

---

Theses and Dissertations

---

Summer 2013

# The Problem of Easy Justification: An Investigation of Evidence, Justification, and Reliability

Samuel Alexander Taylor  
*University of Iowa*

Copyright 2013 Samuel Alexander Taylor

This dissertation is available at Iowa Research Online: <https://ir.uiowa.edu/etd/4917>

---

## Recommended Citation

Taylor, Samuel Alexander. "The Problem of Easy Justification: An Investigation of Evidence, Justification, and Reliability." PhD (Doctor of Philosophy) thesis, University of Iowa, 2013.  
<https://ir.uiowa.edu/etd/4917>.

---

Follow this and additional works at: <https://ir.uiowa.edu/etd>



Part of the [Philosophy Commons](#)

THE PROBLEM OF EASY JUSTIFICATION: AN INVESTIGATION OF  
EVIDENCE, JUSTIFICATION, AND RELIABILITY

by  
Samuel Alexander Taylor

A thesis submitted in partial fulfillment  
of the requirements for the Doctor of  
Philosophy degree in Philosophy  
in the Graduate College of  
The University of Iowa

August 2013

Thesis Supervisor: Professor Richard A. Fumerton

Copyright by  
SAMUEL ALEXANDER TAYLOR  
2013  
All Rights Reserved

Graduate College  
The University of Iowa  
Iowa City, Iowa

CERTIFICATE OF APPROVAL

---

PH.D. THESIS

---

This is to certify that the Ph.D. thesis of

Samuel Alexander Taylor

has been approved by the Examining Committee  
for the thesis requirement for the Doctor of Philosophy  
degree in Philosophy at the August 2013 graduation.

Thesis Committee:

\_\_\_\_\_  
Richard A. Fumerton, Thesis Supervisor

\_\_\_\_\_  
David Cuning

\_\_\_\_\_  
Evan Fales

\_\_\_\_\_  
Ali Hasan

\_\_\_\_\_  
Gregory Landini

To My Parents

A wise man proportions his belief to the evidence.  
David Hume, *An Enquiry Concerning Human Understanding*

## ACKNOWLEDGMENTS

The ideas for this dissertation came rather easily and naturally. However, actually writing the dissertation and attempting to explain these ideas has been one of the most difficult and trying tasks which I've ever undertaken. In completing the project I've accumulated debts of gratitude to many people, and I apologize to those of you who I will inevitably forget to thank.

I want to begin by thanking my family. Throughout my entire life my parents have instilled and encouraged in me an intellectual curiosity. They have challenged me to continually improve and do my best in all of my endeavors. Throughout my graduate career they have been a continual source of encouragement and support. Without them this dissertation wouldn't have been possible. I'd also like to thank my siblings Dana, Siri, Jake, and Matt. I couldn't have asked for a better family. Each of you inspires me in your own unique way. I'd also like to thank my soon to be wife Hannah Caster for being there for me during both the highs and lows of completing this dissertation. Finally, while not technically family, I want to offer special thanks to my friends Dillon and Jon simply for being such great friends for as long as I've known them.

I'd also like to thank the entire philosophy faculty at Iowa. It's because of your courses, the camaraderie you encourage, and your general excitement about philosophy that my time here has been amongst the best years of my life. If I could be a graduate student in this department forever I would. I also owe special thanks to my advisor Richard Fumerton. His influence on my philosophical thought will be readily apparent to anyone familiar with his work. I'm dumbfounded as to how Richard has always had time to help me work through and develop my inchoate ideas and still have the time for his own work and responsibilities. Without his help my mistakes would be much more numerous and the ideas would have been even less clearly expressed; his suggestions and criticism have undoubtedly improved this dissertation and it would likely have been even better had I heeded more of

his suggestions and criticism. I'd also like to explicitly thank the remaining members of my committee: David Cunning, Evan Fales, Ali Hasan, and Gregory Landini. I apologize for making you all read such a long dissertation. I'd like to offer special thanks to Ali. It was in my first seminar with Ali that I first heard of the problem of easy justification. More importantly, he is an amazing teacher whose feedback and guidance has no doubt made me a better philosopher. Everything I've learned and written in his classes has influenced this dissertation in one way or another.

Lastly, I need to thank the other philosophy graduate students that I've met throughout my years at Iowa. When I was applying to graduate schools one of my professors who encouraged me to attend graduate school told me that I would learn as much from my classmates as I did from my professors. He was right. Your conversations, feedback, shouting matches (not to mention any names but, cough, Kris), bar time philosophy, and friendships have been invaluable to me both personally and philosophically. I won't list names of those I would like to thank most simply because I would inevitably mention everyone who has been a philosophy graduate student at Iowa during my time here. I cannot give you all enough thanks.



## ABSTRACT

Our beliefs utilize various sources: perception, memory, induction, etc. We trust these sources to provide reliable information about the world around us. My dissertation investigates how this trust could be justified.

Chapter one introduces background material. I argue that justification rather than knowledge is of primary epistemological importance, discuss the internalism/externalism debate(s), and introduce an evidentialist thesis that provides a starting point/framework for epistemological theorizing.

Chapter two introduces a puzzle concerning justification. Can a belief source provide justification absent prior justification for believing it's reliable? Any answer appears to either make justifying the reliability of a source intellectually unsatisfying or all together impossible.

Chapter three considers and rejects a plethora of proposed solutions to our puzzle. Investigating these solutions illustrates the need to further investigate evidence, evidence possession, and evidential support.

Chapter four discusses the metaphysics of evidence. I argue that evidence always consists of a set of facts and that fact-proposition pairs stand in confirmation relations isomorphic to those holding between pairs of propositions.

Chapter five argues that justification requires what I call actually connected possession of supporting evidence: a subject must be aware of supporting evidence and of the support relation itself.

Chapter six argues that the relation constitutive of a set of facts being justificatory evidence is a *sui generis* and *irreducible* relation that is knowable a priori.

Chapter seven begins by showing how Richard Fumerton's acquaintance theory meets the constraints on a theory of justification laid down in previous chapters. I modify the theory so as to: (i) make room for fallible foundational justification, and (ii) allow inferential justification absent higher-order beliefs about evidential connections.

Chapter eight applies the developed theory of justification to our initial puzzle. I show how my modified acquaintance theory is in a unique position to vindicate the idea that *necessarily* a source provides a person with justification only if she is aware of evidence for the reliability of that source. However, this awareness of evidence for a source's reliability falls short of a justified belief and thereby avoids impalement from our dilemma's skeptical horn.

## TABLE OF CONTENTS

LIST OF FIGURES .....	xi
CHAPTER ONE: SURVEYING THE EPISTEMOLOGICAL TERRAIN .....	1
1.1 Metaepistemology and Normative Epistemology .....	1
1.2 Leaving Knowledge by the Wayside .....	2
1.3 Internalist Theories of Epistemic Justification .....	8
1.3.1 Access Internalism .....	9
1.3.2 Mentalism .....	16
1.3.3 Inferential Internalism .....	20
1.3.4 Internalism and the 1 <sup>st</sup> -Person Perspective .....	23
1.4 Externalist Theories of Epistemic Justification .....	25
1.5 Why Internalism? .....	27
1.6 Why Externalism? .....	31
1.7 An Evidentialist Framework .....	33
1.7.1 The Central Insight .....	33
1.7.2 Propositional and Doxastic Justification .....	35
1.7.3 Treating Evidentialism as a Heuristic .....	36
1.8 A Preview of Things to Come .....	37
CHAPTER TWO: EASY JUSTIFICATION AND THE PROBLEM OF THE CRITERION .....	41
2.1 The Source of Our Puzzle .....	41
2.1.1 The JR Principle .....	41
2.1.2 Contemporary Theories and the Denial of JR .....	47
2.2 The Problem of Easy Justification .....	52
2.2.1 The Problem from Closure .....	52
2.2.2 Easy Justification and Epistemic Bootstrapping .....	57
2.2.3 It Isn't a Compliment When Someone Calls You Easy .....	65
2.2.3.1 Running in Circles? .....	66
2.2.3.2 Failure Is Not an Option .....	69
2.2.3.3 I'm Sorry but I Have a Prior Commitment .....	74
2.3 The Problem of the Criterion .....	75
2.3.1 Justification: From Easy to Impossible! .....	75
2.3.2 Easy Responses to "Easy" Skepticism .....	79
CHAPTER THREE: PROPOSED SOLUTION TO OUR PUZZLE .....	83
3.1 Dialectical vs. Epistemological Failure .....	84
3.2 A Reidian Inspired Rejection of JT .....	88
3.3 The Wright Way of Rejecting JT .....	103
3.4 Sosa, Levels of Justification, and Coherence .....	109
3.5 Restricting Evidential Support .....	122
3.5.1 Vogel on Rule-Circularity .....	122
3.5.2 Weisberg's Probabilistic Approach .....	125
3.6 Lessons Learned .....	128
CHAPTER FOUR: FACTS AND EVIDENCE .....	130

4.1 A Fact Ontology of Evidence .....	131
4.1.1 Advantages of a Fact Ontology of Evidence.....	135
4.1.2 Challenging the Fact Ontology of Evidence .....	137
4.1.2.1 Intuitive Resistance and the Notion of False Evidence .....	139
4.1.2.2 Neta: Same Facts, Different Justification.....	143
4.1.2.3 Williamson: Evidence and Confirmation Relations .....	150
4.2 Factoring Accounts of Justifying Evidence.....	159
4.3 Conclusion .....	162
 CHAPTER FIVE: HAVING EVIDENCE, AWARENESS, AND JUSTIFICATION ....	163
 5.1 Evidence Possession and Awareness .....	164
5.2 Justification and Awareness: Part I.....	168
5.2.1 Appealing to the Cases.....	168
5.2.2 Radical Externalist Attempts at Accommodation.....	170
5.2.2.1 Norman, Nymoon, and Normina .....	175
5.2.3 The Subject's Perspective Objection (SPO) .....	185
5.3 Justification and Awareness: Part II.....	188
5.3.1 Three Forms of Evidence Possession .....	188
5.3.2 Motivating a Strong Evidence Possession Requirement.....	191
5.3.3 An Outline of a Theory of Prima Facie Justification .....	199
 CHAPTER SIX: THE IRREDUCIBILITY OF EPISTEMIC SUPPORT .....	204
 6.1 Examples of Epistemic Support.....	206
6.2 Epistemic Support is Not Subjective .....	210
6.2.1 The Subjectivist View(s) .....	210
6.2.2 Subjectivism as Epistemological Anarchy.....	216
6.2.3 A Blameless Appeal.....	219
6.3 Against Externalist Accounts of Epistemic Support .....	221
6.3.1 The New Evil Demon Objection.....	224
6.3.2 A Blameless Response.....	226
6.3.3 Attempts at Accommodation.....	228
6.3.3.1 Actual World Reliabilism.....	229
6.3.3.2 Home World Reliabilism .....	234
6.3.4 The Really New Evil Demon Problem.....	237
6.4 Epistemic Support: A Keynesian Approach .....	245
6.4.1 Keynesian Probability .....	246
6.4.1.1 Objectivity .....	246
6.4.1.2 Irreducibility .....	247
6.4.1.3 Internal vs. External Relations .....	249
6.4.1.4 A Priori Probability .....	251
6.4.2 A Keynesian Inspired Theory of Epistemic Support.....	252
6.4.3 What the New Evil Demon Really Shows.....	256
6.5 Objections to the Keynesian Approach .....	258
 CHAPTER SEVEN: AN ACQUAINTANCE THEORY .....	265
 7.1 Fumerton's Acquaintance Theory .....	265
7.1.1 Acquaintance .....	266
7.1.2 Correspondence .....	271
7.1.3 The Epistemic Regress Argument for Foundationalism .....	271
7.1.4 The Conceptual Regress Argument for Foundationalism.....	278
7.1.5 Acquaintance, Evidence, and Foundational Justification .....	278

