well as his response to Gettier problems—subsequently arguing that Hetherington's epistemology is simply not viable, especially in light of other proposals now on offer.[64]

Like so many great philosophical projects before it, Hetherington begins *Good Knowledge, Bad Knowledge* with a wholesale rejection of a longstanding and nearly platitudinous assumption in hopes of revitalizing various stale and/or tricky debates.[65] The longstanding assumption that Hetherington wants to call into question is what he calls Epistemic Absolutism:

> **Epistemic Absolutism:** Knowledge is absolute, in the sense that it is impossible for a person to have *better*, or to have *worse*, knowledge of a fact. (Hetherington 2001, 3)

According to Hetherington, most epistemologists assume that knowledge is something that you simply have or lack, that knowledge is something that does not, as he puts it, admit degrees. And while Hetherington grants that there is an absolute point where a given belief will either start or stop being knowledge (i.e., it is not as if Hetherington thinks that all beliefs are some degree of knowledge), he nevertheless thinks that knowledge is the sort of thing that can, once had, become either better or worse.

---

[64] I start with *Good Knowledge, Bad Knowledge* (2001) instead of his paper "Actually Knowing" (1998), because while the latter indeed advocates the relevant view, the former is more complete and is, I think, better able to stand under recent criticisms (cf. Madison 2011).

[65] This is, to be sure, just the sort of way Timothy Williamson begins Williamson 2000. Indeed, this strategy is at the heart of this dissertation.

To be clear, consider the following example:

> **Flood:** One day, at t_1, Bill's uncle, who is known to frequently lie, tells him that there has been a massive flood in Indiana, where Bill is originally from. At t_2, Bill checks the various newspapers and finds that there has indeed been a massive flood in Indiana. Later that day, at t_3, Bill calls various trustworthy friends in Indiana to make sure they are safe and hears their accounts of the flooding.

According to Hetherington, most epistemologists would give Flood the following analysis: while Bill's uncle's testimony may provide some warrant for believing that there has been a flood in Indiana, given the uncle's propensity to lying, Bill does not know that there has been a massive flood in Indiana at t_1; at t_2, after reading the newspapers, Bill is seemingly now warranted enough to know that there has been a massive flood in Indiana; finally, at t_3, Bill now knows, after speaking with his friends, that there has been a massive flood in Indiana w*ith an even greater degree of warrant.* And it is this analysis of Bill's epistemic state at t_3 with which Hetherington takes issue. Granted that Bill does not know that there has been a massive flood in Indiana at t_1, Hetherington wants to say that surely Bill knows about such flooding far better at t_3 than at t_2. In other words, according to Hetherington, at t_3 Bill not only knows with a higher degree of warrant, *Bill knows to a higher degree in accord with the warrant.*

Now, why should we grant that knowledge admits degrees? What is the payoff of rejecting Epistemic Absolutism? It is not, after all, immediately apparent; indeed, the distinction that Hetherington seems to be making can, at first blush, simply seem like a matter of semantics. Whatever reasons Hetherington gives for rejecting Epistemic Absolutism and whatever merit anti-absolutism brings to various perennial epistemic debates, we need only concern ourselves at this point with how Hetherington applies this anti-absolutism to handle the Gettier Problem.

For Hetherington, Gettier problems will be unavoidable so long as "we see knowledge as being categorical and [warrant] as being gradual"; hence, he proposes that once we see both warrant and knowledge as admitting degrees, "[t]he Gettier Problem will disappear" (Hetherington 2001, 72).[66] How so? Consider the following case:

[66] I have replaced "justification" with "warrant" in the above quote, because, given the general way Hetherington is understanding "justification" (cf. Hetherington 2001, 88) and the running definition of "warrant" in this dissertation, such terms are synonymous; indeed, Hetherington expressly admits as much in Hetherington 2001, 109.

> **Dog/Sheep:** Standing outside a field you see what seems to you to be a normal sheep. You do not think the false thought, "That is a sheep"; instead, you think the true thought, "That looks exactly like a sheep." From this you infer that there is a sheep in the field. And you are right. But the real sheep in the field is out of sight, hidden from your gaze, eating peacefully in a distant corner of the field. What you are seeing is a dog, disguised in a sheep's fleece. (Hetherington 2001, 71)

The typical analysis of such a case would be that the given protagonist simply lacks the relevant bit of knowledge; Hetherington, however, does not concur. According to Hetherington, all Gettier cases contain what he calls a "Strange Occurrence"; that is, all Gettier cases contain something that makes the "epistemic subject . . . somewhat lucky to have his well-justified true belief" (Hetherington 2001, 73). And instead of precluding knowledge, a Strange Occurrence only weakens the strength of a given epistemic agent's purported knowledge.[67] Hetherington offers the following general elucidation:

> (1) If an epistemic subject x knows that *p* within a normal (non-Gettier) situation, then none of the close—the similar—accessible worlds, as regards that knowledge of x's, are epistemic-failure worlds for x in relation to *p*. That is, in every close world were at least two of the three traditional components of x's knowing that *p* are present, the third component is present. Because (on our hypothesis), x knows that *p* in a normal way, that normality is also present within any very close worlds where at least two of those traditional components are present, with the result—the normal result, after all—being that the third component is also present. (2) On the other hand, if x knows that *p* within a Gettier situation, then there *are* some close accessible worlds, as regards that knowledge, that are epistemic-failure worlds for x in relation to *p*. This is because the close accessible worlds, as regards x's knowing that *p* within a Gettier case, are not simply worlds where at least two of the three traditional components of x's knowing that *p* are present; they are also worlds containing the Strange Occurrence from x's Gettier situation. And in some of these worlds the Strange Occurrence's presence ensures that the third traditional component of x's knowing that *p* is not present. Relative to a specific case, there need not be worlds like this for each one of the three components, but there will be some for at least one of the three. For example, the luck with which, given the Strange Occurrence, x's good evidence for *p* and his believing that *p* coincided in this world's Gettier situation with *p*'s being true might be modelled by there being some close truth-failure worlds for x in relation to *p*—some close truth-failure worlds for x in relation to *p*, that is, where the Strange Occurrence is present. (Hetherington 2001, 76-77)

Applying his degree-theoretic analysis of knowledge, Hetherington reinterprets Gettier cases like Dog/Sheep as exhibiting weak or "failable" knowledge. To put it roughly, knowledge is failable, according to Hetherington, if it nearly wasn't knowledge; in other

[67] As Hetherington points out, it would be a fallacy to confuse *almost* failing to achieve knowledge with *genuinely* failing to achieve knowledge (Hetherington 1998, 456–459).

words, S's knowledge is failable if in some close possible worlds S does not have it.[68] As such, instead of showing that the warranted true belief is insufficient for knowledge, Gettier cases simply show that knowledge might be largely (though not completely) undermined, by showing how knowledge might be very weak or failable. All this being correct, Hetherington thinks he has dissolved the Gettier Problem.

But the main problem with this proposed solution to Gettier counterexamples is that it misidentifies what is actually doing the work. Despite Hetherington's express claim that the Gettier Problem will dissolve once knowledge is acknowledged to admit degrees, Hetherington's anti-absolutism, his rejection of Epistemic Absolutism, in itself *plays absolutely no role in solving Gettier counterexamples*. Assume for the sake of argument that, while knowledge has an absolute cutoff point (i.e., a point where a given belief will either start or stop being knowledge), knowledge nevertheless admits of degrees. Let us also make the reasonable assumption that warrant, that which bridges the gap between true belief and knowledge, should bear some close but not inviolable relation to truth if a given belief is to be knowledge. With these two assumptions, Hetherington's anti-absolutist epistemology seems every bit as stymied by Gettier cases like Dog/Sheep as standard absolutist epistemologies. The protagonist's warrant for believing that "there is a sheep in the field" is critically disconnected from the belief's truth, and, as such, the relevant belief simply falls short of knowledge given our second assumption that warrant should bear some close though not inviolable relationship to truth; the anti-absolutism does not even come into play.

In accord with Chapter 1, what *is* solving the Gettier Problem for Hetherington is his tacit assumption that warrant need not bear *any* relationship to truth. Consider how he described normal (non-Gettier) epistemic circumstances versus Gettier circumstances above: normal (non-Gettier) epistemic circumstances are roughly those where all three components for knowledge (warrant, truth, belief) are simply "present" in all close possible worlds; Gettier circumstances are roughly those where at least one of the said components is not "present" in some close possible worlds. Knowledge, for Hetherington, is seemingly like a recipe; simply add equal parts truth and belief and a sufficient amount of warrant, stir, and then you will have knowledge (to some degree). Indeed, if warrant can be completely divorced from truth, we may wonder whether warrant was really necessary for knowledge in the first place; as such, it should come as no surprise that this is exactly the

[68] For a full account of what Hetherington means by "failable knowledge," see Hetherington 2001, 40-47.

sort of view that Hetherington eventually endorses. In *Good Knowledge, Bad Knowledge* (2001), "Is there a World where Knowledge has to include Justification? (2007), and "Elusive Epistemological Justification" (2010), Hetherington, rather shockingly, posits that mere true belief can suffice for knowledge.[69] But this is clearly a radical break from the typical understanding of knowledge. In any case, Hetherington assumes that warrant need not bear any relationship to truth. And it is this assumption that helps him solve the Gettier Problem, and it is this assumption that puts him radically at odds with the practices and intuitions of most epistemologists.

Of course, one of the main problems with such an assumed theory of knowledge is that it is *extremely* counterintuitive when it comes to Gettier-type cases. We can, to be sure, question the philosophical merit of such intuitions (and we will consider some philosophers who do just this in Section 3), but insofar as Hetherington wants to somehow honor or account for such Gettier intuitions (as I think any viable epistemology should), his proposed solution to the Gettier Problem is deeply problematic. And Hetherington does seem to try to honor or account for Gettier intuitions in at least two ways. First, Hetherington tries to defuse Gettier intuitions by arguing that to deny knowledge in Gettier scenarios is to commit what he calls *the epistemic counterfactuals fallacy*. Second, Hetherington tries to account for Gettier intuitions through his rejection of Epistemic Absolutism. We will now consider both strategies in turn.

First, in an attempt to address or at least account for the aforementioned Gettier intuitions, Hetherington argues that such intuitions rest on an easily made mistake, namely, the mistake of conflating *almost possessing knowledge* with *genuinely not possessing knowledge*—a mistake Hetherington identifies as *the epistemic counterfactual fallacy*. As Hetherington explains:

---

[69] See Hetherington 2001, chap. 4; Hetherington 2007; Hetherington 2010. In "Is This a World Where Knowledge has to Include Justification?" (2007), for example, Hetherington posits that in a random and shifty world, a world where warrant is simply not possible, having beliefs that are luckily true can stand as knowledge—the luck, according to Hetherington, now bearing the value typically attributed to warrant (Hetherington 2007, 43–46). But this sounds crazy. A world where warrant is simply not possible strikes us, whether we are committed to the analysis of knowledge or not, as a world where knowledge is not possible. Good luck is in no way a substitute for warrant, and to suggest that mere true belief can be knowledge seems to suggest that Hetherington has changed the subject; what he means by "knowledge" simply cannot be what the rest of us seem to mean by "knowledge." Perhaps his alternative "epistemology" is not a relevant alternative after all.

> Within each Gettier case, the epistemic subject knows, because he has a well-[warranted] true belief. But he is very lucky to have all that (and thereby the knowledge). This luck is a mark of how very failably he knows within the Gettier case. He does know—even though he almost failed to do so. He does know—even though he is very close to not doing so. Consider an analogous situation. Maybe there is luck involved in life's ever having come to exist in this universe. There are so many possible and slight differences to the universe's initial conditions that, had they been actual, would not have led to the creation of life. So, by having the initial conditions it had, the universe was close to having different initial conditions—and thereby to not containing life. Nevertheless, it did not have those different initial conditions. Most of us do not say that consequently there are no natural laws within the universe. There are, in spite of the luck involved in their obtaining. And failable knowing within a Gettier case is like that. In each such case, the epistemic subject is very close to failing to know that *p*. He almost fails to know that *p*. And his being so close to failing to know that *p* is what misleads people into thinking that he does fail to know that *p*. Instead, we may say that his being so close to not knowing that *p* is simply part of his knowing very failably that *p*. (Hetherington 2001, 82)

Gettier cases, for Hetherington, are scenarios where at least one of the three general ingredients for knowledge (warrant, truth, and belief) is almost not present. As such, Gettier cases are, for Hetherington, simply cases where a given protagonist almost lacks knowledge: so to deny the said protagonists knowledge is, by Hetherington's light, to commit the epistemic counterfactual fallacy. Just as we would not deny knowledge to someone who was almost struck by lightning, we should not deny knowledge to someone whose belief is almost not true, say.

But this is a confused understanding of Gettier cases.[70] Gettier cases are not simply scenarios where one of the essential ingredients for knowledge is almost missing. As we saw in the previous chapter, luck regarding the existence of a given belief (Doxastic Epistemic Luck), the truth of a given belief (Content Epistemic Luck), or the warrant of a given belief (Evidential Epistemic Luck) are all epistemically benign and manifestly *not* the luck at issue in Gettier problems; Gettier cases are scenarios that exhibit luck regarding the truth of a given belief in relation to the relevant warrant (Veritic Epistemic Luck). As such, scenarios like Hetherington's existence of life example are simply *not* analogous. Given the analysis of luck provided in the previous chapter, Gettier cases are not scenarios where a given protagonist almost does not have knowledge; they are cases where, given the kind of luck at issue, the protagonist genuinely does indeed lack knowledge.

What about the aforementioned second strategy? Hetherington tries to account for Gettier intuitions through his rejection of Epistemic Absolutism by arguing that Gettier

[70] See Madison 2011.

cases only *seem* to lack knowledge because they are, in fact, very bad (i.e., very failable) instances of knowledge that do indeed fail to be knowledge in close possible worlds (Hetherington 2001, 88–92). As Hetherington explains, identifying the knowledge in Gettier scenarios as *very* failable "does justice to the feeling that there must be a difference in the quality of the instances of knowing in, respectively, a normal situation where there is failable knowledge that *p*, and a Gettier situation where there is failable knowledge that *p*" (Hetherington 2001, 76). But given what has been said in the previous chapter, this strategy is now clearly defunct. Gettier cases, it has been argued, can be created for any belief that is not knowledge in all possible worlds; if luck comes in degrees (as it seemingly does), even very "good" instances of knowledge in Hetherington's estimation can be Gettiered. As such, Hetherington cannot invoke a good knowledge/bad knowledge distinction to save face when it comes to Gettier intuitions.

In the preface to *Good Knowledge, Bad Knowledge*, Hetherington says that the "main justification" for his epistemic project is its ability to solve "many of epistemology's traditional puzzles" (Hetherington 2001, vi). But insofar as (i) Hetherington's solution to that traditional puzzle of central importance, the Gettier Problem, leads to an extremely counterintuitive analysis of knowledge (an analysis that says, at best, warrant need not relate to truth at all, and at worst that mere true belief is sufficient for knowledge) and (ii) other more viable alternative solutions (such as Timothy Williamson's) are on offer, it seems Hetherington's "main justification" is undermined. As such, the infeasibility of assuming a divorced relationship between truth and warrant, the third leg of my trilemma, still holds.

Section 3: Two Meta-Objections

Having defended two of the legs of Zagzebski's trilemma (the third leg being vindicated throughout the course of the dissertation), we now turn to consider some meta-objections against taking Gettier counterexamples seriously to begin with. While this dissertation is largely based on the reasonable and widespread assumption that Gettier counterexamples *should* be taken seriously, it is nevertheless worth considering some such meta-objections that can be found in the literature. That said, however, my aim in this section is relatively

weak; I do not intend to wholly disprove the pertinent meta-objections but simply to cast sufficient doubt on them.[71]

In this section, we will consider two such meta-objections—both focusing on the epistemic import of the Gettier intuition (e.g., the intuition that Gettier counterexamples are genuine counterexamples). The first meta-objection comes from the paper, "Normativity and Epistemic Intuitions" (2001), by Jonathan M. Weinberg, Shaun Nichols, and Stephen Stich, in which Weinberg *et al.* claim that the intuitions that have generated the strong response to Gettier counterexamples (e.g., earnestly trying to avoid being vulnerable to Gettier cases) are culturally relative and as such are "seriously undermined" (Weinberg, Nichols, and Stich 2001, 429). The second meta-objection comes from Brian Weatherson's paper, "What Good are Counterexamples?" (2003), in which Weatherson argues that intuitions are regularly unreliable and particularly unreliable when it comes to Gettier counterexamples—insisting that we should champion "a simple, systematic and largely successful theory [of knowledge]" over "respecting" the Gettier intuition (Weatherson 2003, 1). We will now consider each of these objections in turn.

Section 3.1: Weinberg *et al.*'s Objection

In their landmark paper in experimental philosophy, Jonathan M. Weinberg, Shaun Nichols, and Stephen Stich call into question the philosophical import of intuitions (including our intuitions regarding Gettier counterexamples) through a series of empirical observations. They argue, roughly, that if key epistemic intuitions are not shared across cultures or across socioeconomic demographics, then their evidential value in our philosophical theorizing is largely undermined. In this subsection, I first briefly elucidate Weinberg *et al.*'s argument and consider an example of their purported empirical observation regarding intuitions in Gettier counterexamples. I will then argue that 1) we have very good reason to be extremely skeptical of Weinberg *et al.*'s findings, and 2) that their empirical observations, even if generally correct, are not a threat to taking epistemic intuitions very seriously in our philosophical theorizing.

Weinberg *et al.* argue for the following two claims:

> [F]irst . . . that a sizeable group of epistemological projects—a group which includes much of what has been done in epistemology in the analytic tradition—

[71] After all, such a weak response is fitting because I do not think either meta-objection under consideration is meant to be offering a "conclusive argument" (Weinberg, Nichols, and Stich 2001, 429).

> would be seriously undermined if one or more of a cluster of empirical hypotheses about epistemic intuitions turns out to be true . . . . [S]econd . . . that, while the jury is still out, there is now a substantial body of evidence suggesting that some of those empirical hypotheses *are* true. (Weinberg, Nichols, and Stich 2001, 429)

These two claims, of course, can be nicely summarized in a simple modus ponens argument:

> **P1**: If one or more of a cluster of empirical hypotheses are true, a sizable group of epistemic projects would be seriously undermined.
>
> **P2:** One or more of a cluster of empirical hypotheses are true.
>
> **C:** Therefore, a sizable group of epistemic projects are seriously undermined.

While their goal is expressly "not to offer a conclusive argument demonstrating that . . . [such] epistemological projects . . . are untenable," they do want to force those of us who take such projects seriously to offer a defense—to put the burden of proof on those who wish to work within that given "sizable group of epistemological projects" (Weinberg, Nichols, and Stich 2001, 430).

Now, just what is this "sizable group of epistemological projects" that are so threatened by Weinberg *et al.*'s yet unspecified empirical hypotheses? What are the defining characteristics of the epistemic projects Weinberg *et al.* wish to critique? To quote, the epistemic projects Weinberg *et al.* have in mind will meet the following three conditions:

> (i) The [project] must take epistemic intuitions as data or input. (It can also exploit various other sorts of data.)
>
> (ii) It must produce, as output, explicitly or implicitly normative claims or principles about matters epistemic. Explicitly normative claims include regulative claims about how we ought to go about the business of belief formation, claims about the relative merits of various strategies for belief formation, and evaluative claims about the merits of various epistemic situations. Implicitly normative claims include claims to the effect that one or another process of belief formation leads to justified beliefs or to real knowledge or that a doxastic structure of a certain kind amounts to real knowledge.
>
> (iii) The output of the [project] must depend, in part, on the epistemic intuitions it takes as input. If provided with significantly different intuitions, the strategy must yield significantly different output. (Weinberg, Nichols, and Stich 2001, 432)

Weinberg *et al.* call the "sizable group of epistemological projects" that satisfy the above conditions *Intuition-Driven Romanticism* (hereafter IDR).[72] And most importantly given our purposes, one of the projects under the umbrella of IDR is the enterprise of taking the intuitions generated by Gettier-style counterexamples as of normative import in the project of defining knowledge. Subsequently, it is Weinberg *et al.*'s contention that if one or more of the given set of yet unspecified hypotheses is true, the strategy of using the Gettier Problem to motivate various epistemic conclusions (the very thing being done in this dissertation) is "seriously undermined."

Just what are these volatile hypotheses that Weinberg *et al.* mentioned, which, if true, seriously undermine various projects in analytic epistemology (including responses to the Gettier Problem)? They are:

> **Hypothesis 1:** Epistemic intuitions vary from culture to culture.
>
> **Hypothesis 2:** Epistemic intuitions vary from one socioeconomic group to another.
>
> **Hypothesis 3:** Epistemic intuitions vary as a function of how many philosophy courses a person has had.
>
> **Hypothesis 4:** Epistemic intuitions depend, in part, on the order in which cases are presented. (Weinberg, Nichols, and Stich 2001, 437–438)

The general worry, here, is clear enough. If people of other demographics (i.e., people of demographics outside professional analytic epistemology) have significantly different intuitions regarding foundational epistemic issues, then, assuming we have no plausible way to champion our own intuitions over all others, the motivational force behind IDR projects seems to evaporate—IDR will be "significantly undermined."

But why think at least one of the above hypotheses are true? Why, for example, think that people in different cultures or people with different socioeconomic backgrounds have radically divergent epistemic intuitions about key issues in epistemology? Well, straightforwardly, Weinberg *et al.* asked people from different cultures and socioeconomic

[72] "Intuition-Driven Romanticism" is a rhetorically charged term and, regardless of Weinberg *et al.*'s intentions, it is something of a guise. If what we are doing in contemporary epistemology only has roots that go back to nineteenth-century Romanticism, then an "uprooting" proposal (so to speak) like Weinberg *et al.*'s will not seem that implausible. However, if what we are doing in contemporary epistemology has roots that go far deeper (e.g., arguable back to Plato), as Weinberg *et al.* seem to concede, then we should expect an uprooting proposal to be far more cataclysmic and, as such, far less plausible.

backgrounds what their intuitions are regarding key issues in epistemology and found that they do indeed differ starkly from the intuitions of highly educated Westerners (the general demographic of analytic philosophers). In other words, Weinberg *et al.* tested Hypothesis 1 and Hypothesis 2 using empirical experiments (which Weinberg *et al.* called "intuition probes"), and, while they admit their findings are not conclusive, their findings suggest that both hypotheses are quite plausibly true. Using a sociological "ethnic identification questionnaire," Weinberg *et al.* were able to classify people around New Brunswick, New Jersey (usually undergraduates at Rutgers University), as "Western," "East Asian," etc. (Weinberg, Nichols, and Stich 2001, 439, 457). So classified, Weinberg *et al.* gauged the epistemic intuitions of different ethnicities using a simple survey. For example, both East Asians and Westerners were asked whether the protagonist in the following Gettier counterexample either "Really Knows" or "Only Believes":

> **American Car:** Bob has a friend, Jill, who has driven a Buick for many years. Bob therefore thinks that Jill drives an American car. He is not aware, however, that her Buick has recently been stolen, and he is also unaware that Jill has replaced it with a Pontiac, which is a different kind of American car. Does Bob really know that Jill drives an American car, or does he only believe it? (Weinberg, Nichols, and Stich 2001, 443)

Interestingly, Weinberg *et al.* found that there was a stark difference here between the intuitions of Westerners and the intuitions of East Asians. Consider the following graph portraying how Westerners and East Asians responded respectively:

Figure 2.2

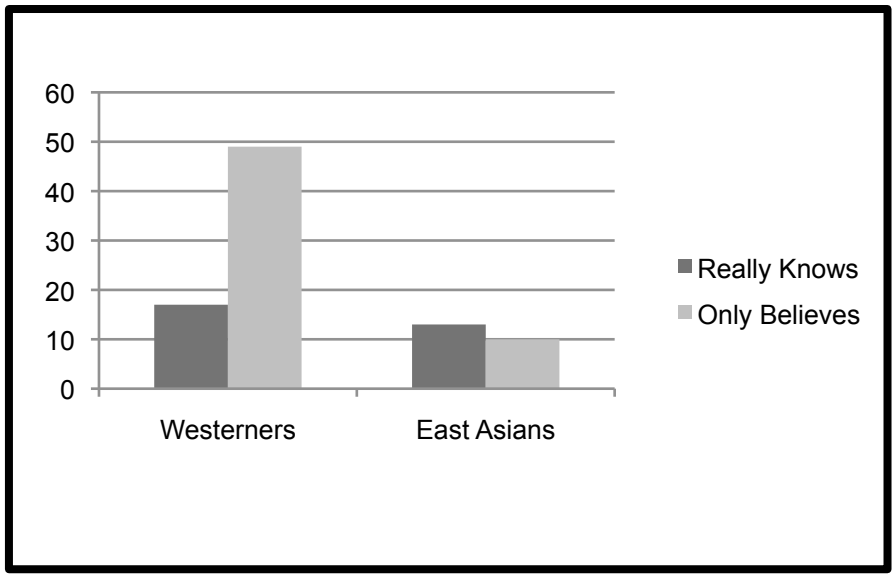

To quote Weinberg *et al.*: "The striking finding in this case is that a large majority of [Westerners] give the standard answer in the philosophical literature, viz., 'Only Believes.' But among [East Asians] this pattern is actually *reversed*! A majority of [East Asians] say that Bob really knows" (Weinberg, Nichols, and Stich 2001, 443). And given that Weinberg *et al.* found similar incongruity between the epistemic intuitions of various ethnic and socioeconomic groups in similar tests, doesn't that give us good reason to suspect the philosophical import of intuitions? Doesn't that undermine their use in philosophical argument?

Not necessarily. To start, we have very good reason to be skeptical of Weinberg *et al.*'s findings—of their empirical experiments or "intuition probes." First of all, it's worth noting that sociological surveys are tricky endeavors—fraught with perils of false data and false conclusions—even when conducted by trained sociologists. When *philosophers* like Weinberg *et al.* undertake a transcultural sociological survey, we should, at the very least, be extremely skeptical regarding any conclusions they prescribe. For example, East Asians may have a radically different conception of "knows" than Westerners do.[73] As such, if they were to reflect on "knowledge" in the way that we mean "knowledge" their intuitions may not vary at all. In other words, if their cultural perspective is so different as to suggest that they may have different intuitions on various cases, why can't it be so different as to suggest that they didn't understand the question as we intended? Weinberg *et al.*'s argument is self-defeating. But what is more, the sampling done by Weinberg *et al.* is incredibly questionable on at least two counts: 1) surveying undergraduates at one of the researcher's (in this case Stich's) own institution probably is not a sociologically respectable way to collect genuine data. 2) Weinberg *et al.*'s sample size in their experiments is abysmally small—sometimes as few as 8 (!) individuals are surveyed to represent a population of over a billion souls. As such, the margin of error in Weinberg *et al.*'s findings is going to be staggeringly huge—making their results very easy to simply dismiss offhand.

[73] Weinberg *et al.* respond to an objection along these lines, where it is suggested that those surveyed were thinking of "know" more informally, as a term expressing a high degree of subjective certainty (see Weinberg, Nichols, and Stich 2001, 449-450). But this is not quite what I mean here; my worry is not that the survey subjects were using "know" informally; my worry is that the East Asians, for example, understand "know" in a fundamentally different way than Westerners do.

Even so, don't Weinberg *et al.*'s findings, suspect though they may be, highlight a legitimate worry? Perhaps so. After all, perhaps where we would have formerly dismissed the suggestion that different cultures or socioeconomic groups might have significantly different epistemic intuitions as either impossible or a fiction, Weinberg *et al.*'s findings might at the very least give us reason now to worry that one or more of the proposed hypotheses might be true.[74] Consider the following two points in response. *First of all,* even if Weinberg *et al.*'s findings do suggest that intuitions will vary according to various demographics, this does not mean we should embrace their conclusion. Weinberg *et al.*'s argument will presumably generalize. If we dismiss large swaths of contemporary epistemology because central intuitions are not always shared transculturally, then we should probably dismiss large portions of metaphysics, ethics, logic, etc., on the same grounds. What is more, there is a real worry that Weinberg *et al.*'s argument will generalize outside of philosophy. It seems reasonable to suspect that intuitions lie at the heart of not only philosophy but also other disciplines across the arts and sciences. If this is right, Weinberg *et al.*'s argument would presumably lead to an intractable form of skepticism. *Secondly*, there is nothing intrinsically wrong with supposing that one group or class is right and everyone else is wrong. This is not, to be sure, anything that is politically incorrect. It may just be that no one has reflected on technical Western conceptions of knowledge (what we're talking about) quite as well as technical Westerners. We do not have any reason to suppose that people of different groups would come to different conclusions once they were sufficiently informed. Intuitions can, after all, be corrected, so the intuitions of the various opposing groups need not be set in stone.

Section 3.2: Weatherson's Objection

In "What Good are Counterexamples?" (2003), Brian Weatherson calls into question the intuitions behind Gettier counterexamples—arguing that they are (at least in the case of the Gettier Problem) unreliable and should not be respected at the loss of the standard analysis of knowledge. While a large portion of the paper (over half) is devoted to elucidating the naturalness of meaning in proposed definitions of knowledge and defending the standard analysis against psychologists (see Rosch and Mervis 1975) and those who would take the standard analysis as woefully out of step with folk epistemology (see Stich 1988), our main concern in this subsection is Weatherson's critique of the Gettier intuition. I will argue 1)

[74] See Stich 1993; Weinberg, Nichols, and Stich 2001, 435.

that Weatherson's reasons for questioning the Gettier intuition are largely misguided and 2) that following Weatherson in discrediting the Gettier intuition generalizes widely so as to lead to "an untenable kind of skepticism" (Weatherson 2003, 4).

Epistemologists, according to Weatherson, generally take intuitions regarding counterexamples far more seriously than ethicists. When an epistemologist encounters a counterexample to a given analysis of knowledge, the classic response has been to reject the given analysis—the counterexample being independently sufficient for undermining the said analysis. Ethicists, according to Weatherson, are far more inhibited regarding the ramifications for counterexamples. When a utilitarian, for example, is presented with a counterexample (perhaps a scenario that seems to suggest that torturing a few individuals to increase the happiness of the majority is, counterintuitively, the moral course of action), the typical response is *not* to dismiss utilitarianism at once. While such counterexamples are taken seriously, ethicists simply do not seem to attribute the same theoretical import to counterexamples as epistemologists do. And at the very least, Weatherson wants to argue that epistemologists should take a cue from ethicists here—considering counterexamples in light of additional factors.

After all, as Weatherson notes, "whether something seems to be true can be independent of whether we believe it to be true" (Weatherson 2003, 3). In other words, just because a given counterexample suggests that such and such, it need not necessitate our believing that such and such. For example, to quote Weatherson:

> Frege's Axiom V seems to be true though we know it is false. It does not seem to be the case, in the relevant sense, that 643 x 721 = 463603. Unless one is rather good at mental arithmetic, there is nothing that 643 x 721 seems to be; it is out of the reach of intuition. . . . One can judge that something seems to be the case while neither believing nor disbelieving it. This is a sensible attitude to take towards the view that one cannot *know* that a particular ticket will lose in a fair lottery. This is despite the fact that it certainly *seems* one cannot know this. (Weatherson 2003, 3)

Such a realization can seemingly help us limit the damage counterexamples do to a given theory. Just as there may be cases where utilitarianism seems to get the wrong result, we may very well go on believing in utilitarianism for other reasons. Accordingly, Gettier counterexamples may very well seem to be cases where warranted true beliefs fail to acquire knowledge, but that need not keep us from believing that the warranted true belief analysis of knowledge is at base correct for other reasons.

But what is more, given that intuitions (such as our intuition that Frege's Axiom V is true) are regularly mistaken, it seems as though we need independent reason to think that the intuitions at play in the relevant cases (i.e., Gettier cases) are right before we can assign them any significant philosophical importance. To quote Weatherson:

> Given . . . that the example of Axiom V shows that seemings can be mistaken, what evidence have we that they are not mistaken in cases we consider here? Arguably, we have very little indeed. Robert Cummins (1998) argues that in general intuition should not be trusted as an evidential source because it cannot be calibrated.[75] We wouldn't have trusted the evidence Galileo's telescope gave us about the moon without an independent reason for thinking his telescope reliable. Fortunately this can be done; we can point the telescope at far away terrestrial mountains, and compare its findings with the findings of examining the mountains up close and personal. There is no comparable way of calibrating intuitions. Clearly we should [be] suspicious of any method that has been tested and found unreliable, but there are tricky questions about the appropriate level of trust in methods that have not been tested. (Weatherson 2003, 4)

People do indeed seem to regularly have faulty intuitions regarding questions of logic (cf. the rate of error in the Wason Selection Task), and people do indeed regularly have conflicting if not contradictory intuitions regarding various moral issues (after all, we can't all be right), so the evidential role of intuitions should seemingly be duly limited until we have reason to think the given intuition is a good (i.e., a reliably truth-conducive) one.

All this is to say we should not necessarily let some pesky counterexamples get in the way of championing "a simple, systematic and largely successful theory" (Weatherson 2003, 1). In other words, given that intuitions are not necessarily belief guiding and given that intuitions are regularly mistaken, we should not, according to Weatherson, abandon the standard analysis of knowledge in the face of Gettier counterexamples. The quality of a given theory should (minus reasons to the contrary) trump counterexamples. The Gettier Problem, according to Weatherson, "is not in itself decisive" (Weatherson 2003, 2).

When determining the robustness of a given theory, we should not simply consider whether or not it is prone to counterexample (important though that may be), but we should also consider whether or not the given theory is *systematic*. According to Weatherson:

> [T]he true theory of knowledge is the one that does best at (a) accounting for as many as possible of our intuitions about knowledge while (b) remaining systematic. A "theory" that simply lists our intuitions is no theory at all, so condition (b) is vital. And it is condition (b), when fully expressed, that will do

[75] See Cummins 1998.

> most of the work in justifying the preservation of the [standard analysis] theory in the face of the counterexamples. (Weatherson 2003, 7)

To elaborate, there are, according to Weatherson, four criteria, four tests, we should use to judge a given philosophical theory. First, as we might expect, a given theory should avoid counterexamples—"counterexamples to a theory," as Weatherson acquiesces, "count against it" (Weatherson 2003, 8). Second, a given philosophical theory should not "have too many *theoretical* consequences which are unacceptable" (Weatherson 2003, 8). Third, a given philosophical theory should analyze concepts that are "theoretically significant" and in such a way that they can be "analysed in other theoretically significant terms" (Weatherson 2003, 9). And fourth, a given philosophical theory should be appraised in accord with its simplicity—a successful philosophical theory "must be simple" (Weatherson 2003, 9).

While it is not entirely clear what these criteria amount to, the conclusion Weatherson draws from them is explicit:

> My main claim is that even once we have accepted that the [standard analysis] theory seems to say the wrong thing about Gettier cases, we should still keep an open mind to the question of whether it is true. The right theory of knowledge, the one that attributes the correct meaning to the word "knows," will do best on balance at these four tests. Granted that the [standard analysis] theory does badly on test one, it seems to do better than its rivals on tests two, three, and four, and this may be enough to make it correct. (Weatherson 2003, 10)

Why think a thing like that? Why, for example, think that the standard analysis is better at avoiding unacceptable theoretical consequences, better at analyzing significant concepts, and simpler than competing definitions of knowledge such as, say, Timothy Williamson's? Weatherson simply does not say. He simply assumes it to be such. But without any clear reason to champion the standard analysis over rival theories using Weatherson's criteria, the standard analysis's shortcomings regarding Gettier counterexamples is uncomfortably conspicuous. Nevertheless, perhaps Weatherson's weaker point stands—perhaps we should "keep an open mind" regarding the truth of the standard analysis even in light of Gettier counterexamples. After all, if the standard analysis is, as Weatherson claims, "simple, systematic, and largely successful," perhaps we should give it more credit.