7.1.6 The SPO, Acquaintance, and Bergmann’s Challenge.....	283
7.1.6.1 “He Has No Idea”.....	284
7.1.6.2 Changing the Scope.....	287
7.2 Moving to Fallible Foundations.....	289
7.2.1 Phenomenological Considerations.....	290
7.2.2 Foundational Justification, Uncertainty, and Fallibility.....	292
7.2.3 Moving from Correspondence to Probabilistic Relations.....	298
7.3 The Inferential Internalist Insight.....	303
7.4 Having Evidence and Justification: Stating the View.....	310
7.4.1 Facts and Evidence: Conditions (i) and (ii).....	311
7.4.2 Evidence Possession: Conditions (i) and (iii).....	312
7.4.3 The No Defeaters Clause: Condition (iv).....	312
7.5 Objections to the Acquaintance Theory.....	314
7.5.1 The Over-Intellectualization Objection.....	314
7.5.2 Etiological Objections.....	316
 CHAPTER EIGHT: EASY JUSTIFICATION AND SKEPTICISM REVISITED.....	 319
8.1 Reviewing Our Puzzle.....	319
8.1.1 Easy Justification.....	319
8.1.2 The Problem of the Criterion.....	323
8.2 Getting Re-Acquainted with the Acquaintance Theory.....	324
8.3 Basic Inferences, Deduction, and Easy Justification.....	324
8.3.1 Basic Deductive Inferences.....	325
8.3.2 Modus Ponens and Easy Justification.....	326
8.3.3 Getting Acquainted with the Way Out.....	330
8.4 Generalizing the Solution.....	337
8.4.1 Irreducible Epistemic Support’s Role in Our Solution.....	337
8.4.2 Defending Defeasible A Priori Evidence for Reliability.....	340
8.4.3 Induction, Perception, and Easy Justification.....	346
8.5 Meta-Level Issues.....	349
 CONCLUSION.....	 355
 BIBLIOGRAPHY.....	 360

## LIST OF FIGURES

### Figure

1.	Figure 1: Isomorphic Explanatory Relations .....	154
2.	Figure 2: Isomorphic Reliability Relations .....	156
3.	Figure 3: Isomorphic “Logical” Probability Relations .....	158
4.	Figure 4: Attending to Dots .....	269
5.	Figure 5: Acquaintance and Foundational Justification .....	309
6.	Figure 6: Acquaintance and Inferential Justification.....	310

## CHAPTER ONE

### SURVEYING THE EPISTEMOLOGICAL TERRAIN

Epistemology investigates philosophical issues pertaining to knowledge, justification, evidence, etc. In this dissertation I develop and defend a controversial set of epistemological views and illustrate how such a view resolves one of the most difficult puzzles for a theory of epistemic justification. Before doing this, however, I must place my work in context.

#### 1.1 Metaepistemology and Normative Epistemology

Following Richard Fumerton we can distinguish two distinct projects within epistemology: metaepistemology and normative epistemology.<sup>1</sup> Metaepistemology aims to provide an understanding of epistemological properties (or, depending on your metaphilosophy, words or concepts); achieving such an understanding involves bringing to light both essential and accidental properties of epistemological properties while also highlighting various interrelations amongst these properties. Normative epistemology, on the other hand, takes metaepistemological principles for granted and attempts to shed light on what we in fact do and do not know, what we do and do not have justification for believing, what we do and do not have evidence for, etc.

Examples of metaepistemological issues include the fallibilism/infallibilism, internalism/externalism, and foundationalism/coherentism/infinitism debates. Each debate concerns the nature of epistemic properties while ignoring the question of which beliefs do or do not have such properties. Alternatively, debates regarding skepticism occur at the normative level. Arguments for external world skepticism, for instance, attempt to show that beliefs about the external world lack a certain epistemic property such as being knowledge by

---

<sup>1</sup> See Fumerton (1995) chapter 1.

presupposing a necessary condition for knowledge and arguing that external world beliefs fail to meet this condition.

A majority of my dissertation is devoted to developing and defending various metaepistemological views but normative issues inevitably arise along the way. Ideally one might hope to develop a metaepistemology in a vacuum isolated from any normative considerations and then subsequently let the normative chips fall where they may. While I'm sympathetic that this would be ideal, I'm skeptical that normative epistemology can be completely excised from metaepistemology. The trickiest issue is: how should the fact that a metaepistemology leads to skepticism affect our assessment of it? That a view avoids skepticism is not always a reason to believe the view. Most externalist epistemologies avoid skepticism but the very ease with which they do so is incredibly implausible and might actually constitute a reason to reject these views. Similar remarks apply to various theories often classified as internalist. Nevertheless, I do hold that a metaepistemology leading to skepticism counts against the plausibility of the theory if only minimally; *ceteris paribus* a metaepistemology that avoids certain skeptical results is preferable. The issues here are incredibly complicated. My goal here is to simply state my views on these issues in a very rough outline. My take on the import skeptical results have for a metaepistemology will be further explored at the end of the next chapter.<sup>2</sup>

### 1.2 Leaving Knowledge by the Wayside

Within metaepistemology there are many properties that you might be interested in: knowledge, justification, rationality, evidence, probability, etc. Knowledge has, without a

---

<sup>2</sup> I've been hinting at Chisholm's (1977) problem of the criterion. See also Amico (1988), (1995); Cling (1994), (1997); and Fumerton (2008) for discussions of the problem of the criterion. The related issues concerning the proper role of skeptical arguments in epistemological inquiry are complicated and require more discussion than I can devote here. For particularly good discussions see Chisholm (1977); Greco (2000); and Fumerton (1995). Further explicit discussion of the issue can be found in 2.3.2 and at the end of my discussion of a Reidian inspired epistemology in section 3.2.



doubt, taken center stage in the history of epistemology. I leave discussion of knowledge by the wayside in order to concentrate on epistemic justification, rationality, and evidence. Despite the fact that this focus is becoming more common in contemporary discussions, such an idiosyncrasy cannot go without comment. I take justification to be the property of primary epistemological interests for two reasons.

My first reason for considering justification to be the central epistemological concern is that most epistemologists agree that justification is a necessary condition in an analysis of knowledge. S's having justification for believing *P* stands in a part-whole relation to the fact that S knows *P*. The reverse, however, does not hold. This asymmetry entails that an understanding of knowledge requires an understanding of justification but not vice versa.<sup>3</sup>

The second reason I consider justification to be of central importance can't be explained as quickly since it relies on a highly controversial metaepistemological view: infallibilism about knowledge (henceforth "infallibilism"). Infallibilism is the view that a subject S knows that *P* only if S's justification for believing *P* guarantees *P*'s truth.<sup>4</sup>

If infallibilism is correct then we likely don't know very much; we likely don't know any truths about the external world, the past, or the future.<sup>5</sup> However, given this strong

---

<sup>3</sup> Williamson (2000) has inspired a new wave of epistemologists who have adopted a kind of "knowledge first" epistemology that attempts to analyze justification in terms of knowledge. Discussing the arguments for the knowledge first approach would take us too far afield. At this point I merely note Williamson's dissent while noting that I find his arguments for taking knowledge as unanalyzable to be unconvincing. For further discussions of Williamson's knowledge-first epistemology see the papers in the volume edited by Greenough and Pritchard (2009).

<sup>4</sup> Defining infallibilism and fallibilism about knowledge is actually quite difficult. See Lehrer (1974) for a good discussion of complications with defining infallible belief.

<sup>5</sup> An infallibilist might attempt to avoid this result by incorporating more into her evidence base. A direct realist about perception that holds that physical objects are actually constituents of our perceptual experiences upon which we base our beliefs about the external world might be able to accommodate infallibilist knowledge of the external world; see Brewer (2002), (2013); Fish (2009); Johnston (2004); McDowell (1982), (1994); Pritchard (2012); and Smith (2002) for more on direct realism. Some of these papers focus on metaphysical direct realism, others such as the work by Pritchard and McDowell focus on a kind of epistemological disjunctivism such that we can have factive support for our perceptual beliefs (where, perhaps, this needn't rely on a kind of *metaphysical*

standard for knowledge, this skepticism is neither very surprising nor very disturbing. As Crispin Wright explains:

[K]nowledge is not really the proper central concern of epistemologico-sceptical enquiry. There is not necessarily any lasting discomfort in the claim that, contrary to our preconceptions, we have no genuine knowledge in some broad area of our thought – say in the area of theoretical science. We can live with the concession that we do not, strictly, know some of the things we believed ourselves to know, provided we can retain the thought that we are fully justified in accepting them.<sup>6</sup>

Such a thought doesn't strike us as a retreat so much as a candid admission of epistemic humility. One of the first reactions many students have when presented with Descartes' skeptical arguments isn't one of worry but rather triviality: "Of course we can't Know (said with special emphasis) that we aren't brains in a vat *but we still have plenty of justification for believing that there is a table in front of us and that we're not actually brains in vats.*" Descartes himself takes much of the bite out of his skeptical arguments by making this same kind of claim toward the end of the first meditation. We don't care much that our beliefs fall short of certainty as long as we're still rational and justified in holding our beliefs.

The previous thought is predicated upon an acceptance of infallibilism. Why accept such a strict requirement on knowledge? Initial motivation comes from linguistic intuitions concerning the fallibilist's commitment to the truth of so called "abominable conjunctions" concerning knowledge attributions:

1. I know my lottery ticket is a loser but I can't rule it out that the ticket is a winner.
2. Sally knows her lottery ticket is a loser but she can't, from her own perspective, rule it out that her ticket is a winner.
3. I know I'll be spending Christmas with my family, but it's possible that I'll get into a car wreck and die before Christmas.
4. Sally knows she will be spending Christmas with her family, but from her own perspective, it's possible that she will get into a car crash and die before Christmas.

---

direct realism and disjunctivism); and others focus on utilizing metaphysical direct realism to build a kind of epistemological direct realism.

<sup>6</sup> Wright (1991), p. 88.

If the fallibilist is correct and a subject can have knowledge despite her justification failing to guarantee the truth of her belief then some conjunctions of this form could be true. When we're explicitly presented with the sentences whose possible truth is *constitutive* of fallibilism the sentences strike us as absurd. And this gives some initial reason to think that our ordinary notion of knowledge isn't the fallibilist notion.<sup>7</sup>

Related are intuitive inconsistencies that arise between *fallibilist* knowledge attributions and practical behavior. Lotteries provide an especially useful illustration.<sup>8</sup> Most lotteries are such that I have justification that makes my belief that my ticket will lose all but certain. Assuming fallibilism, my belief that my ticket will lose is a prime candidate for something that I can know.<sup>9</sup> If I *know* my lottery ticket will lose why buy a ticket? Considering one's possible behavior *after* buying a lottery ticket is even more telling. Even assuming it was irrational to buy the ticket in the first place, if I *know* that my ticket lost why does it appear that throwing the ticket away before the winning numbers are announced

---

<sup>7</sup> David Lewis is the first person I'm aware of who makes this point. He pleads with his fallibilist opponents, "If you are a contented fallibilist, I implore you to be honest, be naïve, hear it afresh. 'He knows, yet he has not eliminated all possibilities of error.' Even if you've numbed your ears, doesn't this overt, explicit fallibilism *still* sound wrong?" (1996), p. 550. Our uneasiness when presented with the kinds of sentences whose truth being a conceptual possibility is constitutive of the fallibilism suggests that our ordinary notion of knowledge is an infallibilist notion. I take it that fallibilism isn't our ordinary notion but is rather a modification that people quickly move to in an attempt to avoid skeptical worries. However, a retreat to a vindication of the rationality of our beliefs seems as good of a response to skepticism that doesn't abandon our ordinary concept. Of course, many philosophers influenced by Craig's (1990) attempt to investigate our concept of knowledge by considering what role such a concept plays in our lives and what features the concept must have in order to serve that purpose will be skeptical that the infallibilist notion of knowledge really is our ordinary notion. What purpose would such a concept have if it so rarely applies? This is a serious concern but responding would take us too far afield. For now I'll just note that assertions using certain concepts might play important roles that don't depend on assertions involving those concepts being *literally* true. Even if most knowledge claims are literally false it might be incredibly useful to make such assertions where these purposes aren't as easily satisfied by assertions involving weaker concepts that accurately depict the situation. *Exaggeration is often incredibly useful.*

<sup>8</sup> See Hawthorne (2004) for an excellent extended discussion of these and many more issues surrounding how we think about knowledge in lottery cases.