While an acquiescence that the standard analysis *may* yet be a successful philosophical theory despite the Gettier Problem is not necessarily incompatible with my proposed thesis, we should nevertheless cast some doubt on Weatherson's optimism. Is it really the case that the standard analysis is "simple, systematic, and largely successful"? Arguably, the

standard analysis as we know it did not take shape until Edmund Gettier's famous paper. While no doubt the thought that knowledge involved warrant, truth, and belief can be found in pre-Gettier literature (e.g., Plato's *Meno*), we should nevertheless be skeptical as to whether or not viewing such concepts as conceptually primitive necessary and the jointly sufficient conditions on knowledge has a long and illustrious history. In many ways, it seems as though the Gettier Problem gave birth to the standard analysis of knowledge, in which case calling it "largely successful" is a bit hasty to say the least. Simple and systematic our contemporary understanding of the standard analysis may well be (though seemingly no *more* simple or *more* systematic than the rival theories), but in calling it "largely successful" when the problem that more or less birthed it has yet to be solved seems uncalled for.

While we certainly may question Weatherson's (unsubstantiated) optimism regarding the virtues of the standard analysis, we can also question his pessimism regarding our intuitions in Gettier counterexamples. Weatherson contrasts the way epistemologists and ethicists typically respond to intuitions in counterexamples. Epistemologists, as shown in their response to Gettier cases, typically seem to take such intuitions as independently sufficient for rejecting the given theory of knowledge. Ethicists, says Weatherson, are more reserved. When utilitarianism faces a hypothetical scenario where it intuitively produces the wrong result (results such as "it's permissible to torture a few individuals if it benefits the majority"), utilitarianism is not at once dismissed. Ethicists, it seems, put less philosophical stock in such intuitions. But it is worth pointing out that Gettier counterexamples have important disanalogies with such counterexamples contra utilitarianism. In the case of the latter, we have conflicting intuitions; under many circumstances what seems like the right thing to do is to increase happiness (or well-being) as much as possible—we have intuitions in favor of utilitarianism qua utilitarianism. As such, when utilitarians face counterexamples they usually already have strong intuitions in their favor—giving them a position from which to try to undermine or dismiss the contrary intuitions generated by the counterexample. Defenders of the standard analysis are seemingly not in the same position. Seemingly we do not have much in the way of strong intuitions in favor of the standard analysis qua the standard analysis. While we surely have the intuition that knowledge involves truth, belief, and something like warrant, none of this distinctly favors the standard analysis over, say, a Williamsonian model. As such, when epistemologists see Gettier counterexamples, it seems perfectly understandable and

reasonable when they dismiss the given analysis—it does not, after all, seem to have any intuitions going specifically in its favor. If an ethical theory was in such circumstances (with counterexamples against it and with nothing else clearly and specifically for it), ethicists seemingly would be every bit as cavalier as epistemologists in dismissing the given theory on those very grounds.

What about Weatherson's claim (via Cummins) that intuitions are generally not a trusted source of evidence since it cannot be calibrated? Someone certainly might worry that such a wholesale besmirching of intuitions may generalize to, say, perception and lead to a form of skepticism; perception, after all, might seem equally unable to be calibrated. Ernest Sosa (1998) levels just this sort of objection against Cummins, and Weatherson has a response in turn.[76] To quote:

> Ernest Sosa (1998) argues in response to Cummins that this kind of reasoning leads to an untenable kind of scepticism. Sosa notes that one can make the same point about perception as Cummins makes about intuition: we have no independent way of calibrating perception as a whole. There is a distinction to be drawn here, since perception divides into natural kinds, visual perception, tactile perception, etc., and we can use each of these to calibrate the others. It is hard to see how intuitions can be so divided in ways that permit us to check some kinds of intuitions against the others. In any case, the situation is probably worse than Cummins suggests, since we know that several intuitions are just false. (Weatherson 2003, 4)

But is this a sufficiently compelling response? Sosa's objection is that besmirching the evidential value of intuition because it cannot be calibrated means that we should seemingly equally besmirch the evidential value of perception *on the whole*. Given that such a conclusion puts us in the grip of skepticism, we shouldn't so besmirch intuition. And it is not entirely clear how it can be of any help to note that we can use some types of perception to gauge the accuracy of other types of perception. We seemingly have a whole host of intuitions that range across a wide range of domains; why can't we use our intuitions in certain domains to gauge the accuracy of intuitions in others? This is, after all, precisely what we sometimes seem to do. Say a contestant is trying to win a car on the Monty Hall show by guessing which of three doors has the car hidden behind it. When Monty reveals that one of the two unchosen doors does not have the car behind it, we may initially have the intuition that the contestant now has a 50/50 chance of winning the car (there are, after all, only two doors left) and would be in no better probabilistic position if

[76] See Sosa 1998.

he were to now switch doors. We can, however, gauge the accuracy of such an intuition by bringing to bear some of our other basic intuitions regarding probability and probabilistic reasoning; we can judge our initial intuition false while working to develop new and more robust intuitions about the contestant's chances at winning the car.[77] And pointing out that intuitions are sometimes wrong, to be sure, does nothing to assuage Sosa's worry because perceptions, clearly, are sometimes wrong too.

There is, however, another more troubling way Weatherson's besmirching (via Cummins) of the evidential value of intuitions seems to lead to an untenable kind of skepticism. There is good reason to think that the vast majority of our philosophical theorizing is critically undergirded by a vast array of intuitions, such that one should worry that sullying the evidential value of intuitions affectively takes the floor out from under us. One pressing example: if we systematically question the value of our intuitions in Gettier counterexamples, we should in the same way question Weatherson's intuition that a good theory of knowledge should, say, be systematic and simple. In questioning the Gettier intuition, Weatherson will likely have no way to champion the standard analysis because any intuition that seems to do so will be equally questionable. Not only, then, do we have a reason to think that Weatherson is leading us to a far-reaching skepticism, but we should also seemingly worry that Weatherson is leading us down a road that is ultimately self-defeating.

Both "Normativity and Epistemic Intuitions" (2001) by Jonathan M. Weinberg, Shaun Nichols, and Stephen Stich and "What Good are Counterexamples?" (2003) by Brian Weatherson call into question the philosophical import of our intuitions, particularly our epistemic intuitions when it comes to Gettier counterexamples. Given 1) that such views are undeniably in the minority (at least for those of us working within analytic epistemology) and 2) that the current thesis is ultimately aimed at those of us who do take Gettier intuitions seriously (the majority), we are not under any great onus to respond to

[77] This raises a point that Weatherson seems to overlook: that we are not stuck with the intuitions we have; our intuitions can be corrected and changed—often thanks to other intuitions. In "What Good are Counterexamples?" Weatherson makes much of the fact that people have at various points in history had intuitions that led to wildly false beliefs—such as whales are fish, the sun is not a star, and that the earth is flat. But what Weatherson does not seem to appreciate is that our intuitions regarding such beliefs seem to have been corrected, thanks to our intuitions in other domains.

such views in detail. Nevertheless, we have seen in this subsection various positive reasons to reject both Weinberg *et al.*'s and Weatherson's arguments—leaving us all the more comfortable (and justified) in working under common post-Gettier assumptions.

Conclusion

At the end of Chapter 1, we expressed some residual worries that questioned the parameters of the proposed diagnosis of the Gettier Problem: worries about whether or not assuming an inviolable relationship between warrant and truth really leads to radical skepticism; worries about whether assuming a close (though not inviolable) relationship between warrant and truth really leads to further Gettier counterexamples; worries about whether assuming that warrant bears no relationship to truth is, in fact, counterintuitive; and worries about the epistemic import of Gettier-styled counterexamples themselves. In this chapter, we have considered the best the contemporary philosophical literature has to offer in promotion of these latter three concerns, and we have found them all lacking. In the first section, we considered Daniel Howard-Snyder, Frances Howard-Snyder, and Neil Feit's paper, "Infallibilism and Gettier's Legacy" (2003), which directly objected to our conclusion in the previous chapter that a close but not inviolable relationship between warrant and truth will always lead to Gettier counterexamples; and I argued that their proposed solution to the Gettier Problem is, ironically, successful only insofar as it conforms to the proposed diagnosis. In the second section, we considered recent work by Stephen Hetherington (particularly his book *Good Knowledge, Bad Knowledge*, 2001), which proposed an analysis of knowledge that assumes a divorced relationship between warrant and truth; and I argue that Hetherington's epistemology is hopelessly counterintuitive and that it need not be advocated in light of viable alternatives. Finally, in Section 3, we considered landmark papers by Brian Weatherson (2003) and Jonathan M. Weinberg, Shaun Nichols, and Stephen Stich (2003), which call into question the epistemic import of Gettier counterexamples; and I argue that such criticisms either lead to intractable forms of skepticism or are ultimately incoherent or self-defeating.

But what about the first worry—the worry that assuming an inviolable relationship between warrant and truth will not really lead to radical skepticism? While it may seem very intuitive that requiring warrant to infallibly track truth will lead to skepticism, and while infallibilistic accounts of warrant indeed have a long history of doing just that, we do not have a systematic reason to think that assuming an inviolable relationship between warrant

and truth will *always* lead to radical skepticism. While the diagnosis of the Gettier Problem proposed in Chapter 1 gives us reason to think that Gettier counterexamples will always occur so long as we assume a close (but not inviolable) relationship between warrant and truth, it gives us no such reason to think that skepticism *always* follows infallibilism about warrant—as intuitive as such an assumption may be. And while the accounts considered thus far (e.g., Pritchard's account, Howard-Snyder *et al.*'s account, etc.) have all confirmed our suspicions, perhaps virtue epistemology has something special to offer that these other accounts do not, something that allows them to endorse an inviolable connection between warrant and truth that does not lead to radical skepticism. In the next two chapters, we will put such a worry to rest. Through an examination of two eminent virtue epistemologies, those of Alvin Plantinga and Ernest Sosa, we will find the proposed diagnosis of Gettier counterexamples vindicated and every instance of infallibilism leading to skepticism.

PART II:
APPLYING THE DIAGNOSIS

3 PLANTINGA'S VIRTUE EPISTEMOLOGY

In Chapter 1, I advocated a diagnosis of the Gettier Problem that predicted that Gettier counterexamples cannot be feasibly avoided within the reductive analysis project. I proposed that the only way to avoid Gettier counterexamples, given the nature of luck, is to assume that warrant is either completely divorced from truth or, conversely, necessarily wedded to it—neither option seeming viable. In other words, in accord with Zagzebski's 1994 diagnosis, I have argued that as long as we reasonably assume that warrant bears a close but not inviolable relationship to truth, the Gettier Problem is unavoidable (Zagzebski 1994, 65).

With this diagnosis in hand, we are now ready to consider what is perhaps the most iconic virtue epistemology of the twentieth century, Alvin Plantinga's analysis of knowledge in terms of properly functioning cognitive faculties—a view that developed throughout his monumental warrant trilogy: *Warrant: the Current Debate* (1993), *Warrant and Proper Function* (1993), and *Warranted Christian Belief* (2000). Plantinga's epistemology stands in stark contrast to our proposed diagnosis and the growing trend in contemporary epistemology to abandon the analysis of knowledge—offering what is meant to be a viable reductive account that is immune to Gettier counterexample. Our goal in this chapter is to apply the diagnosis of Gettier problems developed in Chapter 1 to each iteration of Plantinga's epistemology and argue that, in accord with Zagzebski's diagnosis, Plantinga's reductive analysis of knowledge repeatedly becomes either unpalatable or unable to surmount the Gettier Problem.

We will work toward this goal in two sections, following the chronological development of Plantinga's virtue-theoretic analysis of knowledge. In Section 1, we will elucidate and critique Plantinga's analysis of knowledge as it is found in *Warrant and Proper Function* (1993). In Section 2, we will elucidate and critique the proposed modifications to

Plantinga's original account found, first, in "Respondeo" (1996) and "Warrant and Accidentally True Belief" (1997) and, then, in *Warranted Christian Belief* (2000). In both sections, we will find our proposed diagnosis of Gettier problems vindicated, with each iteration and proposal failing precisely along the lines our diagnosis predicted.

Section 1: Plantinga's 1993 Virtue Epistemology

In *Warrant: The Current Debate* (1993), Plantinga surveys several dominant accounts of warrant (i.e., that which bridges the gap between true belief and knowledge) in the contemporary literature and subsequently argues that each fails due to an inability to track virtue-theoretic intuitions across a range of cases. It is his particular focus on the *properly functioning* human knower that is meant to distinguish his account from all others. Plantinga's notion of proper function is meant to rightly connect a given agent to the facts (cf. René and the Gambler's Fallacy) and, of course, rightly preclude knowledge from cases of malfunction (cf. Brain Lesion). According to Plantinga, other theories of knowledge fail due to their inability to track knowledge ascription in accord with the proper functioning of the relevant cognitive faculties behind a given belief's genesis. Proper function is, for Plantinga, not only the key virtue-theoretic concept in his account, it is meant to be the "rock on which" competing theories of knowledge "founder" (Plantinga 1993b, 4).

Knowledge, for Plantinga, is warranted true belief.[78] Plantinga's 1993 theory of warrant can be approximately summarized as:

> **Plantinga's 1993 Warrant:** A belief B is warranted for S when B is formed by cognitive faculties (of S's) that are functioning properly in the right environment in accord with a good design plan aimed at truth.[79]

In other words, a given belief, B, will be warranted for Plantinga if and only if the following four conditions are met:

> 1) the cognitive faculties involved in the production of B are functioning properly;
>
> 2) [the] cognitive environment is sufficiently similar to the one for which [the agent's] cognitive faculties are designed;
>
> 3) the design plan governing the production of the belief in question involves, as purpose or function, the production of true beliefs.;

[78] Or more accurately, *sufficiently* warranted true belief.

[79] See Plantinga 1993b, 19.

> 4) the design plan is a good one: that is, there is a high statistical or objective probability that a belief produced in accordance with the relevant segment of the design plan in that sort of environment is true. (Plantinga 1993b, 194)[80]

Although such a compendious rendering of Plantinga's theory can be sufficiently understood without further elucidation, we will, for diligence sake, unpack it a bit further.

Again, it is Plantinga's focus on the epistemically virtuous agent (i.e., the agent with mechanically sound cognitive equipment) that is the cornerstone to his epistemology. But what, then, is proper function? If proper function is meant to be the cornerstone of Plantinga's account of warrant, then one may worry that the term "proper function" is every bit as enigmatic as "warrant" and that introducing the former to explain the latter is counterproductive; however, Plantinga thinks that we all have a more or less rough and ready understanding of what it means for something to be functioning properly or malfunctioning (Plantinga 1993b, 5–6). We all know what it means when a car cannot go in reverse because the transmission is not (mechanically) functioning properly. We know that a properly functioning human being should generally be able to walk in a straight line without tripping or swerving and how enough alcohol impairs this proper functioning. As such, Plantinga provisionally stipulates that the relevant cognitive facilities involved in a given belief's formation must be functioning properly if the belief is to have warrant. If John is on a hallucinogenic drug, his perceptual faculties will no longer be functioning as they should; hence, his belief that the sky is melting will not have any epistemic value in terms of warrant.

Although we may indeed have a sufficient rough-and-ready grasp of proper function, it will, nevertheless, be helpful to make some general clarifications. First of all, to function properly is not to function normally (as understood in the general statistical sense). Gottlob Frege may be better at logic than the normal human being, but this does not mean that his prowess in logic is the result of some cognitive malfunction. Likewise, to use one of Plantinga's examples, if due to some disaster most everyone on earth was blinded (such blindness would be a statistically normal condition), the few sighted individuals would not be suffering from malfunctioning perceptual faculties (Plantinga 1993b, 9–10). Secondly, not *all* of one's cognitive faculties need to be functioning properly for them to produce warranted beliefs (Plantinga 1993b, 10). If my cognitive faculties associated with vision are

[80] For other tabulations of Plantinga's 1993 account of warrant (what later iterations refer to as the "nutshell" or "central core" of warrant) see Chignell 2003, 445; Plantinga 2000, 156.

faulty, that would not preclude my producing warranted beliefs with my auditory faculties. What is more, all that must be working are "the faculties (or subfaculties, or modules) involved in the production of the particular belief in question" to be warranted. For example, if I am colorblind, I can still produce all sorts of warranted visual beliefs—beliefs concerning distance, the presence of people and objects, and so on. Third and finally, proper function comes in degrees (Plantinga 1993b, 10–11). My running ability may not be functioning as well as it could (if, say, I chose to pursue athletics instead of academics in life), but that doesn't mean my running ability is somehow malfunctioning. Likewise, my cognitive faculties associated with my abilities at logic may not be as good as they could be (I could spend more time working through complex proofs, memorizing truth-tables, etc.), but I can still produce warranted beliefs based on logical deduction. How proper functioning does a given cognitive faculty need to be in order to be warrant conferring? Here, Plantinga concedes, he has no answer; however, he notes that we independently recognize that knowledge and warrant are vague to some degree—his hope, then, is that the vagueness of his theory in this instance corresponds to the vagueness of knowledge and warrant in general.

While Plantinga's focus on the mechanical proper functioning of an agent's various cognitive faculties is the cornerstone of his account, he concedes that proper function, though necessary for warrant, is not sufficient. To demonstrate, he presents us with the following case:

> **Alpha Centauri Elephants:** You have just had your annual cognitive checkup at MIT; you pass with flying colors and are in splendid epistemic condition. Suddenly and without your knowledge you are transported to an environment wholly different from earth; you awake on a planet revolving around Alpha Centauri. There conditions are quite different; elephants, we may suppose, are invisible to human beings, but emit a sort of radiation unknown on earth, a sort of radiation that causes human beings to form the belief that a trumpet is sounding nearby. An Alpha Centaurian elephant wanders by; you are subjected to the radiation, and form the belief that a trumpet is sounding nearby. (Plantinga 1993b, 6)

Your belief in such a case, though formed by properly functioning cognitive faculties, does not have warrant. What is more, even if it turned out to be true that a trumpet is sounding nearby (one is being played in a soundproof booth out of sight), your belief still wouldn't have much if any warrant.[81] Why? Because the environment on Alpha Centauri's planet is

[81] The same would be true if, say, the radiation from the elephant didn't produce the belief that "a trumpet is playing nearby," but instead that "a big grey object is nearby" (Plantinga 1993b, 7).

not congenial to your human cognitive faculties, properly functioning though they may be. A properly functioning toaster cannot be relied on to make toast underwater, because water is the wrong environment for toasters; likewise, properly functioning human cognitive faculties cannot be relied on to produce warranted beliefs on Alpha Centauri's planet (or in a vat, or under the influence of an evil demon, etc.), because the environment there is the wrong environment for human cognitive faculties.[82] What constitutes the right environment for human cognitive faculties? According to Plantinga, the kind of environment that our cognitive faculties were designed for (by God, evolution, or both) is an environment like earth with "such. . . features as the presence and properties of light and air, the presence of visible objects, of other objects detectable by cognitive systems of our kind, of some objects not so detectable, of the regularities of nature, of the existence and general nature of other people, and so on" (Plantinga 1997, 143).[83]

So, according to Plantinga, in order for the belief B to have warrant for S, it needs to be formed by cognitive faculties (of S's) that are functioning properly in the sort of environment they were designed for; however, this is not yet sufficient for warrant either. Proper function presupposes a design plan (Plantinga 1993b, 21). To assess whether or not a toaster, say, is functioning properly we need to have an idea of how it was designed to work. Likewise, in order to assess whether or not cognitive faculties are functioning properly we need to have some idea as to how they were designed to function (by God, evolution, or both). To be sure, the design plan for human cognitive faculties is something like "a set of specifications for a well-formed, properly functioning human being—an extraordinarily complicated and highly articulated set of specifications, as any first-year medical student could tell you" (Plantinga 1993b, 14). Indeed, our cognitive faculties are so complex that they are designed to occasionally function divorced from the goal of arriving at truth. Consider how women, it seems, are designed to believe in hindsight that the pain of childbirth was less than it actually was. Similarly, people seem to be designed to optimistically believe they will survive a terrible disease far beyond what statistics vindicate.

[82] So in order for a given belief to have warrant it needs to be produced by cognitive faculties that are functioning properly within a congenial environment, but just how congenial? What if I am knowingly in an uncongenial environment and know how I might be misled—would relevant beliefs still be barred from warrant? Such issues, Plantinga admits, highlight a vagueness in the theory. There may not be an answer to such questions. See Plantinga 1993b, 11.

[83] See Plantinga's description of what he meant by environment in 1993b, 6–7. Also see Plantinga 2000, 156–158.

In such cases, the beliefs formed are produced by cognitive faculties that are functioning properly in the right environment according to their design plan, but are nevertheless unwarranted because the cognitive faculties in question are not, at least in this instance, aimed at truth; rather, they are aimed at either propagation of the species or survival. In order for a belief to be warranted, not only does it need to be produced by cognitive faculties that are functioning properly in the right environment, but it also must be produced by cognitive faculties that are functioning properly according to a design plan that is aimed at truth (Plantinga 1993b, 11–14).

Finally, as Plantinga points out, even if S's cognitive faculties are functioning properly in a congenial environment and in accord with a design plan aimed at truth, the given belief is not necessarily warranted. Say an incompetent angel (or one of Hume's infant deities) set out to create a species of rational persons ("capable of thought, belief and knowledge"); however, due to the angel's incompetence the vast majority of the created persons' beliefs turn out to be absurdly false (Plantinga 1993b, 17).[84] In such a case, the given people's beliefs are produced by cognitive faculties that are functioning properly in a congenial environment in accord with a design plan aimed at truth, but the design plan just turns out to be a rubbish one—despite the incompetent angel's best efforts, the subjects he created have wholly unreliable cognitive faculties. Warrant, according to Plantinga, requires reliability; more precisely, in order for a given belief to be warranted "the module of the design plan governing the production of that belief must be such that the statistical or objective probability of a belief's being true, given that it has been produced in accord with that module in a congenial cognitive environment, is high" (Plantinga 1993b, 18).[85] In other words, to use the language of Plantinga's 1993 Warrant specified earlier, the design plan of the relevant cognitive faculties must be a "good" one if a belief produced by those faculties is to have any hope for warrant.

[84] Also see Hume 1948, part 5.

[85] How high must the statistical objective probability of the beliefs being true be? Here again, Plantinga concedes that there is vagueness. For Plantinga, the amount of warrant a given belief enjoys is relative to how firmly it is believed by the given agent. However high the statistical objective probability of the belief's truth needs to be, Plantinga presumes that "the degree of reliability varies as a function of degrees of belief" (Plantinga 1993b, 18).

Now that we have elucidated Plantinga's 1993 account of warrant, we are ready to consider how Plantinga's 1993 theory of knowledge, *warranted true belief*, fares against the Gettier Problem. Given the proposed diagnosis of Chapter 1, this comes down to what relationship Plantinga assumes warrant bears to truth—whether it bears an inviolable, a close, or a divorced relationship to truth. Obviously, given Plantinga's agent reliabilism such that warrant is meant to reliably track truth, Plantinga would not endorse a divorced relationship between warrant and truth. The question then becomes: Is it possible, on Plantinga's account, to form a warranted false belief? If it is possible, then warrant is not bearing an inviolable relationship to truth. If it is not possible, then warrant is bearing an inviolable relationship to truth. Thankfully, Plantinga seems to answer this very question for us:

> On an adequate account of warrant, what counts is not whether my experience somehow *guarantees* the truth of the belief in question (and how *could* it do a thing like that?), but whether I hold it with sufficient confidence and whether it is produced in me by cognitive faculties successfully aimed at the truth and functioning properly in an appropriate environment. If so, it has warrant; and if it is also true it constitutes knowledge. (Plantinga 1993b, 55 – emphasis Plantinga's)[86]

Plantinga is (reasonably) assuming that warrant bears a close but not inviolable relationship to truth. It is, on Plantinga's account, possible to have a warranted false belief. On the proposed diagnosis, then, Plantinga's 1993 theory of warrant, his 1993 epistemology, will be susceptible to Gettier counterexample.

That certainly is not a result Plantinga (in 1993) would have predicted. In *Warrant and Proper Function* (1993), Plantinga interpreted Gettier problems *as problems specifically for internalism*; the lesson to be learned from Gettier problems, he says, is that "internalist accounts of warrant are fundamentally wanting" (1993b, 33).[87] For Plantinga, internalism

---

[86] Quoted in Klein 1996, 106. Notice that Plantinga seems to express one of the guiding intuitions of this thesis: that infallibilism about warrant is ultimately unfeasible. Plantinga also seems to expressly allow for warranted false beliefs in 1993b, 40–41. To be sure, in later iterations of his epistemology, Plantinga denies the possibility of warranted false beliefs in an attempt to surmount Gettier counterexamples (see Plantinga 1996, 311–312, 329). We will consider the success of such a venture in the next section.

[87] We should note that it is simply not true that internalist accounts of warrant simpliciter are "fundamentally wanting" in regard to avoiding Gettier counterexamples. It is, after all, quite commonplace to identify someone as an internalist who sees having "some sort of special cognitive access to whatever confers warrant for a given belief" as a necessary condition on warrant even though such a condition is not, on its own, sufficient for it. After all, one's theory of warrant could consist of an internalistic theory of justification and a

requires a given agent to have some sort of special cognitive access to whatever confers warrant for a given belief if that belief is to be warranted. The point of Gettier cases for Plantinga, then, is that having special cognitive access to whatever confers warrant on my belief that, say, "that is a barn," "Jones is going to get the job," or "Smith owns a Ford" is not sufficient for that belief to be warranted. And this leads Plantinga to a rather bizarre conclusion that stands in opposition to our proposed diagnosis: "We should . . . expect that an externalist account such as the present account will enjoy a certain immunity to Gettier problems" (1993b, 36–37).[88] Thankfully for our diagnosis, however, this was a conclusion Plantinga was forced to abandon. In the Kvanvig edition, *Warrant in Contemporary Epistemology* (1996), and elsewhere, Plantinga comes face to face with Gettier counterexamples to his externalist (1993) epistemology. And in his 1996 essay, "Respondeo," Plantinga again asks "What are Gettier problems?" and though he is still quick to posit that Gettier problems show internalist theories of knowledge to be wanting, he nevertheless acknowledges that Gettier problems are not localized to such theories; instead, simply focusing on the luck involved in Gettier counterexamples—an approach that is more in line with the diagnosis offered in Chapter 1 (1996, 309).

And indeed, there was a flurry of cases that demonstrated the vulnerability of Plantinga's 1993 theory of knowledge to Gettier counterexample. For example, Linda Zagzebski, in "The Inescapability of Gettier Problems" (1994), offered the case of Lucky Mary:

> **Lucky Mary:** Suppose that Mary has very good eyesight, but it is not perfect. It is good enough to allow her to identify her husband sitting in his usual chair in the living room from a distance of fifteen feet in somewhat dim light (the degree of dimness can easily be specified). She has made such an identification in these circumstances many times. . . . There is nothing at all unusual about either her faculties or the environment in these cases. Her faculties may not be functioning perfectly, but they are functioning well enough, so that if she goes on to form the belief "My husband is sitting in the living room," that belief has enough warrant to

---

suitably strong safety theory or a suitably strong defeasibility condition. Perhaps a lesson to be learned from the Gettier Problem is that pure internalism (i.e., a theory of warrant that consists solely of an internalistic justification condition) is "fundamentally wanting," but surely, as I argued in Chapter 1, the more general lesson is that knowledge is incompatible with a certain species of luck, regardless of the theory of warrant employed. For more on Plantinga's objections to internalism see Feldman 1996, 199–209.

[88] This is a bizarre conclusion at least in part because there are numerous counterexamples in the literature that do not rely on an internalistic conception of warrant or justification. For example, consider Fake Barn-type cases.

> constitute knowledge when true and we can assume that it is almost always true. . . . Suppose Mary simply misidentifies the chair-sitter who is, let us suppose, her husband's brother. Her faculties may be working as well as they normally do when the belief is true and when we do not hesitate to say it is warranted in a degree sufficient for knowledge. . . . Her degree of warrant is as high as it usually is when she correctly identifies her husband. . . . We can now easily emend the case as a Gettier example. Mary's husband could be sitting on the other side of the room, unseen by her. (Zagzebski 1994, 67–68)

Similarly, Peter Klein, in his paper "Warrant, Proper Function, Reliabilism, and Defeasibility" (1996), offered the case of Lucky Ms. Jones:

> **Lucky Ms. Jones:** Jones believes that she owns a well-functioning Ford. She forms this belief in perfectly normal circumstances using her cognitive equipment that is functioning just perfectly. But as sometimes normally happens (no deception here), unbeknownst to Jones, her Ford is hit and virtually demolished—let's say while it is parked outside her office. But also unbeknownst to Jones, she has just won a well-functioning Ford in the Well-Functioning Ford Lottery that her company runs once a year. (Klein 1996, 105)

In both of these cases, the protagonist in question seems to be using properly functioning cognitive faculties. Arguably there is nothing wacky about the environment in these cases—the respective environments are perfectly earth-like, the cognitive faculties in question are not operating in a brain in a vat or on a planet of Alpha Centauri, and there are no liars or illusions at work. And presumably the relevant cognitive faculties in both cases are operating in accord with a design plan that is both (i) good and (ii) aimed at truth; unless we can distinguish the cognitive faculties at work in these cases from the everyday cognitive faculties of perception, memory, or credulity, denying that they meet either of these conditions will lead us directly to some unhappy skeptical conclusions.

But before we chalk one up for our proposed diagnosis of the Gettier Problem, we need to consider Plantinga's expressed methods for defending his 1993 account against Gettier cases. In *Warrant and Proper Function* (1993), Plantinga offers two distinct strategies for handling Gettier cases. Consider the second case from Edmund Gettier's original 1963 bombshell, "Is Justified True Belief Knowledge?":

> **Classic Case 2**: Smith comes to Jones bragging about his new Ford—providing manifold justification to this effect (showing Jones the bill of sale, showing Jones the title to the car, taking Jones for a ride in it, etc.). Based on this superb evidence, Jones comes to the belief "Smith owns a Ford." For whatever reason (perhaps Jones simply loves to play with logic), Jones adds to that belief the disjunct "or Brown is in Barcelona," where Brown is an acquaintance of Jones's whom he has no reason to think is in Barcelona. By chance, it turns out that Smith does not own a Ford. Also by chance, it just so happens that Brown is indeed in Barcelona. As such, Jones's belief that "Smith owns a Ford or Brown is in

> Barcelona" turns out to be true and highly justified (based on the aforementioned evidence), though few would call it a case of knowledge. (paraphrased from Gettier 1963, 122–123)

On the face of it, Jones's belief could easily seem to meet all the conditions of Plantinga's theory of warrant. Jones's cognitive faculties seem to be functioning properly and the environment does not seem overly bizarre (after all, Jones is not a brain in a vat, on Alpha Centauri's planet, or anything like that), so the first two conditions on warrant seem to be met. In prudential cases where, say, a patient believes he will survive an illness against the odds, Plantinga would want to say that the relevant segment of the patient's cognitive faculties' design plan governing the production of his belief is aimed at survival and not truth, so in such cases Plantinga's second condition would not be met; however, nothing like this seems to be the case in Classic Case 2, so Plantinga's third condition on warrant seems to be met. And finally, we do not have any reason to think Plantinga's general reliability condition, the "good" design plan condition, is not met either. However, in *Warrant and Proper Function* (1993), Plantinga argues, first, that the environment in Gettier counterexamples is somehow uncongenial for the relevant cognitive faculties—that his second condition on warrant is unmet in Gettier cases like Classic Case 2. And later, Plantinga argues that the relevant segments governing the production of a Gettierized belief are not, in fact, "aimed at truth"—denying that his third condition on warrant is met. Let us consider both strategies in turn.

Consider the strategy of denying that Plantinga's second condition is met, the environmental condition on warrant. To quote Plantinga:

> The locus of infelicity, in these cases too, is not the cognitive faculties of the person forming the justified true belief that lacks warrant; they function just as they should. The locus is instead on the cognitive environment; it deviates, ordinarily to a small degree, from the paradigm situations for which the faculty in question has been designed. . . . What we have in Gettier situations is a belief's being formed in circumstances differing from the paradigm circumstances for which our faculties have been designed. (Plantinga 1993b, 35)[89]

Jones's belief is formed out of his cognitive faculty of *credulity*—"whereby for the most part we believe what our fellows tell us"—and, according to Plantinga, such a faculty is not suited for environments with liars in them (Plantinga 1993b, 33).[90] Again, to quote Plantinga:

[89] Quoted in Feldman 1996, 211.

[90] Also see Feldman 1996, 211.

> Still, credulity is part of our design plan. But it does not work well when our fellows lie to us or deceive us in some other manner. . . . It does not work well in the sense that the result of its proper function in those circumstances does not or need not result in true belief. More exactly, it's not that *credulity* does not work well in these cases—after all, it may be working precisely according to its specifications in the design plan. It is rather that credulity is designed, we might say, to work in a certain kind of situation (one in which our fellows do not intend to mislead us), and designed to produce a given result (our learning something from our fellows) in that situation. But when our fellows tell us what they think is false, then credulity fails to achieve the aimed at result. (Plantinga 1993b, 34)[91]

So while we may grant that Jones's cognitive faculties are functioning properly (including the specific faculty of *credulity*) and that the relevant cognitive faculties at work are aimed at truth (even though they miss their mark in this case), Plantinga will seemingly want to deny that the environment is suitable for Jones's faculties—Plantinga's second condition for warrant.

To be sure, the environment in which Jones forms his belief must be *generally* suitable for his faculties. For example, surely Jones knows that Smith is talking to him, that Barcelona is a real place, that he is being appeared to in such and such a fashion, etc. As such, Plantinga would not want to say that the environment is bad *simpliciter* because in so doing he would be overly limiting what we can know. Instead Plantinga will want to say that the environment is bad for the cognitive faculty of *credulity* in particular. While the cognitive faculty of credulity is not meant for an environment with liars, such as Smith, such an environment does not prohibit perceptual beliefs, say. Thus, Plantinga can sustain the intuition that Jones knows he is talking to Smith while denying that Jones's belief that "Smith owns a Ford or Brown is in Barcelona" meets the conditions for warrant.

However, as Richard Feldman notes in "Plantinga, Gettier, and Warrant," such a reply to Classic Case 2 "turns on untestable and implausible assumptions about the circumstances for which our cognitive faculties were designed and on equally questionable assumptions about the individuation of our cognitive faculties" (1996, 212). Let us say that instead of being duped by Smith, Jones is an acute cynic and able to aptly detect most liars like Smith. In other words, let us say that Jones could tell Smith was lying about owning a Ford. After all, Plantinga himself recognizes that people can "learn to believe some people under some circumstances and disbelieve others under others" (Plantinga 1993b, 33).[92] Indeed, I think we can reasonably posit that cynic-Jones is so good at detecting liars that in

[91] Quoted in Feldman 1996, 211.