<sup>9</sup> I wouldn't be surprised if it's statistically more likely that odd lighting conditions cause you to have a deceptive color experience than it is that you win the Powerball.

would be irrational? Why would it be irrational to sell a ticket that cost \$10 for \$0.25 to my friend? If I *know* my ticket is a loser I'd be crazy not to take that deal! But it isn't implausible to suppose that it's actually *irrational* to sell a \$10 lottery ticket for \$0.25 and this is indicative of fallibilism's *prima facie* implausibility. These first two worries amount to an appeal to intuition; the next two worries for fallibilism are more technical.

Fallibilism encounters a problem of arbitrariness in setting a degree of justification *both* necessary and sufficient for meeting the justification condition on knowledge.<sup>10</sup> Infallibilism amounts to the view that a subject has met the justification condition on knowledge iff  $\Pr(p|e) = 1$  ( $e$  represents the subjects evidence or justification). If we reject infallibilism, how high do we set a standard that is both necessary and sufficient for meeting the justification condition on knowledge?  $\Pr(p|e) \geq .99$ ?  $\Pr(p|e) \geq .75$ ?  $\Pr(p|e) \geq .5873$ ? Making  $\Pr(p|e) > .5$  necessary is non-arbitrary since surely knowledge requires that one's justification make the proposition more likely to be true than false, but we need a principled threshold such that meeting it is both necessary and sufficient. Any such threshold appears utterly arbitrary.<sup>11</sup>

Finally, in setting the threshold sufficient for meeting the justification condition on knowledge less than 1, fallibilism is committed to denying that knowledge is closed under

---

<sup>10</sup> In presenting this worry I assume that the reader is familiar with probabilistic notation. The basic idea is that probabilities can be represented on an interval from 0-1. An assignment of "0" means that the proposition has absolutely no chance of being true: the proposition is guaranteed to be false. An assignment of "1" means that the proposition has 100% chance of being true: it's guaranteed to be true. And an assignment of ".5" means that the proposition is equally likely to be true as it is to be false. The notation " $\Pr(p|e) = N$ " is used to represent the idea that the probability of the proposition  $p$  relative to some evidence or justification  $e$  is the number  $N$  on the 0-1 scale.

<sup>11</sup> Someone might think that this is just a Sorites problem and apply whatever solution they advocate to Sorites problems more generally to this issue. Two more recent attempts to provide a non-arbitrary way of setting the degree of justification necessary and sufficient for meeting the justification condition on knowledge are an appeal to pragmatic encroachment and contextualism. See Fantl and McGrath (2002), (2007), (2009) for sustained defenses of pragmatic encroachment. See Cohen (1988), (1998), (1999), (2005), (2008); DeRose (1992), (1995), (1996), (2009), and Lewis (1996) for sustained defenses of contextualism.

known entailment. A majority of epistemologists accept a principle akin to the following: if  $S$  knows that  $P$  and knows that  $P$  entails  $Q$  then  $S$  is *in a position* to know that  $Q$ . Principles of this kind are “closure principles” (you can have closure principles for both knowledge and justification). Philosophers who accept this principle (or some modification of it)—a majority of epistemologists<sup>12</sup>—often speak of knowledge as being “closed under known entailment”. Many have even suggested that closure is so plausible that an epistemological theory  $T$  being committed to a denial of knowledge closure is tantamount to a *reductio ad absurdum* of  $T$ .

Given the widespread acceptance of closure, it’s curious that fallibilists often remain silent on the consistency of closure and fallibilism. That fallibilism must deny closure follows from elementary facts about probability theory. Fallibilism is committed to the existence of a threshold  $N$  such that  $1 > N > .5$  and  $\Pr(p|e) \geq N$  is necessary and sufficient for meeting the justification condition on knowledge. But we need only imagine that  $S$ ’s justification for probabilistically independent true propositions  $P$  and  $Q$  is right at the threshold. If this is the case then  $\Pr(p \& q|e) < N$  and so  $S$  isn’t in a position to know the conjunction  $P \& Q$  even if  $S$  infallibly knows that  $P$  and  $Q$  entails  $P \& Q$ .<sup>13</sup> Therefore, the basic tenets of fallibilism are tantamount to a rejection of knowledge closure.<sup>14</sup>

---

<sup>12</sup> Deniers of closure include Audi (1988); (1991); Drestke (1970); (2005); Nozick (1981); and Schechter (2013). See Feldman (1995) and Hawthorne (2004) and (2005) for good defenses of closure principles.

<sup>13</sup> Similar instances will arise even in cases where  $P$  and  $Q$  aren’t right at the threshold; the assumption just makes the point easier to illustrate. Similarly, the assumption that the subject *infallibly* knows the entailment relation is also inessential. Allowing fallibilist knowledge of an entailment just creates more potential instances of knowledge closure.

<sup>14</sup> Even if contextualism or pragmatic encroachment solve the arbitrariness problem they don’t help solve the closure problem. Failures of closure will arise any time that context or pragmatic features set the threshold below 1. Once the threshold is set the problem arises. It’s therefore curious that one of the claimed advantages of contextualist theories as opposed to traditional relevant alternative theories was that the former but not the later can accept closure.

These arguments for infallibilism are underdeveloped<sup>15</sup>, and I want to stress that my goal isn't to give a vigorous defense of this metaepistemological presupposition. My goal is to give some initial motivation so as to explain my focus in this dissertation on justification rather than knowledge. If my commitment to infallibilism is correct, justification is clearly the epistemological property of most significance to our lives.

However, even if infallibilism turned out to be incorrect I'd be willing to let the motivation for focusing on justification rest solely on two claims. 1) Understanding knowledge requires an understanding of justification and not vice versa. 2) Justification is a concept that is interesting enough in and of itself so as to warrant a dedicated investigation.

### 1.3 Internalist Theories of Epistemic Justification

One metaepistemological debate concerning justification has taken center stage over the past 50 years: the internalism/externalism debate. An extended discussion of justification that didn't situate itself within this controversy would be seriously lacking and so I must take some care in laying out the issues. Another reason that it's important to understand what is at issue between internalists and externalists is that in chapter two I explain—following Stewart Cohen (2002)—that a problem of easy justification often associated with externalism afflicts paradigmatic versions of internalism as well; a solution to our puzzle requires more finesse than simply rejecting externalism.

Throughout my dissertation the issues underlying the internalism/externalism debates, while often not explicit, are lurking in the background. In my development and defense of my preferred theory of justification it will become apparent that I align myself very closely with the arguments and concerns traditionally motivating internalism. However,

---

<sup>15</sup> See Bonjour (2010); Dodd (2011); and Unger (1971) and (1975) for defenses of infallibilism.

the internalism/externalism distinction has been made in many different ways and the theory I defend might be a version of internalism on some characterizations but not others.

What exactly is internalism? Internalist theories hold that justification must in some sense be internal to the believer. Externalist theories make justification depend (to some extent) on factors external to the believer. There will be as many different interpretations of internalism as there are of what it means for justification to be internal to a believer. Thus, it's probably best to talk in terms of multiple internalism/externalism debates corresponding to each of the ways philosophers attempt to characterize the notion of justification being internal rather than external to the believer.

### 1.3.1 Access Internalism

A traditional way of distinguishing internalist and externalist theories of justification is in terms of “access” requirements. Let us stipulate that a “J-factor” of a subject S’s belief that *P* is anything that is relevant to S’s having justification for believing that *P*. Accessibilism construes what it is for justification to be internal to the believer in terms of the subject’s having access to a belief’s J-factors or, perhaps, to a limited set of J-factors. Accessibilist theories of justification are those that impose the following access requirement:

**AR** A subject S’s belief *B<sub>p</sub>* is justified *iff* (i) there is some J-factor X that contributes to *B<sub>p</sub>*’s justification and (ii) S has access to the fact that X obtains.<sup>16</sup>

Troubles quickly arise, however, when we consider what it means to have “access” to a J-factor. Often enough it’s clear that accessibilists think access is epistemic; to have access to X is to know or have justification for believing X obtains. More specifically, however, the accessibilist construes the required access in terms of having *a priori* or *introspective* knowledge/justification for believing X obtains.

---

<sup>16</sup> My formulation of access internalism in terms of an acceptance of AR is inspired by and deeply indebted to Michael Bergmann’s characterization of internalism more generally in his (2006) chapter 1.

AR is proposed as a requirement applying to both foundational and to inferential justification. Thus, accessibilism faces a threat of conceptual circularity. Our hope is to give an analysis of justification and identify conditions constitutive of justification. Imposing an access requirement in such an *analysis* makes S's meeting condition X and S's having access to X's obtaining part of this analysis. Once we make the epistemological nature of access explicit, it's clear that one is analyzing justification in terms of justification!

Accessibilists must propose AR as a requirement that any adequate theory of justification must vindicate but that isn't part of the analysis. In other words, in order to avoid vicious conceptual circularity, access internalism must construe AR as a *synthetic* necessary truth. Accessibilists hold that any adequate analysis of justification in terms of a belief's meeting conditions X, Y, and Z must be such that S's having justification for Bp entails that S has justification for believing that X, Y, and Z obtain. Having this justification, however, isn't to be construed as constitutive of S's justification for believing the first-order belief Bp. Similarly, "*P* is true" entails "*P* is true is true" but no one would dare suggest that this ought to be included as a part of an analysis of truth.

Versions of accessibilism can be separated along two fault lines. First, we can distinguish stronger and weaker versions on the basis of whether they require access to all or only some J-factors. Does justification require global or merely local access to J-factors? And, second, we can distinguish stronger and weaker versions depending on whether the required access is current/actual or merely potential (i.e. the ability to access the J-factor).

The most common objections to global accessibilism are that it's too demanding, likely leads to skepticism, and engenders vicious regress. Let's consider the first two related worries. A J-factor is anything that is in any way relevant to the justification of a belief. Every fact that played a causal role in the production of my belief is relevant to the beliefs being justified. Surely we don't want to demand that a person have knowledge of all the causal facts that led to her belief in order to have justification. And surely we don't want to demand that a subject have *introspective* or *a priori* justification for beliefs about these causal facts. It



would seem that hardly anyone (if anyone) ever has introspective or a priori knowledge of all the causal facts leading to any of their beliefs; global access internalism would therefore provide too easy a victory for global skepticism.

It's tempting to think that what global accessibilists really want isn't access to every fact that played any kind of role in the subject's having a justified belief but rather as advocating a kind of JJ thesis such that having justification for a belief B entails that one has justification for believing that one has justification for believing B. This, however, quickly leads to regress. In order to have justification for believing B I must have justification for the belief B\* that I have justification for believing B. In order to have justification for believing B\* I must have justification for the belief B\*\* that I have justification believing B\*. In order to have justification for believing B\*\* I must have justification for the belief B\*\*\* that I have justification for believing B\*\*. This continues *ad infinitum*.

Consider what the belief B\*\*\* actually amounts to. B\*\*\* is the belief that *I have justification for believing that I have justification for believing that I have justification for believing B*. I already have a difficult time entertaining let alone believing this proposition at this third level of complexity. If global access internalism is true then justification for any belief will require that one has justification for believing an infinite set of beliefs of infinitely increasing complexity. It's difficult to convince yourself that you could possess justification for a belief that occurs at the 139<sup>th</sup> iteration when, psychologically, the proposition is too complex for you to even entertain. The regress produced by requiring global access is vicious and leads to the view that justification is humanly impossible... a hard pill to swallow.

Note that this skeptical result also threatens a global accessibilism that appeals to *potential* access.<sup>17</sup> The difficulty isn't that global access requires that people have an infinite

---

<sup>17</sup> This depends on what interpretation of "potential" we're using. How are we to understand the modal status involved in the "potentiality"? If potential access merely involves *logical possibility* then potential access surely would help avoid the viciousness of the regress. Such a view, however, would seem trivial. We could identify justification with almost anything—including reliability, proper function, etc—and having justification for a belief would entail that we have

number of beliefs but that it requires that people have justification for an infinite number of beliefs of *infinitely increasing complexity!* It's psychologically impossible for us to entertain most of these beliefs and so it doesn't appear that we even have potential justification for these beliefs in any interesting sense.<sup>18</sup>

Global accessibilism is often interpreted as advocating a JJ principle. I'm not sure, however, that this is an accurate reading of some global accessibilists. Global accessibilism requires that a person have access to all of a belief's J-factors. There is an ambiguity here. On a strong reading justification requires access to the fact that J-factor X obtains (i.e. the subject is justified in believing *that the J-factor X obtains*). On a weaker reading justification requires that any J-factors X is such that the subject has access to the fact that X obtains

---

potential access to the fact that we have justification for that belief. There is no *contradiction* in supposing that we have justification for believing a belief has the relevant features.