[92] Referenced in Feldman 1996, 213.

Classic Case 2 he would form a warranted belief that "Smith does not own a Ford." As such, Feldman levels the following dilemma-prompting question: "What faculty produced *that* warranted belief and what [environment] is it designed for?" (Feldman 1996, 212 - emphasis mine). If we say that cynic-Jones's belief is formed using his faculty of *credulity*, then it looks like Plantinga is wrong to say that credulity cannot operate in an environment with liars. Alternatively, if we say that cynic-Jones formed his belief about Smith not owning a Ford using another faculty, the faculty of *incredulity*, say, then Plantinga's view begins to look more and more ad hoc.

Other cases pose an even greater challenge for this strategy—the strategy of faulting the general cognitive environment in Gettier scenarios. Consider the pre-Gettier Gettier case developed originally by Bertrand Russell:

> **Russell's Stopped Clock:** I glance at a clock, forming the opinion that it is 3:43 pm; as luck would have it, the clocked stopped precisely twenty-four hours ago. The belief I form is indeed true; again, however, it is true "just by accident" (the clock could just as well have stopped an hour earlier or later); it does not constitute knowledge. (Plantinga 2000, 157)

Naturally, Plantinga will want to deny that the protagonist in this example has a warranted belief, and, as in Classic Case 2, he will do this by denying that the environment is congenial to the formation of the belief that "it is 3:43 pm." Again, Plantinga cannot say stopped clocks are epistemically incompatible with our cognitive faculties simpliciter; after all, we can seemingly know all sorts of things in such environments (i.e., that there is a clock in the room, what time the clock says, etc.) that we would not want to deny. So, as in Classic Case 2, Plantinga will want to say that stopped clocks are epistemically incompatible with a specific cognitive faculty. But what could that specific faculty be? Our time-telling faculty? Seemingly, as Feldman points out, whatever answer we choose it will come across as "entirely arbitrary and unsupportable" (1996, 213). Even if we allowed for ad hoc specific cognitive faculties, it is unclear how helpful they would be. If this is indeed the state Plantinga's account of warrant has come to, "there's little by way of a general theory and nothing that is illuminating or informative about why there is no warrant in Gettier cases" (Feldman 1996, 213).

In the second chapter of *Warrant and Proper Function,* Plantinga goes on to develop a completely different kind of Gettier-defense. Now, instead of denying the congeniality of the relevant general cognitive environments (i.e., the strategy of denying that the second condition on warrant is met), Plantinga denies that "the segment of the design plan

governing the production of . . . [Gettierized] belief is aimed at the production of true beliefs" (i.e., the strategy of denying that the third condition on warrant is met). According to Plantinga, though our cognitive faculties are generally apt at acquiring true beliefs, due to limitations in our design there are times in which our beliefs are formed as a result of what Plantinga calls a "compromise" or "trade-off." For example:

> [S]traight sticks look bent in water; it can falsely appear that there is an oasis just a mile away; a dry North Dakota road looks wet on a hot summer day; an artificial apple among the real apples in the bin can deceive almost anyone (at least so long as she doesn't touch it); a hologram of a ball looks just like a ball and can confuse the unwary perceiver; in those famous Müller-Lyer illusions, the shorter line looks the longer because of the direction of the arrow heads. (Plantinga 1993b, 38)

According to Plantinga, the beliefs formed as a result of such tricks and illusions, though formed by cognitive faculties operating in accord with their design plan in the appropriate environment, fail to be warrant simply because the segment governing the production of said belief is not aimed at truth. Plantinga offers the following explanation:

> [C]onsider these perceptual illusion cases, and for definiteness imagine that our faculties have actually been designed; and then think about these matters from an engineering and design point of view. The designer aims at a cognitive system that delivers truth (true beliefs), of course; but he also has other constraints he means or needs to work within. He wants the system to be realized within . . . a humanoid body . . . in a certain kind of world, with certain kinds of natural laws or regularities. . . . He also wants the cognitive system to draw an amount of energy consonant with our general type of body, and to require a brain of only modest size. . . . So the designer's overall aim is at truth, but within the constraints imposed by these other factors; and this may require trade-offs. (Plantinga 1993b, 38–39)[93]

In Classic Case 2, even though Jones's cognitive faculties are generally aimed at truth, other goals in Jones's design qua Jones such as efficiency and modest brain size sometimes take precedence. As such, employing the second strategy, Plantinga denies that Jones's belief that "Smith owns a Ford or Brown is in Barcelona" is warranted because the segment of his cognitive faculties' design plan that produced it was not aimed at truth but rather something else (perhaps efficiency).

As Richard Feldman notes, such a response to Classic Case 2 is a bit "puzzling" (1996, 214). As we just discussed, Plantinga has already argued that a belief like Jones's fails to be warrant because the cognitive faculty that produced it, namely credulity, was operating in an improper environment; however, Plantinga now seems to be ignoring the environment

[93] Quoted in Feldman 1996, 213–214.

proviso.[94] What is more, although Plantinga had just said that *credulity* as a faculty is aimed at truth, now it seems as though Plantinga is saying that this particular exercise of *credulity* is not aimed at truth; however, as Feldman clarifies, seemingly "what Plantinga has in mind is that the specific "segment" of the design plan involved in producing the belief that Smith owns a Ford is not something so general as *credulity*" (Feldman 1996, 214). So though *credulity* in *general* may be aimed at truth, perhaps the particular exercise of it found in Classic Case 2 is not because the *specific* design segment at play in the given belief's genesis.

Plantinga describes these "segments" of a given faculty's design plan in terms of triplets—composed of a circumstance, a belief, and a purpose (Plantinga 1993b, 22).[95] So in a standard (i.e., non-Gettierized) case of a belief based on testimony, the segment triplet governing its production, as Feldman notes, might be something like this:

> <honest testifier, belief that *p*, truth> (Feldman 1996, 215)

In Classic Case 2, however, the segment governing the production of Jones's belief could perhaps be construed quite differently:

> <dishonest testifier, belief that Jones owns a Ford, efficiency> (Feldman 1996, 214)

As such, given that the segment of Jones's cognitive faculties that produced his belief is not aimed at truth, the belief in this instance is not warranted—thus defusing the case.

To be sure, there is something to this strategy. Surely it is the case that many of the Gettier counterexamples would dissipate if only our cognitive faculties were not so limited. If only we had faculties that were not tricked by the various deceptions we currently find ourselves tricked by, the protagonist in many Gettier cases would not be fooled into a false belief (more accurately, a luckily true belief). Indeed, as the proposed diagnosis of Gettier problems in Chapter 1 would predict, Gettier counterexamples could be avoided altogether if only our cognitive faculties infallibly (inviolably) tracked truth. However, it also seems feasible to say that if our cognitive faculties were not so limited, we would not have been able to survive as a species because, perhaps, our heads would have been too big and we would not have been able to hold them up, let alone escape from predators. However, as Feldman highlights, trying to preclude Gettier counterexamples by claiming that "the

[94] See Plantinga 1993b, 31–32, 37, 39.

[95] Also see Feldman 1996, 214.

segments of our design plan that get us knowledge are . . . aimed at truth while the segments responsible for the beliefs in Gettier cases have other aims is rather implausible" (1996, 215).[96]

To see this, consider a case exactly like Classic Case 2 except that Smith is actually telling Jones the truth about owning a Ford—a case where we would typically attribute knowledge to Jones's belief that "Smith owns a Ford."[97] Seemingly there is absolutely nothing different from Jones's perspective between this standard testimony case and Classic Case 2; however, Plantinga would have to say that while the segment governing the production of Jones's belief is aimed at truth in a case where Smith is telling the truth, in Classic Case 2, the segment governing the production of his belief must be aimed at something else. Somehow, the segment governing the production of Jones's belief in the former case is markedly different from the segment governing the production of Jones's belief in the latter case even though nothing has changed from Jones's perspective. With the ad hoc multiplication of design segments and the mystery as to what is regulating their aims, such a conclusion seems at the very least far-fetched.

To further accentuate the implausibility of the claim that "the segments of our design plan that get us knowledge are . . . aimed at truth while the segments responsible for the beliefs in Gettier cases have other aims," Feldman provides another case:

> **Newspaper:** You are reading the sports section of the newspaper. There is a list of the scores of last night's basketball games. You believe that each game did have the outcome reported in the paper. Suppose there is a misprint in one of the reports. (Feldman 1996, 216)

In such a case, you form any number of true beliefs, but you do form one false belief due to the misprint (a belief that could easily be Gettierized). According to the Gettier-defense currently under consideration, as you read through the scores the segment of your cognitive faculties' design plan that is governing the production of your beliefs is aimed at truth up until (and immediately after) you read the misprint. For the instant you read the

[96] Feldman expressly argues for the implausibility and not the refutation of such a strategy because "there are so few constraints on what one might say about the 'direct aims' of various segments of our design plan . . . there is so little to restrict what one says about the aims of the segments of the design plan that govern the production of our beliefs, it is possible for a defender of Plantinga's theory to say what's necessary to defend the theory from potential counterexamples. However, what one has to say becomes increasingly implausible" (1996, 215).

[97] See Feldman 1996, 216.

misprinted score, the segment changes and is no longer aimed at truth. As Feldman notes, "this seems wildly ad hoc" (Feldman 1996, 216). Although such cases do not decidedly refute Plantinga's second strategy, they do nevertheless condemn it as decidedly meriting rejection (Feldman 1996, 217).[98]

To defuse cases like Classic Case 2, Plantinga offered two independent strategies: (i) deny that the given cognitive environment is suitably congenial or (ii) deny that the segment of the design plan governing the production of the Gettierized belief is aimed at truth. Unfortunately for Plantinga, such strategies simply fall flat—being intractable and unable to systematically distinguish Gettier environments from normal environments, Gettier design plans from normal design plans. Plantinga's 1993 account of warrant, Plantinga's 1993 analysis of knowledge, simply lacks the material for precluding Gettier environments and design plans without precluding from knowledge our everyday beliefs formed in normal environments and according to normal design plans aimed at truth. Both of Plantinga's proposed strategies fail because they lead to skeptical conclusions, on pain of adhocery.

Plantinga's 1993 account of warrant, his reductive virtue epistemology, takes warrant as bearing a close but not inviolable relationship to truth. As such, our proposed diagnosis of the Gettier Problem predicted that Plantinga's 1993 virtue epistemology would be susceptible to Gettier counterexample. And that is precisely what the literature has shown, and that is precisely what we have seen. Plantinga's reductive account of knowledge easily lends itself to Gettier counterexample—his proposed strategies for avoiding Gettier cases being ultimately intractable. But before we count this a victory for our proposed diagnosis of the Gettier Problem, let's briefly consider some other strategies Plantinga could have tried to invoke in his defense against cases like Lucky Mary, Lucky Ms. Jones, and now Classic Case 2.

For starters, Plantinga could have tried to deny that warrant can be transferred from a false belief. But, as Peter Klein points out in "Warrant, Proper Function, Reliabilism, and

[98] Feldman goes on to muse over Plantinga's theistic commitments and how that might explain aspects of his theory; such considerations, however, do nothing to stave off Feldman's "wildly ad hoc" indictment. That is to say, the claim that "the segments of our design plan that get us knowledge are 'directly' aimed at truth while the segments responsible for the beliefs in Gettier cases have other aims" is going to lead to ad hoc conclusions whether or not you are a theist.

Defeasibility" (1996), such a strategy is generally accepted to be a dead end—being both too strong and too weak.[99] It is too strong because it seems to prohibit some legitimate instances of knowledge . . .

> For example, suppose that my belief that I have an appointment at 3:00 p.m. on April 10th is based on my warranted but false belief that my secretary told me on April 6th that I had such an appointment. If my secretary told me that I had such an appointment, but she told me that on April 7th (not the 6th) my belief that I have an appointment at 3:00 p.m. can still be knowledge, even though the belief that supports it is false. (Klein 1996, 106)

. . . And too weak because further Gettier cases can be produced even if it is adopted. Someone could easily manipulate cases like Lucky Ms. Jones and Lucky Mary so that there is no clear inference from a false belief being made.[100]

Similarly, denying the transferability of warrant full stop would not be any more fruitful as a strategy for Plantinga to pursue. While such a strategy would indeed handle Classic Case 2, it does not seem as effective against Lucky Mary or, at least, Lucky Ms. Jones. As Klein points out, though someone might read the case of Lucky Ms. Jones as involving warrant transmission—e.g., thinking Ms. Jones inferred the belief that she owns a well-functioning Ford from the prior belief that she drove a well-functioning Ford to work that morning—the case need not be read that way. Ms. Jones could simply form the given belief without drawing it from an antecedent. Just as Henry in Fake Barns sees the barn and then simply believes "there's a barn," perhaps Jones in Lucky Ms. Jones simply remembers and subsequently simply believes that she owns a well-functioning Ford. The transfer of warrant does not seem to be a necessary component of Gettier counterexamples.

What if Plantinga were to deny that the cognitive faculties at issue in the relevant cases were properly functioning to a sufficient degree—denying that condition 1 of his account of warrant is met? In other words, what if Plantinga denied that Mary's faculties of perception (in the case of Lucky Mary), or Ms. Jones's faculties of memory (in the case of Lucky Ms. Jones), or Jones's faculties of credulity (in Classic Case 2) . . . were functioning as properly as they should if the corresponding beliefs they produce were to count as knowledge? The problem with such a response to these Gettier problems is that there is absolutely no indication that the relevant faculties at issue are not already functioning extremely well. Indeed, it seems like the only way this strategy could do Plantinga any good

[99] See Klein 1996, 106; Clark 1963; Shope 1983, 21–26.

[100] See Klein 1996, 106–107.

is if he required *cognitive faculty perfection* for warrant and correspondingly for knowledge. As Linda Zagzebski rightly points out, however, Plantinga expressly does not require perfection, nor should he (Zagzebski 1994, 68). Such a demand would simply be unreasonable—leaving us with only a precious few beliefs that we know. Without requiring perfection (which would be tantamount to requiring an inviolable relationship between warrant and truth), however, there is no way that denying that the cognitive faculties at issue in the relevant cases were properly functioning to a sufficient degree will produce any lasting success against Gettier counterexamples.

Could Plantinga instead deny that the relevant design plans governing the production of Gettierized beliefs are good? Could Plantinga deny that his fourth condition on warrant is met? Seemingly not. There does not seem to be any ground for making such a claim. There is absolutely no indication that the design plans governing the production of Gettierized beliefs (as in Lucky Mary, Lucky Ms. Jones, and Classic Case 2) are not already quite good. Indeed, it seems like the only way this strategy could do Plantinga any good is if he required not just good design plans but *perfect* (infallible) design plans for warrant and correspondingly for knowledge. We can stipulate that the design plans are as good as we like, but again it seems like so long as they are anything less than (statistically) infallible, then they will produce a false belief—a false belief that can be Gettierized. Without requiring design plan *perfection* (which would be tantamount to requiring an inviolable relationship between warrant and truth and presumably lead directly to radical skepticism), there is no way that denying that the relevant design plans governing the production of Gettierized beliefs are good will produce any lasting success against Gettier counterexamples.

According to Plantinga's reductive virtue epistemology, a belief is warranted if it is produced by properly functioning cognitive faculties that are operating in a congenial environment and in accord with a good design plan aimed at truth. And given Plantinga's 1993 reading of such a definition, warrant bears a close but not inviolable relationship to truth. And in accord with our proposed diagnosis of the Gettier Problem, Plantinga's account of knowledge fell prey to Gettier counterexample. Plantinga's expressed strategies for defusing Gettier counterexamples were found to be intractable—lacking material for sufficiently (i.e., without adhocery) distinguishing between Gettier scenarios and normal, everyday circumstances. We considered some other possible arguments Plantinga could

invoke to try to save his 1993 theory, but they were ultimately no more helpful—at best, avoiding Gettier counterexamples by assuming an inviolable relationship between warrant and truth, and only then at the cost of radical skepticism. Thus far, the diagnosis of Gettier counterexamples proposed in Chapter 1 seems completely on target. However, we are far from finished with Plantinga's reductive virtue epistemology. After accepting that his 1993 account of knowledge was susceptible to counterexample, Plantinga admits that cases like Classic Case 2, Russell's Stopped Clock, Lucky Mary, and Lucky Ms. Jones show that a given belief can be produced by cognitive faculties functioning properly in the right kind of environment in accord with a good design plan aimed at truth and nevertheless fail to have warrant.[101] In the next section, we will explore the series of two modifications Plantinga proposes to his environment proviso—starting with "Respondeo" (1996) and "Warrant and Accidentally True Belief" (1997) and working our way to Plantinga's final modifications in *Warranted Christian Belief* (2000)—and discover that neither of them is any more successful at precluding Gettier counterexamples than his original 1993 account, failing precisely along the lines predicted in Chapter 1.

Section 2: Plantinga's Modifications

So, to reiterate, according to Plantinga's original 1993 conception of warrant, a belief B will be warranted for S if and only if the following four conditions are met:

> 1) the cognitive faculties involved in the production of B are functioning properly;
>
> 2) [the] cognitive environment is sufficiently similar to the one for which [the agent's] cognitive faculties are designed;
>
> 3) the design plan governing the production of the belief in question involves, as purpose or function, the production of true beliefs;
>
> 4) the design plan is a good one: that is, there is a high statistical or objective probability that a belief produced in accordance with the relevant segment of the design plan in that sort of environment is true. (Plantinga 1993b, 194)[102]

Or as I summarized earlier: "A belief B is warranted for S when B is formed by cognitive faculties (of S's) that are functioning properly in the right environment in accord with a

[101] See Plantinga 1996, 308–329.

[102] For other tabulations of Plantinga's 1993 account of warrant (what later iterations refer to as the "nutshell" or "central core" of warrant) see Chignell 2003, 445; Plantinga 2000, 156.

good design plan aimed at truth."[103] As a number of critics pointed out, however, such conditions are not sufficient for warrant—which, given that Plantinga's 1993 account assumes a close but not inviolable relationship between warrant and truth, is precisely what our proposed diagnosis of Gettier problems would predict. In any case, Plantinga acquiesces as much in "Respondeo" (1996) and "Warrant and Accidentally True Belief" (1997)—admitting that cases like Classic Case 2, Russell's Stopped Clock, and Lucky Ms. Jones show that a given belief can be produced by cognitive faculties functioning properly in the right kind of environment in accord with a good design plan aimed at truth and nevertheless fail to have warrant. According to Plantinga, the problem with his 1993 account of warrant lies in his environment proviso. Though the beliefs formed in the aforementioned cases are formed by cognitive faculties operating within the right sort of environment, there is nevertheless some "cognitive environmental pollution," an abiding "lack of resolution" between the proper function of the given cognitive faculties and the environment in which they operate (Plantinga 1996, 309–310, 316, 327). So naturally, it is the environment condition that Plantinga seeks to expand on in his revised accounts of warrant. In his proposed modifications, Plantinga's environment condition, his second condition for warrant, is his anti-luck condition, his condition that is meant to preclude Gettier counterexamples.

Plantinga notes that in his original account the environment proviso was more or less a general environmental condition—a environmental condition that required warranted beliefs to be formed in an earth-like environment, an environment with "the presence and properties of light and air, the presence of visible objects, of other objects detectable by our kind of cognitive system, of some objects not so detectable, of regularities of nature, the existence of other people, and so on" (Plantinga 1996, 313). Plantinga has us call this notion of environment the "maxi-environment." However, as shown by the aforementioned cases, having a congenial maxi-environment alone is not sufficient for avoiding Gettier counterexamples. So Plantinga introduces the concept of what he calls a cognitive "mini-environment." Plantinga explains:

> We can think of a cognitive mini-environment of a given exercise of cognitive powers E as a *state of affairs* (or propositions)—one that includes all the relevant epistemic circumstances obtaining when that belief is formed. Consider any current belief B I hold and the exercise E of cognitive powers that produced it: the mini-environment M for E (call it "MBE") includes the state of affairs specified by

[103] See Plantinga 1993b, 19.

> my cognitive maxi-environment, but also much more specific features of my epistemic situation. (Plantinga 1996, 314)

According to Plantinga, the aforementioned "cognitive environmental pollution" and the abiding "lack of resolution" highlighted by Gettier cases like Classic Case 2, Russell's Stopped Clock, Lucky Mary, and Lucky Ms. Jones are not found in their respective maxi-environments but rather in the mini-environments of the relevant exercises of cognitive powers. In other words, the cognitive pollution and lack of resolution come from a given MBE not being favorable for said E.

As such, Plantinga adds the following resolution condition to his *Warrant and Proper Function* (1993) account of warrant:

> **Resolution Condition:** A belief B produced by an exercise E of cognitive powers has warrant only if MBE is favorable for E. (Plantinga 1996, 328)

Not only does a given belief's cognitive maxi-environment need to be "sufficiently similar to the one for which [its] cognitive faculties are designed" (the original second condition on warrant), its *mini*-environment needs to be favorable for the exercise of cognitive faculties that produced it. To be sure, Plantinga goes on to say that we can make a given MBE as full and detailed as we please, with the exception of truth or falsehood of the given belief—lest mini-environments where a given belief is only luckily true, as in Gettier cases, be deemed favorable (Plantinga 1996, 314–315). As such, the relevant MBEs of cases like Classic Case 2, Russell's Stopped Clock, Lucky Mary, and Lucky Ms. Jones will, respectively, include the presence of liars, stopped clocks, visiting brothers-in-law, and Fords being unforeseeably destroyed. Plantinga's hope, then, is to explicate favorability in such a way that deems the relevant MBEs in such cases as unfavorable—precluding the corresponding beliefs from being warranted and therefore precluding their (mistakenly) being deemed knowledge by his account. At the heart of Plantinga's anti-Gettier strategy and his proposed modifications is his explication of "favorability."

Section 2.1: 1996/1997 Favorability

Just what does it mean for a mini-environment to be "favorable" for a given exercise of cognitive powers? While Plantinga is initially skeptical as to whether this kind of detail is "attainable or necessary here," nevertheless, in "Respondeo" (1996) and "Warrant and Accidentally True Belief" (1997), Plantinga posits the following definition (1996, 327):

> **1996/1997 Mini-Favorability**: MBE is favorable for E, if and only if, if S were to form a belief by way of E, S would form a true belief. (Plantinga 1997, 144; Plantinga 1996, 328)

At first blush, it may seem as though Mini-Favorability is completely ineffectual. It would seem as though the mini-environment of the particular exercise of cognitive powers that produced Ms. Jones's belief that she owns a Ford is indeed favorable as the belief is true. However, Plantinga avoids this hurdle by specifying that Mini-Favorability's counterfactual semantics are non-standard (i.e., non-Lewisian, non-Stalnakerian)—"the truth of *p* and *q* is not sufficient for the truth of the counterfactual *if p then* q" (Plantinga 1996, 328).[104] Mini-Favorability is "a point where the usual semantics for counterfactuals is inadequate" (Plantinga 1996, 328–329).[105] The counterfactual semantics that Plantinga instead stipulates is one where "the counterfactual is true only if there is no *sufficiently close* possible world in which *p* is true but *q* is not" (Plantinga 1996, 329 - emphasis Plantinga's). In other words, according to Plantinga, a given MBE is favorable for E, if and only if, if S were to form a belief by way of E in MBE, S would form a true belief in *all close possible worlds*. As such, the MBEs of the relevant cases are not meant to be favorable for their corresponding exercises of cognitive powers; seemingly, so it goes, in many close possible worlds, the protagonist's belief in question would be false. Subsequently, if the MBE is not favorable in such cases, then Plantinga's resolution condition on warrant is not met.

Perhaps realizing that Gettier cases will be unavoidable so long as warrant is not infallibly connected to truth, Plantinga's 1996/1997 account of warrant, where his new Resolution Condition is understood in terms of 1996/1997 Mini-Favorability, is meant to "[guarantee] that no false belief has warrant" (1996, 329). It is meant to be impossible for a warranted belief to be false, and as such it is supposed to be impossible for a warranted belief to only be true for reasons not captured by the warrant. As such, according to our proposed diagnosis of the Gettier Problem, Plantinga's 1996/1997 account of warrant and corresponding virtue epistemology should be immune to Gettier counterexamples but unfeasible, leading to some skeptical conclusion. At first blush, however, no such

[104] Also see Plantinga 1997, 144n.

[105] Plantinga goes on to reference some other points where the standard counterfactual semantics are inadequate. Quantum effects: "perhaps in fact the photon went through the right slit rather than the left; that is not enough to entail that if it had gone through either slit, it would have gone through the right" (1996, 328–329). Die tossing: "I toss the die; it comes up 5. That is not sufficient to entail that if I had tossed the die, it would have come up 5" (1996, 329).

conclusion seems to be at hand. Does this mean that Plantinga has found an exception to our diagnosis of Gettier counterexamples, finding a viable reductive, virtue-theoretic solution to the Gettier Problem?

Unfortunately for Plantinga (and fortunately for the proposed diagnosis), the answer is no. Plantinga was simply wrong about his 1996/1997 account's infallibilism. Despite what Plantinga says, his 1996/1997 account of warrant simply does not "guarantee that no false belief has warrant." Consider the following familiar scenario: Ms. Jones believes that she owns a well-functioning Ford, and she forms this belief under good circumstances using properly functioning cognitive equipment. But, as it happens, a freak event occurs; Ms. Jones's car is utterly destroyed by a stray meteor. In such a scenario, Ms. Jones's belief seems to meet all of Plantinga's original 1993 conditions on warrant and it seems to satisfy his Resolution Condition as understood in terms of 1996/1997 Mini-Favorability—presumably, in all close possible worlds, Ms. Jones's Ford is not hit by a roaming meteor. As such, Plantinga's 1996/1997 account of knowledge *does* allow for warranted false beliefs; and as such, it should, according to the proposed diagnosis, be susceptible to Gettier counterexample. And this is precisely what we find. To produce such a counterexample, simply stipulate that Ms. Jones wins that "Well-Functioning Ford Lottery" in all close possible worlds (perhaps one of her friends rigged the lottery without Ms. Jones knowing about it), and then we have produced a Lucky Ms. Jones-like counterexample to Plantinga's revised account of warrant. Again, our proposed diagnosis of Gettier problems seems to get it exactly right.

And as much is vindicated in the relevant literature. In "Gettier and Plantinga's Revised Account of Warrant" (2000), Thomas Crisp has leveled a similar Gettier counterexample against Plantinga's 1996/1997 account:

> **Prune Guessing:** Suppose your uncle runs the town's annual guess-the-number-of-prunes-in-the jar contests. Your prankish friend takes it on good authority that the jar contains 138 prunes and lets you in on the secret. Unbeknownst to both you and your friend, though, the number he is given is incorrect. Now, suppose further that your uncle has taken ill with an unusual brain fever and has come to believe that the fate of the nation hangs on your winning the contest. Since he can't remember how many prunes were in the jar to begin with, he empties it and refills it with the exact number of prunes indicated on your contest entry card. The day of the contest arrives and the town gathers for the beloved counting of the prunes. You believe firmly that the jar contains 138 prunes. And indeed it does. But your belief is true by accident: had your uncle not taken ill with the fever, your belief would have been false. (2000, 47)

Again, the protagonist's belief in Crisp's case seemingly meets not only the original conditions for warrant but also the conditions of favorability set out in 1996/1997 Mini-Environment Favorability. The full and detailed state of affairs in the relevant MBE would include the presence of the uncle's fever and his delusional conviction about national security such that in no close possible worlds is the belief formed by the pertinent exercise of cognitive powers false; though, again, clearly the belief is still only luckily true. And indeed, by the time Plantinga wrote *Warranted Christian Belief* (2000), he seems to have realized this problem with his 1996/1997 modification. [106] According to Plantinga, his 1996/1997 account of mini-environment favorability is simply insufficient because "the relevant counterfactual itself can be true 'just by accident'"—a feature that should be simply impossible if his 1996/1997 account of warrant really was infallibilistic regarding truth (2000, 159). To illustrate as much, Plantinga provided his own Gettier counterexample to his 1996/1997 account:

> **Foggy Fake Barns:** Suppose I am driving through [fake barn territory] on an early September morning when there is a good deal of mist and fog. I glance to the right and see a real barn; as it happens, all the nearby fake barns (which outnumber the real ones) are obscured by the morning mist; I say to myself, "Now that is a fine barn!" The belief I form is true; the relevant counterfactual is also true because of the way the fake barns are obscured by mist; but the belief does not have warrant sufficient for knowledge. (Plantinga 2000, 159–160)

Once again, the protagonist's belief in Foggy Fake Barns seems to meet not only the original 1993 conditions for warrant but also the conditions of favorability set out in 1996/1997 Mini-Environment Favorability. The full and detailed state of affairs in the relevant MBE would include the presence of the fog such that in no close possible worlds is the belief formed by the pertinent exercise of cognitive powers false; though clearly the belief is still only luckily true. [107]

[106] Though Crisp's paper is not explicitly mentioned, Plantinga credits Crisp with prompting his later 2000 modification. See Plantinga 2000, 159n. Presumably if Crisp's paper did not explicitly prompt the change in Plantinga's account, we can assume that informal ancillary conversations between Crisp and Plantinga on account of said paper did.

[107] For another worry on the 1996/1997 conception of "favorability" concerning mini-environments, see Chignell 2003, 448.

Stymied once again, Plantinga went back to the drawing board, so to speak, and proposed a different account of favorability in *Warranted Christian Belief* (2000).[108] If Plantinga's environmental condition on warrant, in particular his Resolution Condition on mini-environments, is meant to be his anti-Gettier condition, his 1996/1997 modification is simply insufficient. As such, in a final attempt to surmount Gettier counterexamples via reductive analysis, Plantinga tries once more to explicate mini-environment favorability.

Section 2.2: 2000 Favorability

While Plantinga is again unsure as to whether "we can say anything more definite" as to what it means for a mini-environment to be favorable for a given exercise of cognitive powers—beyond our intuition that a given mini-environment will be favorable for a given exercise insofar as that "exercise can be counted on to produce a true belief" in that mini-environment—he is, in *Warranted Christian Belief* (2000), nevertheless compelled to try (Plantinga 2000, 159). Again, Plantinga attributes the aforementioned "lack of resolution"—the lack of fit between the proper function of the given cognitive faculties and the environment in which they operate. And Plantinga characterizes this "lack of fit" in *Warranted Christian Belief* (2000) as a discrepancy between a given mini-environment as *perceived* (or detected) by the pertinent agent and the mini-environment *in full* (Plantinga 2000, 160). In other words, the discrepancy causing the "lack of resolution" is between a given MBE and the conjunct of states of affairs that are *detectible* via the said agent's exercise of cognitive powers E, or DMBE (Plantinga 2000, 160). Hence, to preclude this

[108] Other aspects of Plantinga's 1996/1997 account were critiqued as well, though failing to motivate a change in the 2000 account. For example, in "Gettier and Plantinga's Revised Account of Warrant" (2000), Thomas Crisp argued that no state of affairs fulfills Plantinga's definition of mini-environments. In sum, Crisp's worry is that "for any exercise E of one's cognitive power in maximally specific circumstances C, there will be . . . reasons for doubting that there is a *closest* state of affairs to C that neither includes nor precludes the proposition that E yields true belief" (2000, 45). The "reasons for doubting" that Crisp mentions are that for any given state of affairs that is taken to be closest to C, we can think of one that is closer. In other words, there may be an infinite series of states of affairs that get closer and closer to C indeterminate in regard to the truth of the given belief such that one never arrives at a "closest" point. Though I have not fully elucidated Crisp's argument by any means, let me just note that by my lights Plantinga is not putting nearly the same amount of weight on there being a "closest state of affairs . . ." as Crisp assumes. After all, Plantinga invites us to make a given mini-environment as full and as detailed *as we please* (1996, 314)—seemingly implying that getting to the absolute "closest state of affairs . . ." is not centrally important to Plantinga's account.

"lack of resolution" so understood, Plantinga proposes a new conception of mini-environment favorability that aims to track warrant ascriptions along with the absence of this sort of discrepancy:

> **2000 Mini-Favorability:** MBE is favorable just if there is no state of affairs S included in MBE but not in DMBE such that the objective probability of B with respect to the conjunction of DMBE and S falls below r, where r is some real number representing a reasonably high probability. (Plantinga 2000, 160)[109]

Plantinga's hope, then, is that something like 2000 Mini-Favorability will provide not only a clearer understanding as to what is required in his Resolution Condition but also a successful defense against Gettier-style counterexamples.

According to Plantinga's analysis of knowledge, the original 1993 conditions for warrant are sufficient for warrant once the environmental condition is amended with his Resolution Condition, now understood in terms of 2000 Mini-Favorability; and warrant, so understood, is with truth and belief meant to be both necessary and jointly sufficient for knowledge.[110] In support of his second attempt at specifying MBE favorability, Plantinga has us consider another Gettier case:

> **Peter and Paul:** I am not aware that Paul's look-alike brother Peter is staying at his house; if I'm across the street, take a quick look, and form the belief that Paul is emerging from his house, I don't know that it's Paul, even if in fact it is (it could just as well have been Peter emerging); . . . if Peter would not have been in the neighborhood, I would have known. (Plantinga 2000, 157)

As Plantinga notes, the resolution problem in this case "arises because I can't (for example) distinguish Paul from Peter from across the street just by looking" (Plantinga 2000, 160).

---

[109] One may worry that Mini-Favorability 2000) sneaks knowledge into the definition of warrant, since, as Thad Botham notes, having a proposition *p* be detectable to S could perhaps be seen as simply akin to having S know *p* (2003, 434). If this were the case then Plantinga's account of knowledge would appear to be viciously circular. However, according to Botham, Plantinga presumably has a viable response to this worry. He says: "A merchant's scale detects its being the case that the goods weigh 400 pounds, but the scale cannot detect its being the case that a handful of chocolates weighs three ounces. A thermostat detects its being the case that the room's ambient temperature is eighty degrees Fahrenheit but cannot detect the temperature to four significant digits. A speedometer detects the automobile's moving 50 miles per hour but cannot detect the precise rate of 50.2 miles per hour. Even though the scale, thermostat, and speedometer detect various states of affairs only when they function properly, they do not have knowledge. Neither do they have anything close to justified or warranted beliefs. Thus it's not clear that detection entails anything as strong as knowledge or justified belief" (Botham 2003, 434).

[110] See Plantinga 2000, 161.

By taking a glance across the street toward Paul's house the protagonist can detect all sorts of things ("the appearance of a person, of a man, of someone across the road, and the like"), but he cannot with such a glance tell the difference between Paul and Peter from across the street (Plantinga 2000, 160). This discrepancy between the DMBE and the MBE in the case of Peter and Paul is such that B is seemingly not probable in respect to DMBE. Likewise, cases like Russell's Stopped Clock, Lucky Ms. Jones, Lucky Mary, and Classic Case 2 seem to be defused as well. Though the protagonist in Russell's Stopped Clock can detect through his pertinent exercise of cognitive powers that there is a clock on the wall and that it says 3:43 pm, he cannot detect that the clock is broken. Though Ms. Jones in the case of Lucky Ms. Jones can detect through her pertinent exercise of cognitive powers that she had earlier driven her Ford to the office, she could not detect that the Ford she drove to the office was destroyed. Though Mary in the case of Lucky Mary can detect through her pertinent exercise of cognitive faculty that a husband-like figure is sitting in the living room, she could not detect that the figure is actually her husband's brother. Though Jones in Classic Case 2 can detect that Smith has manifold evidence to the effect that he owns a Ford, Jones could not detect that Smith was lying. In accord with Plantinga's analysis, in Russell's Stopped Clock, Lucky Ms. Jones, Lucky Mary, and Classic Case 2, just like in the case of Peter and Paul, there is a discrepancy between the respective DMBE and the MBE such that the relevant beliefs are not objectively probable in respect to their given DMBEs.