<sup>18</sup> I'm only introducing some traditional worries for global accessibilism. I don't wish to commit myself to the claim that global accessibilism is committed to a vicious regress (at least not in this introduction). Three avenues for responding to the regress seem promising. First, one might make the common distinction between propositional and doxastic justification. Propositional justification amounts to a subject *having* justification for believing a certain proposition. The subject needn't actually *believe* the relevant proposition. Doxastic justification, on the other hand, amounts to a subject having a belief that *is* justified. This requires both that the subject has the relevant belief and that the belief is *based* on that which propositionally justifies it (see 1.7.2 for more on this distinction). Applying the global accessibilists requirement solely to propositional justification isn't as worrisome since having propositional justification doesn't require a subject actually believes the relevant proposition; in which case, *prima facie*, the fact that the proposition is too complex for the subject to entertain isn't *as* worrisome. Second, an advocate of global accessibilism might argue that all of the higher order beliefs are justified by the same fact F and that a subject needn't be able to entertain all of the higher-order propositions they need only see that F would justify them. Such a subject might have justified *de re* beliefs that the higher-order propositions are true despite lacking the corresponding *de dicto* beliefs. Finally, one might provide an interpretation of the modal status involved in *potential* access that makes it more interesting than mere logical consistency. One might attempt to understand potential access in terms of its being true that *if one were able to entertain the higher-order propositions then one would have justification for believing them*. The truth of this counterfactual is assessed by looking at the possible world most similar to the actual world but where the antecedent is true; if the consequent is true in this possible world then the counterfactual is true (talk of "possible worlds" is metaphorical). If the counterfactual is true then it would seem that everything relevant to the justification of the high-order propositions was already present in the actual world *accept* the ability to entertain the proposition. I actually find this last interpretation of global access internalism quite attractive. However, difficulties arise for this way of understanding the JJ thesis if we require for justification that a subject be aware of the evidential relation that holds between her evidence and a proposition. And I do (see chapters 5 and 7).

without necessarily having access to X's status as a J-factor. The weak reading doesn't commit the global access internalist to a JJ thesis. However, regress still threatens since the access to the J-factor is *itself* a J-factor. And so global accessibilism entails that if a subject has justification for believing B then there is a J-factor X and S has access to X. This access is a J-factor and so S must have access to the fact that S has access to X. Again, this is a J-factor and S must have access to the fact that S has access to the fact that S has access to X. The regress rears its ugly head once again.<sup>19</sup>

The threat of regress may move an accessibilist to a more modest view that only requires access to some J-factors. You might hold that a subject must have access to the evidence for her belief but not to the adequacy of this evidence. Dana's having justification for believing that her nephew Maxwell is playing with blocks might require that she has justification for believing her perceptual evidence (i.e. her experience as of Max's playing with blocks) exists without requiring access to this evidence making her belief probable. She would then need evidence for her introspective belief that she has this perceptual experience and she must have access to this evidence. This, however, isn't likely to result in a vicious regress since there is no guarantee that these evidential states will be of infinitely increasing complexity. Actually, the perceptual and introspective beliefs will likely be justified by the same evidence. Dana's evidence for her perceptual belief is the perceptual experience. As such, this local access internalism entails that Dana only has justification for her perceptual belief if she has justification for believing the perceptual experience is present. Her evidence for her introspective belief might again be her perceptual experience. That this perceptual

---

<sup>19</sup> One could avoid this wrinkle by redefining J-factors not as *anything* that is relevant to the justification of one's belief but as conditions *constitutive* of a belief's justification. And as I explained earlier, the access internalist ought to understand AR as a *synthetic* necessary truth so that access isn't a *constituent* of one's justification for the first-order proposition. However, the regress is likely to arise for even this modified version of global access internalism since part of what constitutes justification isn't just your evidence for the belief but also the adequacy of this evidence, i.e. the fact that your evidence supports your belief. And this quickly leads to regress. The resulting regress is the regress presented by Bergmann (2006) chapter 1 for a view he calls Strong Awareness Internalism.

experience justifies Dana in believing she has this perceptual experience *entails* that Dana has justification for believing such evidence is present. That one is justified in believing *P* entails that one is justified in believing *P*. No regress arises.

At the end of chapter 5 I argue that theories that require access to evidence without requiring any sensitivity to its adequacy fail to provide a subject with the kind of assurance paradigmatic of a philosophically interesting concept of justification. For now I only note that requiring so little access is unlikely to get at the heart of the internalism/externalism debate. As Richard Fumerton notes in his discussion of access requirements:

[A]n externalist who accepted such a view would stay an externalist provided that the access referred to was still given a paradigmatic externalist understanding.<sup>20</sup>

Two incredibly influential externalist views appear to do just this, i.e. put access requirements on justification where this access is understood in externalist terms.

The local version of accessibilism presented earlier is directly inspired by William Alston's (1988) indicator reliabilism. According to this view a subject *S* must have adequate evidence *E* for believing a proposition, *S* must have access to *E*, but *S* needn't have access to the fact that *E* is adequate evidence for *P*. Alston's view, however, is clearly a version of externalism since he gives externalist analyses of the nature of one's access to evidence and of the *adequacy* of evidence. Access to evidence *E* is understood in terms of having a process that takes a reliable indicator as input and outputs the belief that *E* is present. And the adequacy of evidence is understood in terms of there being a reliable connection between the presence of *E* and the truth of *P*.

Another possible example of externalism with access requirements is Ernest Sosa's influential externalist theory which posits a bifurcation of justification into two levels:

---

<sup>20</sup> Fumerton (1995), p. 65.

animal-level justification and reflective-level justification.<sup>21</sup> At the animal-level Sosa's view is a virtue reliabilism without access requirements. S's belief that *P* has animal-level justification iff the belief is *adroit* (i.e. manifests S's intellectual competence where this is given a reliabilist treatment); animal justification doesn't require that a subject have access to the fact that this belief has its source in a reliable intellectual competence. Sosa, however, can't help but think that there is something of distinctive epistemic value missing if one considers the question "Do I know *P*?" and must either answer in the negative or withhold belief.

Suppose that, while consciously confident that *p*, one also considers, at that same time, whether one not only believes but knows that *p*. Exactly three options open up: one might say either (a) "No, I don't know that," or (b) "Who knows whether I know it or not; maybe I do, maybe I don't," or (c) "Yes, that is something I do know." ... [O]nly answer (c), of the three, entirely avoids disharmony within that consciousness at that time. If one has to give answer (a), or even answer (b), one thereby falls short, and one's belief that *p* itself falls short. That belief is then not all that it could be. One is not as well justified as one might be, epistemically.<sup>22</sup>

This thought leads Sosa to claim that *adroit* belief is only sufficient for a primitive kind of justification. Animal-justification merely represents the same kind of achievement as a thermometer's being sensitive to the temperature, it doesn't represent the kind of intellectual achievement that philosophers, and intellectually curious people more generally, seek.

However, from this primitive animal justification Sosa thinks we can pull ourselves up so as to gain a perspective on our beliefs that endorses our beliefs as the products of reliable intellectual competence. At which point we arrive at a kind of reflective justification more worthy of the title and that provides us with assurance that our beliefs are true. Reflective justification involves *adroit* belief *adroitly* noted. Gaining such a perspective on our beliefs, however, is understood on the externalist model presented at the animal-level.

---

<sup>21</sup> Sosa's discussion usually proceeds in terms of knowledge rather than justification. My discussion of his view here is rather brief. See chapter 3 section 3.4 for a more detailed discussion of Sosa's views.

<sup>22</sup> Sosa (2007), pp. 115-116.

We have reflective justification for *P* only when we have animal justification for *P* and animal justification for believing that *P* has met the conditions constitutive of animal justification (i.e. that *P* has its source in reliable faculties).<sup>23</sup> Nevertheless, Sosa's notion of reflective justification involves accepting a kind of access to the conditions constitutive of animal justification. I won't assess Sosa's incredibly sophisticated theory here—I discuss it further in chapter 3. Here I'm merely using it to illustrate that access requirements fail to fully capture the heart of the internalism/externalism debate.

### 1.3.2 Mentalism

An alternative way of carving the internalism/externalism divide has been proposed by Earl Conee and Richard Feldman. On their view internalism is identified with the thesis that they label “mentalism”: justification is fully determined by the subject's mental life

A mentalist theory may assert that justification is determined entirely by occurrent mental factors, or by dispositional ones as well. As long as the things that are said to contribute to justification are in the person's mind, the view qualifies as a version of mentalism.<sup>24</sup>

The justificatory status of a person's doxastic attitudes strongly supervenes on the person's occurrent and dispositional mental states, events, and conditions.<sup>25</sup>

If any two possible individuals are exactly alike mentally, then they are alike justificationaly, e.g., the same beliefs are justified for them to the same extent.<sup>26</sup>

Mentalism initially appears to makes good sense of the idea that justification is internal to the believer in that it's internal to the believer's mental life. Contingent external facts relating the

---

<sup>23</sup> Sosa's discussion of reflective knowledge/justification often fluctuates between a characterization of reflective knowledge as a kind of coherentism and at other times as merely apt belief aptly noted (i.e. simply adding more animal knowledge/justification at a higher-level). See chapter 3 section 3.4 for a more detailed discussion of Sosa's view.

<sup>24</sup> Conee and Feldman (2004), p. 55.

<sup>25</sup> Conee and Feldman (2004), p. 56.

<sup>26</sup> Conee and Feldman (2004), p. 56.

believer's mental states to the environment are irrelevant to the justification of her beliefs—below I discuss worries that arise in light of the recent popularity of content externalism.

Mentalism also appears to classify paradigmatic externalist theories correctly. Paradigmatic externalist theories make contingent facts about the subject's environment and causal history—facts independent of the subject's mind—responsible for justificatory differences. Two mental duplicates might differ justificationally on account of differences in contingent external facts. On externalist views mental facts may play a role in determining facts about justification but they do not *fully* determine these facts (further facts about relative frequencies, reliability, causation, etc. are relevant as well).

Mentalism also shows why an externalist could add accessibility requirements and remains a form of externalism. William Alston's indicator reliabilism puts a potential access requirement such that a believer must have access to psychological grounds that make her belief objectively probable. Despite imposing this access requirement, the view is still externalist since contingent facts (external to the subject's mind) about relative frequencies determine the "objective probability" and thereby affect facts about justification. Similar remarks apply *mutatis mutandis* to Ernest Sosa's notion of reflective justification.

However, it isn't clear that a mentalist characterization of internalism can go much further in capturing natural characterizations of views as internalist or externalist. Surely the epistemology presented in Russell's (1912) *The Problems of Philosophy* is a classic example of internalism. According to Russell all knowledge by description (i.e. knowledge of truths) is to be grounded in knowledge by acquaintance. For instance, my knowledge of the truth that I'm having an experience as of a red round object is due to my being directly acquainted with a red round sense datum. But consider someone like G.E. Moore who toyed with the idea that sense data might be mind-independent. Or consider the new wave of direct realists like Bayer (2012), Brewer (2002) and (2013), and Hobson (*forthcoming*) who hold that one might be directly acquainted with mind-independent physical objects! Or consider that Russell famously thought that one could be acquainted with mind-independent *universals*. Surely any

foundationalism grounded in direct acquaintance with such items ought to be classified as internalist. But should we characterize these *relational* states of being acquainted with a *mind-independent* sense datum, physical object, or universal as mental states? If we give a negative answer, these theories don't qualify as internalist on the mentalist construal. Alternatively, you might give a positive answer since these relational states include at least one mental relata. At which point we risk classifying too many theories as internalist. If acquaintance with an external fact constitutes a mental state in virtue of having a mental relata, it appears arbitrary to deny that a belief's being causally related to an external object or even being reliably produced wouldn't constitute a mental state since these also include mental constituents. Now we risk characterizing paradigm forms of externalism as internalist.