Has Plantinga finally developed a reductive virtue epistemology that avoids Gettier counterexample? According to our proposed diagnosis, this will depend on whether or not it is possible for a warranted belief, as elucidated in *Warranted Christian Belief* (2000), to be false. If it is possible, then, according to the proposed diagnosis, Plantinga's 2000 account of knowledge should still be susceptible to Gettier counterexample. If it is not possible to have a warranted false belief, then, according to the proposed diagnosis, Plantinga's 2000 account of knowledge will be Gettier-proof but unpalatable, leading to radical skepticism.

Although in *Warrant and Proper Function* (1993) Plantinga seems to expressly allow for false belief having warrant, this certainly does not seem to be his position in his revised accounts.[111] As we already noted, Plantinga fully intended for his 1996/1997 account of warrant to "[guarantee] that no false belief has warrant" (Plantinga 1996, 329). And while

[111] See Plantinga 1996, 311–312, 329.

Plantinga is nowhere near as explicit in *Warranted Christian Belief* (2000) on this score, he nevertheless seems to imply as much. Consider the following passage:

> [T]here can be mini-environments for a given exercise of our faculties, in which it is just by accident, dumb luck, that a true belief is formed, if one is indeed formed. A true belief formed in such a mini-environment does not have warrant sufficient for knowledge, even if it has some degree of warrant. To achieve that more exalted degree of warrant, the belief must be formed in a mini-environment such that the exercise of cognitive powers producing it can be counted on to produce a true belief. (Plantinga 2000, 161)[112]

Given that "being counted on to produce a true belief" is Plantinga's synonymous lingo for mini-environment favorability, Plantinga is saying in this passage that mini-environment favorability precludes the possibility of a luckily true belief. And given that luckily true beliefs are lucky because they could just have easily been false, Plantinga seems to be saying that mini-environment favorability precludes the possibility of a false belief. In other words, Plantinga in his most recent work seems to take warrant as inviolably related to truth.[113]

And even if this is not the correct way to read such passages, we might have nevertheless *guessed* that Plantinga is assuming that warrant bears an inviolable relationship to truth simply on the grounds that it is extremely difficult to think of a false belief that could satisfy 2000 Mini-Favorability, let alone the other conditions on warrant. One seems utterly at a loss for thinking of any false belief that would not have *some* epistemically relevant state of affairs that is not detectible by the given exercise of cognitive powers that would be included in the respective mini-environment. It seems like for any given false belief there will be some state of affairs not included in the given DMBE that undermines whatever reason(s) we may have to believe said belief.

If there is not a false belief that meets Plantinga's conditions for warrant, then it will not be possible for a belief to be so warranted but true for other reasons. In other words, if there is not a false belief that meets Plantinga's conditions for warrant, then his account should be immune to Gettier counterexamples. As such, in accord with our proposed diagnosis of Gettier problems, Plantinga's account should be unpalatable, leading to

---

[112] Plantinga is not clear in such passages as to how a given belief can have any degree of warrant without meeting all of the conditions of warrant. It is reasonable to assume that this relates to the other woes afflicting Plantinga's understanding of degrees of warrant. See Markie 1996.

[113] This understanding of Plantinga's 2000 account of warrant was affirmed during my personal correspondence with Plantinga while presenting at the University of Notre Dame in 2009.

skeptical conclusions. Fortunately for the proposed diagnosis (though unfortunately for Plantinga's account), this is exactly what has been noted in the literature. As critics have repeatedly pointed out, 2000 Mini-Favorability is simply too strong—for many instances that we would intuitively deem knowledge there is *some* state of affairs that would be included in the given MBE but not the DMBE that would significantly reduce the given belief's objective probability in respect to DMBE and that state of affairs. [114] Thad Botham (2003) provides the following three cases:

> **Only Paul:** Consider a version of the [Paul and Peter case] where we stipulate that Paul is an only child, thereby removing Peter from the scene. In addition, the moment before you believe the proposition "There's Paul," his uncle—unbeknownst to you—lies to a friend, telling her that Paul has an identical twin brother who's visiting Paul at that very moment and that he just spoke with each of them on the telephone. Paul's uncle asserts this falsehood while in London, thousands of miles away. (Botham 2003, 435–436)

> **Tiny Fake Barns:** Consider . . . [the case of Fake Barns] with the following alteration. Rather than constructing life-sized barn facades, the locals manufacture model barns so tiny that standard passersby cannot view them. Perhaps the tiny barns are only two inches in height, built to suit ant communities. As you drive through the heart of this anomalous territory, you see one of the only real barns and believe the proposition "That's a fine barn." (Botham 2003, 436)

> **Dalmatian:** Suppose you observe a Dalmatian by looking through a window into the backyard. You believe the proposition "There's a Dalmatian." However, there is a state of affairs [S]—being such that there are ten mechanical Dalmatians in the backyard, each of which appears like a real Dalmatian—that together with [the given DMBE] renders the objective probability of your belief less than reasonably high. Nonetheless, the owner locked all of the robotic Dalmatians in a shed in his backyard to prevent them from rusting. Indeed, they've resided in the shed for about ten years, the lock is rusted shut, and no one has known the key's whereabouts for at least two years. (Botham 2003, 436)

Andrew Chignell (2003) provides a similar case:

> **Song Sparrow:** Johnson . . . has acquired his ornithological training by reading some birding books and by listening to recordings of birdcalls. Johnson hasn't seen a song sparrow (Melospiza melodia) before, but he is familiar with their plumage, body shape, and calls. Unbeknownst to him, however, song sparrows have suffered from a devastating virus in recent weeks, and there are now only two of them left on the entire continent. The Lincoln's sparrow (Melospiza lincolnii), which looks and sounds very similar to the song sparrow, has been unaffected by the virus and there are still quite a few of them living in Johnson's region. As Johnson strolls through the forest, he hears what seems to him to be the call of a song sparrow. He approaches the relevant tree to get an up-close look at the bird. The bird is, in fact, one of the two remaining song sparrows on the continent.

[114] See Botham 2003; Chignell 2003.

> Johnson studies the bird for some time: it looks to him like the song sparrow pictured in his books, and its call sounds like the recordings he has heard. On this basis these observations, Johnson assents, with a degree of strength that is just enough to put him over the threshold required for knowledge, to the proposition that the bird he is observing is a song sparrow. (Chignell 2003, 449)

In all four of these cases we intuitively think that the given protagonists possess knowledge—knowledge that "there's Paul," "that's a fine barn," "there's a Dalmatian," "that's a song sparrow," respectively—however, in all four cases there is a state of affairs that "together with [the respective DMBE] makes the objective probability of [the protagonist's] belief less than reasonably high" (Botham 2003, 436). Though the protagonists in such cases seem to know their respective beliefs, the conditions for warrant in Plantinga's 2000 account are not met, given 2000 Mini-Favorability.

And to be sure, the problem is general. Plantinga's 2000 account of warrant does not just fail to track knowledge across a limited range of cases, it seems to lead us directly into radical skepticism. Not only is it extremely difficult to think of a false belief that would satisfy Plantinga's 2000 Mini-Favorability, it is extremely difficult to think of a true belief that satisfies it. For almost any given belief there is going to be *some* state of affairs that would reduce the objective probability of said belief in light of the perceived mini-environment (DMBE). Take my belief that grass is green, for example—a belief that is presumably quite epistemically secure. There may be any number of facts that might (either individually or in conjunction) reduce the objective probability of such a belief in light of the mini-environment as I perceive it: it might be the case that color blindness runs in my family; it might be the case that lots of people use the terms "green" and "grass" in ways that I do not;[115] I might have a personal history of getting some colors a bit mixed up, etc. As such, according to Plantinga's 2000 account of warrant, I presumably do not know that grass is green! Take my belief that I am now in pain. Again, there seem to be any number of facts that might (either individually or in conjunction) reduce the objective probability of such a belief in light of the mini-environment as I perceive it: it may be the case that I have a history of hypochondria; it may be the case that I am prone to exaggeration, etc. Again, according to Plantinga's 2000 account of warrant, I presumably do not know that I am now in pain—another radical conclusion. Given Plantinga's infallibilism about warrant, his

[115] This state of affairs would presumably apply even if I was completely orthodox in my use of such terms. In other words, there may be lots of people who misuse such terms; but as far as the state of affairs as stated above is concerned, I would not know if I was one of them.

assumption that his modified account of warrant "guarantees" truth, he can seemingly avoid Gettier counterexample, but only at the cost of skepticism—precisely as our proposed diagnosis would predict.

In response to the Gettier problems afflicting his 1993 account of warrant, Plantinga amended his environmental condition to preclude not only ill-suited maxi-environments but also unfavorable mini-environments. And in his 1996/1997 and his 2000 modifications, Plantinga tentatively tried to say a bit more about what it means for a mini-environment to be favorable for a given exercise of cognitive competence. Both attempts, as we have now seen, seem to fall flat—either leading to further Gettier counterexamples or leading to unsavory skeptical conclusions. But maybe the lesson to be learned from all this is that Plantinga should not try to explain what does not warrant or need explanation. After all, as we have already said, in "Respondeo" (1996) Plantinga worries that perhaps elucidation of mini-environment is neither "attainable or necessary" (1996, 327). In *Warranted Christian Belief* (2000) Plantinga is again unsure as to whether "we can say anything more definite" as to what it means for a mini-environment to be favorable for a given exercise of cognitive powers beyond our intuition that a given mini-environment will be favorable for a given exercise insofar as that "exercise can be counted on to produce a true belief" in that mini-environment (Plantinga 2000, 159). Perhaps Plantinga should just leave his account of warrant somewhat underdefined—omitting any explication of what precisely it means for a mini-environment to be favorable for a given exercise of cognitive faculty.

Sadly, I do not think this is viable strategy for Plantinga to take for at least two reasons. First of all, leaving his account of warrant underdefined in this way would seem negligent. For nearly fifty years, people have been wrestling with the Gettier Problem, with little success; so to simply say that Gettier counterexamples can be avoided so long as one's environment is suitably favorable, there will be deserved protests and outcries if no explication of favorability is given. Secondly, leaving his account of warrant underdefined in this way would simply mean that the problems that afflict it would be likewise underdefined. Even if he does not try to explain what it means for a given mini-environment to be favorable, our proposed diagnosis of Gettier problems would still apply. Assuming that Plantinga wants warrant to bear some relationship to truth, the question we need to ask is whether that relationship is inviolable or not. If it is not, if it is possible to

have a warranted false belief, then presumably Plantinga's account of warrant, even left underdefined, is going to be susceptible to Gettier cases of the following form:

> **Under-Defined Plantinga Counterexample:** S forms the belief B as a result of an exercise of cognitive faculties E. S's cognitive faculties are functioning properly in accord with a good design plan aimed at truth, and the maxi and mini environments are favorable for E. By some accident B is not true for reasons related to E (since we are agreeing for the time being that this is possible), but what is more, by some other accident, B still happens to be true (divorced from E). As such, though S has a warranted true belief that B, S does not know that B.

And if warrant, left underdefined, *does* inviolably track truth, then skeptical worries will continue to loom large. Presumably, if warrant needs to guarantee the truth of the belief in question, then Plantinga's environmental proviso (his anti-luck/anti-Gettier condition) needs to preclude any possibility for a false belief, which will be tantamount to requiring a perfect cognitive environment—a condition that presumably few of even our most secure beliefs can meet.

After realizing that his fallibilistic 1993 conception of warrant was susceptible to Gettier counterexample, Plantinga proposed a series of two modifications that sought to strengthen his account of warrant so as to preclude Gettier problems. Unfortunately for Plantinga, both proposals (like his original account) failed precisely along lines predicted by our proposed diagnosis of Gettier counterexamples. Plantinga's 1996/1997 modification offered a strengthened account of warrant that was, nevertheless, ultimately still fallibilistic. As such, it fell prey to strengthened Gettier counterexamples. Plantinga's 2000 modification produced an account of warrant that seemed genuinely infallibilistic, an account of warrant where warrant inviolably tracks truth. As predicted, this led to unsavory skeptical conclusions. And in all of this, the lines taken by Plantinga's critics in the contemporary literature time and time again affirmed our verdicts and our proposed diagnosis of the Gettier Problem.

Conclusion

Alvin Plantinga's reductive account of knowledge is perhaps the most iconic virtue epistemology of the twentieth century—inspiring an entire generation of epistemologists with similar theories of knowledge that are no less intricate and no less multifaceted. However, as we have just seen, Plantinga's reductive analysis of knowledge in terms of

virtue, in terms of properly functioning cognitive faculties, is simply unable to viably surmount the Gettier Problem. From *Warrant and Proper Function* (1993), to "Respondeo" (1996) and "Warrant and Accidentally True Belief" (1997), to *Warranted Christian Belief* (2000), Plantinga's virtue epistemology, in its various stages and iterations, fails in exactly the way our diagnosis of the Gettier Problem in Chapter 1 would predict. Whenever Plantinga's account assumed a close but not inviolable relationship between warrant and truth, he faced Gettier counterexamples. Whenever it assumed an inviolable relationship, he faced unsavory skeptical conclusions. And what is more, our proposed diagnosis explains and informs not only the development of Plantinga's virtue epistemology, but also the critiques leveled against it. Time and time again, there was a revealing correspondence between our proposed diagnosis of Gettier counterexamples and the shortcomings of Plantinga's virtue epistemology as explicated in the relevant literature.[116]

Alvin Plantinga's theory of knowledge, a pillar of reductive virtue epistemology, has fallen—being ineffectual against the Gettier Problem. But there are, to be sure, other, quite different approaches to reductive virtue epistemology—approaches that are simpler and perhaps far more elegant. Ernest Sosa's recent virtue epistemology, which embodies over 30 years of experience on the subject, is perhaps the most prominent virtue epistemology on offer in the recent literature. And instead of trying to satisfy the high demands of reductive analysis by producing an overly complex, intricate, (perhaps) ad hoc account of knowledge, Sosa proposes in *A Virtue Epistemology* (2007) and *Reflective Knowledge* (2009) a relatively straightforward and stylish analysis of knowledge; a belief will be knowledge by Sosa's lights if it is *apt*, if it is "true because competent," if its truth is somehow "attributable to" or "derived from" relevant cognitive faculties (2007, 1:23, 33). In the next chapter, we will elucidate Sosa's analysis of knowledge and consider it in light of our proposed diagnosis of the Gettier Problem, exploring whether or not it can be any more successful at precluding Gettier counterexample.

[116] Within such literature there were a number of diverse proposals trying to salvage Plantinga's reductive virtue epistemology. See, for example, Botham 2003; Chignell 2003; Crisp 2000. Unfortunately, because of our limited goals and space, we are not able to consider such proposals here; however, given the proposed diagnosis of Gettier problems and the light of experience, I think we have excellent prima facie reason to suspect that such proposals are ultimately no more successful.

4 SOSA'S VIRTUE EPISTEMOLOGY

While Alvin Plantinga's reductive analysis of knowledge is perhaps the iconic virtue epistemology of the twentieth century, Ernest Sosa's recent work in *A Virtue Epistemology* (2007) and *Reflective Knowledge* (2009) is arguably the seminal virtue epistemology of the last decade. Where Plantinga tried to satisfy the high demands of reductive analysis through an increasingly complex and intricate account of warrant, Sosa's virtue epistemology, in contrast, is simple and elegant. Knowledge, for Sosa, is *apt belief*, belief that is "true because competent" (2007, 1:23). And what is more, this straightforward analysis of knowledge is iconic in its own right—being representative (at least in form) of an approach taken by other prominent virtue epistemologists, not least Linda Zagzebski, in *Virtues of the Mind* (1996) and "What is Knowledge?" (1999), and John Greco, in "The Nature of Ability and the Purpose of Knowledge" (2007a), "Knowledge and Success from Ability" (2009), and *Achieving Knowledge* (2010).[117]

In this chapter, we will elucidate Sosa's analysis of knowledge and consider it in light of the diagnosis of the Gettier Problem proposed in Chapter 1. And insofar as Sosa's virtue epistemology has not succumbed to a series of modifications and iterations, our work in this chapter can be relatively straightforward. First, in Section 1, we will elucidate Sosa's theory of knowledge—unpacking its virtues and its various dimensions. Second, in Section 2, we will consider it in light of our proposed diagnosis, arguing that the simplicity of Sosa's analysis of knowledge is deceptive—hiding dubious ambiguity. I argue that as soon as we try to elucidate *knowledge as apt belief*, we find that under any plausible interpretation, Sosa's analysis, like Plantinga's, either succumbs to Gettier counterexamples or yields a radical and untenable form of skepticism. This being in keeping with our proposed diagnosis of Gettier counterexamples.

[117] Similar formulations of knowledge can also be found in Greco 2003b, 116; Riggs 2002, 93–94; Sosa 1991, 277.

Section 1: Sosa's Current Virtue Epistemology

Sosa uses the example of an archer to help us understand his theory of knowledge. In assessing an archer's shot, we will be interested in three things: 1) its accuracy—whether or not the shot hits its target; 2) its adroitness—whether or not it manifests the competence of the archer; 3) its aptness—whether or not the shot's accuracy is a result of its adroitness. The question: When is the accuracy of the archer's shot attributable to him? Naturally, if the accuracy of the shot had nothing to do with the archer's competence such that it was not adroit, we would not credit the shot to him; we would not, for example, credit the accuracy of a difficult shot to an inept archer who just so happened to clumsily fire the arrow in the right direction. What is more, if a given shot is accurate and adroit though not apt (i.e., not accurate *because* adroit), we would not credit it to the archer either. Imagine a scenario where an expert archer takes a shot at a relatively easy target, a shot that in normal circumstances would have hit the target without difficulty; imagine, then, that while the arrow is still in the air it is blown off target by a sudden burst of strong wind such that it was then on an inaccurate trajectory and only lucky blown back on target by another sudden burst of strong wind in the opposite direction. While the shot indeed hits the target and seemingly manifests the archer's competence (after all, in normal circumstances it easily hits the target), it is not accurate *because of* the archer's competence—instead being accurate because of the second burst of wind. As such, we would not credit such a shot to the archer. An archer's shot, then, *is* creditable to him if and only if it is accurate as a result of its being adroit. In other words, an archer's shot is creditable to him if and only if it is apt.

Beliefs, like an archer's shots, can also be evaluated according to accuracy (i.e., truth), adroitness (whether or not it manifests the agent's epistemic competence), and aptness (whether or not the belief is true as a result of its adroitness). A given true belief is knowledge, according to Sosa, if and only if it hits upon the truth (so to speak) as a result of the given agent's epistemic competence. In other words, knowledge, for Sosa, is apt belief—belief that is "true *because* competent" (Sosa 2007, 1:23). Parallel to the archery example, if the accuracy of a given belief had nothing to do with the agent's cognitive competence such that it was not adroit, we obviously would not credit the belief to him as knowledge; if, for example, I suddenly and without reason started believing that the pope is my fourth cousin, we would not credit me with knowledge even if by some chance it

turned out to be true. What is more, if a given belief is accurate and adroit though not apt (i.e., not accurate *because* adroit), we would not credit the pertinent agent with knowledge either. Imagine a Gettier-style scenario where I do a detailed investigation into my family history and come to the calculated belief that the pope is my fourth cousin; imagine, then, that while as a result of some bad luck my familial calculations were fundamentally misguided, my belief is nevertheless true—true for reasons my calculations would never have predicted. While my belief in such a case is true and adroit, the truth has nothing to do with the adroitness—it is not true because adroit, it is not apt—and, therefore, we would not credit my belief to me as knowledge.[118]

According to Sosa, S knows that *p* if and only if (i) S believes that *p* and (ii) that belief is apt, true *because of* S's cognitive competency.[119] With this initial analysis in mind, Sosa distinguishes two tiers of knowledge. The first tier of knowledge is what Sosa calls "animal knowledge" and it consists simply of apt belief; a given agent possesses animal knowledge if and only if he possesses an apt belief. The second tier of knowledge, what Sosa calls "reflective knowledge," on the other hand, is more demanding—it is, according to Sosa, not only apt belief but "defensibly apt belief," second order apt belief that a given belief is apt (Sosa 2007, 1:24, 32). I may have apt belief that "that's a zebra," say, and therefore have animal knowledge. To have reflective knowledge, on the other hand, I need to aptly believe that my belief that "that's a zebra" is apt—I might be an expert zoologist who was able to inspect the zebra closely, for example.

Sosa's distinction between animal knowledge and reflective knowledge can be quite handy. Consider, for example, the notorious chicken sexer case:

> **Chicken Sexer:** Naïve and Reflective are both chicken sexers. Their job is to look at baby chickens, determine their genders, and then segregate the chickens accordingly—putting male chicks in one box and female chicks in another. Both Naïve and Reflective are equally good at their job—both are highly reliable at determining the gender of the baby chickens. There is, however, one important

[118] For a big picture of how Sosa sees his account in light of the Gettier Problem see Sosa 2009a, 2:185–189.

[119] For Sosa, (i) and (ii) are seemingly necessary and jointly sufficient conditions on knowledge. What is more, Sosa presumably takes *belief* and *aptness* to be conceptually prior to *knowledge*—*knowledge*, after all, is understood in terms of *belief* and *aptness* and not vice versa. While Sosa's simple analysis looks substantially different from more classic and more complex reductive analyses like Plantinga's, Sosa's virtue epistemology is a reductive analysis, nevertheless. See Sosa 2009b; Williamson 2009b. To see Williamson's critique of Sosa's earlier work, see Williamson 2004 in Greco's edition *Sosa and His Critics* (2004); unfortunately, Sosa's "Replies" (2004) contains no response.

> difference between them: Naïve has no idea how he is able to correctly determine the gender of the chicks; he "just does it"; Reflective, on the other hand, is very much aware of how he makes such a judgment, by looking for a certain pattern in the chick's feathers. Does Naïve's ignorance affect his ability to know "that's a male chick," "that's a female chick," etc.?

The notoriety of this case comes from its divisiveness. Externalists about warrant are typically inclined to say that Naïve's ignorance is no hurdle for knowledge. Internalists, on the other hand, are likely to balk at such a suggestion because Naïve seemingly does not have the right sort of reflective access to what is grounding the pertinent belief. One way to characterize this differing of opinion is through what look like conflicting intuitions. On the one hand, Naïve obviously seems to be doing something right—he is, after all, a highly reliable judge. On the other, surely Reflective's epistemic state is preferable to Naïve's—surely they are not on equal epistemic footing. One of the nice things about Sosa's analysis is that it can account for both of these intuitions. Naïve, it seems, has apt beliefs about the genders of chicks and as such has animal knowledge; Reflective, in contrast, seems to be aptly aware of his apt beliefs concerning the gender of the chicks and as such he enjoys the more exalted reflective knowledge. As such, Sosa's epistemology, with its distinction between animal and reflective knowledge, seems to be able to satisfy some basic internalist intuitions while nevertheless remaining decidedly externalist—it recognizes Naïve's competence all the while appreciating the superiority of Reflective's epistemic state.

Despite such merits, Sosa's analysis of knowledge as apt belief seems to run into some immediate problems. If, for example, apt belief is sufficient for knowledge then seemingly the protagonist in the Fake Barns case has knowledge—a seemingly counterintuitive conclusion. If all it takes for a given belief to be knowledge is aptness, then it seemingly will not matter how lucky that aptness may be. Imagine an archer who is unknowingly in an absolutely terrible environment for archery—an environment that has been plagued by wild and irregular winds for thousands of years. If when he takes his shot a quasi-miracle occurs, the winds cease, and his arrow is able to fly true and hit the target, the shot is, according to Sosa, creditable to the archer—after all, it seemingly hit the target as a result of the archer's competence. And if the archer's shot hit its target because of the archer's competence (i.e., if the shot is apt), the shot is creditable to him no matter how incredibly lucky it may be that the winds died down when they did. Likewise, when a given credulous subject happens to believe "that's a barn" correctly even amid a sea of barn facades, he possesses knowledge. If his belief is true as a result of his epistemic competence, he has

knowledge no matter how incredibly lucky it may be that he picked a real barn. Sosa has us consider the following case, which bears obvious similarities to Fake Barn-type scenarios just referenced:

> **Kaleidoscope**: You see a surface that looks red in ostensibly normal conditions. But it is a kaleidoscope surface controlled by a jokester who also controls the ambient light, and might as easily have presented you with a red-light+white-surface combination as with the actual white-light+red-surface combination. (Sosa 2007, 1:31)

According to Sosa, when presented with the white-light+red-surface combination, the protagonist of Kaleidoscope knows that the surface is red. In other words, so long as the protagonist's belief is true as a result of his epistemic competence, he possesses knowledge no matter how incredibly lucky it may be that the jokester was not at that moment presenting him with a red-light+white-object combination. Environmental luck of that sort does not, according to Sosa, block knowledge.

Given the broad consensus that the protagonist in cases like Kaleidoscope (including cases like Fake Barns) lacks knowledge—that his belief is at the very least somehow epistemically deficient—is not such a conclusion problematic? According to Sosa, not if we make use of the *animal knowledge/reflective knowledge* distinction. By Sosa's lights, the protagonists in such cases do indeed have first order (or first tier) knowledge, animal knowledge, but that is not to say that their epistemic state is not somehow deficient. They may very well have apt beliefs, but surely they lack defensible apt beliefs—in other words, surely they lack higher order reflective knowledge. According to Sosa, when the protagonists in such cases rightly identify barns or the color of surfaces using their epistemic competencies, they should be credited for doing as much even if they could have very easily gotten it wrong. Indeed, where most accounts will simply say that something is epistemically *wrong* in cases like Kaleidoscope, perhaps it is a virtue of Sosa's account that it can also tell us what is epistemically *right*.

That said, it is worth noting that Sosa's diagnosis of cases like Kaleidoscope and Fake Barns may have some serious drawbacks. In his recent paper "Knowledge as Aptness" (2009), Stewart Cohen raises some particularly pressing concerns for Sosa's analysis of knowledge as apt belief that center on examples like the Kaleidoscope case. First of all, consider a case where a jokester makes it such that the given kaleidoscope surface looks red but is actually white illuminated by red lighting for all but one second in a given hour where the surface is indeed red and illuminated by white lighting. If someone who was unaware of

the jokester's presence were to stare at the kaleidoscope surface during that hour, he would, according to Sosa's account, know the surface was red during that one second when the lighting was favorable. But that seems simply wrong. Sosa may try to invoke his distinction between animal knowledge and reflective knowledge to soften the blow, but appealing to our intuitions here would be a mistake since animal knowledge and reflective knowledge are technical terms (Cohen 2009, 122). What is more, invoking such a distinction for these cases only makes sense if we felt pulled in two different directions in the first place; however, in the case just set out, it is not at all clear that we feel such a pull—there is no waffling as to whether or not the protagonist knows.[120] If this is right, then "there is no basis" for introducing the animal/reflective distinction in the first place (Cohen 2009, 122). Consider another worry: contrary to Sosa's account, it seems like once someone was told about the jokester operating the kaleidoscope surface, they would not know its surface was red even in favorable conditions. Cohen points out that Sosa could amend his account to include a "no conflicting competencies" clause such that if someone is aware of the jokester's presence he would not know the kaleidoscope surface was red even in favorable lighting because their competencies are reporting conflicting data—the "proper use of evidence" competency is conflicting the perception competency (Cohen 2009, 122). The problem with such a solution, however, is that it seems strange to require someone to have a proper use of evidence competency in order to make use of a perception competency; one could advocate such a solution but it seems to leave us in the uncomfortable position of denying perceptual knowledge to someone who does not have a competency with evidence.

But putting aside such worries, why does the perceiver in the Kaleidoscope case fail to have reflective knowledge? If he has animal knowledge, why should he not also have reflective knowledge? Putting aside questions as to whether or not Sosa's animal/reflective distinction solution is palatable in principle when applied to such Gettier counterexamples, let us explore the details of Sosa's solution a bit further. Consider Sosa's own explanation:

> The [Kaleidoscope] perceiver would there be said to have apt belief and animal knowledge that the seen surface is red. What he lacks . . . is *reflective* knowledge, since this requires apt belief that he aptly believes the surface to be red (or at least it requires that he aptly take this for granted or assume it or presuppose it, a qualification implicit in what follows). Why should it be any less plausible to think

[120] Sosa, to be sure, does have the intuition that the protagonist in such a case knows that the surface is red for that brief instance. See Sosa 2009c, 140-143.

> that he aptly believes that he aptly believes than to think that he aptly believes *simpliciter*? Well, what competence might he exercise in believing that he aptly so believes, and how plausible might it be to attribute to that competence his being right in believing that he aptly believes? . . . It seems a kind of default competence, whereby one automatically takes the light to be normal absent some special indication to the contrary. And that is presumably what the kaleidoscope perceiver does, absent any indication of a jokester in control. So we may suppose him to retain that competence unimpaired *and* to exercise it in taking for granted the adequacy of the ambient light, so that he can aptly take the surface to be red. Since the belief that he believes aptly is a true belief, and since it is *owed* to the exercise of a competence, how then can we suppose it not to be itself an apt belief? Well, recall: the requirement for aptly believing is not just that one's belief is true, and derived from a competence. The requirement is rather that one believes *correctly* (with truth) through the exercise of a competence in its proper conditions. What must be attributed to the competence is not just the belief's existence but its correctness. (Sosa 2007, 1:32–33)

While the protagonist in Kaleidoscope arguably has animal knowledge, he lacks reflective knowledge because the epistemic competence (that "default competence") by which he might believe his belief that the surface is red to be apt is not itself apt. Why? Because while his belief that his belief that the surface is red is apt is true and causally "owed" to the exercise of the relevant "default" competence, it is not true *because* or *through* the relevant competence. Or, as Sosa says later, the Kaleidoscope protagonist's belief that his belief is apt is not apt because the relevant competence "might . . . too easily have issued a false belief that the lights are normal" (Sosa 2007, 1:33).[121]

Even if Sosa is right that the protagonists in Kaleidoscope cases possess animal knowledge but lack reflective knowledge, his explanation as to *why* this is the case (which we just considered) seems, by his own admission, to lead to a second problem. As Sosa notes, if the main character of the Kaleidoscope case fails to have reflective knowledge because the relevant competence (the default competence by which he might believe his belief that the surface is red to be apt) might too easily have produced a false belief given the close possibility of the jokester's mischief, we might worry that the close possibility of the skeptic's dream hypothesis might prevent us from *ever* having perception-based reflective knowledge. In other words, if the Kaleidoscope protagonist does not have reflective knowledge because the jokester could have easily been playing a trick on him,

[121] We should note that it is very confusing at this point what exactly aptness requires. On the one hand, Sosa seems to say that the protagonist's higher-order belief fails to be apt because its truth is somehow not appropriately contingent on the relevant epistemic competence. On the other hand, Sosa seems to be saying that the higher-order belief fails to be apt because the relevant competence could have very easily led to a false conclusion. This serious discrepancy will be explored further in the next section.

perhaps for similar reasons we cannot possess reflective knowledge of perceptions because we seemingly could easily be so perceiving while in a lucid dream.[122] It looks as though if we apply Sosa's response to cases like Kaleidoscope more generally, we are in the grip of radical skepticism concerning higher-order knowledge.

Before we consider how Sosa avoids such a conclusion, let us consider the Kaleidoscope case a bit further. The reason the protagonist possesses animal knowledge (i.e., apt belief) in such a case is, according to Sosa, that the protagonist's "faculty of color vision" is being exercised "in normal conditions for lighting, distance, size of surface, etc., in conditions generally appropriate for the exercise of color vision" (Sosa 2007, 1:31). Our hero in the Kaleidoscope case has a belief that is apt, a belief that is "true because competent," simply because the competence governing its production is operating in the right sort of conditions such that the given belief could not easily be false and so based (that is, based in the given competence operating in the relevant conditions).[123] In other words, in order for a given belief to be apt, Sosa seems to require that it be produced by a cognitive competence that is operating in good conditions where "operating in good conditions" means at the very least something like "reliably based." Likewise, the reason the protagonist in the Kaleidoscope case does *not* have reflective knowledge is seemingly because the relevant competence, the "default competence" being exercised is *not* operating under normal conditions, in conditions generally appropriate for its exercise such that it would be reliably counted on to produce a true belief. In other words, the protagonist in the Kaleidoscope case does not possess reflective knowledge because the conditions under which the given competence is operating is not a reliable basis for forming true belief given the presence of the jokester.

Now, does the possibility of the skeptic's dream hypothesis usurp reflective knowledge in the standard case in the same way the jokester's presence usurps reflective knowledge in the Kaleidoscope case? Given what has just been said, arguably not. If 1) having an apt belief means having a belief that is produced by a cognitive competence that is operating in

[122] While Sosa explicitly wants aptness to be distinct from safety (Sosa 2007, 1:29), it is at points like this where it looks like a safety principle is somehow built into Sosa's understanding of aptness. For example, it seems like Sosa is saying that the protagonist in the Kaleidoscope case does not have reflective knowledge because the relevant competence could have led to a false belief in close possible worlds—because the truth of the belief, given the relevant competence, is not safe. Contrast Sosa's definition of safety (Sosa 2007, 1:25) with his explanation why the Kaleidoscope perceiver fails to have reflective knowledge (Sosa 2007, 1:33).

[123] See Sosa 1996, 33-35.

appropriate conditions such that the belief could not easily be so based and false, and 2) the skeptic's dream hypothesis did indeed threaten the aptness required for reflective knowledge in the standard cases, then seemingly the conditions under which our relevant competency (i.e., that "default" competency we use when we assume ourselves awake) is operating are not appropriate, that the conditions for our given cognitive competency is not a good basis for reliable true belief. But why think a thing like that? Why would the possibility of a dream-hypothesis affect the appropriateness of the conditions under which my everyday cognitive competencies operate? We can arguably resist such a conclusion by simply assuming that the conditions our "default competence" is operating under when we are awake is markedly different from when we are dreaming. If we want to "appeal, with Austin, to the vividness and richness of wakeful experience, and with Descartes to its coherence, as part of the basis for our belief that we are awake," we can deny that the relevant conditions under which our cognitive faculties are operating when we are awake are the same as when we are dreaming (Sosa 2007, 1:38).[124] Given a difference in conditions, it is simply not true that my relevant default competence would be so operating in the dream scenario—it is not true that my default competence could be so based and easily lead to a corresponding false belief.

Just as we would only credit an archer with hitting the target if he hit it due to his competence as an archer, according to Sosa, we should only credit someone with knowledge if he "hit upon" the given true belief as a result of his cognitive competence. Knowledge, for Sosa, is apt belief—belief that is true because competent. And from this simple definition, Sosa specifies two tiers of knowledge. The first tier, "animal knowledge," is simply apt belief; the second tier, "reflective knowledge," is apt belief aptly believed—an apt belief that one's (first-order) belief is apt. Amongst its various assets, this analysis of knowledge lets us, when faced with Chicken Sexer-style cases, to satisfy both internalist and externalist intuitions; with it, we can say why the reflective chicken sexer's cognitive state is superior to his naïve counterpart, while upholding that there is something the latter is indeed doing right. No doubt, however, when faced with cases like Kaleidoscope someone may worry about the conclusion that the given protagonist knows under such circumstances, but such worries may be assuaged by the animal/reflective distinction—by pointing out that while the Kaleidoscope protagonist may indeed have animal knowledge

[124] See Austin 1962, 48–49 and the last paragraph of Descartes' *Meditations* in 1955.

he surely lacks reflective knowledge. What is more, despite appearances, such a move does not (as we've just seen) lead to widespread higher-order skepticism given closeness of the skeptic's dreaming hypothesis. In fact, Sosa's analysis of knowledge seems to allow for a new and innovative virtue-theoretic solution to certain forms of dream skepticism.

Nevertheless, putting away initial worries regarding Kaleidoscope cases and dreaming, some unresolved problems remain. Most pressingly, we are not yet entirely sure how to interpret Sosa's understand of aptness as "true *because* competent" (Sosa 2007, 1:23).[125] After all, taken literally, a casual reader may think that knowledge, according to Sosa, requires truth to be contingent on cognitive abilities; and given that truth is rarely so contingent, such a reading of Sosa would surely significantly limit what we can know.[126] If *not* taken literally, however, the question remains how we *should* take it. In the next section, we will explore this worry further and eventually argue that Sosa's account, like every other account considered thus far, will be forced to either embrace an unpalatable conclusion (like radical skepticism) or remain susceptible to Gettier counterexamples.

Section 2: Sosa's Virtue Epistemology and the Gettier Problem

When Sosa defines knowledge as apt belief where aptness is understood as "true *because* competent," a casual reader may think that Sosa is saying that knowledge literally requires truth to be contingent on cognitive competencies—a view that seems overtly mistaken. Surely, for example, I can know that the earth orbits the sun even though such a fact has nothing to do with my epistemic competencies—even though the earth's movement around the sun is in no way contingent on my faculties. Such a view aside, we are nevertheless left uncertain how exactly to understand Sosa's conception of aptness. If we should not read aptness as literally "true because competent," how *should* we read it? Our

[125] This is a worry that seems to plague many recent accounts of virtue epistemology. For example, in *Achieving Knowledge* (2010), John Greco defines knowledge in such a way that "*S* knows that *p if and only if* (i) *p* is true; (ii) *S* believes that *p*; and (iii) *S* believes the truth *because S*'s belief is produced by intellectual ability" (2010, 12 - emphasis mine). Greco goes on to admit that "the account is still not maximally informative. This is because it depends on both semantic and pragmatic considerations that are not well understood. Specifically, we would like a better understanding of causal explanations and causal explanation language." However, the worry here, for Sosa *and* Greco, is not just that such "because" terminology is less than "maximally informative" but that it hides dubious ambiguity.

[126] Indeed, we would only know those beliefs that assert, self-referentially, that they were so produced by their given cognitive competence.

focus in this section will be the difficulties that surround various interpretations of aptness, particularly in relation to the Gettier Problem; I argue that no matter how we understand aptness in Sosa's account, it is going to either run into Gettier counterexamples or be forced to embrace an intuitively infeasible conclusion, namely, radical skepticism.

Before we consider how we might interpret Sosa's conception of aptness, it is worth noting that the ambiguity regarding how truth relates to warrant, how truth relates to virtue, is not a new problem for virtue epistemology. In *Virtues of the Mind* (1996) and "What is Knowledge?" (1999), Linda Zagzebski defines knowledge as "belief arising out of acts of intellectual virtue" where a belief "arises out of" an intellectual virtue if and only if the belief is true "because of" or "due to" the said virtue (Zagzebski 1996, 270–271; Zagzebski 1999, 108–109)—an analysis of knowledge that has come under some harsh criticism because it is unclear how to understand the "because of" or "due to" relation. To be sure, while Zagzebski acknowledges that she knows "of no account of the *because of* relation that fully captures [the relevant connection]" and "this concept is in need of further analysis and I do not know of one that is adequate," the problem is bigger than she supposes (Zagzebski 1999, 108, 111). The problem is not just that the "because of," "due to," etc. relation is underdescribed but that there simply is no feasible way to cash out such a relationship. To quote Michael Levin in his critique of Zagzebski, "The problem . . . is not that a gap is left in her account. In this respect the account fares no worse than any analysis of anything that takes the notion of explanation as primitive. The trouble is that there is *no* notion of explanation to fill the gap" (Levin 2004, 401).

As we search for clarity regarding Sosa's understanding of aptness, then, one may worry that Sosa's account is heading down the same path as Zagzebski's. After all, parallels between the two accounts are apparent. Where Zagzebski, a neo-Aristotelian virtue epistemologist, defines knowledge as a belief "arising from," "due to," "because of" acts of intellectual virtue (that is, "true because virtuous"), Sosa, an agent-reliabilist virtue epistemologist, defines knowledge as belief that is true "because" competent. Where Zagzebski might (loosely) define knowledge as something like "virtuous belief," Sosa defines it as "apt belief." While there will, no doubt, be some important differences between the accounts (for example, Sosa's agent-reliabilist account will not be constrained in precisely the same ways Zagzebski's neo-Aristotelian account is), we might nevertheless have a better perspective and a better appreciation as to what is at stake once we turn to

consider Sosa's understanding of knowledge as apt belief (belief that is "true because competent") if we first summarize the worry for Zagzebski's account.[127]

In the paper "Virtue Epistemology: No New Cures" (2004), Michael Levin critiques Zagzebski's analysis of knowledge along these lines—providing us with a good summary of what is at issue in Zagzebski's account.[128] As we saw in Chapter 1, Zagzebski proposed that so long as we assume that truth bears a close but inviolable connection with warrant, Gettier counterexamples are inescapable.[129] In keeping with this diagnosis, Zagzebski, in

[127] For more on the differences between agent-reliabilism virtue epistemology and neo-Aristotelian virtue epistemology, see Greco 2000.

[128] Interestingly, where I am wanting to point to the Gettier Problem to highlight a trend to abandon the standard analysis of knowledge, Levin begins his paper "Virtue Epistemology: No New Cures" by pointing to the Gettier Problem to highlight a trend to abandon the notion of justification—where I would see virtue epistemology as bucking trends Levin sees virtue epistemology as a trend-setter. In this paper Levin argues that, despite being a *response* to the Gettier Problem, virtue epistemology's solution to it is inadequate. He does this by focusing on the work of Linda Zagzebski, who, according to Levin, was the only virtue epistemologist to tackle the Gettier Problem—arguing that if Zagzebski's solution is found inadequate we have prima facie reason to think "other, weaker, versions" of virtue epistemology will be so too (Levin 2004, 397).

[129] In accord with what I have said in Chapter 1, Levin describes Zagzebski's diagnosis of Gettier problems as turning "on bad luck canceled by good"—that so long as a given theory of warrant is does not take warrant as only violably connected to truth (or is not divorced from truth) it "*must entail truth*" because otherwise the said theory would "be liable to such double accidents" (Levin 2004, 398 - emphasis mine). Strangely, however, immediately after providing some quick examples that lend credence to such a diagnosis, Levin concludes that "[a]ccording to Zagzebski . . . a Gettier-proof condition for knowledge must (*not only entail* but) be responsible for truth" (Levin 2004, 399 - emphasis mine). On the face of it, Levin seems to be saying conflicting things about Zagzebski's diagnosis; to avoid Gettier counterexamples is it enough for warrant to entail truth or must it also "be responsible for truth"? To be sure, while Zagzebski's virtue epistemology may be one such account where warrant not only entails truth but is also responsible for it (Zagzebski 1996, 270; Zagzebski 1999, 108–109), I do not take Zagzebski as saying that as much is the threshold for avoiding Gettier counterexamples. For example, consider Fake Barn-type cases. Regarding such a case Zagzebski notes, "Gettier cases are based on situations in which the belief is true, but it might just as well have been false," but if warrant entails truth then the possibility that Henry picks a barn façade is "excluded from the class of justified (warranted) beliefs." According to Zagzebski one could simply assume that warrant entails truth to avoid cases like Fake Barns because in so doing one would be concluding that the barn facsimiles are epistemically ineffectual just so long as the protagonist in question is looking at a real barn. More generally, it looks like someone *could indeed* avoid any Gettier counterexample by (just) assuming that warrant entails truth

Virtues of the Mind (1996) and "What is Knowledge?" (1999), proposes an analysis of knowledge where truth is built into the warrant (where truth and warrant are inviolably connected) and in such a way that it does not, at least at first blush, seem to lead to radical skepticism (as opposed to, say, when safety accounts establish an inviolable connection between truth and warrant). As we have already noted, Zagzebski defines knowledge as "belief arising out of acts of intellectual virtue" such that the truth of the belief is, as she says elsewhere, "due to," "because of," or "explained by" the given virtue (Zagzebski 1999, 109). That said, however, this leaves unexplained how the truth of a belief is (in any way) "explained by" acts of intellectual virtue; the problem, as Levin argues, is that simply no feasible explanation is available.[130]

Zagzebski's official definition of knowledge is "belief arising out of acts of intellectual virtue" or "cognitive contact with reality arising out of acts of intellectual virtue" (Zagzebski 1999, 109). To help make such a definition clearer, Levin provides us with the following notation for Zagzebski's understanding of virtue:

> Let $T(V)$ be the end or telos of virtue V. Then, for Zagzebski, act A, is an "act of V" iff (i) A expresses V-ish motives, (ii) A is the sort of thing V-ish people do, and (iii) A brings about $T(V)$ because of (i) and (ii). (Levin 2004, 399)

For example, the action of saving a child from drowning is an act of courage, say, if and only if (i) such an action expresses courageous motives, (ii) it is the sort of action that courageous people do, and (iii) if, in so doing, the child is actually saved. To be sure, according to Zagzebski, if I try to so save a child yet fail, I may very well have acted courageously but I did not *perform an act of courage*. To use an epistemic case, the action of thinking carefully about a given inquiry is an act of scrupulousness, say, if and only if such an action (i) expresses scrupulous motives, (ii) is the sort of thing that scrupulous people would do, and (iii) if, in so doing, the given telos (which would include the truth) is

---

because it is simply impossible for a belief to be so warranted and false. To be sure, such a move would seemingly be very ill-advised—surely the barn facsimiles *are* epistemically effectual—but one seemingly could, nevertheless, avoid Gettier problems in such a way, unpalatable though it may be. Regardless, this is presumably why Zagzebski, in her own account of warrant, wishes to do something more than simply entail truth—she wants warrant to "be responsible for truth." This is precisely the feature of Zagzebski's account that Levin wishes to criticize. Regardless, this is presumably why Zagzebski, in her own account of warrant, wishes to do something more than simply entail truth—she wants warrant to "be responsible for truth." This is precisely the feature of Zagzebski's account that Levin wishes to criticize.

[130] To see additional worries for *true "because" virtue*-type accounts of knowledge, see Shope 2008.

reached. Again, if I am so cognizing and yet nevertheless form a false belief, I may very well have cognized scrupulously without performing an act of "cognitive scrupulousness" (Levin 2004, 399–400). As such, a given belief is knowledge for Zagzebski if and only if it arises out of a given *A* that (i) expresses *V*-ish motives, (ii) is the sort of thing that *V*-ish people do, and (iii) the given telos (which would include truth) is met.

Rephrased in this way, however, it looks like Zagzebski's official definition is insufficient. Seemingly there is nothing different in the cognitive character between Henry in a "Fake Barn territory" Gettier case and Henry in a "Real Barn territory" good case. In both scenarios, by hypothesis, Henry is performing identical cognitive actions—actions that have *V*-ish motives and that meet their telos. Seemingly, then, according to Zagzebski's official definition, Henry knows that "that's a barn" in Fake Barns—a conclusion that most of us will intuitively find repugnant. To be sure, other problematic cases abound; consider (once again) one of Zagzebski's own Gettier cases.

> **Lucky Mary:** Suppose that Mary has very good eyesight, but it is not perfect. It is good enough to allow her to identify her husband sitting in his usual chair in the living room from a distance of fifteen feet in somewhat dim light (the degree of dimness can easily be specified). She has made such an identification in these circumstances many times. . . . There is nothing at all unusual about either her faculties or the environment in these cases. Her faculties may not be functioning perfectly, but they are functioning well enough, so that if she goes on to form the belief "My husband is sitting in the living room," that belief has enough warrant to constitute knowledge when true and we can assume that it is almost always true. . . . Suppose Mary simply misidentifies the chair-sitter who is, let us suppose, her husband's brother. Her faculties may be working as well as they normally do when the belief is true and when we do not hesitate to say it is warranted in a degree sufficient for knowledge. . . . Her degree of warrant is as high as it usually is when she correctly identifies her husband. . . . We can now easily emend the case as a Gettier example. Mary's husband could be sitting on the other side of the room, unseen by her. (Zagzebski 1994, 67–68)

Here again it looks like there is nothing different in the cognitive character of Mary in the above case and Mary in the nearly identical (good) case where the only difference is that she really does see her husband (call them Mary and Twin-Mary respectively). By hypothesis, Mary and Twin-Mary are performing identical cognitive actions—actions that have *V*-ish motives and that meet their telos. Seemingly, then, according to Zagzebski's

definition Mary knows that "My husband is sitting in the living room" in the case of Lucky Mary—again, an intuitively unacceptable conclusion.[131]

These sorts of problems arise because, according to Levin, Zagzebski's "official definition is too elaborate for its purposes" (Levin 2004, 401). On the face of it, defining knowledge as "belief arising out of acts of intellectual virtue" does not seem to incorporate truth in the right sort of way (according to Zagzebski's own diagnosis of Gettier problems)—such that what is doing the warranting in Zagzebski's account (the virtuous action) entails the truth of the target belief. Repeatedly, Zagzebski describes knowledge as belief that is true "because of" or "due to" an intellectual virtue—a noncausal "acqui[sition of] true belief through the virtues that brought it about"—and it is this feature that it seems is missing in Zagzebski's official definitions (Zagzebski 1996, 270; referenced in Levin 2004, 399). However, it is precisely this feature of her account that she invokes to surmount cases like Mary and Her Husband. To quote Zagzebski:

> [I]n the case of Mary's belief that her husband is in the living room, she may exhibit all the relevant intellectual virtues and no intellectual vices in the process of forming the belief, but she is not led to the truth *through* those virtuous processes or motives. So even though Mary has the belief she has because of her virtues and the belief is true, she does not have the truth because of her virtues. (Zagzebski 1996, 297 - emphasis mine)[132]

As Levin points out, requiring a given belief, *B*, to "arise from some act of virtue *A*, where *A* attains *A*'s end, truth, via good motives *A* expresses" in order for it to be knowledge "interposes *A* between the belief and truth; neither *B* nor *B*'s truth arises from motives good or bad—a result at variance with Zagzebski's plain intent that it is the *truth of the relevant belief* that must be reached via good motives (and acts) for there to be knowledge" (Levin 2004, 401 - emphasis Levin's). To incorporate this feature into Zagzebski's official definition, according to Levin, we need to "bypass the middleman" and identify the given "*A* with the formation of *B*," thus restating Zagzebski's definition of knowledge as "belief

[131] See Pritchard 2005, 194-199. To be sure, Zagzebski, unlike Sosa, does not seem to have the conceptual resources for defusing (or purportedly defusing) Fake Barn-type cases—i.e., the animal knowledge/reflective knowledge distinction.

[132] Also quoted in Pritchard 2005, 196-197.

whose attainment of truth arises from, or is explained by, intellectually good motives" (Levin 2004, 401).[133]

The question then becomes what this "arising from," "due to," or "through" feature of Zagzebski's account consists of. According to Zagzebski, it is not enough to form a true belief via an act of intellectual virtue—the true belief must *arise out of* such an act. As some commentators have noted, however, "[t]his distinction is obscure . . . since it is not at all clear what it involves"—after all, it is not at all obvious what distinguishes Mary's cognitive act in the case where she is "Gettierized" from the case where she's not (Pritchard 2005, 197). To be clear, as Levin points out, no causal explanation is available to Zagzebski—"all causal paths from even the surest motives to belief permit double accidents" and, what is more, such a move would fundamentally change Zagzebski's account from neo-Aristotelian virtue epistemology into motive reliabilism (Levin 2004, 401).

Given Zagzebski's understanding of knowledge as "belief whose attainment of truth arises from, or is explained by, intellectually good motives," how, then, might truth "'arise noncausally from,' or be 'due noncausally to,' or 'explained noncausally by'" virtuous motives (Levin 2004, 401)? We are not told. As we have said earlier, Zagzebski admits that she "know[s] of no account of the *because of* relation that fully captures" the link between truth and virtuous motives—adding that "[w]e all have intuitions about what it means for something to happen because of something else, but this concept is in need of further analysis and I do not know of one that is adequate" (Zagzebski 1999, 111; also quoted in Levin 2004, 401). According to Levin, however, the problem with Zagzebski's account of knowledge is not just that it contains an explanatory gap (after all, as Levin notes, "any analysis of anything that takes the notion of explanation as primitive" will suffer from such problems) but rather "the trouble is that *no* notion of explanation is able to fill the gap"—that "there is no logical room for one" (Levin 2004, 401, 405).

As Levin points out, explaining a given belief will involve either 1) explaining how the belief was formed or 2) explaining why the belief is true (what makes it true). Using Levin's example, if we wanted to explain Homer's belief that he weighs 239 lbs. we might have "one of two explananda in mind"—either we will want to know how he formed such a

[133] It is worth noting that even if we are skeptical as to whether this is the definition of knowledge that Zagzebski intends, it is nevertheless very similar to Sosa's definition—apt belief, that is belief whose attainment of truth is because of or due to a cognitive competence. As such, even if Levin is wrong about how we should read Zagzebski, his reading nevertheless will inform our understanding of Sosa.

belief (e.g., by standing on a scale) or we will want to know whether the said belief is true (i.e., does Homer really weigh 239 lbs.?) (Levin 2004, 401). Take *B*(*p*) to mean "belief that *p*." To rephrase the current point, when asked to explain *B*(*p*), we will want to provide an account of either how *B*(*p*) was formed or of the truth of *B*(*p*), namely, *p*.

Consider the first possible explanandum—explaining how the given belief was formed—what Levin calls a "presence-explanation" (Levin 2004, 401). In seeking a presence-explanation of someone's belief, we are wanting to know why he/she has such a belief—how it got into the person's head. Presence-explanations, to be sure, "are causal . . . citing factors in the environment and within the subject" (Levin 2004, 402). So when we are seeking a presence-explanation for Homer's belief concerning his weight for example, we are looking for an account of how he came to believe such a thing. Levin grants that the causal factors that led Homer to mount the scale may include virtuous motives such as "a burst of no-holds-barred honesty" for self-knowledge; however, "[t]he point to notice is that a presence-explanation of a belief *B*(*p*), the sort of explanation in which mental habits generally appear, is neutral with respect to the truth of *B*(*p*), i.e., of *p*" (Levin 2004, 402). While Homer's belief that he weighs 239 lbs. may indeed be true, Homer could have had the exact same cognitive character with exactly the same relevant stimuli and nevertheless produced a false belief. In other words, the cognitive character and stimuli that led Homer to believe that he weighs 239 lbs. would be the same in a 239 lb. Homer or a 300 lb. Homer. What is more, it seems as though any presence-explanation "ostensibly citing the accuracy of stimuli or the truth-conduciveness of motives" can be consistently diminished to an "internal presence-explanation that omits them" (Levin 2004, 402). As Levin points out, the presence-explanation, "Homer believes he weighs 239 lbs because the scale revealed this," for example, can be consistently diminished to "Homer believes he weighs 239 lbs. because the display read 239 lbs.," which itself can be diminished to "Homer believes [etc.] because he looked at the scale closely enough to see the display," which is also diminishable to "He looked at the scale closely enough to form the belief that it said 239 lbs." (Levin 2004, 402). This all being the case, Levin arrives at his first lemma: "While there is a sense of 'explain,' namely, presence-explain, in which motives explain belief, and these beliefs may incidentally be true, motives do not in this sense explain truth" (Levin 2004, 402–403).

Now consider the second possible explanandum—explaining what makes a given belief true—what we will call a "truth explanation." As Levin points out, it is clearly this

explanandum that Zagzebski (and other virtue epistemologists) has in mind when she describes knowledge as a true belief that is "explained by" (or "due to" or "arising out of" or "because of," etc.) an intellectual virtue. Consider once again Mary and Twin-Mary. Both Mary and Twin-Mary have identical beliefs (or at least qualitatively equivalent beliefs based on "equivalent evidence from identical motives" (Levin 2004, 403). As such, Mary and Twin-Mary have identical cognitive characters concerning the belief in question. The difference between their beliefs, then, is not some psychological feature in either Mary or Twin-Mary. The one difference that Zagzebski is picking up on in her account is that, unlike Twin-Mary's belief, the motives of Mary's belief leave the truth of the belief unexplained (Levin 2004, 403). As Levin notes, however, "Shifting the explanandum to truth . . . pushes the explanans out of the believer's mind" (Levin 2004, 403). Levin explains:

> Why was Homer's belief about his weight true? Not because of whatever motives led him to form this belief, nor the merits of those motives, but because Homer does weigh 239 lbs. Whence did the truth of this belief arise? From Homer's weight, or perhaps the diet, genes, and lifestyle responsible for it. (Levin 2004, 403)

Virtuous motives, being psychological in character, in no way "explain" or "give rise to" the truth of (most) beliefs—"Insofar as truth can be explained, motives do not explain it" (Levin 2004, 403).[134] A second lemma, then: "while there is a sense in which the truth of beliefs can be explained, the truth of beliefs cannot be explained in the manner Zagzebski requires" (Levin 2004, 403).

To be sure, there may be cases where a presence-explanation seems to coincide with a truth-explanation. Consider the following:

A) "Homer was right about his weighing 239 lbs. because he forthrightly mounted a scale." (Levin 2004, 403)

In such a statement, the truth of Homer's belief seems to be directly dependent on his actions—the presence-explanation of Homer's belief seems to encompass its truth-explanation. In other words, the true belief in such a statement indeed seems to be "due to," "because of," or "explained by" the protagonist's cognitive character. As Levin points out, however, appearances deceive—"[i]n effect, (A) telescopes a nonevaluative presence-

[134] As Levin points out in a footnote, the exception would be when "*p* asserts, self-referentially, that *B*(*p*) was produced by that very motive" (Levin 2004, 403).

explanation of Homer's belief into a nonexplanatory characterization of that belief as true" (Levin 2004, 403–404). Once the presence-explanation and truth-explanation are isolated, however, A amounts to the following:

B) "Homer weighs 239 lbs. & he believes he does because he forthrightly mounted a scale." (Levin 2004, 404)

Leaving the two explananda telescoped, however, A amounts to this:

C) "The truth of Homer's belief that he weighs 239 lbs. is explained by his willingness to mount a scale." (Levin 2004, 404)

As Levin points out, however, while B seems plausible (certainly no less plausible than A), C seems extremely odd if not mistaken. We should not, therefore, conclude from statements like A—statements where a presence-explanation seems to coincide with a truth-explanation—that there is some middle ground where presence-explanations and truth-explanations overlap. In accord with Levin's argument, presence-explanations and truth-explanations are well and truly distinct.

Presence-explanations do not explain the truth of a given belief (lemma 1), and in the vast majority of cases truth-explanations in no way derive from the given agent's cognitive character (lemma 2). From these two lemmas, Levin summarizes the above reasoning with the following—what he calls a "quick a priori argument" (Levin 2004, 405):

> Beliefs have two sorts of properties: psychological, in virtue of occurring in minds, and semantic, in virtue of possessing content. Presence-explanations cover properties of the first sort, [truth-explanations] cover properties of the second. (The absence of mention of these two properties should not be mistaken for mention of some further one.) The two lemmas taken together, then, rule out the kind of explanation, the sort of "because" Zagzebski desires. There are, then, no other senses in which a true belief can be explained. There is no logical room for one (Levin 2004, 405).

As such, if we, like Zagzebski, are to describe knowledge as true belief "arising out of," "due to," "because," or "explained by" virtuous motives, then we must embrace the uncomfortable conclusion of radical skepticism. To the extent that it is impossible to satisfy the "explaining" condition, it is impossible to satisfy the conditions of knowledge so defined.

Zagzebski wants to define knowledge as something like "belief whose attainment of truth arises from, or is explained by, acts of intellectual virtues"; such a definition, however, leaves us wondering how to fill in the "explained by" relation. As Levin points out, there are three ways we might explain a given true belief—*reliabilistically*: explaining causally how

a given agent arrived at said true belief; *in terms of belief:* explaining why the agent believes said belief; *in terms of truth:* explaining why the given belief is true. According to Levin, we should not read Zagzebski's definition of knowledge *reliabilistically* because 1) "all causal paths from even the surest motives to belief permit double accidents" and 2) doing so would be to fundamentally change the neo-Aristotelian character of Zagzebski's account (Levin 2004, 401). We should not read Zagzebski's definition of knowledge *in terms of belief* either, because doing so is incompatible with the given definition. What is more, we should not read Zagzebski's definition of knowledge *in terms of truth* because doing so would radically limit what we know—seemingly, we would know very few things if knowledge required truth to be contingent on our intellectual virtues. This being the case, how should we understand Zagzebski's definition? Lacking any further possible filling-in, it looks as though Zagzebski's account of knowledge is fundamentally incomplete—bearing a critical gap that simply cannot be filled.[135]

So what does all of this have to do with Sosa? Let's quickly take stock. Clearly, Sosa's analysis of knowledge parallels Zagzebski's in many important respects. As we have already

[135] As Levin points out, some people may find it "implausible that Zagzebski could have gone wrong in the fundamental way suggested" (Levin 2004, 408). This is, it seems to me, a very legitimate worry to have; nevertheless, in response, Levin highlights a way Zagzebski (and perhaps other, like-minded virtue epistemologists) could have easily been so confused:

"Let *C* be a condition reasonably thought to be by definition necessary (and with truth sufficient) for knowledge. Some traditional epistemologists allegedly confused analyses of the form

(E) □(*p* is known ≡ *p* is true & *C*)

with

(F) *p* is known ≡ □(*p* is true & *C*),

and inferred from that what is known is necessarily true, or, worse, that there is a necessary connection between the truth of *p* and *C*. On its face Zagzebski's definition,

(G) □(knowledge = truth via motives from which truth arises),

gives □ properly wide scope, hence requires nothing from which truth arises necessarily. But 'arises' as Zagzebski understands it carries necessity already built in. Recall her central argument that motive (or any other factor) merely contingently connected to truth is Gettierable, *hence insufficient for knowledge*. It is natural to conclude (and the suggestion is that Zagzebski has concluded) that knowledge calls for something close enough to truth to exclude the possibility of error and accident, in effect pushing the □ in (F) right next to 'arises,' or, more accurately, turning 'arises' into 'necessarily issues.' . . . The idea that the knowledge-creating condition is motivational plus the ideal that it be necessary would explain the quest for a motivational state capable by itself of attaining truth" (Levin 2004, 408–409).

noted, where Zagzebski defines knowledge as a belief "arising from," "due to," "because of" acts of intellectual virtue (that is, "true because virtuous"), Sosa similarly defines knowledge as belief that is true "because" competent—belief where the truth is "derive[d] from" or "attributable to" a competence (Sosa 2007, 1:33). Likewise, where Zagzebski might loosely define knowledge as something like "virtuous belief," Sosa defines it as "apt belief." Indeed, just as we might wonder how a belief can be true "because" virtuous in Zagzebski's account, we wonder how a belief can be true "because" competent in Sosa's. Seemingly, then, as in Zagzebski's case, we can try to fill in Sosa's "because" terminology in any of three ways—*reliabilistically*: explaining causally how a given agent arrived at said true belief; *in terms of belief:* explaining why the agent believes said belief; *in terms of truth:* explaining why the given belief is true. Seemingly, we should not read Sosa's definition of knowledge *in terms of belief* because doing so is simply incompatible with the given definition—there is no way to explain the "because" relation in terms of belief and account for truth. Likewise, just as in Zagzebski's case, we should not try to read Sosa's definition of knowledge *in terms of truth* either because doing so would radically limit what we know; we would seemingly know very few things if knowledge required truth to be contingent on our competencies. *Unlike* Zagzebski's account, however, Sosa's account is, as we would expect, perfectly compatible with a *reliabilistic* reading—doing so would by no means be a fundamental departure from the brand of virtue epistemology Sosa wishes to endorse.[136]

And thankfully, this seems to be precisely the view that Sosa is assuming when he defines knowledge as apt belief—as belief that is "true because competent." Instead of interpreting "true because competent" literally so as to take truth contingent on epistemic competencies, Sosa is seemingly saying something far more commonplace, if not far more sensible. When Sosa says that in order for a belief to be knowledge it must be "true because competent," perhaps he is simply saying something like in order for a belief to be knowledge it must be produced by an epistemic competency that is very likely to hit upon the truth. After all, Sosa sometimes seem to say just that. When Sosa explains why the protagonist in Kaleidoscope Case does not aptly believe he has an apt belief, he says it is because, given the jokester, "the exercise of that competence might then too easily have issued a false belief that the lights are normal" (Sosa 2007, 1:33). In other words, Sosa

[136] Such a reading is, of course, all in keeping with Sosa's broader philosophical commitment to agent-reliabilism virtue epistemology. After all, in Sosa's more recent work he occasionally even identifies his own position as "virtue reliabilism" (Sosa 2009a, 2:138).

seems to be denying the protagonist in Kaleidoscope Case the aptness required for second-order knowledge because the relevant epistemic competence at issue is not sufficiently reliable under the given circumstances. Elsewhere, Sosa seems to worry about what the closeness of the skeptic's dream hypothesis might do to the reliability of our everyday epistemic competencies, saying:

> If common sense is to prevail, based on our virtue epistemology, we must see how, in ordinary perceptual belief, one can aptly presuppose, or take for granted, that the relevant competence and conditions are in place. But the aptness of any such presupposition would require that it be correct because of a competence exercised in the conditions in which it is exercised. And the relevant competence seems nothing more than a default competence of assuming ourselves awake whenever conscious, absent any specific indication to the contrary. But the ease with which we might have gone wrong by so presupposing on such a basis is proportional to the proximity of the dream possibility, and that is really too close for comfort. So we would have to conclude that our getting it right when we ordinarily believe ourselves awake is not attributable (sufficiently) to the exercise of our default competence. . . . We do not get it right through competence in presupposing ourselves awake, since the supposed competence that we exercise, in its proper conditions, *might too easily lead us astray*. (Sosa 2007, 1:35 - emphasis mine)

Here again, Sosa seems to assume that in order for a belief to be apt the relevant epistemic competencies must reliably lead to the truth, since a belief won't be apt if the relevant competencies "might too easily lead . . . astray." Perhaps, then, in accord with Levin's diagnosis, there is good reason to read Sosa as a reliabilist of sorts, and perhaps he would agree that having truth somehow contingent on epistemic competencies is crazy; the question then becomes: Is such a view, more sensible though it may be, successful against Gettier counterexamples?

The answer, unfortunately, is a resounding "no." Given the diagnosis of the Gettier Problem found in Chapter 1 with its account of luck, the shortcomings of such a reliabilist theory when it comes to Gettier counterexamples are apparent. Reliabilist theories have notoriously been unable to answer the Gettier Problem because demanding anything shy of complete reliability lends itself to Gettier double-accidents. Given 1) that luck admits degrees such that even the slightest degree (of the relevant sort) is of epistemic consequence and 2) that degrees of luck can be mapped onto degrees of reliability, then the only way such a reliabilist theory can entirely avoid Gettier counterexamples is if it requires perfect reliability, radically limiting what we know.

Examples demonstrating as much abound; simply take a belief that is produced by an incredibly reliable (though fallible) cognitive competence and Gettierize it—stipulating that

1) by some accident the competence in question was in error and 2) by some additional accident the belief is true for reasons not (or at least not entirely) predicted by the given competence. Indeed most any Gettier case can be used to demonstrate as much. Consider once again a classic Gettier case:

> **Classic Case**: Smith and Jones are applying for the same job. Smith has very strong evidence for thinking that Jones will get the job (e.g., the employer tells Smith that he will hire Jones, etc.), and for thinking that Jones has 10 coins in his pocket (e.g., Jones emptied his pockets in front of Smith and then clearly, slowly, in good lighting, and perhaps even counting out loud, placed 10 coins in his pocket). As such, Smith forms a belief in the general proposition that "the man who gets the job has 10 coins in his pocket." As it turns out, however, Smith gets the job and he happens to also have 10 coins in his pocket. (paraphrased from Gettier 1963, 122)

So long as we stipulate that the cognitive faculties responsible for Smith's various beliefs regarding who is going to get the job and how many coins are in Jones's pocket are incredibly reliable, then Sosa's account of knowledge (read reliabilistically) will give us the wrong result—it will predict that Smith knows that "the man who gets the job has 10 coins in his pocket."

If we define knowledge as apt belief, and aptness is roughly understood to mean reliably (though fallibly) truth-conducive, then no small number of Gettier counterexamples will be at hand. Sosa could, of course, always object that the cognitive competencies in such cases are not reliable *enough*, but, as the post-Gettier literature has shown time and time again, strengthened cases can be easily produced. Given that we are talking about cognitive competency types and not tokens, we can (for at least ease of comprehension) put matters probabilistically. Let us say that in order for a given belief to be knowledge it needs to be produced by a cognitive competency that is so reliably truth-conducive that it only produces a false belief once every 1,000,000 times it is exercised; a Gettier case can nevertheless be produced out of that one time the given cognitive faculty is imperfect by making the relevant belief true for reasons divorced from the competence. To be sure, the same thing can be done with a belief that is produced by a cognitive faculty that is so reliably truth-conducive that it only produces a false belief once every 10^{30} times it is exercised. As such, the only way Sosa's account (read reliabilistically) can completely avoid Gettier counterexamples is if it demands absolutely *perfect* reliability from cognitive competencies. And given that few (if any) of our beliefs can meet such a demand, one may

worry that Sosa's account would then be in the grip of radical skepticism—precisely as our proposed diagnosis of Gettier counterexamples would predict.