More worries for the mentalist construal arise from the increasingly popular idea that mental states "*ain't in the head*".<sup>27</sup> Philosophers have become increasingly enamored with the idea that the content of a subject's intentional states cannot be individuated without essential reference to the kind of environment the subject inhabits. A number of philosophers working on perception are now even suggesting that sensations themselves are to be construed as intentional states whose content is determined in an externalist fashion. Combining mentalism with this kind of externalism about mental content qualifies as a form of internalism on the mentalist construal but such a view would be a far cry from traditional internalist views.

A mentalist construal of internalism also appears to classify Timothy Williamson's (2000) epistemology as internalist, surely an odd result. Williamson's theory is content externalism on steroids. He suggests that knowledge itself is a mental state and that a subject's evidence consists of her knowledge (Williamson's E=K thesis). If a subject's evidence determines facts about justification, such a view holds that justification is fully

---

<sup>27</sup> See Putnam (1975).



determined by her mental states. If two subjects are exactly alike mentally then they know the same things, have the same evidence, and have the same justification for their beliefs. Such a view is a version of mentalism but it stretches credulity to characterize it as internalist.

Lastly, Sawyers and Majors (2005) have intertwined content externalism and process reliabilism in a way that satisfies the mentalist thesis. They hold that a subject's justification for believing *P* is fully determined by the mental process' home world reliability. A subject's home world is the set of possible worlds with content determining properties such that: were the subject placed in the environment it would have determined a mental life with the intentional content that the subject actually has. This definition of a home world guarantees that two subject's with the same mental life will be equally justified in their beliefs since the relevant reliability will *necessarily* be assessed relative to the same set of worlds. Such a version of process reliabilism is an externalist view *despite* satisfying the mentalist theses.<sup>28</sup>

There is a more fundamental problem for the mentalist construal. Not all mental states are relevant to justification. A mentalist theory must specify which mental states are relevant. Inclusive accounts might claim that all mental states play a justifying role but such views are incredibly implausible. Some desires, memories, fears, etc. might be buried so deep in the unconscious that they cannot become conscious even with extreme time and effort. Surely we don't want to suggest that these mental states confer justification. So which mental states are relevant to justification? Feldman suggests a rather restrictive view where justification is determined only by mental states of which one is currently thinking, where "thinking of" is construed broadly enough to include conscious feelings or appearances. But it now appears that the mentalist (insofar as he or she is an internalist) is attempting to

---

<sup>28</sup> Comesaña's (2010) evidentialist reliabilism merges a kind of indicator reliabilism with a view that indexes reliability to the actual world. Such a view might also satisfy mentalist supervenience despite being a clear example of externalism. Goldman (2008) suggests making this same move so as to satisfy internalist insistence that epistemic support relations are necessary.

delimit the mental states relevant to justification by identifying a feature  $\Phi$  such that, in some sense, having  $\Phi$  makes a mental state *available* to the subject so as to provide justification. It's this availability to the subject that is really important to *internalist* justification. The mentalist thesis doesn't really get at what the internalist thinks is important. Mentalism might look like a good characterization of internalism since fixing the mental facts also fixes which mental states are *available* to a subject. But what makes a version of mentalism internalist isn't the idea that justification is determined by mental states but rather by states that are in some sense available to the subject. According to mentalist versions of internalism the states that are available to a subject are always mental states. Nevertheless, the more a theory intuitively strikes us as allowing aspects of mental life (such as content determined in an externalist fashion) that aren't available to the subject to affect justification the more it strikes us as a version of externalism. This suggests that it isn't the mentality but the availability that is important to internalism.

### 1.3.3 Inferential Internalism

Richard Fumerton (1995), (2004), and (2006) has been a principle defender of a principle of inferential justification supposed to divide paradigmatic internalists and externalists. And so Fumerton distinguishes the following two views: inferential internalism and inferential externalism.

It's natural to think that many of our justified beliefs are based on and "owe" their justification to our competently inferring them from other beliefs. My justification for believing that the global climate is changing due to human activity is partly constituted by my competently inferring this from my belief that there is wide consensus that this is so among scientific experts. Part of my justification for believing that it rained last night arises from inferring this from my belief that the roads and my lawn are wet. Each case illustrates a notion of *inferential* justification where justification for a belief B is partly constituted by my having a further belief that acts as its basis.

Of course no one thinks I have justification for believing humans have contributed to climate change or that it rained last night *merely* in virtue of having these further beliefs. My beliefs that there is wide consensus about human contribution to climate change among scientific experts and that the roads are wet might be completely irrational and unjustified. Perhaps these beliefs are the result of a wild guess, self-deception, a reading on a magic 8-ball, or some other epistemically pernicious basis. Surely the irrationality of these beliefs infects the beliefs based upon them. Irrational beliefs can't inferentially justify further beliefs.

An inference from  $P$  to  $Q$  only provides inferential justification for  $Q$  if one can *competently* infer  $Q$  from  $P$ . Our discussion shows that competent inference minimally requires that a subject have justification for the basis of the inference, i.e. the belief that  $P$ :

**IJ1** S has inferential justification for believing  $Q$  on the basis of S's further belief that  $P$  only if S has justification for believing  $P$ .

IJ1 is a relatively uncontroversial principle accepted by epistemologists with wildly different views. Process reliabilists, proper functionalist, evidentialists, mentalists, access internalists, coherentists, infinitists, etc. will all likely accept IJ1.

More is required of competent inference than that specified by IJ1. Making an inference from a justified belief isn't sufficient for inferential justification, there must also be an appropriate connection between  $P$  and  $Q$ . Now the question arises: does inferential justification also require that the subject have *access* to the connection? It's here that the controversy arises between inferential internalists and externalists: the former give an affirmative answer and the latter a negative answer. The inferential internalist accepts (whereas the externalists rejects) a more demanding principle of inferential justification IJ2:

**IJ2** S has inferential justification for believing  $Q$  on the basis of S's further belief that  $P$  only if (a) S has justification for believing  $P$  and (b) *S has justification for believing that P makes probable Q.*<sup>29</sup>

---

<sup>29</sup> One will formulate the defining principle of inferential internalism differently depending on whether one takes probability relations or epistemic rules to be more fundamental.

We'll return to a discussion of inferential internalism in chapter 7. Right now I want to note the controversy and consider how the debate between inferential internalism and externalism relates to the internalism/externalism controversy more generally.

It's clear that paradigmatic versions of externalism reject clause (b) of IJ2. On process reliabilist accounts of inferential justification what matters is that the input belief to an inferential process be justified and that the process be conditionally reliable. Traditionally, reliabilists place no requirement on inferential justification that a subject has a justified belief about the legitimacy of the inferential process (i.e. to its conditional reliability). Similar remarks apply *mutatis mutandis* to the other most influential forms of externalism.

Unfortunately, it isn't clear that IJ2 captures a sharp distinction between internalist and externalist theories even in the case of inferential justification. Yet again, there is nothing preventing an externalist from adding access requirements to her account of inferential justification while remaining an externalist *provided that the access to the inferential connection is understood along externalist lines*. Such an externalist can even insist that some beliefs about inferential connections are foundationally justified. As long as justification for the relevant belief is given a reliabilist analysis, such a view would still be paradigmatically externalist.<sup>30</sup>

IJ2 is even controversial among internalists. It's natural to characterize the epistemological views of Tim and Lydia McGrew, Conee and Feldman, and Michael Huemer as versions of internalism but all reject the 2<sup>nd</sup> clause of IJ2. The debate concerning inferential internalism and externalism is a debate above and beyond the traditional internalism/externalism debate. However, what one says concerning the traditional internalism/externalism debate may play a significant role in determining what one says concerning inferential internalism and externalism.<sup>31</sup>

---

<sup>30</sup> Jack Lyons (2009) has proposed a process reliabilist account of inferential justification along these lines.

<sup>31</sup> Fumerton uses IJ2 to argue for a Keynesian view of a priori epistemic probability modeled on entailment. This view of probability underlying Fumerton's inferential internalism

### 1.3.4 Internalism and the 1<sup>st</sup>-Person Perspective

Each of these proposals for characterizing the internalism/externalism debate are very different. Is there anything unifying these distinct internalist/externalist debates? I don't think we can give a precise definition of internalism but there is a common theme stressed by internalists: the 1<sup>st</sup>-person perspective. Internalist epistemologies are entirely egocentric.

Consider the skeptical predicament as it arises in Descartes' First Meditation. Descartes takes on a 1<sup>st</sup>-person evaluation of his beliefs as an example for us to follow. When I follow Descartes' lead I enter into the following kind of 1<sup>st</sup>-person reasoning: "I've encountered much disagreement with others and I've even found out that many of the claims I previously believed to be true have turned out to be false. Might I not unknowingly be in the same situation regarding my current beliefs? I take each of my beliefs to be true, but what reason do I have to accept their truth? Why should I think it unlikely that, as with my previous beliefs, unbeknownst to me, they are false? For instance, what reason do I have for thinking I'm not currently dreaming that I have a body with two hands?"

Descartes' skeptical worries arise from reflecting on his own beliefs and reasons. He then embarks on a project of deciding which of his beliefs he epistemically ought to maintain and which he ought to abandon. The internalist is convinced that all one has to go on in response to such reflection are the materials included in one's own 1<sup>st</sup>-person perspective. When deciding what to believe on any given matter I can't escape the egocentric perspective and it's these materials that *necessarily* determine the epistemic status of my beliefs. Externalists, on the other hand, hold that facts about the rationality of my beliefs aren't fully determined by my egocentric perspective. The rationality of my beliefs depends on contingent facts external to my perspective, e.g. the *de facto* reliability of my faculties.

---

actually does play a significant role in separating internalist and externalist theories of justification. See chapter 6.

Internalists want to restrict the materials that can make a difference to justification to those that make up a subject's 1<sup>st</sup>-person perspective. This is what characterizes internalist views. A sharp distinction, however, would require distinguishing what is and is not included in a 1<sup>st</sup>-person perspective. Thus, we get different internalism/externalism distinctions when different philosophers attempt to characterize the 1<sup>st</sup>-person perspective differently.

Accessibilists understand the egocentric perspective as consisting of facts, states, properties, events, etc. to which the subject has a priori or introspective access. Why privilege these kinds of access? Once we understand the connection between internalism and the 1<sup>st</sup>-person perspective this restriction intuitively makes sense. One intuitive way of trying to distinguish my perspective from yours is in terms of what I have privileged access to. Moreover, this also explains why, despite the access requirements, Alston's and Sosa's theories shouldn't be thought of as internalist. Even if the materials making up a subject's egocentric perspective play a role in determining justification on these theories, such materials don't *fully* determine facts about justification. Whether these materials justify believing *P* depends on whether there is a *contingent* reliable connection between having such a perspective and *P*'s truth.

Alternatively, you also get a clear rationale for a mentalist characterization of internalism. It's initially plausible that my perspective consists of my private mental states. Non-mental items such as physical objects, other persons, the frequencies with which certain things are related, etc. only make a difference to my perspective if they make a difference to my mental states. Moreover, once one sees that internalism really concerns the connection between justification and the egocentric perspective it becomes clear why mentalist views that incorporate elements of content externalism (e.g. Williamson's knowledge first view) fail to qualify as internalist. If we admit that mental states have content determined externally and we allow this broad content to play a justificatory role, it intuitively appears that justification is no longer fully determined by materials within an egocentric perspective.

The connection between internalism and the 1<sup>st</sup>-person perspective also makes sense of why inferential internalism is a kind of “internalism.” The motivations for this view arise from a consideration of a subject’s egocentric perspective. Fumerton explains that IJ2:

[P]lays no favorites, recognizes no special burdens of proof. The astrologer and the astronomer, the gypsy fortune-teller and the economic forecaster, the druid examining entrails and the physicist looking at tracks in cloud chambers are all expected to have reason to believe that their respective evidence makes probable their conclusions if the conclusions are to be rational.<sup>32</sup>

If the physicist infers that certain particles are behaving in a certain way on the basis of various tracks left in a cloud chamber but lacks justification for believing the latter makes probable the former then her perspective on the latter belief appears no better than the astrologer’s perspective on her predictions. And if justification is determined by a subject’s egocentric perspective then we should treat the rationality of the physicist’s and astrologer’s beliefs as on a par epistemically. The “internalist” view is motivated by an appeal to what is going on in a subject’s 1<sup>st</sup>-person perspective. Internalism is egocentric epistemology.