To be sure, Sosa might respond by saying that "the requirement for aptly believing is not just that one's belief be true, and derive from a competence . . . [it is] rather that one believe *correctly* (with truth) through the exercise of a competence in its proper conditions" (Sosa 2007, 1:33). In other words, while Sosa might grant that the true belief in the above Gettier cases is "owed to the exercise of competence," but deny that it is appropriately "derived from" or "attributable to" the competence—after all, the belief must be "true *because* competent" (Sosa 2007, 1:23, 33). While there are certainly many places where Sosa seems to think about aptness reliabilistically (as referenced earlier), perhaps that is not the right way to read him. But now, of course, we are back to where we started, unsure as to how to understand this "derived from," "attributable to," "because" relation. If we read such a relation literally instead of reliabilistically, then Sosa's aptness requirement on knowledge seems to demand that the truth of a given belief be contingent on the relevant agent's cognitive faculties. And unless we concoct a nonstandard understanding of truth (one that will surely be at odds with common sense), then such a literal reading of Sosa's requirement simply seems crazy. Again, I know that the earth orbits the sun even though such a fact has nothing to do with my cognitive competencies—the truth of such a belief is in absolutely no clear way "derived from," "attributable to," or "because" of my cognitive competencies. So if we don't read Sosa's aptness requirement reliabilistically, Sosa seems to be in the same position as Zagzebski—left trying to fill in the "because" terminology either *in terms of belief* or *in terms of truth*. Unfortunately, however, neither one of these options seems viable.[137]

[137] When I leveled this objection against Sosa at the Edinburgh Epistemology Graduate Conference (2011), part of his reply was to invoke the notion of manifestation—arguing that what he means by "true because competent" is that the given competence *manifests* the truth of the belief. Interestingly, a similar line has recently been taken in John Turri's "Manifest Failure: The Gettier Problem Solved" (2011). Unfortunately, however, I don't think this sort of response makes significant headway against the above critique, and I certainly don't think it is of any help whatsoever against the Gettier Problem. Explicating manifestations with any degree of precision is a notoriously sticky issue in metaphysics, so trying to clarify "true because competent" by invoking the notion of manifestation seems to only trade obscurity for obscurity in another subdiscipline. What is more, manifestation suits degrees, so, given our current diagnosis, the only way an account of knowledge as *true belief that manifests competence* will be able to surmount Gettier counterexample is if

Even if we allow that there is some way to fill out the "because" relation (in some way that is not *reliablistic* or *in terms of belief* or *in terms of truth*), Sosa's account still runs into Gettier-trouble. Even if we leave Sosa's account ambiguous as to how we should interpret the "because" relation, Sosa's theory of knowledge seems to succumb to the same problems that have plagued every other account considered thus far.[138] According to the diagnosis of the Gettier Problem offered in Chapter 1, the only way to avoid Gettier counterexamples is to assume that either warrant is divorced from truth or inseparable from it, but such strategies, so the worry goes, lead to unfeasible results. Denying that warrant bears any relationship to truth, of course, seems wrong as it results in baldly denying that the protagonists in Gettier cases fail to know—going against the near-universal intuition that warrant has something to do with truth. Assuming that warrant is inviolably connected to truth, on the other hand, seems to readily lead to radical skepticism (as we saw with Plantinga's and Pritchard's respective accounts). Sosa's account, as I will now argue, is no different—in trying to avoid Gettier counterexamples, Sosa has constructed a theory of knowledge that leads to an infeasible result *even if aptness is left ambiguous*.

As we have already discussed, a given belief will be knowledge, according to Sosa, if it is accurate because adroit, if it is true because competent, if it is apt. What is more, given that truth is built into the warrant of Sosa's account ("*true because competent*"), he is assuming for his epistemology an inviolable connection between warrant and truth, and in so doing he is eliminating the possibility of Gettier counterexample (in accord with the diagnosis of Chapter 1) since it is simply impossible for a warranted belief in Sosa's account to be false—it is impossible for a belief to be so warranted but true for other reasons. *By putting pressure on Sosa's notion of warrant and pushing the "true because competent" condition to its logical conclusions, however, we will see how in avoiding Gettier problems Sosa's account, even left ambiguous, is of the sort of shape that could still easily be taken as leading to either Gettier counterexample or radical skepticism.*

---

it demands that the true belief *perfectly* or *completely* manifest the relevant competence—a requirement that presumably very few of even our most secure beliefs can meet.

[138] This strikes me as a very uncomfortable position, but perhaps, for example, there is some way to read "true because competent" literally but in such a way that is palatable. In other words, perhaps there is a way to use a nonstandard definition of truth or extend our commonsense understanding of cognitive competencies in such a way that allows for a literal reading of "true because competent" while preserving the core of Sosa's epistemology.

As Sosa freely admits, adroitness and aptitude both come in degrees—an archer's shot may be more or less skillful and subsequently may hit its target as a result of that skill to a greater or lesser extent (Sosa 2007, 1:22).[139] As such we can imagine situations where an apt belief fails to be knowledge because the relevant conditions are only weakly met. For example, with a difficult target, an expert archer may need to do a hundred different things to guarantee success (e.g., regulate breathing, hold the arrow *just so*, etc.) such that when a novice, only doing three of those things, hits the target we would not credit the shot to him. While the shot was to a small degree accurate because of those three things the novice did (i.e., marginally apt), he was, nevertheless, wildly lucky in hitting his mark. Likewise, an expert horticulturalist may need to do a hundred different things to make sure she has rightly identified a rare species of plant; as such, when I, a novice horticulturalist at best, only clumsily do a couple of those things, we would not credit me with knowledge even if I happen upon the right answer. While my answer was to a small degree correct because of my (limited) competence as a horticulturalist, I was, nevertheless, wildly lucky in hitting the epistemic mark.

This is not, to be sure, a problem in and of itself; it's not, in any case, a problem particular to Sosa's analysis of knowledge. Every analysis of knowledge that accounts for degrees of warrant will run into similar circumstances—circumstances where a weakly warranted true belief does not amount to knowledge. As such, the tacit qualification in such analyses is something like "knowledge is *sufficiently* warranted true belief." Needless to say, then, in defining knowledge as apt belief, Sosa surely tacitly means *sufficiently* apt belief.[140] There are, to be sure, difficult questions as to just *how* justified, *how* safe, *how* apt, *how* warranted a given belief needs to be in order for it to qualify as knowledge, but such is simply endemic of the vagueness that plagues epistemology and is certainly not particular to Sosa's analysis.

Highlighting degrees of competence and degrees of aptitude does, however, help us see just what Sosa's conditions for knowledge must demand if they are to rightly handle Gettier counterexamples. Consider the case of Dr. Jones and the Virus once again:

---

[139] Sosa also acquiesces that accuracy comes in degrees. In the epistemic case, however, we can put this aside; we need only concern ourselves with whether or not the belief in question is true or false—not "how true" or "how false" the belief may be.

[140] Sosa says as much in a footnote (Sosa 2007, 1:22).

> **Dr. Jones and the Virus:** Smith is ill and exhibits a unique set of symptoms, S. Given these symptoms, Dr. Jones forms the belief that "Smith has Virus X," which she deduces from the true proposition that "Virus X is the only known virus to exhibit S." What is more, Dr. Jones does a blood test which verifies that Smith's body contains antibodies for Virus X, further justifying Jones's belief. Based on the evidence, it is extremely feasible that Smith has Virus X. As it happens, however, Smith's symptoms are in fact due to an unknown virus, Virus Y, which exhibits identical symptoms to Virus X; Smith only exhibits antibodies for Virus X due to an idiosyncratic feature of Smith's particular biochemistry which causes his immune system to maintain high levels of antibodies long past a given infection. Nevertheless, Dr. Jones's belief turns out to be true divorced from Smith's symptoms or his bloodwork, because Smith was infected with Virus X just before meeting with Dr. Jones—the infection being so recent that bloodwork cannot detect it and it is causing no symptoms. (paraphrased from Zagzebski 1994, 71)

Surely it would be wrong to say that Dr. Jones's hitting on a true belief is completely divorced from her epistemic competences as a doctor. Surely, for example, Dr. Jones's hitting on a true belief exhibits far more adroitness than, say, someone with no medical training's wild guess. As such, to rightly handle such a case Sosa must stipulate that Dr. Jones's belief is not nearly apt enough—its truth does not depend *enough* on Dr. Jones's adroitness. Given that degrees of luck can be mapped onto the degrees of aptitude, however, stronger cases can (once again) be produced. Consider the following similar case:

> **Dr. Virtuoso and the Virus:** Smith is ill and exhibits an extremely unusual set of symptoms, S. Given these symptoms along with a myriad of advanced, thorough, and highly technical medical tests, Dr. Virtuoso, the leading medical mind of her day, forms the belief that "Smith has Virus X," which she deduces from the true proposition that "Virus X is the only known virus to exhibit S and to so respond to the tests." Based on this superb evidence, it is extremely feasible that Smith has Virus X. As it happens, however, Smith's symptoms and test results are due to an unknown virus, Virus Y, which, for various bizarre and quasi-miraculous reasons, exhibits extremely X-like symptoms and produces extremely X-like test results in Smith. Dr. Virtuoso's belief, nevertheless, happens to be true, because just before meeting with Dr. Virtuoso, Smith contracted Virus X as a result of an extremely bizarre and quasi-miraculous state of affairs—Smith has contracted Virus X so recently, in fact, that it is neither producing symptoms nor affecting the aforementioned tests.

We can stipulate that no one other than Dr. Virtuoso could have so skillfully conducted and read the "myriad of advanced, thorough, and highly technical . . . tests" or so wisely remembered that Virus X is the only known virus to produce S. Surely, then, Dr. Virtuoso's correct diagnosis is largely due to her medical expertise and exhibits a great deal of adroitness. In other words, Dr. Virtuoso's belief is surely *to a large extent* true because competent (no matter how we understand such a "because" relation); however, as it is not

completely true because competent, there is room for luck and room for Gettier counterexample.

In Chapter 1, I lent credence to Linda Zagzebski's diagnosis of the Gettier Problem through an analysis of luck—arguing that Gettier cases can arise out of even the slightest degree of (the relevant sort of) luck. We now see this thesis vindicated in Sosa's analysis. In a footnote, Sosa admits that while aptness is a matter of degree, "a performance is apt only if its success is *sufficiently* attributable to the performer's competence" (Sosa 2007, 1:22 - emphasis Sosa's); unfortunately, however, the only way for Sosa's analysis of knowledge to sufficiently avoid Gettier counterexamples is if aptness requires success to be *completely* attributable to the performer's competence. As we saw with the case of Dr. Virtuoso and the Virus, anything less than "*completely* true because competent" will not do. If Sosa is to rightly build truth into his account of warrant so as to make Gettier counterexamples impossible, "a performance is apt only if its success is *sufficiently* attributable to the performer's competence" simply must mean "a performance is apt only if its success is *completely* attributable to the performer's competence."

Unfortunately, however, while "a performance is apt only if its success is *sufficiently* attributable to the performer's competence" may sound plausible, "a performance is apt only if its success is *completely* attributable to the performer's competence" does not. If knowledge requires apt belief and if apt belief requires the truth of the belief to be completely (and perfectly) attributable to the agent's epistemic competencies (whatever that might mean), surely we know very little. Surely there are only a precious few beliefs we may have whose truth is fully attributable to our competencies—even animal knowledge would be extremely difficult to come by (in humans let alone other animals). If knowledge is apt belief, then, even left ambiguous as to what that means, it looks as though we might be (once again) in the grips of radical skepticism—precisely as our diagnosis of Gettier problems in Chapter 1 would predict.

Ernest Sosa's virtue epistemology is innovative and interesting. However, even if we are willing to accept Sosa's nonstandard diagnosis of Kaleidoscope and Fake Barns-type cases, his analysis of knowledge as *apt belief*, belief that is "true because competent," seems to face some serious problems. First of all, it is not at all clear how such an analysis should be understood. As Levin's critique of Linda Zagzebski showed us, every possible interpretation of "because" leads to either radical skepticism (when knowledge requires the

truth of a given belief to be literally caused of a cognitive competence) or Gettier counterexample (when read reliabilistically). But what is more, no matter how Sosa's analysis of knowledge is understood, it looks as though the only way Sosa's analysis of knowledge as apt belief will be able to avoid Gettier counterexamples altogether is if it requires *perfect* aptness (whatever that might mean). Such a requirement seems a tall order; surely very few of even our most epistemically secure beliefs are candidates for perfect aptness. Radical skepticism, then, seems to be the only way Sosa's account can enjoy immunity from the Gettier Problem. Our proposed diagnosis of Gettier problems is, yet again, vindicated.

Conclusion

The theory of knowledge advanced in *A Virtue Epistemology* (2007) and *Reflective Knowledge* (2009) is truly monumental—arguably the seminal virtue theoretic account to date. Indeed, as some commentators have observed, "*A Virtue Epistemology* . . . is arguably the single-most important monograph to be published in analytic epistemology in the last 10 years" (Battaly 2009, 382). Nevertheless, Sosa's analysis of knowledge falls on the horns of the dilemma explicated in Chapter 1—the same dilemma that plagued Plantinga's analysis in Chapter 3.[141] Again, it seems as though the only way Sosa's analysis of knowledge as apt belief (under any plausible interpretation) will be able to sufficiently avoid Gettier counterexamples is if it requires *perfect* aptness, however "aptness" is understood. And again, such a requirement seems a tall order; surely very few of even our most epistemically secure beliefs are candidates for perfect aptness. Radical skepticism, then, seems to be the only way Sosa's account can enjoy immunity from the Gettier Problem, once again confirming our diagnosis of the Gettier Problem.

Ernest Sosa's seminal and iconic virtue epistemology is no more successful at avoiding Gettier counterexamples than Alvin Plantinga's. While Sosa and Plantinga are both champions of virtue epistemology and the reductive analysis project more generally, neither of them is able to break the series of analyses followed by critique and counterexample that has been the trademark of the reductive analysis project for nearly fifty years running. As

[141] I am again assuming here that Zagzebski's trilemma diagnosis of Gettier counterexamples is functionally only a dilemma for philosophers like Plantinga and Sosa. In other words, I simply assume that Planting and Sosa would find a divorced relationship between warrant and truth (the third leg of Zagzebski's trilemma) unpalatable.

such, one may begin to wonder when enough is enough; one may wonder whether the only way to preserve virtue-theoretic epistemology is to adapt it to suit an alternative non-reductive model of knowledge. In the next chapter, we motivate and explore by way of prolegomena the possibility of non-reductive virtue epistemology.

PART III:
EXPLORING ALTERNATIVES

5 PROLEGOMENA TO NON-REDUCTIVE VIRTUE EPISTEMOLOGY

In the introduction of this thesis, I painted a rather dim picture of the project of pursuing a reductive analysis of knowledge—the project of defining knowledge in terms of its necessary and jointly sufficient conditions. While the post-Gettier literature saw an explosion of energy and interest in epistemology, producing a dizzying array of analyses with diverse and multifaceted conceptions of warrant, no lasting success was found—every proposed analysis either falling victim to further Gettier-style counterexamples or leading to unpalatable results (e.g., radical skepticism).

Despite this grim and seemingly hopeless history, the majority of work done in contemporary epistemology toward defining knowledge is nevertheless still pursued via reductive analysis.[142] The reductive analysis project is still highly prevalent; its ghost has by no means yet been exorcised. While there is, no doubt, a growing trend to abandon the reductive analysis project, advocates of such a move are plainly in the minority. One of the primary and simple goals of this thesis has been to lend credence to this trend by criticizing some of the chief proponents of what is perhaps the most popular species of reductive analysis to date, virtue epistemology.

To this end, I proposed a diagnosis of the Gettier Problem in Chapter 1, which accounted for when and why Gettier counterexamples occur. In Chapter 2, I defended this diagnosis against possible objections found in the contemporary literature. Then, in Chapters 3 and 4, I applied my diagnosis of the Gettier Problem to the recent analyses of knowledge of two eminent virtue epistemologists—Alvin Plantinga and Ernest Sosa—whose collective works have been enormously influential and largely responsible for virtue

[142] I am assuming that there is a difference between "defining" knowledge and providing an "analysis" of knowledge, that abandoning the reductive analysis project does not mean the abandonment of defining knowledge. Just as we might still define the concept "red" without defining it in terms of a reductive analysis, I take it that we can define the concept "know" without defining it in terms of a reductive knowledge.

epistemology's widespread popularity.[143] Applying the diagnosis proposed in Chapter 1, I argued in Chapter 3 that Plantinga's epistemology is, as predicted, either susceptible to Gettier counterexamples or unfeasible, leading to radical skepticism. And in Chapter 4, I argued that Sosa's epistemology ultimately faces the same dilemma.[144]

And given the pattern that has been established in this thesis, there is no reason to think that the problems afflicting Plantinga and Sosa are localized to their respective accounts. In other words, we have no prima facie reason to think that other prominent virtue-theoretic epistemologies on offer in the contemporary literature will fair any better against the Gettier Problem as understood in Chapter 1. First of all, consider Linda Zagzebski's neo-Aristotelian virtue epistemology. Despite having convincingly argued that Gettier counterexamples are inescapable in 1994, Linda Zagzebski proposed in *Virtues of the Mind* (1996) and "What is Knowledge?" (1999) a virtue-theoretic reductive analysis of knowledge that is meant to surmount the Gettier Problem. Unfortunately for Zagzebski (and somewhat ironically), however, her analysis will seemingly fail precisely along the lines "The Inescapability of Gettier Problems" (1994) predicts. Knowledge, for Zagzebski, is roughly defined as "true belief arising out of intellectual virtue" (Zagzebski 1996, 270–271; Zagzebski 1999, 108–109); given that a belief can "arise out of" an intellectual virtue to a greater or lesser extent, however, degrees of luck (as explicated in Chapter 1) can seemingly be mapped onto the degrees in which a true belief can more or less "arise out of" an intellectual virtue—leading to the now-familiar dilemma of forcing the analysis to either place an unreasonably high demand on knowledge (e.g., "true belief that has arisen *completely* from an intellectual virtue") or fall victim to Gettier counterexample.[145] A similar story can be told regarding John Greco's most recent virtue-theoretic account in "The Nature of Ability and the Purpose of Knowledge" (2007), "Knowledge and Success from Ability" (2009), and *Achieving Knowledge* (2010), which, unlike some of his previous accounts,

[143] To be sure, the influence of Plantinga and Sosa extends beyond epistemology and into a number of other subdisciplines such as philosophy of religion, metaphysics, and ethics.

[144] I am assuming that the third leg of Zagzebski's trilemma, that warrant is completely divorced from truth, is not a viable alternative for these virtue epistemologists.

[145] And this is of course to ignore the more pressing worries for Zagzebski's definition of knowledge as explicated in Chapter 4 (by way of a critique of Sosa) and in Michael Levin's "Virtue Epistemology: No New Cures" (2004).

are meant to directly handle the Gettier Problem.[146] Unfortunately for Greco, however, his account (another form of agent-reliabilism virtue epistemology) will seemingly succumb to the same sort of critique I have already leveled against Plantinga and Sosa. Greco defines knowledge as (roughly) an intellectual achievement where a given agent's epistemic success is "due to" his or her cognitive ability; given that a successful belief can be "due to" a cognitive ability to a greater or lesser extent, however, degrees of luck (as explicated in Chapter 1) can seemingly be mapped onto the degrees in which a belief can be more or less "due to" the said cognitive ability—leading again to the familiar dilemma of forcing the analysis to either place an unreasonably high demand on knowledge (e.g., "successful belief that is *completely* due to a cognitive ability") or fall victim to Gettier counterexample.[147]

While I advocate the trend to abandon the reductive analysis of knowledge in favor of alternative epistemic models by criticizing dominant forms of virtue epistemology, it is nevertheless my conviction that accounting for virtue within epistemology is ultimately meritorious—that there are significant advantages to pursing virtue-theoretic epistemology even if we must renounce the analyzability of knowledge. Having (i) championed contemporary virtue epistemology as one of the most sophisticated and popular species of reductive analysis to date and (ii) argued that, given the proposed diagnosis of the Gettier Problem, we have very good reason to doubt that any reductive analysis (virtue-theoretic or otherwise) can feasibly surmount Gettier counterexamples (as demonstrated in the work of Plantinga and Sosa), I nevertheless want to explore in this chapter the possibility of adapting virtue-theoretic epistemology to suit an alternative epistemic model. In Section 1, I briefly motivate such a project by explicating some of what is so attractive about virtue-theoretic epistemology, giving us reason to preserve it via an alternative model. In Section 2, I elucidate and outline the seminal alternative epistemic model to date, namely, the non-reductive model of knowledge developed by Timothy Williamson in *Knowledge and its Limits* (2000). Third, in Section 3, I explore how virtue might be incorporated within a non-reductive model, generally. Finally, in Section 4, I reconsider the possibility of non-

[146] In "Virtues and Vices of Virtue Epistemology," for example, Greco simply added a "Gettier Problems aside" caveat to his analysis of knowledge (1993, 413).

[147] Greco's *success* (which surely includes truth) *due to ability* analysis of knowledge is similar to Sosa's (2007); as such, we have reason to suspect that all of the problems afflicting Sosa's analysis of knowledge that we detailed in Chapter 4 will apply equally well to Greco's account—including the problem of making sense of the "due to" terminology.

reductive virtue epistemology in light of Williamson's specific non-reductive epistemology and the respective virtue-theoretic concepts of Plantinga and Sosa.

Section 1: Why Virtue Epistemology?

So why should we care about virtue epistemology? If virtue epistemology really is in the dire straits that I have suggested as a result of its commitment to reductive analysis, why not simply cut our losses and continue with our philosophical lives? What is so attractive about virtue-theoretic epistemology in the first place that warrants attempts toward its preservation? There are, no doubt, many things someone might find appealing about virtue epistemology—its ability to withstand arguments for "naturalizing" epistemology, for one,[148] and its unique approach to debates surrounding the structure of knowledge, for another.[149] And any of these reasons might sufficiently motivate a project that sought to explore the possibility of non-reductive virtue-theoretic epistemology. But the two reasons I want to focus on now, the reasons I happen to be most interested in personally, are (i) virtue epistemology's special fit for interesting and exciting intra-disciplinary dialogue, particularly between epistemology and moral philosophy and between epistemology and philosophy of religion, and (ii) the fact that we seem to have virtue-theoretic intuitions across a range of cases. Let's briefly explore each of these in turn.

Everything else being equal, I take it as a given that an epistemology that complements and informs other subdisciplines (and vice versa) is superior to an epistemology that does not. In other words, if the only relevant difference between epistemology A and epistemology B is that A allows for robust dialogue with other subdisciplines (e.g., ethics, philosophy of religion, etc.) and B does not, then A should be championed over B. To put it yet another way, I take it that a synthesized and holistic philosophical worldview is preferable over a disjointed and static philosophical worldview, everything else being equal.[150] One of the chief virtues of virtue-theoretic epistemology, as I see it, is its ability to complement and inform (and be informed by) subdisciplines such as ethics and philosophy

[148] See Quine 1969 and McDowell 1994, 132-134; Sosa 1991, 100-105; Zagzebski 1996, 334-338.

[149] See Sosa 1980.

[150] It is tempting to say that a synthesized and holistic worldview should always to be preferred over one that is not, but this would surely be a mistake. Surely an absurd worldview should not be championed over a plausible worldview, even if the former is more synthesized and holistic than the latter. It is for this reason that I included the "everything else being equal" caveat.

of religion in ways that other epistemologies simply cannot. Virtue epistemology has the potential to produce a more synthesized and holistic philosophical worldview; for this reason, I think there is a strong onus to preserve virtue-theoretic epistemology, if possible.

Regarding moral philosophy: Given the robust normative virtue-theoretic concepts at the heart of virtue epistemology, some rich parallels between virtue epistemology and moral philosophy (especially virtue ethics) can be easily identified. How we incorporate virtue within epistemology, for example, can easily affect how we incorporate virtue within ethics (and vice versa). Virtue epistemology also allows for discussions regarding epistemic value to uniquely parallel discussions regarding moral value. Additionally, explications as to the nature of virtue within epistemology can easily dovetail with explications as to the nature of virtue within ethics. Indeed, with ethics lying at the foundation of knowledge (as some virtue epistemologists have suggested), virtue epistemology can even inform (and be informed by) discussions regarding the scope of the moral.[151]

Regarding philosophy of religion: It is no coincidence that so many epistemologists with interest in philosophy of religion, such as Alvin Plantinga and Linda Zagzebski, have endorsed virtue epistemology. Character and virtue are at the heart of many major religions, allowing for virtue epistemology to uniquely dovetail to an array of projects central to these faiths. Theistic proofs, apologetic methodologies, the epistemic value of sacred texts, knowledge of God and his attributes, etc., can all uniquely inform and be informed by virtue-theoretic epistemology. And what is more, virtue epistemology can even be tied to various projects in biblical hermeneutics—offering explanatory power to passages such as Proverbs 1:7 ("the fear of the Lord is the beginning of knowledge") where there is an apparent link between one's epistemic prowess and his/her moral character.

In addition to fostering a unique relationship between epistemology and other subdisciplines, virtue epistemology tracks our intuitions across a range of cases. And even though we will not be pursuing a reductive analysis of knowledge, this is a feature that any epistemology should nevertheless strive to be sensitive to. Consider two cases first discussed in the introduction of this thesis:

> **René and the Gambler's Fallacy:** René thinks he can beat the roulette tables with a system he has devised. Reasoning according to the Gambler's Fallacy, he believes that numbers which have not come up for long strings are more likely to come up next. However, unlike Descartes' demon victim, our René has a demon helper. Acting as a kind of epistemic guardian, the demon arranges reality so as to

[151] See Zagzebski 1996.

> make the belief come out as true. Given the ever present interventions of the helpful demon, René's belief forming process is highly reliable. But this is because the world is made to conform to René's beliefs, rather than because René's beliefs conform to the world. (Greco 1999, 286)[152]

> **Brain Lesion:** Suppose . . . that S suffers from this sort of disorder and accordingly believes that he suffers from a brain lesion. Add that he has no evidence at all for this belief: no symptoms of which he is aware, no testimony on the part of physicians or other expert witnesses, nothing. (Add if you like, that he has much evidence against it, but then add also that the malfunction induced by the lesion makes it impossible for him to take appropriate account of this evidence.) Then the relevant [process] will certainly be reliable but the resulting belief—that he has a brain lesion—will have little by way of warrant for S. (Plantinga 1993a, 199)[153]

Such cases not only seem to generate intuitions *against* process reliabilism, they seem to generate intuitions *for* virtue epistemology generally. I think our intuitions in such cases not only tell us that reliable process is not enough for knowledge, they tell us that knowledge fundamentally demands the involvement of the relevant agent. What sort of involvement knowledge requires is open to dispute, but the sheer fact that knowledge requires *some* such involvement from the agent is seemingly enough to call for an adaptation in any non-reductive epistemology we might pursue to suit such virtue-theoretic intuitions.

While I have argued in this dissertation that contemporary iterations of virtue epistemology, given their commitment to the reductive analysis of knowledge project, are in dire straits, this does not mean that we should simply give up on virtue-theoretic epistemology. There are plenty of reasons to try to preserve virtue epistemology; and indeed, there are plenty of virtues of virtue epistemology that are particular to it. We have now briefly noted two of them. No doubt, a full account would take great measures to elucidate and explore far more and in greater detail; however, our current aims are humble, and our space is limited.[154] Autobiographically, virtue epistemology's ability to widely engage with robust intra-disciplinary dialogue, and its ability to track our intuitions across a range of cases, are the two merits of virtue epistemology that *I* find most interesting. So these merits in particular I would hope to preserve via an alternative epistemic model. And insofar as these are general virtues of virtue epistemology, I will assume that they can be

[152] Also quoted in Pritchard 2005, 187.

[153] Also quoted in Pritchard 2005, 188.

[154] For more on the virtues (and vices) of virtue epistemology, see Greco 1993.

minimally preserved (to at least some degrees) by simply incorporating a virtue-theoretic concept in a given theory of knowledge.

Section 2: The Williamsonian Model

So we have good reason to be attracted to virtue epistemology; however, as I have argued throughout this dissertation, we also have good reason to be pessimistic regarding the reductive analysis model of knowledge—the assumed model for every prominent virtue epistemology. So if we are to preserve the virtues of virtue epistemology, it is my suggestion that we do so via an alternative epistemic model. In this section, I elucidate and outline the seminal alternative epistemic model to date, namely, the non-reductive model of knowledge developed by Timothy Williamson in *Knowledge and its Limits* (2000).[155] In the sections that follow, we will explore how a non-reductive *virtue epistemology* might be developed along the same lines.

Recall Immanuel Kant's momentous claim in *The Critique of Pure Reason* to a Copernican revolution within metaphysics:

> Thus far it has been assumed that all our cognition must conform to objects. On that presupposition, however, all our attempts to establish something about them *a priori*, by means of concepts through which our cognition would be expanded, have come to nothing. Let us, therefore, try to find out whether we shall not make better progress in the problems of metaphysics if we assume that objects must conform to our cognition. This assumption already agrees better with the demanded possibility of an *a priori* cognition of objects—i.e., a cognition that is to ascertain something about them before they are given to us. The situation here is the same as was that of Copernicus when he first thought of explaining the motions of celestial bodies. Having found it difficult to make progress there when he assumed that the entire host of stars revolved around the spectator, he tried to find out whether he might not be more successful if he had the spectator revolve around and the stars remain at rest. Now we can try something similar in metaphysics, with regard to our intuition of objects. (1998, 638).

We can see Timothy Williamson in *Knowledge and its Limits* (2000) as initiating a similar revolution in epistemology. After all, we have "thus far . . . assumed" that knowledge should be defined in terms of a reductive analysis, but as we are now seeing such a project has seemingly "come to nothing" (or at least very little). As such, "having found it difficult to make progress," perhaps, as Williamson will suggest, it is time we do away with that

[155] While similar non-reductive accounts of knowledge can be found in the literature (cf., Michael Polanyi's *Personal Knowledge*, 1974; John McDowell's "Criteria, Defeasibility, and Knowledge," 1988), Williamson's is unquestionably *the* pivotal account within contemporary analytic epistemology.

assumption and "[try] to find out whether [we] might not be more successful" exploring other options.[156]

As was explained in the Introduction, the standard practice in contemporary epistemology is to provide an analysis of knowledge in terms of truth, belief, and warrant. Given (i) that truth and belief are necessary for knowledge and (ii) the widespread assumption that belief is conceptually prior to knowledge, the contemporary project of defining knowledge has for nearly fifty years taken the shape of a reductive analysis of trying to explicate what, with true belief, is sufficient for knowledge. Now, we may agree that truth and belief are indeed necessary for knowledge, but why should we assume that belief is conceptually prior to knowledge? Such an assumption is critical for the reductive analysis project, yet, as Williamson points out, it is often just "taken for granted" and "rarely supported by argument" (Williamson 2000, 2). To be sure, we might think the simple fact that "knowledge entails belief but not vice versa" is enough to get the aforementioned assumption off the ground, but this would be a mistake (Williamson 2000, 3). As Williamson explains:

> Given that knowledge entails belief, it is trivial that one knows *p* if and only if (i) one believes *p*; (ii) *p* is true; and (iii) if one believes *p* and *p* is true, then one knows *p*. But that equivalence is useless for establishing that belief is conceptually prior to knowledge, for it is circular: "know" occurs in (iii). The received idea is that we can conceptualize the factors whose conjunction with belief is necessary and sufficient for knowledge independently of knowledge; we can think of the former without already thinking of the latter, even implicitly. But the argument does not show that such independent conceptualization is possible, for a necessary but insufficient condition need not be a conjunct of a non-circular necessary and sufficient condition. Although being coloured is a necessary but insufficient condition for being red, we cannot state a necessary and sufficient condition for being red by conjoining being coloured with other properties specified without reference to red. Neither the equation "Red = coloured + X" nor the equation "Knowledge = true belief + X" need have a non-circular solution . . . belief can be a necessary but insufficient condition of knowledge even if we do not implicitly conceptualize knowledge as the conjunction of belief with that which must be added to belief to yield knowledge." (Williamson 2000, 3)

As such, the brute fact that knowledge seems to entail belief but not vice versa does not provide suitable ground on which to assume that belief is somehow conceptually prior to

[156] As Williamson notes, "[t]he programme of analysis had its origin in great philosophical visions [e.g., logical positivism]. . . . Now the philosophical visions which gave it a point are no longer serious options. Yet philosophers continue to pursue the programme long after the original motivation has gone. Correct deep analyses would doubtless still be interesting if they existed; what has gone is the reason to believe that they do exist" (2000, 32).

knowledge. Besides, the fact that knowledge seems to entail belief could just as easily support a hypothesis that sees knowledge as conceptually prior to belief. As Williamson explains, "[i]f believing *p* is conceptualized as being in a state sufficiently like knowing *p* 'from the inside' in the relevant aspects, then belief is necessary for knowledge since knowing *p* is sufficiently like itself in every respect, even though knowledge is conceptually prior to belief" (Williamson 2000, 3). How else, then, might someone defend the widespread assumption that belief is conceptually prior to knowledge? To be sure, if we *do not* have any good reason to suppose that belief is conceptually prior to knowledge, and if we *do*, as I've tried to argue in this thesis, have a good reason to be pessimistic regarding any viable success against the Gettier Problem from a reductive analysis, perhaps this is simply all the more reason to give up such a project.[157]

Now someone might agree that the reductive analysis project has (at best) a lackluster history but nevertheless advocate the assumption that belief is conceptually prior to knowledge by holding out hope that a more viable analysis (after nearly 50 years) is finally on the horizon—after all, Gettier problems aside, many analyses have been at least decent approximations of knowledge. But the mere fact that *p* is a close approximation of *q* in no way entails that *q* can be analyzed in terms of *p* or that *q* is somehow conceptually prior to *p* or that *q* cannot be understood divorced from *p*. Consider Williamson's example of parenthood and ancestry. As Williamson points out, "*x* is an ancestor of *y* and *x* is not an ancestor of an ancestor of *y*" is a very close approximation of "x is a parent of y," but not perfectly so (Williamson 2000, 4). There are insurmountable counterexamples involving rare cases of incest. If a daughter bears a son by her father, that father is both a parent of

[157] The merit of a Williamsonian model of knowledge is not strictly contingent on every analysis somehow failing. While we may have very good reason to doubt the viability of the reductive analysis project, Williamson's epistemology can still get off the ground even if a successful analysis of knowledge in terms of necessary and jointly sufficient conditions *is* found. While "[t]he present account does not strictly entail that no analysis of the traditional kind provides correct necessary and sufficient conditions for knowing . . . ," it does call into question the reasoning for pursuing such an analysis in the first place (Williamson 2000, 30). To quote Williamson, "[e]ven if some sufficiently complex analysis never succumbed to [Gettier] counterexamples, that would not entail the identity of the analyzing concept with the concept *knows*. Indeed, the equation of the concepts might well lead to more puzzlement rather than less. For knowing matters; the difference between knowing and not knowing is very important to us. Even unsophisticated curiosity is a desire to *know*. This importance would be hard to understand if the concept knows were the more or less ad hoc sprawl that analyses have had to become; why should we care so much about *that*?" (2000, 30–31).

the aforementioned son while also being his ancestor's ancestor. As Williamson explains, "[s]ince the father and the mother of his daughter are symmetrically related to the daughter and son in terms of ancestry but not in terms of parenthood, parenthood cannot be defined in terms of ancestry without extra conceptual resources" (Williamson 2000, 4). Even though parenthood can be closely approximated via ancestry, this in no way entails that "parent" can be analyzed in terms of ancestry or that the latter is in any way critical to our understanding of the former or that ancestry is somehow conceptually prior to parenthood. According to Williamson, such examples show us that "[t]he possibility of approximating knowledge in terms of belief and other concepts is not good evidence for the conceptual priority of belief over knowledge" (Williamson 2000, 4). It's not clear, then, that we have any robust theoretical reason to accept the widespread assumption that belief is somehow conceptually prior to knowledge, that knowledge should be analyzed in terms of true belief plus whatever makes true belief knowledge. Besides, as Williamson notes, generally "one would not expect the concept *knows* to have a non-trivial analysis in somehow more basic terms. . . . 'Bachelor' is a peculiarity not a prototype" (2000, 31).[158] If we can dispel such an assumption, then seemingly we have all the more reason to explore alternative epistemic models, specifically those that see knowledge as, at the very least, on conceptually equal footing with belief.

Williamson's is a knowledge-first epistemology. Given (i) that solving the Gettier Problem via reductive analysis seems like a Sisyphean endeavor and (ii) that we have no robust theoretical reason to assume that such an approach to knowledge is, in fact, appropriate, Williamson proposes an epistemic model that takes knowledge to be an unanalyzable primitive, where knowledge is conceptually prior to belief.[159] Instead of trying to build up (via warrant) to knowledge from true belief, Williamson starts with knowledge in hopes of exploring notions of belief, justification, and evidence in turn. Knowledge, for

[158] For example, "[a]ttempts to analyse the concepts of *means* and *causes* . . . have been no more successful than attempts to analyse the concept *knows*, succumbing to the same pattern of counterexamples and epicycles. The analysing concept does not merely fail to be the same as the concept to be analysed; it fails even to provide a necessary and sufficient condition for the latter" (Williamson 2000, 31). "The pursuit of analyses," it seems, "is a degenerating research programme" (Williamson 2000, 31).

[159] Williamson does not give us any reason for thinking that knowledge is not simply equi-primitive with related concepts like belief and justification, for thinking that knowledge is not on conceptually equal footing with belief and justification. Indeed, most of Williamson's arguments in favor of his primitive-knows hypothesis would apply equally well to an equi-primitive-knows hypothesis.

Williamson, is "a state of mind"; it will be our task in the rest of this section to unpack in broad strokes a bit of what that means.

While Williamson gives relatively little explanation as to what exactly a mental state is, we are presumably meant to have an intuitive enough grasp of such terminology so as to proceed without hindrance.[160] After all, it seems as though we can easily identify paradigmatic examples of mental states ("love, hate, pleasure, pain . . . believing that it is so, conceiving that it is so, hoping or fearing that it is so, wondering whether it is so, intending or desiring it to be so"), so all we need to do is extrapolate from these to get a rough and ready understanding of how Williamson is wanting to think of knowledge (Williamson 2000, 21). An initial hurdle to doing so, however, is that, unlike paradigmatic examples of mental states, knowledge is factive; where paradigmatic examples of mental states such as "believing that *p*" and "hoping that *p*" remain neutral as to whether or not *p* is in fact true; "knowing that *p*" surely entails that "*p*."[161] To be sure, we could of course retreat somewhat and say that knowledge is only a mental state insofar as it necessarily involves believing, a paradigmatic mental state; such a move, however, would seemingly be a retreat into the reductive analysis project, defining knowledge as having a mental state of a certain sort that is also true. As Williamson explains:

> Someone might expect knowing to be a state of mind simply on the grounds that knowing *p* involves the paradigmatic mental state of believing *p*. If those grounds were adequate, the claim that knowing is a state of mind would be banal. However, those grounds imply only that there is a mental state being in which is *necessary* for knowing *p*. By contrast the claim that knowing is a state of mind is to be understood as the claim that there is a mental state being in which is necessary *and sufficient* for knowing *p*. (Williamson 2000, 21 - emphasis Williamson's)

When Williamson says that knowledge is a mental state, he means that knowledge is *merely* a mental state (Williamson 2000, 21). What, then, are we to make of the non-factive/factive disparity between paradigmatic mental states and knowledge? Given that knowledge

[160] A similar reading of Williamson can be found in Cassim 2009, 13, in the edited volume by Greenough and Pritchard (2009).

[161] Of course, as is typical in philosophy, there is not universal agreement here. While the thought that *knowledge requires truth* is very nearly a ubiquitous platitude, there are those very few who would digress (cf., Rosenberg 2003). Regardless, such views are so far in the minority, we will not give them further consideration here.

requires truth, a non-mental component, how can knowledge be *merely* a state of mind? Williamson says the following:

> Our initial presumption should be that knowing is a mental state. Prior to philosophical theory-building, we learn the concept of the mental by examples. Our paradigms should include propositional attitudes such as believing and desiring, if our conception of the mental is not to be radically impoverished. But factive attitudes have so many similarities to the non-factive attitudes that we should expect them to constitute mental states too; we expect a concept to apply to whatever sufficiently resembles its paradigms. It would be strange if there were a mental state of fearing but no mental state of regretting, or a mental state of imagining but no mental state of remembering. Indeed, it is not clear that there are any pretheoretic grounds for omitting factive attitudes from the list of *paradigmatic* mental states. That the mental includes knowing and other factive attitudes is built into the natural understanding of the procedure by which the concept of the mental is acquired. Of course, that does not exclude the subsequent discovery of theoretical reasons for drawing the line between the mental and the non-mental somewhere else. But the theory behind those reasons had better be a good one. (Williamson 2000, 22)

Williamson, then, has effectively shifted the explanatory onus onto anyone who would like to deny that factivity (regarding external propositions) precludes the mental. When we ask "Why should we think that knowledge, being factive, is purely mental?" Williamson simply seems to reply "Why not?"

To see one reason why we might think the factivity knowledge requires precludes its being mental, contrast knowledge with believing truly. *Believing truly* (when the proposition in question is in the external environment) is not a mental state; while *believing* is a mental state, believing *truly* typically extends beyond the boundaries of the mental. And if (i) everything within the state of believing that *p* is contained within the state of believing truly that *p* and (ii) everything within believing truly that *p* is within the state of knowing that *p*, then it seems as though once we assume that knowledge is indeed mental we seem forced to oddly sandwich a non-mental state (believing truly that *p*) between two mental states (believing that *p* and knowing that *p*). But, to quote Williamson:

> That something sandwiched between two mental states need not itself be a mental state is not as paradoxical as it may sound. Consider an analogy: the notion of a geometric property. For these purposes, we can understand geometrical properties to be properties possessed by particulars in physical space. Let π_1 be the property of being an equilateral triangle, π_2 the property of being a triangle whose sides are indiscriminable in length to the naked human eye, and π_3 the property of being a triangle. Necessarily, everything that has π_1 has π_2, because lines of the same length cannot be discriminated in length; necessarily, everything that has π_2 has π_3. Nevertheless, although π_1 and π_3 are geometrical properties, π_2 is not a geometrical property, because it varies with variations in human eyesight. Something

> sandwiched between two geometrical properties need not itself be a geometrical property. Similarly, there is no structural reason why something sandwiched between two mental states should itself be a mental state. . . If S is a mental state and C a non-mental condition, there need be no mental state S* such that, necessarily, one is in S* if and only if one is in S and C obtains. The non-existence of such an S* is quite consistent with the existence of a mental state S** such that, necessarily, one is in S** only if (but not: if) one is in S and C is met. A mental state can guarantee that conjunction only by guaranteeing more than that conjunction. (Williamson 2000, 27–28)

Just because knowledge seemingly requires both truth (a non-mental condition) and belief (a mental state) does not mean there needs to be a mental state *believing truly*. There is no inconsistency in taking *believing truly* to be non-mental while simultaneously taking "knowledge that *p*" to be a mental state that includes both "belief that *p*" and "*p*" (the truth that *p*). While it may initially seem strange that *believing truly* is, as it were, sandwiched between the mental states of *believing* and *knowing*, this in itself is no reason to doubt that knowledge is a mental state—the aforementioned strangeness seemingly dissipates upon reflection.[162]

There are, to be sure, a number of factive mental states: "seeing that *p*," "remembering that *p*," "hearing that *p*," etc.—such that for any given factive mental state, **Φ**, if "S **Φ** that *p*" then "*p*." Knowledge, for Williamson, is the most general factive mental state; "seeing that *p*," "remembering that *p*," "hearing that *p*," etc., are all ways of "knowing that *p*"—if "S **Φ** that *p*" then "S knows that *p*." Consider the following explanation:

> To picture the proposal, compare the state of knowing with the property of being coloured, the colour property which something has if it has any colour property at all. If something is coloured, then it has a more specific colour property; it is red or green or. . . . Although that specific colour may happen to lack a name in our language, we could always introduce such a name, perhaps pointing to the thing as a paradigm. We may say that being coloured is being red or green or . . ., if the list is understood as open-ended, and the concept *is coloured* is not identified with the disjunctive concept. One can grasp the concept *is coloured* without grasping the concept *is green*, therefore without grasping the disjunctive concept. Similarly, if one knows that A, then there is a specific way in which one knows; one can see or

[162] Someone might respond to this strangeness by trying to find some liberal sense in which *believing truly* might be construed as a mental state; however, as Williamson points out, even if there is some such liberal understanding of "mental state," "there is also a more restrictive but still reasonable sense in which believing truly is not a mental state but the combination of a mental state with a non-mental condition" (Williamson 2000, 28). When Williamson defines knowledge as a state of mind he means that "knowing is a mental state in *every* reasonable sense of that term: there is no more restrictive but still reasonable sense of "mental" in which knowing can be factored, like believing truly, into a combination of mental states with non-mental conditions" (Williamson 2000, 28).

> remember or . . . that A. Although that specific way may happen to lack a name in our language, we could always introduce such a name, perhaps pointing to the case as a paradigm. We may say that knowing that A is seeing or remembering or . . . that A, if the list is understood as open-ended, and the concept *knows* is not identified with the disjunctive concept. One can grasp the concept *knows* without grasping the concept *sees*, therefore without grasping the disjunctive concept. (2000, 34)

Lacking any robust theoretical reason to deny that mental states can be factive, Williamson not only describes knowledge as a factive state of mind, he describes it as *the most general* factive state of mind. In other words, "knowing," for Williamson, "is the most general factive stative attitude, which one has to a proposition if one has any factive stative attitude to it at all" (2000, 34).

Of course, for any given factive mental state there is going to be a difference between "**Φ** that *p*" and "**Φ** a situation that *p*"; as Williamson points out, "only the former requires the [agent] to grasp the proposition [*p*]" (Williamson 2000, 38). As such, using Williamson's example, there is going to be a deference between seeing people play chess and seeing a situation where people are playing chess without knowing what they are doing. "A normal observer in normal conditions who has no concept of chess can see a situation in which Olga is playing chess," says Williamson, "by looking in the right direction, but cannot see *that* Olga is playing chess because he does not know what he sees to be a situation in which Olga is playing chess" (Williamson 2000, 38 - emphasis Williamson's).

To be sure, Williamson's understanding of knowledge as a factive mental state needs to be somewhat environmentally sensitive—it needs to track the truth of a given proposition such that "**Φ** that *p*" requires the truth that *p* to be secure. For example, "seeing that *p*" needs to be understood in such a way that the protagonist in Fake Barn-type cases does not, in a real sense, *see* that "that's a barn"—*seeing*, so understood, requires a more stable environment. Thankfully, however, such amends can easily be made. We could, with Alan Millar, simply stipulate that knowledge (or factive mental states) requires a given agent to possess *a way of telling*—given the proximity of fake barns, the protagonist in Fake Barns simply does not possess the ability to tell whether or not a given object is a barn.[163] Or we could, with Williamson himself, stipulate that knowledge requires safety, such that a given agent can "**Φ** that *p*" only when *p* is safe.[164] In any case, once some such caveat has been

[163] See Millar 2008, Millar 2009, and Pritchard, Millar, and Haddock 2010, chaps. 5-7.

[164] See Williamson 2000, chaps. 5-7. Strangely, Williamson does not address the Gettier Problem in any detail; indeed, he only invokes the problem as indicative of his diagnosis of the reductive analysis project. Seemingly,

made, we need no longer worry about Gettier-like counterexamples where the truth of the relevant proposition is not sufficiently secure.

But why does not this sort of solution fail for precisely the same reasons all of the previous accounts (virtue theoretic and otherwise) fail? Why, for example, is Williamson's use of safety within a *non-reductive account* of knowledge any more viable than anyone else's use of safety within a *reductive account* of knowledge? In both instances, safety comes in degrees. And accordingly, degrees of luck could be mapped onto the degrees of Williamson's account of safety just as easily as degrees of luck can be mapped onto any traditional/reductive account of safety. So why think non-reductive accounts are going to be any more successful? How does Williamson avoid the dilemma that has plagued so many other accounts? In other words, why is not Williamson's account forced either to place an unreasonably high demand on knowledge (e.g., *perfect* safety) or fall victim to Gettier-like counterexample?

By way of an answer, we should first note that non-reductive accounts are simply not nearly as demanding. As we saw in Chapter 1, in order for any reductive analysis of knowledge (approximately) in terms of *warranted true belief* to surmount the Gettier Problem, whatever is doing the warranting must be either (i) strongly infallibilistic (i.e., bearing a Inviolable WT-relation) such that it necessarily entails the truth of the belief in question or (ii) assumed to be completely divorced from the truth (e.g., bearing a Divorced WT-relation). These are high demands, asking an epistemic price very few are willing to pay. And given (i) that safety theorists want to assume that warrant bears *some* relationship to truth and (ii) that safety is meant to be a key component of that warrant, safety must infallibilistically guarantee the truth of the belief in question—a requirement that demands *perfect* safety and leads to radical skepticism. Not so for non-reductive accounts like Williamson's. Having given up the project of trying to provide a reductive analysis of knowledge in terms of necessary and jointly sufficient conditions, the constraints on such non-reductive accounts have been substantially reduced. Williamson is only interested in explicating the necessary conditions for knowledge; sufficiency conditions are not his aim. And as such, what drives safety conditions within reductive analysis to such extremes is simply missing in the Williamsonian model.

---

then, a fuller account of Williamson's epistemology would give the Gettier Problem more extensive treatment, diagnosing when and why such problems occur (as we have done in this thesis).

But what is more, if we are to see why Williamson's non-reductive account of knowledge does not fall into the same dilemma that has plagued all of the reductive accounts considered in this thesis (virtue theoretic and otherwise), we need to be clear about the sort of cases that threaten it. Consider the following two cases, which mirror the features of popular Gettier counterexamples in the literature considered earlier:

> **Sheep 1:** Bill lives in Scotland, where sheep are fairly common. He looks out his window and sees what clearly looks like a sheep in a field. As such, Bill believes "there is a sheep in that field." Unfortunately, however, what Bill saw was a sheep façade, a wolf in sheep's clothing. As it happens, however, Bill's belief is nevertheless true; as luck would have it, there was a sheep in that field; it was just hidden behind a gorse bush.[165]

> **Sheep 2:** Bob lives in Scotland, where sheep are fairly common. He looks out his window and sees what clearly looks like a field with hundreds of sheep in it. Looking at one in particular, Bob believes "that's a sheep." Unbeknownst to Bob, however, farmers in his area have recently constructed hundreds of sheep façades, such that imitations outnumber real sheep 100 to 1. Nevertheless, the object Bob identified as a sheep was one of the real sheep. Bob's belief was true, but only luckily so.

Williamson's non-reductive account can easily handle cases like Sheep 1 without any reference to safety. Straightforwardly, Bill does not really *see* (as a factive mental stative attitude) that there is a sheep in the field, so he does not *know* that there is a sheep in the field. However, if we were to try to invoke safety within a *reductive* account to handle Sheep 1, the dilemma between Gettier-vulnerability and radical skepticism would immediately emerge. Simply put, features could be built into Sheep 1 such that it would be extremely likely that there would be a sheep hiding behind the gorse when Bill looks out his window so as to minimize the relevant epistemic luck at issue and maximize the safety of Bill's belief while nevertheless preserving our intuition that Bill lacks knowledge. In other words, Sheep 1 can be strengthened such that the only way a safety condition (within a reductive analysis) can avoid it is to demand *perfect* safety—a condition few of even our most secure beliefs can meet. It is in cases like Sheep 2, on the other hand, where Williamson might need to invoke an environmental condition like safety in order to get the right result. But unlike Sheep 1, there does not seem to be any obvious way to strengthen Sheep 2 without losing our intuition that Bob clearly lacks knowledge. We could, of course, reduce the number of sheep façades (say to a 1:100 ratio with real sheep) so as to minimize the relevant epistemic luck at issue and to maximize the safety of Bob's belief, but then it is

[165] This case is a close paraphrase of Chisholm's classic case. See Chisholm 1977, 105.

simply no longer apparent that Bob genuinely lacks knowledge. So perhaps Williamson's non-reductive account of knowledge does not need to make extreme demands on knowledge in order to avoid counterexamples, because the potential counterexamples don't demand it. Conditions like safety can be more viably invoked within a Williamsonian model simply because the counterexamples that threaten it do not seem to yield the familiar regression of strengthened cases.[166]

Williamson's positive account of knowledge can be summarized in the following three points: for any given factive mental state, **Φ**:

1) "S **Φ** that *p*" entails that "*p*."
2) "Know" is (merely) a factive mental state.
3) "S **Φ** that *p*" entails "S knows that *p*."[167]

While this account of knowledge may seem relatively thin compared to the elaborate and nuanced analysis of Plantinga, say, perhaps (given the dizzying array of analyses with diverse and multifaceted conceptions of warrant) such a simple account of knowledge is a virtue. As Williamson explains, in his account "the importance of knowing to us becomes as intelligible as the importance of truth. Factive mental states are important to us as states whose essence include a matching between mind and world, and knowing is important to us as the most general factive stative attitude" (2000, 39–40). "This importance," to be sure, "would be hard to understand if the concept *knows* were the more or less ad hoc sprawl that analyses have had to become" (Williamson 2000, 31).

Section 3: Non-Reductive Virtue Epistemology

In recent years, virtue epistemology has enjoyed enormous popularity—offering new and exciting solutions to a variety of problems that have plagued other accounts of knowledge. Even so, we have seen that prominent renditions of virtue epistemology, given their commitment to the reductive analysis project, simply cannot viably surmount the Gettier Problem. But thankfully, there is yet hope. Timothy Williamson's penetrating work offers a new and exciting alternative to the standard analysis of knowledge—an alternative that

[166] Williamson could, of course, always bite the bullet here and say that the protagonists in cases like Fake Barns and Sheep 2, contra our typical intuitions, genuinely know—subsequently disarming them. Of course, such a sacrifice is simply unnecessary if the explanations above are feasible, as I think they are.

[167] See Williamson 2000, 29, 34.

might suit a virtue-theoretic adaptation. Having elucidated the seminal alternative to the reductive analysis of knowledge, Williamson's non-reductive model, we are now in a better position to see how virtue-theoretic concepts might be inserted within such a framework—to see how virtue epistemology might be preserved via a non-reductive model. By way of prolegomena, it is our task in this section to map out the various ways this might be done. Our work in this section will be quite general—leaving the relevant epistemic concepts abstract and uncommitted to any specific account of knowledge (e.g., knowledge as mental states) or any specific account of virtue (e.g., proper functioning cognitive faculties, cognitive competencies, etc.). We will return to such detail in Section 4.

But where to start? As we turn to consider how virtue might be incorporated within a non-reductive model of knowledge like Williamson's, there are two axes worth tracking: (i) whether or not virtue and knowledge are primitive and (ii) whether knowledge entails or is entailed by virtue. Across these axes there are 16 possible ways virtue might relate to knowledge.[168] And given that "K" stands for knowledge, that "V" stands for virtue, and that arrows and underscores signify entailments and conceptual priority, respectively, these 16 possibilities can be tabulated in the following way:

Table 5.1

		Conceptual Priority				
		K & V	K & V	K & V	K & V	
Entailments	K→V	K→V	K→V	K→V	K→V	A
	K←→V	K←→V	K←→V	K←→V	K←→V	B
	V→K	V→K	V→K	V→K	V→K	C
	K & V	K & V	K & V	K & V	K & V	D
		1	2	3	4	

Of course, most of these options can be dismissed out of hand—being either incompatible with non-reductive virtue epistemology or independently unfeasible. First of all, given that our goal is to develop a *virtue epistemology* (an epistemology that sees virtue as intimately tied

[168] Given that our goal here is to explore how virtue might be incorporated into a non-reductive model, how virtue might relate to a non-reductive account of knowledge, we will satisfy ourselves with merely exploring the various ways knowledge might entail (and/or be entailed by) virtue, ways virtue might be a necessary condition for knowledge (and vice versa), and the conceptual ordering of the two. The complexities that arise from considering other possible epistemic concepts and facets (i.e., safety, belief, etc.) will be put aside.

to knowledge), row D can be discarded. And, given that the metaphysical priority established by entailment minimally requires *some* conceptual priority, we can discard (the rest of) column 4 as dubious, if not simply unattainable.[169] On our map of the possible relationships virtue might bear to knowledge (and vice versa), B3 designates the possibility of a reductive analysis of knowledge in terms of necessary and sufficient conditions—a project that I have argued we should abandon. B1, interestingly, designates the possibility of a reductive analysis *of virtue in terms of knowledge*—a project that, lacking independent reason for thinking that virtue yields such an analysis, we presumably should be just as weary of. Under the reasonable assumption that B1 will lead to Gettier-like counterexamples about *virtue* ("Vettier" counterexamples) and similar problems parallel to those afflicting the reductive analysis of knowledge project, B1, too, can be discarded. Finally, I think we can also do away with C1, C2, and C3. In all three instances, virtue is merely sufficient for knowledge, not necessary; knowledge may well be necessary for virtue on these models, but nothing about virtue is necessary for knowledge. As such, whatever the virtues of epistemic virtue, C1, C2, and C3 do not allow them to be brought to bear on knowledge.[170]

If (i) we are giving up on the reductive analysis of knowledge and (ii) we still want to pursue virtue epistemology, what are our alternatives? The remaining four possibilities (A1, A2, A3, and B2) give us a map. And we can classify these remaining possibilities into three general strategies for incorporating virtue into a non-reductive model of knowledge. First, there is what I will call the *knowledge plus virtue* strategy—the strategy of adding a virtue-theoretic condition to a non-reductive epistemic model as a necessary condition on a primitive conception of knowledge (A1, A2). Secondly, there is what I will call the *knowledge as virtue* strategy—the strategy of explicating knowledge itself in terms of a given account of epistemic virtue (B2). Third and finally, there is what I will call the *knowledge within virtue* strategy—the strategy of adding a conceptually primitive account of virtue as a necessary

[169] For any given entailment there is always a question regarding the order of determination, of conceptual priority (this is clearest in the case of biconditionals). So when we consider whether knowledge entails virtue and vice versa, we will be interested in who is taking the conceptual lead, so to speak. Is knowledge conceptually more primitive than virtue (column 1)? Is virtue conceptually more primitive than knowledge (column 3)? Or are they conceptually co-primitive (column 2)? There does not seem to be space for a fourth option.

[170] What is more, C3 designates the possibility of a reductive analysis of knowledge *solely in terms of sufficiency conditions* and will presumably be no more viable (and probably less so) than regular reductive analysis.

condition on a conceptually subsequent account of knowledge, the strategy of explicating virtue as the richest epistemic state (A3).

Our goal in this section is to briefly explore each of these general strategies. Working from the most modest proposals (knowledge plus virtue) to the most extreme (knowledge within virtue), we will be guided by an interest for exploring just how central virtue might be within a non-reductive model of knowledge. I hope to show how all three strategies might allow for the incorporation of virtue-theoretic concepts into a non-reductive framework so as to preserve the merits of virtue epistemology—in particular those merits elucidated in Section 1—without committing to the analyzability of knowledge. Once this is done, we will turn, in Section 4, to consider each strategy in light of Williamson's particular non-reductive model of knowledge and the respective virtue-theoretic concepts of Plantinga and Sosa.

Section 3.1: Knowledge plus Virtue

One straightforward way to incorporate a virtue-theoretic condition into a non-reductive model is to simply identify it as a necessary condition for a primitive account of knowledge. And this is what we see happening in A1 and A2 in Table 5.1. Even as primitive, knowledge can have necessary conditions such as belief, truth, and perhaps even safety. And just as safety might be key to explaining and tracking knowledge ascriptions across cases like Fake Barns and Sheep 2 (as we saw earlier), perhaps a virtue-theoretic condition is needed to explain and track knowledge ascriptions across cases like René and the Gambler's Fallacy and Brain Lesion. Perhaps we can develop a non-reductive virtue epistemology by simply amending a non-reductive model with the caveat that a particular virtue condition (of our favorite stripe) is necessary for a given primitive account of knowledge. There are, to be sure, a couple ways this might be done, contingent on whether or not the said virtue condition is conceptually subsequent to or co-primitive with knowledge.

Figure 5.1

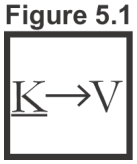

The weakest form of non-reductive virtue epistemology—that is, the weakest incorporation of virtue within a non-reductive model—is the one that holds a conceptually primitive account of knowledge as entailing a conceptually subsequent account of virtue but not vice versa (Figure 5.1, or A1 on Table 5.1). On such an account, knowledge is sufficient for virtue and virtue is necessary for knowledge; but what is more, virtue, being conceptually subsequent to knowledge, is

understood in terms of knowledge. In other words, according to this account, virtue cannot be properly comprehended apart from our (independent) understanding of knowledge.

Figure 5.2

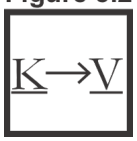

Another form of virtue epistemology can be incorporated by simply making virtue *co-primitive* with knowledge, while keeping everything else the same (Figure 5.2, or A2 on Table 5.1). On such an account, knowledge still entails virtue and not vice versa, but virtue is no longer conceptually dependent on knowledge.[171] This is, to be sure, a stronger use of virtue within a non-reductive model—allowing for an independent account of virtue alongside an independent account of knowledge. That said, however, this ultimately makes the subsequent *virtue epistemology* seem somewhat weaker—as independent concepts, virtue no longer seems as closely tied to knowledge.

The knowledge plus virtue strategy allows us to create a virtue-theoretic epistemology that tracks our intuitions across a wide range of cases, and by adopting a virtue-theoretic condition as a necessary condition on a primitive account of knowledge we will seemingly preserve (at least some) important avenues for dialogue between epistemology and moral philosophy and between epistemology and philosophy of religion. But is this general strategy ultimately going to be satisfying for the lover of virtue epistemology? While the knowledge plus virtue proposals do indeed preserve the relevant virtue-theoretic concepts, the virtue epistemology they produce may seem a bit limp wristed, spiritless, or shallow. The fact that the virtue-theoretic portion of the epistemology is reduced to a mere addendum to, at best, a conceptually co-primitive account of knowledge, marks its departure from the spirit of the virtue epistemologies considered thus far. Yes, such a strategy can account for our relevant virtue-theoretic intuitions across the given cases. And yes, such a strategy can preserve some avenues for dialogue with moral philosophy and philosophy of religion. And this may be all we are looking to do. But there might be a price to pay. We may lose the heart and soul of virtue epistemology, namely the centrality of virtue for epistemology. For many virtue epistemologists, their notion of virtue is meant to be absolutely fundamental for epistemology, and it is hard to see how this could still be the

[171] To be sure, someone like Williamson may wonder whether necessary conditions like virtue can really be understood apart from our concept of knowledge. After all, this seems to be Williamson's take on other necessary conditions like belief. Belief for Williamson is sometimes seen as nothing but a failed attempt at knowledge. See Williamson 2000, 44-48.

case if the virtue condition is ultimately a caveat to a dominant account of knowledge. And indeed, if virtue is not central to epistemology, then perhaps the *richest* dialogue with moral philosophy and philosophy of religion may no longer be open to us.[172]

Section 3.2: Knowledge as Virtue

Figure 5.3

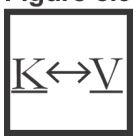

For the lover of robust virtue epistemology, the more radical proposal of *knowledge as virtue* may satisfy (Figure 5.3 or B2 on Table 5.1). The basic concept is straightforward. Simply define knowledge in terms of epistemic virtue. Instead of defining knowledge separate from virtue and subsequently adding a virtue-theoretic condition as a caveat (as we did above), define knowledge itself in terms of your favorite virtue-theoretic concept by explicating a primitive account of knowledge as both necessary and sufficient for a co-primitive account of epistemic virtue. Where the previous two accounts might be called *doubly non-reductive*—both abandoning sufficiency conditions for knowledge and holding knowledge as conceptually primitive—this account is what we might call *singularly non-reductive*. While knowledge is still held primitive on this strategy, a co-primitive conception of virtue is taken to be both necessary and *sufficient* for knowledge. The subsequent account is still non-reductive; knowledge, as a primitive, is simply defined in terms of epistemic virtue.[173]

But what is an *epistemic virtue*? Without advocating any particular virtue-theoretic concept, we will not be able to go into any great detail in answering such a question; however, a general sketch can certainly be provided. An act of virtue, roughly speaking, is

[172] For example, some virtue epistemologists have done some extremely interesting work integrating epistemology and ethics; some even going so far as to say ethics is foundational for knowledge (cf. Zagzebski's *Virtues of the Mind*). If we were to reduce virtue epistemology to a Williamsonian caveat, such projects no longer seem possible.

[173] Given the simplicity and the elegance of Sosa's account of knowledge as *apt belief*, one may worry that Sosa's virtue epistemology is an instantiation of something very much like the knowledge as virtue strategy. This would, however, be a mistake. While Sosa is far from explicit, Sosa's virtue epistemology is well and truly an analysis—a description Sosa has not resisted in personal conversations (Church 2011). Seemingly, S knows that *p*, according to Sosa, if and only if (i) S believes that *p* and (ii) S's belief that *p* is apt, that is "true because competent." For Sosa, (i) and (ii) are necessary and jointly sufficient for knowledge, and *knowledge*, here, is presumably understood in terms of *belief* and *aptness* and not vice versa. Besides, if Sosa *was* truly breaking fellowship with the traditional reductive project, one would hope such a revolutionary move would have been made plain. See Sosa 2009b; Williamson 2009b. That said, however, I will soon argue that Sosa's virtue epistemology is indeed easily adaptable so as to suit a Williamsonian model.

an act that is virtuously motivated and successful in reaching its virtuous telos. So if I am to commit an act of courage by saving a drowning child from a turbulent river, say, then I need to not only be courageously motivated but also successful in reaching my goal. If I save the child out of greed, say (hoping the child's rich parents will give me a monetary reward), then I have not committed an act of courage; I may have reached a virtuous telos, but I was not virtuously motivated. Conversely, if I courageously try to save the child but fail, I still have not committed an act of courage because the telos, saving the child, was unmet; one may act courageously without committing an act of courage. As such, an act of *epistemic* virtue, again roughly speaking, is an act that has epistemically virtuous motivation and is successful in reaching its epistemically virtuous telos (which at the very least includes truth). For example, if I believe that there is a thief in the living room as a result of drug-induced paranoia, then even if such a belief is true (the epistemically virtuous telos) it is not an act of epistemic virtue; drug-induced paranoia is presumably not epistemically virtuous motivation. Conversely, I may falsely believe that there is a thief in the living room after carefully reading reports that thieves were in my area and (having accidentally left the radio on) hearing what happens to sound very much like a thief in my living room, yet not commit an act of epistemic virtue; I may have had epistemically virtuous motives, but since I did not arrive at the truth I did not commit an act of epistemic virtue.[174]

Unlike the *knowledge plus virtue* strategy, the *knowledge as virtue* strategy allows for robust non-reductive virtue epistemologies—virtue epistemologies that not only track our intuitions across the relevant cases but also allowing for richer intra-disciplinary dialogue while capturing the heart of virtue epistemology by placing virtue at the heart of epistemology. That said, however, being more radical, the knowledge as virtue proposal is decidedly more ambitious. Unlike the knowledge plus virtue proposal that simply attaches a virtue-theoretic condition onto a self-standing account of knowledge, the knowledge as virtue proposal forces us to carefully reconsider knowledge itself, conscientiously exploring just how a given concept of virtue can be identified with it.

[174] To be clear, the goal here is manifestly *not* to provide a reductive analysis of "epistemic virtue" in terms necessary and jointly sufficient conditions. If knowledge is a mental state and if mental states are understood roughly as what I am calling a virtuous state, then to pursue a reductive analysis of "virtuous state" would place us well on our way to a reductive analysis of knowledge, a project we now have reason to despair of.

Section 3.3: Knowledge within Virtue

Figure 5.4

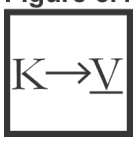

Thus far, we have assumed that knowledge is the conceptual foundation for epistemology. Whether attaching a virtue-rich concept onto knowledge as a necessary condition (the knowledge plus virtue strategy) or redefining knowledge in terms of virtue (the knowledge as virtue strategy), we have assumed that knowledge is the richest epistemic state, that knowledge is our conceptual cornerstone. As such, another family of strategies can be developed by doing away with this assumption and placing virtue *under* knowledge, so to speak—by identifying virtue itself as epistemic bedrock, of which knowledge is merely a facet (Figure 5.4 or A3 on Table 5.1). On this account, a conceptually non-primitive account of knowledge entails a primitive account of virtue; knowledge is sufficient for virtue and virtue is necessary for knowledge, but knowledge is understood in terms of virtue. Virtue is the broader, richer epistemic concept. In other words, knowledge, on this account, cannot be properly comprehended apart from our understanding of virtue. And like the knowledge as virtue proposal, the knowledge within virtue proposal is singularly non-reductive. But unlike the knowledge as virtue proposal, which was non-reductive insofar as knowledge was deemed primitive, the knowledge within virtue proposal is non-reductive insofar as the project of pursuing sufficiency conditions for knowledge has been abandoned.

Consider the following case from Williamson's *Knowledge and Its Limits* (2000), which goes toward demonstrating the richness of knowledge over other concepts like "believing truly":

> **Burglar:** A burglar spends all night ransacking a house, risking discovery by staying so long. We ask what features of the situation when he entered the house led to that result. A reasonable answer is that he knew that there was a diamond in the house. To say just that he believed truly that there was a diamond in the house would be to give a worse explanation, one whose explanans and explanandum are less closely connected. For one possibility consistent with the new explanans is that the burglar entered the house with a true belief that there was a diamond in it derived from false premises. For example, his only reason for believing that there was a diamond in the house might have been that someone told him that there was a diamond under the bed, when in fact the only diamond was in a drawer. He would then very likely have given up his true belief that there was a diamond in the house on discovering the falsity of his belief that there was a diamond under the bed, and abandoned the search. In contrast, if he *knew* that there was a diamond in the house, his knowledge was not essentially based on a false premise. Given suitable background conditions, the probability of his ransacking the house all night, conditional on his having entered it believing truly but not knowing that there was a diamond in it, will be lower than the probability of his ransacking it all

> night, conditional on his having entered it knowing that there was a diamond in it. (2000, 62)

While the knowledge within virtue proposal may seem bizarre and initially inscrutable, cases like Burglar help us see how such a proposal might get off the ground. Insofar as Burglar is a good argument for establishing the conceptual priority of knowledge over other concepts like believing truly, perhaps a similar story could be told to motivate the conceptual priority of virtue over knowledge. If we could come up with a case in which an agent's actions are best explained by epistemic virtue and not merely by knowledge, perhaps we could promote the richness and primacy of the former over the latter. Consider the following real-life example from Tom Wright's book, *Virtue Reborn* (2010):

> **Virtuous Pilot:** Thursday, January 15, 2009, was another ordinary day in New York City. Or so it seemed. . . . Flight 1549, a regular US Airways trip from LaGuardia Airport took off at 15:26 local time, bound for Charlotte, North Carolina. The captain, Chesley Sullenberger III, known as "Sully," did all the usual checks. Everything was fine in the Airbus A320. Fine until, two minutes after takeoff, the aircraft ran straight into a flock of Canada geese. . . . Almost at once both the engines were severely damaged and lost their power. The plane was at that point heading north over the Bronx, one of the most densely populated parts of the city. [Sully] and his co-pilot had to make several major decisions instantly if they were going to save the lives of people not only on board but also on the ground. They could see one or two small local airports in the distance, but quickly realized that they couldn't be sure of making it that far. If they attempted it, they might well crash-land in a built-up area on the way. Likewise, the option of putting the plane down on the New Jersey Turnpike, a busy main road leading in and out of the city, would present huge problems and dangers for the plane and its occupants, let alone for cars and their drivers on the road. That left one option: the Hudson River. It's difficult to crash-land on water: one small mistake—catch the nose or one of the wings in the river, say—and the plane will turn over and over like a gymnast before breaking up and sinking. In the two or three minutes they had before landing, [Sully] and his copilot had to do [several] vital things (along with plenty of other tasks that we amateurs wouldn't understand) [which needed to be done to safely land]. . . . And they did it! Everyone got off safely, with [Sully] himself walking up and down the aisle a couple of times to check that everyone had escaped before leaving himself. (Wright 2010, 7–8)

What best explains Sully's successful landing on the Hudson River? No doubt, he knew what he was doing, but such an explanation does not seem to capture it. Presumably, many pilots would superficially know what to do in the same circumstances but would nevertheless fail at the task. I think the better explanation for Sully's successful landing in that emergency situation and in those extreme circumstances is that he was (is) a very good pilot, a virtuous pilot.[175] As such, if Burglar shows that knowledge is conceptually prior to

[175] This is, to be sure, Wright's conclusion as well. See 2010, 8.

true belief, then presumably cases like Virtuous Pilot can show that virtue is conceptually prior to knowledge, that virtue is the richest epistemic state—a state that, on this model, is both metaphysically necessary for and conceptually prior to knowledge.