Again, I’m not attempting to give a precise definition of internalism. I’m pointing to a general theme closely linked with the internalism/externalism debate. Each of the characterizations of internalism discussed earlier is attempting to capture a sense in which justification is fully determined by the egocentric perspective. However, I throw up my arms and admit that don’t have a characterization of what the 1<sup>st</sup>-person perspective is that might improve upon the intuitive notion. Nevertheless, I think it’s clear that the attempt to develop a fully egocentric epistemology is a unifying theme amongst internalist theories.

#### 1.4 Externalist Theories of Epistemic Justification

We can also distinguish various different kinds of externalism. Access externalism holds that a subject *S*’s having justification for believing *P* doesn’t entail that *S* has

---

<sup>32</sup> Fumerton (1995), p. 41.

justification for believing that some or all J-factors for *P* obtain. There is anti-mentalism which allows that various non-mental states can make a difference to the justificatory status of a belief. And there is inferential externalism which rejects clause (b) of IJ2.

Yet again, however, the unifying theme of externalist theories is the idea that facts about justification are not fully determined by a subject's 1<sup>st</sup>-person perspective. On most versions of externalism what is most important to justification is the existence of various nomological or causal relations which will determine facts about justification whether or not these relations make it into a subject's 1<sup>st</sup>-person perspective. It will help us get a grip on this idea if we briefly consider a few examples of traditional externalist theories.

Process reliabilism is still considered a prototypical externalism. Crudely, reliabilism holds that a belief is justified iff the belief is the output of a reliable process.<sup>33</sup> The *de facto* reliability of perception is (absent defeaters) sufficient for the justification of the resulting beliefs about various physical objects. A similar theory is Alston's indicator reliabilism that I've already discussed. Virtue reliabilism has also become very popular these days but, in the end, I'm unsure whether such a view is any different than process reliabilism. On such a view a belief is justified if it manifests an intellectual virtue. Intellectual virtues are character traits involved in forming beliefs and the virtues as opposed to vices are defined in terms of their reliability.<sup>34</sup>

---

<sup>33</sup> This is a bare bones sketch of process reliabilism. A fully developed version of the view would need to modify the view so as to distinguish between foundational and inferential justification, explain how defeaters work, address the "generality" problem, etc. For further developments and defenses of reliabilism more generally (including process reliabilism, indicatory reliabilism, virtue reliabilism, and attempts to respond to various objections) see Alston (1986), (1988), (1995), (2006); Bach (1985); Beebe (2004); Comesaña (2002), (2006), (2010); Greco (1999), (2000), (2006), (2010); Goldman (1979), (1986), (1988), (1993), (2009), (2011); Heller (1995); Henderson and Horgan (2006), (2011); Leplin (2007); Levin (1997); Littlejohn (2009); Lyons (2009); Riggs (2002); Schmitt (1992); Sosa (1991), (2007), (2009); and Swain (1981). For critical discussion of reliabilist theories of justification and knowledge see Bonjour (1980), (1985); Cohen (1984); Conee (2013); Conee and Feldman (1998), (2004); Feldman (1985); Foley (1985); Fumerton (1995), (2006a), (2006b); McGrew & McGrew (2007); Moser (1989); Vogel (2000), (2008).

<sup>34</sup> See Greco (2000), (2006), (2010) and Sosa (1991), (2007), (2009).



Another popular externalist theory is proper functionalism. This view is in the spirit of reliabilism but proposes a few modifications. Proper functionalism holds that a belief is justified iff (i) the belief is produced by a faculty that was designed for the purpose of forming true beliefs and avoiding false beliefs, (ii) the faculty is reliable relative to the environment for which it was designed, and (iii) the faculty is functioning properly (i.e. in the way it was designed to function).<sup>35</sup>

Safety and sensitivity principles have also become extremely popular among externalists. Crudely we might explain the principles as follows: (SAFETY) *S*'s belief that *P* is safe just when *if S* were to believe *P* then *P* wouldn't be false and; (SENSITIVITY) *S*'s belief that *P* is sensitive just when *if P* were false then *S* wouldn't believe *P*. Such principles are usually used in externalist analyses of knowledge but one could utilize these notions in one's theory of justification as well. Justification might be a result of how safe or sensitive a subject's belief is (where, metaphorically, this might be measured by how far out in the space of possible worlds one would have to go in order for the relevant subjunctive conditional's antecedent to be true and consequent false).<sup>36</sup>

### 1.5 Why Internalism?

I've taken care to spell out a number of internalist/externalist controversies. The issues discussed here are lurking in the background throughout my dissertation. However, while I often make remarks pertaining to specific versions of internalism or externalism, I won't consider how to classify my own theory until the conclusion of the dissertation. As

---

<sup>35</sup> See Bergmann (2004), (2006), (2008) and Plantinga (1993). Plantinga's discussion is framed around what he labels "warrant" rather than justification.

<sup>36</sup> See Pritchard (2005), (2007), (2009); Sosa (1999a), (1999b) for defenses of a safety requirement on knowledge. See Dretske (1971); Nozick (1981); and Roush (2005) for defenses of sensitivity conditions on knowledge and see the edited volume by Becker and Black (2012) for a number of articles dedicated to a discussion of sensitivity principles in epistemology (many of these pieces argue against sensitivity in favor of safety). As I indicated in the main text these modal notions could easily be used in an account of justification.

much as I'm concerned with the issues and arguments surrounding the internalist/externalist debate, I'm not concerned with vindicating a characterization of my view as falling on one side or the other. I reject many paradigmatic externalist theories of justification and I utilize many standard internalist arguments. However, certain features of my view and certain things I say might at times look externalist. Speaking autobiographically, I think of my view as highly internalist but I'm not going to argue if someone wants characterize my view as externalist in certain respects. I'm more concerned with the considerations that have been brought forward as motivation for the views of self-proclaimed internalists and externalists than I am with whether my view falls squarely on one or the other side of the divide. So what are the motivations at play in internalism/externalism debates? I'll outline three historically prominent motivations for favoring internalist over externalist views.

One thought experiment that has figured crucially in the internalism/externalism controversy concerns the justificatory status of our beliefs were it to turn out that (unbeknownst to us) we are the victim of a Cartesian demon. Imagine a subject who is exactly like me, forms the same beliefs, has the same bases for his beliefs, but has the unfortunate luck to be the victim of a Cartesian demon. What it's like to be me is identical to what it's like to be the demon victim. Surely if the demon victim and I form the same beliefs on the same exact basis then the demon has no less justification for his beliefs than I do. Whatever verdict applies regarding the justification I have for my beliefs ought to apply equally to my counterpart in a demon world. It should be clear how such a thought experiment is supposed to support internalism. If my counterpart and I have the same justification for our beliefs then justification cannot depend on external factors such as reliability since these obtain in the actual world but not in the demon world.<sup>37</sup>

---

<sup>37</sup> See Cohen (1984) for the canonical presentation of this thought experiment.

Intuitively my counterpart and I have the same justification. Both the mentalist and the access internalist have a clear explanation for this intuition. It's intuitive to maintain that my counterpart and I have the same mental states and therefore, according to mentalism, the same justification. Similarly, the facts to which both my counterpart and I have a priori or introspective access would intuitively be the same and so, according to (most versions of access internalism), our justification would be the same.

Chapter six considers this thought experiment in much more detail. In the end I think the import of the thought experiment fails to support either mentalism or access internalism. The real import of a proper investigation of the new evil demon thought experiment is that evidential support is a necessary *internal* relation that holds between a set of evidence and a proposition.

A second thought experiment that has also been extremely important to the internalism/externalism debate is Bonjour's case of Norman the Clairvoyant:

Norman, under certain conditions which usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons for or against the general possibility of such a cognitive power or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power under circumstances in which it is completely reliable.<sup>38</sup>

This and related thought experiments have been used as a kind of argumentative foil for arguing against externalism and in favor of internalism. When we consider what Norman's belief would be like from his own perspective we have the intuition that he lacks justification for his belief that the President is in New York City. His belief is no better than an arbitrary hunch of wild guess from his own perspective. Thus, even though his belief might meet various external criteria such reliability, safety, sensitivity, etc. it isn't justified.

---

<sup>38</sup> Bonjour (1985), p. 41.

The access internalist explains Norman's lack of justification by appealing to Norman's lacking access to anything that might plausibly be construed as evidence for his belief. Alternatively, the mentalist can appeal to a lack of mental states that might provide a proper basis for Norman's belief. Norman lacks any conscious experience, beliefs about the existence of clairvoyance, beliefs about the reliability of clairvoyance, beliefs concerning a track-record of previous beliefs that have struck him this way, etc. Discussion of this kind of motivation for internalism will take center stage in chapter five.

Finally, Richard Fumerton has stressed the connection between having a philosophically interesting notion of justification and having a kind of assurance that one's belief is true. Consider again the Cartesian reflection discussed above in connection with the internalist's stress on the 1<sup>st</sup>-person perspective. We notice disagreements we have with others and that we can't both be right. We notice that people's beliefs are often caused by things that appear to be epistemically irrelevant; people are prone to wishful thinking, various biases, cultural influences, etc. Having noticed this we start to reflect on our own beliefs and hope to put them on some solid ground so as to arrive at some kind of assurance that our beliefs are at least likely to be true. My beliefs might be reliably formed, safe, sensitive, the products of properly functioning faculties, etc.:

But should I possess such justification, would it do me any good *at all* in satisfying my intellectual curiosity? Should I possess such justification, would it do me any good *at all* in giving myself the assurance that was shaken by my brief excursion into philosophy? Internalists think that possessing the sort of justification defined by externalists would be utterly irrelevant to possessing the kind of justification we seek when we try to put our beliefs on a secure footing – the kind of justification that gives us assurance.<sup>39</sup>

The new evil demon and Norman the Clairvoyant thought experiments might be construed as particular ways of illustrating this disassociation of justification and assurance on externalist theories.

---

<sup>39</sup> Fumerton (2006b), p. 98.

### 1.6 Why Externalism?

There is much to recommend externalist theories as well. One main concern of externalists is the idea that internalist theories appear to “over-intellectualize” justification. We often characterize “unsophisticated” subjects such as children as having justified beliefs. It’s natural to think that even young children are rational in believing that a toy is red, that there is a juice box on the kitchen table, that their Dad will pick them up after school, etc. But do children have a priori or introspective access to facts that might be used to justify such claims? If I were to ask my nephew what reasons he has to believe he isn’t dreaming could he give me any half-way plausible argument? It isn’t even clear whether my nephew could conceptualize the distinction between seeming to see a cup on the table from the state of seeing a cup on the table. However, it initially strikes us as a kind of intellectual hubris to suggest that my nephew has less reason or justification for believing that there is a cup on the table than I do.

Or consider a case from the movie *Rain Man*. Raymond is autistic but has an incredible ability to count hundreds of objects on sight. When 246 toothpicks fall on the floor Raymond looks at the toothpicks and immediately forms the correct belief. The beliefs he forms via his visual experience are highly reliable. After seeing his track-record counting such large quantities of objects some people will be tempted to characterize Raymond’s belief about numbers of objects formed on the basis of his visual experience as rational. Surely, our believing that there are three toothpicks on the floor on the basis of a visual experience as of three toothpicks is rational. And it’s tempting to think that Raymond’s autism has provided him with this same ability to form justified beliefs about greater numbers of objects. Again, however, it’s hard to imagine that Raymond has introspective or a priori access to facts that he could cite in support of his belief. And even a mentalist who rejects access requirements might have a hard time with this kind of case. Couldn’t I have the same visual experience as Raymond and lack justification for the same belief? The externalist has an incredibly natural suggestion here: Raymond’s belief is rational because

certain external virtues (such as reliability) obtain. When I have the same visual experience as Raymond, however, I don't have justification for this belief since my vision isn't reliable at identifying such large numbers of objects.

The previous worry amounts to the idea that unsophisticated subjects have justification even though they fail to meet the internalist requirements. A related worry is that even cognitively sophisticated subject's beliefs will turn out to be irrational by internalist lights. Many externalists are worried that internalism inevitably leads to skepticism. The worry is just that there aren't enough internal justifiers available to make our beliefs about the external world, other minds, the future, etc. rational.

A final motivation for externalism is that only it appears to give us a strong connection between justification and truth. Why do we care about forming justified rather than unjustified beliefs? Presumably, much of the value of justification derives from the value of having true beliefs. We take it that there is some connection between forming justified beliefs and forming true beliefs, and this connection at least partially explains the value of justification. But how are we to make sense of this connection on internalism?

Internalist theories make justification a function of the subject's perspective but there doesn't appear to be any necessary connection between a subject's perspective and the truth of her beliefs. The new evil demon thought experiment comes back to haunt the internalist. According to the internalist my counterpart in a demon world and I have the same justification for our beliefs. If the internalist is to avoid the charge of skepticism discussed in the previous paragraph then she'll admit that my demon world counterpart world has justified perceptual beliefs and this appears to sever any connection between having justified and having true beliefs. I can have justified perceptual beliefs that are all false!

The connection between justification and truth appears much less problematic on externalist theories. There is a notion of objective likelihood built into the externalist analysis

of justification. As such, externalism appears to vindicate the idea that there is a necessary conceptual tie between having justified beliefs and true beliefs.

### 1.7 An Evidentialist Framework

So far we've focused on the internalist/externalist debate concerning justification. Another prominent debate has been between evidentialist and non-evidentialist theories of epistemic justification, and some philosophers take this distinction to cut across internalist and externalist theories. My own view is that a distinction between evidentialist and non-evidentialist theses isn't a substantive distinction. Nevertheless, I'll suggest treating the evidentialist framework as a useful heuristic for developing a theory of justification, a heuristic I'll make use of in developing and defending a theory of justification.

#### 1.7.1 The Central Insight

Earl Conee and Richard Feldman have spent a large portion of their careers as the principal defenders of "evidentialism."<sup>40</sup> Sometimes they present the view as a thesis which *identifies* justification with a doxastic attitude's standing in a relation of *fit* with a subject's evidence. At other times they present the view as the thesis that justification *supervenes* on a subject's evidence. Each of these theses attempts to capture a general insight that I refer to as the evidentialist maxim or EM:

**EM** The epistemic justification of any doxastic attitude holds in virtue of the subject's possession of evidence and what this evidence supports.

My concern is with this general insight. I won't come back to a discussion of Conee and Feldman's more formal theses until the conclusion of this dissertation.

EM is strikingly plausible. It almost seems to be a mere truism! Imagine that Sally expresses a belief that *P*. We take *P* to be controversial and begin to wonder whether Sally is actually justified in believing *P*. Naturally, we ask for Sally's evidence or reasons. If we find

---

<sup>40</sup> See the collection of essays in Conee and Feldman (2004).

that her evidence supports  $P$ 's truth, we judge that her belief is justified. If we find that her evidence is neutral between  $P$  and  $\neg P$ , we judge that her belief is unjustified and that Sally has justification for withholding belief. Finally, if we find that her evidence actually supports  $\neg P$ , we again judge that her belief is unjustified and that Sally has justification for *disbelieving*  $P$ . As Hume exclaims, "a wise man... proportions his belief to the evidence."<sup>41</sup>

### 1.7.2 Propositional and Doxastic Justification

EM's plausibility is limited to a specific domain. In order to appreciate the import of EM we need to make a crucial distinction that I've been rather loose with so far: the distinction between a person *having justification* for a belief and a person's belief actually *being justified*. EM's proper application is only to the former notion.

Imagine a trial where the prosecution presents damning evidence that the defendant is guilty of murder. Marshall is the defendant's father and has been following the trial closely. Marshall is presented with all of the evidence and he recognizes that it provides overwhelming support for his son's guilt. However, the fact that Marshall loves his son dearly prevents him from actually forming the belief that his son is guilty. Marshall *has* good epistemic reasons for believing that his son is guilty despite the fact that he doesn't actually form this belief; Marshall *has justification for believing* that his son is guilty but, since there is no belief, Marshall doesn't have a *justified belief*. Epistemologists often distinguish these two notions by referring to the former notion as *propositional* justification and the latter notion as *doxastic* justification. The belief that his son is guilty is propositionally justified but not doxastically justified for Marshall.

EM is best construed as a thesis pertaining to propositional rather than doxastic justification. Possessing supporting evidence for a certain proposition plausibly provides the

---

<sup>41</sup> Hume ([1748] 1975), p. 110.



subject with propositional justification but doxastic justification requires that the subject also *base* her belief on this evidence.

Consider Lilly. Lilly is a juror in the aforementioned trial. She too has been exposed to all of the evidence and recognizes that this evidence provides overwhelming support for the defendant's guilt. Lilly had decided in advance that she would form a belief about the defendant's guilt solely on the basis of the reading provided by her Ouija board. When Lilly consulted her Ouija board it "told" her that the defendant is guilty and she thereby formed the corresponding belief. Lilly *has* good reasons to believe the defendant is guilty but, as a matter of fact, she has formed her belief *on the basis of* bad reasons. In such a situation, Lilly has propositional justification for her belief but the belief is still doxastically unjustified.

It's incredibly tempting to think that the concept of propositional justification is more fundamental than that of doxastic justification. Doxastic justification requires that a belief be formed on the basis of whatever provides the subject with propositional justification for her belief. This suggests that an analysis of propositional justification is a prerequisite for an analysis of doxastic justification. My focus throughout the dissertation will therefore be on propositional justification unless otherwise noted.<sup>42</sup> Thus, EM is a natural starting point for developing a theory of justification (though I don't begin this development until chapter four).

### 1.7.3 Treating Evidentialism as a Heuristic

EM is an attractive thesis when limited to the domain of propositional justification. Unfortunately, it's as uninformative as it is attractive. EM is, at best, a mere sketch of a

---

<sup>42</sup> I will only have a few things to say about the basing relation and doxastic justification, and in the rare occasions where I do discuss these notions most of my comments will be non-committal. However, I will note that I'm doubtful that a discussion of propositional justification can be as cleanly separated from discussions of doxastic justification as some epistemologists might wish. Understanding propositional justification will require us to consider the conditions that would enable something to be a *possible* basis of a subject's belief. All of this will become more clear in my discussion of evidence possession and awareness in chapter 5.

theory. A full evidentialist theory of justification requires detailed accounts of evidence, evidence possession, and evidential support.

More worrisome, however, is that a bare evidentialism (i.e. one that is detached from specific accounts of the relevant evidential concepts) doesn't rule out *any* substantial theory of justification! Consider the reliabilist and virtue theories that are often contrasted with evidentialism. Reliabilists hold that justification is a function of the reliability of a subject's cognitive mechanisms. Virtue theories hold that a necessary condition for epistemic justification is that a person engages in responsible inquiry. Neither of these views *need* be interpreted as denying EM. An advocate of such theories can simply build these conditions into her account of evidence, evidence possession, or evidential support!<sup>43</sup> Nevertheless, EM can play an important role in epistemological theorizing.

First, EM pre-theoretically distinguishes epistemic justification from other species of justification such as pragmatic, moral, economic, etc. In some sense of "justification", the fact that the consequences of believing *P* would be extremely valuable might provide me with justification for believing *P*. Such justification, however, would be pragmatic rather than epistemic. That the consequences of a belief would be especially valuable has no bearing on the belief's truth. Epistemic justification is a function of the evidence for the truth or falsity of the belief rather than the value of the consequences of forming the belief.

---

<sup>43</sup> This may be surprising to anyone familiar with the contemporary literature. When Conee and Feldman introduce evidentialism they explicitly contrast it with these reliabilist and virtue theories that make justification dependent upon "the cognitive capacities of people, or the cognitive processes or information gathering practices that led to the attitude" (2004), p. 84. The implication being that by making justification dependent on these factors, these theories are making justification dependent on something *in addition to* the support it receives from the subject's evidence. But this only follows if we're presupposing a certain theory of the relevant evidential concepts. I imagine that the reason Conee and Feldman insist on this contrast is because they are tacitly assuming their "mentalism" which really is inconsistent with these theories. However, an advocate of the reliabilist and virtue theories can accept EM but reject Conee and Feldman's mentalist version of evidentialism. Failure to see that the bare evidentialist thesis isn't incompatible with these views of justification has been exacerbated by the fact that many advocates of reliabilism and virtue theories often attack a view they label "evidentialism" without explaining (or, perhaps, even noticing) that they are also presupposing a certain theory of evidence, evidence possession, and evidential support.

Second, we can treat EM as a helpful heuristic. When considering hypothetical cases of justification we can ask: What is acting as evidence? What explains why these items are included in the subject's evidence? How does this evidence support this belief? This is a kind of top-down strategy for investigating the nature of justification. Alternatively, one could also adopt a bottom-up strategy. We can investigate the evidential concepts independently of our considerations of justification and *subsequently* apply these results (via EM) to our theory of justification.<sup>44</sup> I utilize both strategies in developing a theory of justification—in practice the two approaches often bleed together. EM suggests that there are three evidential notions we must investigate: evidence, evidence possession (i.e. a subject *S*'s *having* evidence), and evidential support.

### 1.8 A Preview of Things to Come

Having laid the conceptual groundwork so as to situate the dissertation within the contemporary literature it's time to turn to the meat of the dissertation.

In chapter 2 I introduce a puzzle for any theory of justification. Our puzzle revolves around the question of whether a belief source has the ability to confer justification absent *prior* justification for believing that source is reliable. If we give a negative answer to this question then we seem to allow subjects to use a source to illegitimately bootstrap their way to justification for believing the source is reliable. Even most internalist theories reject the principle that appears to allow for the illegitimate forms of reasoning and so avoiding the problematic reasoning isn't as simple as rejecting externalism. On the other side of the

---

<sup>44</sup> When we do this we might find that applying our investigation of the evidential concepts to our theory of justification (by filling them into EM) leads to various counterintuitive results. If so, there are two possible reactions: (i) revise one's account of the evidential concepts or (ii) abandon/modify the evidentialist maxim. I imagine that in many cases there won't be a principled reason to prefer one reaction over the other. Rather, in the end the choice will often hinge on personal preference and pragmatic considerations.

puzzle, if one requires prior justification for believing a source is reliable in order for that source to confer justification then a global skepticism threatens.

Chapter 3 considers various proposed resolutions of our dilemma. I find all of the proposals considered in chapter three wanting. But my discussion of these proposals plays an important role in the dissertation. Despite their failure, these proposals teach us some important lessons for the development and defense of a solution to our puzzle. Moreover, my discussion will motivate the idea that an independent investigation of the evidentialist concepts and their relation to justification will be crucial to resolving our puzzle.

In chapter 4 I put our puzzle on the backburner in order to begin an independent development of a theory of justification that I'll eventually argue resolves our puzzle. As I've indicated, I think the evidentialist maxim provides us with a nice heuristic for developing a theory of justification. So in chapter four I consider the ontology of evidence. At this point I'm concerned with evidence in general where this includes various different categories: scientific evidence, legal evidence, historical evidence, justificatory evidence, etc. The chapter develops and defends the thesis that evidence for or against a certain claim always consists of facts. An important part of the chapter develops a controversial metaphysics such that evidential relations hold between fact-proposition pairs that are isomorphic to similar relations that hold between proposition-proposition pairs. One of the main upshots of this metaphysics is that it can vindicate a theory of justification that retains much of the motivation for classical foundationalist views while avoiding many of the associated costs (though this isn't illustrated until chapter 7).

In chapter 5 I turn to a consideration of evidence possession and its relation to justification. The main upshot of this chapter is that I use cases such as Norman the Clairvoyant to argue for a general requirement on justification: S has prima facie justification for believing *P* iff S has what I call actually connected possession of supporting evidence for *P*. S has actually connected possession of evidence *E* for *P* iff S is aware of *E* and S appreciates the evidential connection between *E* and *P*. After arguing for this general

requirement I explain that any plausible externalist theory of justification can and ought to satisfy this requirement.