This entirely pioneering approach to non-reductive virtue epistemology would have little trouble preserving the noted merits of virtue epistemology; in a radical sense, virtue would be at the heart of epistemology and inter-disciplinary dialogue would not only be possible but essential (especially given that the given concept of virtue would presumably extend beyond epistemology). That said, however, the knowledge within virtue proposal is radical and extremely ambitious, requiring a substantial amount of work to adequately develop. Pursuing the knowledge within virtue proposal would not only force us to carefully reconsider knowledge itself, conscientiously exploring just how it fits within a broader concept of virtue, it would force us to develop a thoroughgoing account of virtue to match—a formidable project that would presumably take us well away from epistemology. If it can be viably developed, the *knowledge within virtue* proposal would be an interesting strategy for establishing a non-reductive virtue epistemology and well worth consideration.

Section 4: Rethinking Williamson, Plantinga, and Sosa

Having mapped the various ways in which virtue might be incorporated within a non-reductive model of knowledge, we were able to explicate three general strategies for producing virtue epistemologies that are not committed to the analyzability of knowledge—knowledge plus virtue, knowledge as virtue, and knowledge within virtue. But in all of this, we left the relevant epistemic concepts abstract and uncommitted to any specific account of knowledge (e.g., knowledge as mental states) or any specific account of virtue (e.g., proper functioning cognitive faculties, cognitive competencies, etc.). In this section, we return to the work of Williamson, Plantinga, and Sosa to consider the aforementioned strategies first in light of Williamson's particular account of knowledge and then in light of Plantinga's and Sosa's respective accounts of virtue.

On the face of it, all three strategies are compatible with Williamson's non-reductive account of knowledge. This is not to say, of course, that Williamson himself would approve of any of these strategies, just that there is nothing inherent in his epistemology that precludes them. First, consider the knowledge plus virtue proposal. As Williamson says himself, "[t]he present account of knowing makes no use of such concepts as *justified*,

caused, and *reliable*. Yet knowing seems to be highly sensitive to such factors over a wide range of cases. Any adequate account of knowing should enable one to understand these connections" (2000, 41). Williamson acquiesces that he only "adumbrate[s] a strategy without carrying it out" in *Knowledge and its Limits*—a project that future iterations of Williamson's epistemology should seemingly pursue (2000, 41). Supposedly, then, one set of intuitions that Williamson's epistemology should somehow explain or "be highly sensitive to" are those springing from "the wide range of cases" that help motivate virtue epistemology. In the Section 2, we saw how Williamson ornaments his theory of knowledge with a safety condition. Once he recognizes that "knowing seems to be highly sensitive to such factors over a wide range of cases" (e.g., Fake Barn cases), Williamson is happy to amend his account accordingly—dubbing safety a necessary condition for knowledge (Williamson 2000, 41). The knowledge plus virtue proposal is simply the proposal to likewise ornament Williamson's epistemology with another condition, a virtue-theoretic condition. Again, just as safety might be key to explaining and tracking knowledge ascriptions across cases like Fake Barns and Sheep 2, perhaps a virtue-theoretic condition is needed to explain and track knowledge ascriptions across cases like René and the Gambler's Fallacy and Brain Lesion. Through his incorporation of safety, Williamson paved the way for our favorite epistemic concepts to piggyback on his non-reductive model via the knowledge plus virtue proposal.

Applying the knowledge as virtue proposal to Williamson's account is no less straightforward; instead of defining knowledge as a mental state, simply define it as a virtuous state—explicating *mental states* in terms of *epistemic virtue*, epistemically virtuous motivation that reaches its given telos.[176] Despite being a blatant theoretical "smash and grab" job, such a move does not seem to do any violence to Williamson's epistemic framework. As was noted, Williamson says very little as to what it means for something to be a mental state; and as such, there does not seem to be any theoretical hurdles keeping us from filling in the relevant details with our favorite virtue-theoretic concepts—proper function, cognitive competencies, etc. And to be sure, we only need the *epistemic* mental

[176] Insofar as Williamson wants to explicate *believing something is so* as a failed attempt to *knowing something is so* (2000, 41–48), applying the knowledge as virtue proposal to mental states is both helpful and elucidating. Within a virtue-theoretic framework, perhaps *believing something is so* can be described as an example of an epistemic act that is sufficiently (virtuously) motivated while falling short of its virtuous telos (in this case, truth).

states (i.e., mental states that correlate to knowledge—the *factive mental states*) to be epistemic virtues or virtuous states; we do not need to try to describe *all* mental states in virtue-theoretic terminology.[177] And what is more, in describing knowledge as a mental state by drawing parallels between knowledge and quintessential mental states, as Williamson does, we are not precluded from further elucidating epistemic mental states in terms of virtuous states. Under this proposal, knowledge remains primitive, factive, and indeed the most general factive mental state; the critical difference being that *factive mental state* is now understood using virtue-theoretic concepts. To be sure, someone might resist explicating mental states in terms of virtuous states on the grounds that, unlike quintessential mental states like *believing something is so*, virtuous states typically consist of both an internal and an external component—internal epistemically virtuous motivation and an epistemically virtuous telos that typically contains an external truth component.[178] But insofar as we lack any robust theoretical grounds to assume an overly internalistic conception of the mental, such a worry possesses no real risk.[179] Insofar as we are willing to allow for factive mental states, there is no hurdle to explicating (factive) mental states in terms of epistemic virtue, in terms of virtuous states.[180]

Finally, consider Williamson's account of knowledge in conjunction with the knowledge within virtue proposal. Insofar as we can redefine mental states in terms of epistemic virtue (as in our application of the previous proposal), there are seemingly no intrinsic obstacles within the Williamsonian model to expanding a given account of virtue to be the richest epistemic state, our conceptual cornerstone. In other words, if mental states can be reconceived in terms of virtue, they could easily be planted within a rich and broad virtue-laden context. However, insofar as Williamson's model for knowledge can be given a virtue-theoretic reading, the problem with the knowledge within virtue proposal

[177] Though if we *can* explicate quintessential mental states (e.g., love, hate, imagining that something is so, believing that something is so, etc.) within a virtue-theoretic framework, all the better.

[178] The exceptions being virtuous states where the given telos includes a truth that is internal to the relevant agent.

[179] To put it one way, virtuous states are not in people, people are in virtuous states.

[180] This only gestures toward an outline, without filling in much detail. Williamson's epistemology is revolutionary, with far-reaching ramifications that touch on many subjects in contemporary epistemology. A full account would carefully track the overall effect this virtue-theoretic adaptation has on Williamson's broader epistemic project—making amendments and modifications as need be.

will not come from any incompatibility with Williamson's epistemology, it will come from finding an account of virtue that will viably suit such an expansion.

Now consider Plantinga's virtue epistemology. In keeping with the traditional suppositions, Plantinga's account took the form of a reductive analysis of knowledge in terms of necessary and jointly sufficient conditions—starting with truth and belief and trying to work his way up to knowledge via warrant. And in an attempt to satisfy the high demands of such a project, Plantinga proposed an increasingly complex and intricate account of warrant. To put it roughly, a true belief will be knowledge by Plantinga's lights if it is formed by cognitive faculties that are functioning properly in the right environment (both mini-environment and maxi-environment) and in accord with a good design plan aimed at truth. Despite all of the finely tuned "bells and whistles," the demands of the reductive analysis project were simply too high; as we saw in Chapter 3, Plantinga's analysis of knowledge simply could not surmount the Gettier Problem—either falling victim to Gettier counterexamples or leading to radical and untenable skepticism.

We can preserve the heart of Plantinga's epistemology (i.e., his notion of proper function), however, by adapting it to suit the Williamsonian model using any of the general strategies explicated above. It is his particular focus on the *properly functioning* human knower that is meant to distinguish his account from all others. Plantinga's notion of proper function is meant to rightly connect a given agent to the facts (cf. René and the Gambler's Fallacy) and, of course, rightly preclude knowledge from cases of malfunction (cf. Brain Lesion). According to Plantinga, other theories of knowledge fail due to their inability to track knowledge ascription in accord with the proper functioning of the relevant cognitive faculties behind a given belief's genesis. Proper function is, for Plantinga, not only the key virtue-theoretic concept in his account, it is meant to be the "rock on which" competing theories of knowledge "founder" (Plantinga 1993b, 4). Epistemic virtue for Plantinga is the mechanical proper functioning of cognitive faculties. *Given the knowledge plus virtue proposal,* a non-reductive virtue epistemology that is (largely) in keeping with the heart of Plantinga's theory of knowledge can be developed by simply amending Williamson's model with a caveat that proper functioning is a necessary condition for knowledge.[181] *Given the knowledge as virtue proposal*, knowledge, under a Williamsonian adaptation of Plantinga's epistemology,

[181] As we noted earlier, the worry with such a strategy is that you lose some of the heart and soul of virtue epistemology. Plantinga's notion of proper function is meant to be absolutely fundamental for knowledge, a feature that cannot be preserved when proper functioning is degraded to a caveat.

simply is properly functioning cognitive faculties or at least the state thereof. And of course, the relevant cognitive faculties at issue here are those that are factive, veritic, or (to use Plantinga's terminology) "aimed at truth." Being in a virtuous epistemic state that *p* as a result of cognitive faculty, *F*, will only entail knowledge that *p* if *F* is factive, if *F* is veritic, if *F* is aimed at truth. There are, no doubt, many veritic epistemic states—seeing that *p* (with properly functioning visual faculties), hearing that *p* (with properly functioning auditory faculties), etc.—but being in any such state is sufficient for knowledge. In other words, knowledge is the most general epistemic virtue. And of course, knowledge that *p*, being in a virtuous state that *p*, entails that *p*. Finally, *given the knowledge within virtue proposal*, if we can re-describe Williamson's epistemology in terms of proper function, there will be no intrinsic hurdle to planting such an account within a broader and richer virtue-laden context. In other words, there is nothing about Plantinga's notion of epistemic virtue that would preclude a project that would establish virtue as the richest epistemic state, being metaphysically necessary for knowledge and conceptually prior to knowledge.

Consider once again Sosa's virtue-theoretic account of knowledge. In keeping with the traditional suppositions, Sosa's account, like Plantinga's, also took the form of an analysis of knowledge in terms of necessary and jointly sufficient conditions. And in an attempt to satisfy the high demands of such a project without producing an overly complex, intricate, and (perhaps) *ad hoc* account, Sosa proposed a relatively simple and elegant theory of knowledge; a belief will be knowledge by Sosa's lights if it is *apt*, if it is "true because competent," if its truth is somehow "attributable to" or "derived from" relevant cognitive faculties. From this simple definition, Sosa was able to unfold a two-tiered classification of knowledge: animal knowledge (apt belief) and reflective knowledge (apt belief aptly held). And by invoking this distinction, Sosa is able to provide new and interesting approaches to various philosophical debates (e.g., the internalism/externalism debate). Unfortunately, however, the simplicity of Sosa's analysis of knowledge was deceptive—hiding dubious ambiguity. As soon as we tried to elucidate *knowledge as apt belief*, we found that under any plausible interpretation, Sosa's analysis, like Plantinga's, either succumbs to Gettier counterexamples or yields a radical and untenable form of skepticism.

At the heart of Sosa's definition of knowledge is his virtue-theoretic concept—cognitive competence, more specifically those cognitive competencies from which truths can somehow be "attributed to" or "derived from." And parallel to our treatment of Plantinga's epistemology above, we can preserve this feature of Sosa's account by adapting

it to suit the Williamsonian model using any of the general strategies explicated above. Just like Plantinga's notion of proper function, Sosa's virtue theoretic concept of cognitive competency (his definition of knowledge as apt belief) is meant to rightly connect a given agent to the facts (cf. René and the Gambler's Fallacy) and rightly preclude knowledge from cases of malfunction (cf. Brain Lesion). Epistemic virtue for Sosa is the veritic exercise of cognitive competencies. *Given the knowledge plus virtue proposal*, a non-reductive virtue epistemology that is (largely) in keeping with the heart of Sosa's epistemology can be produced by simply amending Williamson's model of knowledge with the caveat that the veritic exercise of cognitive competencies is a necessary condition for knowledge.[182] *Given the knowledge as virtue proposal*, knowledge, under a Williamsonian adaptation of Sosa's epistemology, is veritic exercise of a cognitive competence or at least the given state thereof. Of course, being veritic, the relevant cognitive competencies are factive. In other words, using my visual cognitive competency to judge that *p*, given that such a competency is veritic, entails *p*. Additionally, being in a virtuous epistemic state that *p* as a result of cognitive competency, *C*, will only entail knowledge that *p* if *C* is factive. There are, no doubt, many veritic epistemic states—seeing that *p* (via visual cognitive competencies), hearing that *p* (with auditory cognitive competencies) etc.—but being in any such state is sufficient for knowledge. Knowledge, again, is the most general epistemic virtue. Finally, *given the knowledge within virtue proposal*, if we can re-describe Williamson's epistemology in terms of veritic cognitive competencies, there will be no intrinsic hurdle to planting such an account within a broad and rich virtue-laden context. In other words, there is nothing about Sosa's notion of epistemic virtue that would preclude a project that would establish virtue as the richest epistemic state, being metaphysically necessary for knowledge and conceptually prior to knowledge.[183]

By way of prolegomena, we have seen how the general strategies of *knowledge plus virtue*, *knowledge as virtue*, and *knowledge within virtue* can be used to develop non-reductive virtue

[182] Again, while this strategy does indeed preserve many of the merits of virtue-theoretic epistemology, the lover of robust virtue epistemology may not find it ultimately satisfying. For example, the heart and soul of Sosa's virtue epistemology is the veritic exercise of cognitive competencies, so reducing such a concept to a mere addendum significantly changes the account.

[183] And of course, none of these recapitulations of Sosa's epistemology does anything to prohibit his innovative approaches to topics like the internalism/externalism debate; a similar two-tiered account of knowledge can still be provided along first-order versus second-order lines, which will allow for just the sort of arguments Sosa is wanting to make.

epistemologies that will seemingly preserve the merits of virtue-theoretic epistemology (cf. Section 3.1 through Section 3.3) without committing to the analyzability of knowledge. All of these strategies are compatible with Williamson's seminal epistemic model (cf. Section 2), and any of them could be used to insert the respective virtue-theoretic concepts of Plantinga and Sosa into a non-reductive framework. These discoveries open the door to a world of valuable research—comparing and developing these respective strategies in new ways, trying additional non-reductive models, and exploring other accounts of virtue. Nevertheless, the humble goals for this discursive introduction have been met. Our goal was to briefly explore some ways we might preserve virtue epistemology using the Williamsonian model, and this is precisely what we have done.

Final Conclusion

Epistemology is on the move. For nearly 50 years philosophers have tried to provide a reductive analysis of knowledge in terms of necessary and jointly sufficient conditions—a project that has, despite heroic efforts, yielded little lasting success. And for a minority of epistemologists, the situation is at its breaking point; for them it is time to abandon the reductive analysis endeavor in favor of alternative epistemic models. In this dissertation I have tried to lend credence to this minority position by examining the nature of luck and systematically criticizing some of the chief proponents of what is perhaps the most popular species of reductive analysis to date, virtue epistemology. And while I have not definitively argued that the Gettier Problem is inescapable, the proposed diagnosis of the Gettier Problem along with the light of experience nevertheless provide excellent *prima facie* reason to be pessimistic regarding any prospects of providing a viable definition of knowledge (virtue-theoretic or otherwise) styled after the standard analysis. That said, however, virtue epistemology nevertheless has a lot to offer. In addition to providing innovative and exciting approaches to a variety of issues that have plagued other accounts of knowledge, virtue epistemology (i) allows for special inroads for dialogue between epistemology and other subdisciplines, not least moral philosophy and philosophy of religion and (ii) tracks our intuitions across a range of cases. And in this final chapter, I have tried to gesture toward a way in which we might preserve the heart of virtue epistemology by importing the relevant virtue-theoretic concepts into an alternative epistemic model—by elucidating Timothy Williamson's seminal account of knowledge as an unanalyzable primitive and adapting it to suit the work of those preeminent virtue epistemologists, Alvin Plantinga and

Ernest Sosa. And as such, it seems as though what attracts us to virtue epistemology need not be necessarily wed to the analyzability of knowledge.

Of course, there is a lot more work that needs to be done. We have only considered in broad strokes how virtue-theoretic concepts might be imported into an alternative epistemic model like Timothy Williamson's—largely leaving absent details and overlooked objections to the reader's forbearance. Building from the work done in this dissertation, however, I now hope to spearhead an exciting and pioneering three-stage research project focused on the establishment of a non-reductive virtue epistemology. In the first stage, I hope to draw from a range of key debates in epistemology, moral philosophy, and philosophy of religion in order to decide on (i) a specific virtue-theoretic concept (proper function? cognitive competencies? intellectual character traits?) and (ii) such a concept's rightful place within a given non-reductive epistemology (virtue as an addendum? epistemic virtue as knowledge?). In the second stage, I hope to take this information and fully elucidate, develop, and defend a robust non-reductive virtue epistemology; this being a first major step toward a positive contribution to the philosophical literature. This non-reductive account may be styled after the Williamsonian model, but not necessarily so; there are, after all, other (less seminal) non-reductive models on offer in the contemporary literature that will be worth surveying. Finally, in the third stage, I hope to explore the broader ramifications of such a theory of knowledge; exploring its effects on other debates in epistemology as well as various issues in moral philosophy and philosophy of religion. It is my ambition that through this research I will be able to establish a virtue epistemology that is able not only to endure what I see as inevitable developments in twenty-first-century epistemology but also contribute positively to a number of debates and discussions across philosophy.

BIBLIOGRAPHY

Anscombe, G. E. M. 1958. "Modern Moral Philosophy." *Philosophy* 33 (124): 1-19.

Austin, J. L. 1962. *Sense and Sensibilia*. Oxford, UK: Oxford University Press.

Axtell, Guy, ed. 2000. *Knowledge, Belief, and Character: Readings in Contemporary Virtue Epistemology*. London, UK: Rowman & Littlefield Publishers.

Battaly, Heather. 2009. "A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I— by Ernest Sosa." *Analysis* 69 (2): 382-385.

Botham, Thad M. 2003. "Plantinga and Favorable Mini-Environments." *Synthese* 135 (3): 431-441.

Brady, Michael, and Duncan Pritchard, eds. 2003. *Moral and Epistemic Virtues*. Oxford, UK: Blackwell Publishing.

Cassim, Quassim. 2009. "Can the Concept of Knowledge be Analysed?" In *Williamson on Knowledge*. Patrick Greenough and Duncan Pritchard, eds. Oxford, UK: Oxford University Press.

Chignell, Andrew. 2003. "Accidentally True Belief and Warrant." *Synthese* 137 (3): 445-458.

Chisholm, Roderick M. 1977. *Theory of Knowledge*. 2nd ed. Englewood Cliffs, NJ: Prentice-Hall.

———. 1982. *The Foundations of Knowing*. Brighton, UK: The Harvester Press.

Church, Ian M. Forthcoming. "Getting 'Lucky' with Gettier." *European Journal of Philosophy*. doi:10.1111/j.1468-0378.2010.00433.x.

———. 2009. "'Entitlement of Cognitive Project' and the Possibility of Gettier-Proof Basic Beliefs." Presented at the University of Notre Dame.

———. 2011. "Accidentally Apt Belief and Skepticism." Presented at the Edinburgh Epistemology Graduate Conference, The University of Edinburgh.

Clark, Michael. 1963. "Knowledge and Grounds: A Comment on Mr. Gettier's Paper." *Analysis* 24 (2): 46-48.

BIBLIOGRAPHY

Code, Lorraine. 1987. *Epistemic Responsibility*. Hanover, NH: Published for Brown University Press by University Press of New England.

Cohen, Stewart. 2009. "Knowledge as Aptness." *Philosophical Studies* 144 (1): 121-125.

Crisp, Thomas M. 2000. "Gettier and Plantinga's Revised Account of Warrant." *Analysis* 60 (265): 42–50.

Cummins, Robert E. 1998. "Reflections on Reflective Equilibrium." In *Rethinking Intuition*. Michael DePaul and William Ramsey, eds. London, UK: Rowman & Littlefield Publishers.

Dancy, Jonathan. 1985. *An Introduction to Contemporary Epistemology*. Oxford, UK: Blackwell Publishing.

DePaul, Michael, and Linda Zagzebski, eds. 2003. *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford, UK: Oxford University Press.

Descartes, René. 1955. *The Philosophical Works of Descartes*. Trans. Elizabeth S. Haldane and G. R. T. Ross. Vol. 1. 2 vols. New York, NY: Dover Publications.

Dretske, Fred I. 1978. "Conclusive Reasons." In *Essays on Knowledge and Justification*. George Sotiros Pappas and Marshall Swain, eds. Ithaca, NY: Cornell University Press.

Fairweather, Abrol, and Linda Trinkaus Zagzebski, eds. 2001. *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Oxford, UK: Oxford University Press.

Feldman, Richard. 1996. "Plantinga, Gettier, and Warrant." In *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge*. Jonathan L. Kvanvig, ed. London, UK: Rowman & Littlefield Publishers.

Floridi, L. 2004. "On the Logical Unsolvability of the Gettier Problem." *Synthese* 142 (1): 61-79.

Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23 (6): 121-123.

Goldman, Alvin I. 1967. "A Causal Theory of Knowing." *Journal of Philosophy* 64 (12): 357-372.

———. 1976. "Discrimination and Perceptual Knowledge." *The Journal of Philosophy* 73 (20): 771-791.

———. 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

Greco, John. 1993. "Virtues and Vices of Virtue Epistemology." *Canadian Journal of Philosophy* 23 (3): 413-432.

———. 1995. "A Second Paradox Concerning Responsibility and Luck." *Metaphilosophy* 26 (1-2): 81-96.

———. 1999. "Agent Reliabilism." *Philosophical Perspectives* 13: 273-296.

———. 2000. "Two Kinds of Intellectual Virtue." *Philosophy and Phenomenological Research* 60 (1): 179-184.

———. 2002. "Virtues in Epistemology." In *The Oxford Handbook of Epistemology*. Paul K. Moser, ed. New York, NY: Oxford University Press.

———. 2003a. "Virtue and Luck, Epistemic and Otherwise." *Metaphilosophy* 34 (3): 353-366.

———. 2003b. "Knowledge as Credit for True Belief." In *Intellectual Virtue: Perspectives From Ethics and Epistemology*. Michael DePaul and Linda Zagzebski, eds. Oxford, UK: Oxford University Press.

———. 2007a. "The Nature of Ability and the Purpose of Knowledge." *Philosophical Issues* 17 (1): 57-69.

———. 2007b. "Worries about Pritchard's Safety." *Synthese* 158 (3): 299-302.

———. 2009. "Knowledge and Success from Ability." *Philosophical Studies* 142 (1): 17-26.

———. 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge, UK: Cambridge University Press.

Greco, John, and John Turri. 2011. "Virtue Epistemology." http://plato.stanford.edu/entries/epistemology-virtue/.

Greco, John, ed. 2004. *Ernest Sosa and His Critics*. Philosophers and their Critics 12. Oxford, UK: Blackwell Publishing.

Greenough, Patrick, and Duncan Pritchard, eds. 2009. *Williamson on Knowledge*. Oxford, UK: Oxford University Press.

Hetherington, Stephen. 1998. "Actually Knowing." *Philosophical Quarterly* 48 (193): 453-469.

———. 2001. *Good Knowledge, Bad Knowledge: On Two Dogmas of Epistemology*. Oxford, UK: Oxford University Press.

———. 2005. "Gettier Problems." http://www.iep.utm.edu/gettier/.

———. 2007. "Is This a World Where Knowledge has to Include Justification?" *Philosophy and Phenomenological Research* 75 (1): 41-69.

———. 2010. "Elusive Epistemological Justification." *Synthese* 174 (3): 315-330.

Hookway, Christopher. 2000. *Truth, Rationality, and Pragmatism: Themes From Peirce*. Oxford, UK: Oxford University Press.

Howard-Snyder, Daniel, Frances Howard-Snyder, and Neil Feit. 2003. "Infallibilism and Gettier's Legacy." *Philosophy and Phenomenological Research* 66 (2): 304-327.

Hume, David. 1748. *An Enquiry Concerning Human Understanding*.

———. 1948. *Dialogues Concerning Natural Religion*. The Hafner Library of Classics. Henry D. Aiken, ed. London, UK: Hafner Press.

Jackson, Frank. 2009. "Thought Experiments and Possibilities." *Analysis* 69 (1): 100-109.

Kant, Immanuel. 1998. "Critique of Pure Reason." In *Modern Philosophy: An Anthology of Primary Sources*. Roger Ariew and Eric Watkins, eds. Cambridge, UK: Hackett Publishing Company.

Klein, Peter. 1996. "Warrant, Proper Function, Reliabilism, and Defeasibility." In *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge*. Jonathan L. Kvanvig, ed. London, UK: Rowman & Littlefield Publishers.

Kvanvig, Jonathan L., ed. 1996. *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Epistemology*. London, UK: Rowman & Littlefield Publishers.

Lackey, Jennifer. 2006. "Pritchard's Epistemic Luck." *Philosophical Quarterly* 56 (223): 284-289.

———. 2008. "What Luck is Not." *Australasian Journal of Philosophy* 86 (2): 255-267.

Latus, Andrew. 2000. "Moral and Epistemic Luck." *Journal of Philosophical Research* 25 (2000): 149-172.

Levin, Michael. 2004. "Virtue Epistemology: No New Cures." *Philosophy and Phenomenological Research* 69 (2): 397-410.

Lewis, David K. 2000. *Counterfactuals*. 2nd ed. Oxford, UK: Blackwell Publishing.

Lycan, William G. 2006. "On the Gettier Problem Problem." In *Epistemology Futures*. Stephen Hetherington, ed. Oxford, UK: Oxford University Press.

Madison, B. J. C. 2011. "Combating Anti Anti-Luck Epistemology." *Australasian Journal of Philosophy* 89 (1): 47-58.

Markie, Peter. 1996. "Degrees of Warrant." In *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge*. Jonathan L. Kvanvig, ed. London, UK: Rowman & Littlefield Publishers.

McDowell, John Henry. 1988. "Criteria, Defeasibility, and Knowledge." In *Perceptual Knowledge*, Oxford Readings in Philosophy. Jonathan Dancy, ed. Oxford, UK: Oxford University Press.

———. 1994. *Mind and World*. Cambridge, MA: Harvard University Press.

Millar, Alan. 2008. "Perceptual-Recognitional Abilities and Perceptual Knowledge." In *Disjunctivism: Perception, Action, Knowledge*. Adrian Haddock and Fiona Macpherson, eds. Oxford, UK: Oxford University Press.

———. 2009. "What Is It that Cognitive Abilities are Abilities To Do?" *Acta Analytica* 24 (4): 223-236.

Montmarquet, James A. 1993. *Epistemic Virtue and Doxastic Responsibility*. Lanham, MD: Rowman & Littlefield Publishers.

Nagel, Thomas. 1979. *Mortal Questions*. Cambridge, UK: Cambridge University Press.

Nozick, Robert. 1981. *Philosophical Explanations*. Oxford, UK: Oxford University Press.

Plantinga, Alvin. 1993a. *Warrant: The Current Debate*. New York, NY: Oxford University Press.

———. 1993b. *Warrant and Proper Function*. New York, NY: Oxford University Press.

———. 1996. "Respondeo." In *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge*. Jonathan L. Kvanvig, ed. New York, NY: Rowman & Littlefield Publishers.

———. 1997. "Warrant and Accidentally True Belief." *Analysis* 57 (2): 140-145.

———. 2000. *Warranted Christian Belief*. New York, NY: Oxford University Press.

Plato. 2002. *Five Dialogues*. Trans. G. M. A. Grube. 2nd ed. Cambridge, UK: Hackett Publishing Company.

Polanyi, Michael. 1974. *Personal Knowledge: Towards a Post-Critical Philosophy*. London, UK: Routledge & Kegan Paul.

Pritchard, Duncan. 2005. *Epistemic Luck*. Oxford, UK: Oxford University Press.

———. 2007. "Anti-Luck Epistemology." *Synthese* 158 (3): 277-297.

Pritchard, Duncan, Alan Millar, and Adrian Haddock. 2010. *The Nature and Value of Knowledge: Three Investigations*. Oxford, UK: Oxford University Press.

Quine, W. V. 1969. "Epistemology Naturalized." In *Ontological Relativity and other Essays*. 6th ed. New York, NY: Columbia University Press.

Reid, Thomas. 1785. *Essays on the Intellectual Powers of Man*.

Riggs, Wayne. 2002. "Reliability and the Value of Knowledge." *Philosophy and Phenomenological Research* 64 (1): 79-96.

———. 2007. "Why Epistemologists Are So Down on Their Luck." *Synthese* 158 (3): 329-344.

Rosch, Eleanor, and Carolyn B. Mervis. 1975. "Family Resemblances: Studies in the Internal Structure of Categories." *Cognitive Psychology* 7 (4): 573-605.

Rosenberg, Jay F. 2003. *Thinking about Knowing*. Oxford, UK: Oxford University Press.

Russell, Bertrand. 1948. *Human Knowledge: Its Scope and Limits*. London, UK: Routledge.

Schaffer, Jonathan. 2009. "Modalities and Methodologies." Presented at the Arché Methodology Workshop, The University of St Andrews.

Sellars, Wilfrid. 1956. "Empiricism and the philosophy of mind." *Minnesota Studies in the Philosophy of Science* 1: 253-329.

Shope, Robert K. 2008. "Abnormality, Cognitive Virtues, and Knowledge." *Synthese* 163 (1): 99-118.

———. 1983. *The Analysis of Knowing: A Decade of Research*. Princeton, NJ: Princeton University Press.

Smith, Martin. 2009. "Transmission Failure Explained." *Philosophy and Phenomenological Research* 79 (1) : 164-189.

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: Selected Essays in Epistemology*. Cambridge, UK: Cambridge University Press.

———. 1996. Proper Functionalism and Virtue Epistemology. In *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge*. Jonathan L. Kvanvig, ed. London, UK: Rowman & Littlefield Publishers.

———. 1998. "Minimal Intuition." In *Rethinking Intuition*. Michael DePaul and William Ramsey, eds. London, UK: Rowman & Littlefield Publishers.

———. 2004. "Replies." In *Ernest Sosa and His Critics*. John Greco, ed. Oxford, UK: Blackwell Publishing.

———. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Vol. 1. Oxford, UK: Oxford University Press.

———. 2009a. *Reflective Knowledge: Apt Belief and Reflective Knowledge*. Vol. 2. Oxford, UK: Oxford University Press.

———. 2009b. "Timothy Williamson's Knowledge and its Limits." In *Williamson on Knowledge*. Patrick Greenough and Duncan Pritchard, eds. 203-216. Oxford, UK: Oxford University Press.

———. 2009c. "Replies to Commentators on A Virtue Epistemology." *Philosophical Studies* 144 (1): 137-147.

Statman, Daniel. 1991. "Moral and Epistemic Luck." *Ratio* 4 (2): 146-156.

Steup, Matthias. 2006. "The Analysis of Knowledge." http://plato.stanford.edu/entries/knowledge-analysis/.

Steup, Matthias, ed. 2001. *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford, UK: Oxford University Press.

Stich, Stephen. 1988. "Reflective Equilibrium, Analytic Epistemology and the Problem of Cognitive Diversity." *Synthese* 74 (3): 391-413.

———. 1993. *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. Cambridge, MA: MIT Press.

Sturgeon, Scott. 1993. "The Gettier Problem." *Analysis* 53 (3): 156-164.

Turri, John. 2011. "Manifest Failure: The Gettier Problem Solved." *Philosophers' Imprint* 11 (8).

Unger, Peter. 1968. "An Analysis of Factual Knowledge." *The Journal of Philosophy* 65 (6): 157-170.

Weatherson, Brian. 2003. "What Good Are Counterexamples?" *Philosophical Studies* 115 (1): 1-31.

Weinberg, Jonathan M., Shaun Nichols, and Stephen Stich. 2001. "Normativity and Epistemic Intuitions." *Philosophical Topics,* 29 (1-2): 429-460.

Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford, UK: Oxford University Press.

———. 2004. "Sosa on Abilities, Concepts, and Externalism." In *Ernest Sosa and His Critics*. John Greco, ed. Oxford, UK: Blackwell Publishing.

———. 2007. *The Philosophy of Philosophy*. Oxford, UK: Blackwell Publishing.

———. 2009a. "Replies to Kornblith, Jackson and Moore." *Analysis* 69 (1): 125-135.

———. 2009b. "Replies to Critics." In *Williamson on Knowledge*. Patrick Greenough and Duncan Pritchard, eds. Oxford, UK: Oxford University Press.

Wittgenstein, Ludwig. 1980. *Culture and Value*. Amended 2nd ed. with English translation. Oxford, UK: Blackwell Publishing.

Wood, W. Jay, and Robert C. Roberts. 2007. *Intellectual Virtues : An Essay in Regulative Epistemology*. Oxford, UK: Oxford University Press.

Wright, Tom. 2010. *Virtue Reborn*. London, UK: SPCK Publishing.

Zagzebski, Linda. 1994. "The Inescapability of Gettier Problems." *The 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, UK: Cambridge University Press.

———. 1999. "What is Knowledge?" In *The Blackwell Guide to Epistemology*. Ernest Sosa and John Greco, eds. Oxford, UK: Blackwell Publishing.

Zimmerman, M. J. 1993. "Luck and Moral Responsibility." In *Moral Luck*. Daniel Statman, ed. Albany, NY: State University of New York Press.