Chapter 6 turns to a consideration of evidential support. What is the nature of the relation  $R$  between a set of facts  $F$  and proposition  $P$  such that  $S$ 's awareness of  $F$  and an appreciation of  $R$  provides  $S$  with prima facie justification for  $P$ ? I begin the chapter by arguing against subjectivist and externalist analyses of epistemic support. I then argue for a non-reductive theory of evidential support which construes evidential support as a *sui generis* and *unanalyzable* relation. My development of the view models the theory of epistemic support on a Keynesian theory of probability where this relation is an internal relation knowable a priori. After articulating the view I respond to the following three objections: (i) some evidential relations are species relative; (ii) some evidential relations can only be known a posteriori; and (iii) such a view severs the truth-connection (the solution to this problem is intimately tied to the solution to the problem of easy justification).

In chapters 4-6 I consider the notions of evidence, evidence possession, and evidential support so as to lay down constraints on an adequate theory of epistemic justification. Chapter 7 brings these insights a full theory of justification. I begin by presenting Richard Fumerton's acquaintance theory of foundational justification, his inferential internalism, and illustrating how these satisfy my constraints from previous chapters. I then argue that both aspects of Fumerton's theory need to be weakened. First, his acquaintance theory of foundational justification needs to be modified so as to allow for fallible foundations.<sup>45</sup> Second, Fumerton's principle of inferential internalism needs to be weakened so that inferential justification requires only acquaintance with a probability relation rather than justification for believing that such a probability relation exists. I close the chapter with by responding to objections to the developed acquaintance theory.

---

<sup>45</sup> Fumerton himself flirts with allowing fallible foundational justification when one is acquainted with a fact similar to the truth-maker for one's belief. In chapter seven I will reject this attempt to allow for fallible foundations.

Finally, in chapter eight we return to our main puzzle. The two key features of the developed acquaintance theory for resolving our puzzle are: (i) epistemic support is a *sui generis* and *unanalyzable* relation such that we can know various truths about epistemic support a priori, and (ii) justification always requires an awareness of the evidential support relation but (importantly) this awareness is weaker than justified belief. The basic idea behind the solution is that one needn't have justification for believing that a source is reliable in order for that source to confer justification. And so we avoid the skeptical worries involved in our puzzle. Nevertheless, my acquaintance theory entails that whenever a source provides justification for a belief a subject will necessarily be aware of evidence for the reliability of her source. Vindicating this claim allows my acquaintance theory to avoid the problem of easy justification. When a subject goes through easy justification reasoning all it does is reveal a priori evidence for the reliability of her source, evidence of which she was *necessarily* aware when she formed the initial object-level belief via that source.

## CHAPTER TWO

### EASY JUSTIFICATION AND THE PROBLEM OF THE CRITERION

Our beliefs come from a variety of sources: perception, memory, introspection, testimony, reasoning, etc.<sup>46</sup> We naturally trust that these sources provide reliable information about the world around us and our place in it. A puzzle arises, however, when considering the following question: can these sources provide justification without our possessing *prior* justification for believing them to be reliable? A positive answer seems to make justification for believing that these sources are reliable trivial. Alternatively, a negative answer appears to make justification all together impossible. My dissertation dances a thin line and develops a theory of justification that makes justification neither trivial nor impossible. My goals in this chapter are to motivate the puzzle and explain why various “easy” answers to the puzzle are philosophically unsatisfying.

#### 2.1 The Source of Our Puzzle

##### 2.1.1 The JR Principle

Stewart Cohen has argued that any epistemology that allows a subject *S* to know *P* without prior knowledge that her source is reliable faces a problem of “easy knowledge.”<sup>47</sup> I develop the problem in terms of justification since I take it to be the concept of primary epistemological interest (see chapter 1).<sup>48</sup> Consider the following principle:

---

<sup>46</sup> Each of these sources can be broken down into more specific categories. Perception can be broken down into: vision, audition, tactition, olfaction, gustation, and possibly proprioception (perception of the motion and position of one’s body) and nociception (perception of pain).

<sup>47</sup> See Cohen (2002) and (2005) for his original presentation of the problem.

<sup>48</sup> Another reason for discussing the problem at the level of justification is that most epistemologists accept that justification is a necessary condition for, and constitutes part of an analysis of, knowledge. See Williamson (2000) for a dissenting view. This suggests that the problem of easy knowledge is ultimately reducible to the analogous problem that arises concerning the structure of justification.

**JR** Necessarily, a belief source  $\Phi$  provides S with justification for believing a particular deliverance (Bp) of  $\Phi$  only if S already has justification for believing that  $\Phi$  is reliable.<sup>49</sup>

Let's use the phrase "basic justification source" to refer to any belief source which can provide a subject with justification prior to the subject's having justification for believing that the source is reliable. Admitting even one basic justification source is tantamount to denying JR, and this supposedly leads to the problem of easy justification.<sup>50</sup> Appreciating the problem requires grasping JR's content which requires, in turn, a basic grasp of JR's two central notions: (i) belief sources and (ii) reliability.

Crudely, belief sources are those things that we can (and often do) use to form beliefs.<sup>51</sup> We all intuitively recognize perception, memory, introspection, testimony, induction, etc. as belief sources in the relevant sense. However, in each of these cases we might be interested in two distinct kinds of belief sources: (i) processes or (ii) grounds. A process can be thought of as a kind of program implemented in forming a belief. A subject S's ground for her belief is S's putative evidence, reason (I'm not using a purely causal sense of "reason" here), or basis for her belief.

Consider the *process* of visual perception. In visual perception an object's surface reflects wavelengths of light, the light hits your eyes, an image is projected on your retina, various rods and cones are stimulated, a complicated neural process produces a visual experience, and we usually form a belief about something in the nearby physical

---

<sup>49</sup> "Already" indicates *epistemic* rather than temporal priority.

<sup>50</sup> Basic justification sources *are not identical to sources of foundational justification*. Moreover, it's helpful to keep in mind that this notion is an epistemological rather than a psychological notion. That a source, as a matter of psychological fact, produces beliefs while a subject lacks beliefs about the source's reliability doesn't entail that the source can provide justification for those beliefs prior to the subject's possessing justification for believing that the source is reliable.

<sup>51</sup> This isn't intended to be a *definition* of belief sources. This is a necessary but insufficient condition for something's being a belief source. We "use" our brains in forming all our beliefs but we wouldn't say that our brain is a belief source in any relevant sense. My aim here is only to gesture at the notion of a belief source.



environment. When this process takes place and results in my believing that there is a computer in front of me, in some sense, I *use* this visual process in forming my belief. We can think of the process as a belief source. One might wish to restrict visual perception as a belief source to the parts of this process that are “seated in the subject.” Do I use the wavelengths of light in forming my belief in any relevant sense? Many philosophers would wish to identify only the parts of this larger process that take place “within the skin” (e.g. the image projected on the retina, the stimulation of my rods and cones, the activity in the primary visual cortex and the extrastriate visual cortical areas, the visual experience, etc.) as the process that constitutes the subject’s belief source. Reasons for and against identifying the narrow rather than the broad construal of the process as the subject’s belief source needn’t concern us here. What is important is that we can contrast this process notion of a belief source with another notion.

A visual process may be a belief source in some sense but it doesn’t constitute the subject’s *grounds* for her belief. The process isn’t the subject’s *reason* for her belief, where a reason is taken to be that which a subject *bases* her belief on.<sup>52</sup> This process can be defined as a functional relation and clearly a subject doesn’t *base* her belief on a functional relation. Moreover, even though each part of the process plays a causal role in the subject’s forming her belief we wouldn’t be tempted to say that a subject forms her belief on the basis of each part. I don’t base my belief that there is a computer in front of me on the stimulation of my rods and cones. I don’t base this belief on the activity in the V1 area of my brain. If anything, I base my belief on my conscious perceptual experience as of a computer in front of me. What’s important is that a process, and even some parts of a process, can constitute a belief source without being what that belief is based on. Alternatively, however, we sometimes talk about that which the belief is based on (assuming there is a basis) as *the source*

---

<sup>52</sup> My discussion of grounds is deeply indebted to discussion in Alston (1988) and Lyons (2009) chapter 2.

of belief. Both notions of a belief source are important. If one begins with an evidentialist framework as outlined in chapter one then one will tend to focus on the *grounds* notion.<sup>53</sup>

Some processes might not even include anything that would intuitively be described as a ground for, reason for, or basis of belief. This is probably the best way to think of cases like Bonjour's example of Norman the Clairvoyant.<sup>54</sup> We can imagine some computational process taking place in Norman's brain that produces a belief but I take it that we're supposed to imagine the example as one where such a process takes place entirely at the sub-personal level such that, from Norman's perspective, the belief appears to just spontaneously pop into his head randomly. If we imagine the formation of Norman's clairvoyant beliefs then intuitively we would claim that Norman doesn't *base* his belief on anything. His clairvoyant belief that the president is in New York City lacks any grounds. Indicating a belief's lack of a ground, basis, or *putative* evidence is part of what we're doing when we describe a belief like Norman's as similar to a mere guess or random hunch.

Whether processes that don't utilize items upon which the subject bases her belief actually exist is controversial. However, this is the way that most people appear to think of the process of wishful thinking. Sosa (2009) takes a priori belief formation to be a process of this kind such that a priori beliefs are delivered by a brute psychological mechanism that spontaneously produces beliefs without utilizing a basis or reason in the current sense.<sup>55</sup>

---

<sup>53</sup> While all of the intrinsic features of an experience might be taken to be *potential* grounds for belief, it seems as if a causal influence on belief will be a necessary feature of something's being a subject's *actual* (as opposed to potential) ground. As such, it seems as if process talk will slide back in when characterizing a belief source even when we take this to be the grounds notion of a belief source. However, this distinction is still relevant since some processes might not include a candidate for a *basis* of belief (as I discuss in the next paragraph). Moreover, the reliability of grounds and processes might come apart. For more extensive discussion of the notion of a ground for a belief see Alston (1988) and Lyons (2009).

<sup>54</sup> See Bonjour (1985), p. 21.

<sup>55</sup> There are of course reason-based views of a priori justification. According to these views, what is distinctive about a priori justification is the *kind* of reason or ground that provides justification. A priori justification is provided by reasons that are not ultimately grounded in sensory experience. On some views the reason is a non-sensory experiential state. Examples include Plantinga

However, it needn't concern us presently whether any processes that actually exist are baseless in this way.

Even more controversial is the question of whether a process that doesn't provide a subject with reasons upon which to base her belief can be justification conferring. I'll eventually argue that only processes that provide *reasons* could provide justification but, yet again, this needn't concern us presently. My purpose in distinguishing processes and grounds as belief sources is simply that it allows us to raise our puzzle in the idiolect that best suits a particular philosopher's theory of justification. Moreover, the reliability of a source will itself be characterized differently depending on which notion of a source one is working with.

Reliability is a fundamentally *modal* notion. A process  $\Phi$  is reliable iff  $\Phi$  *would* produce mostly true beliefs. A ground  $\Phi$  is a reliable indication of a proposition  $P$  iff  $P$  would usually be true when  $\Phi$  occurs.<sup>56</sup> This latter locution is rather cumbersome and I'll usually use the former process based locution.

---

(1993) and Bonjour (1998). Plantinga suggests a kind of reason based theory of a priori justification/knowledge where the status of reasons as reasons is understood in an externalist fashion. However, given Plantinga's proper functionalism he cannot hold that a priori justification *necessarily* involves basing one's belief on a reason. On other views the reasons involved are conceptual relations with which a subject is directly acquainted. See for instance McGrew and McGrew's (2007) chapter 5.

<sup>56</sup> There are two things to note about this characterization of a source's reliability. First, following Goldman (1979) and (1988), we will have to distinguish between processes/faculties that are *unconditionally* and *conditionally* reliable. A process is conditionally reliable when it takes beliefs as inputs and the process' reliability essentially depends upon the truth of these input beliefs. A process is unconditionally reliable when it either (i) doesn't take beliefs as inputs or (ii) its reliability doesn't essentially depend on the truth of the beliefs that it takes as input. Alternatively, in the case of considering reasons for/bases for/grounds of a subject's belief, one will want to distinguish between the fact that an occurrence of an experience(s) or belief(s) is a reliable indicator of  $P$  and the fact that *the truth of a belief(s)* is a reliable indicator of  $P$  (this would be the grounds analogue of Goldman's conditional/unconditional reliability for processes). Second, this counterfactual characterization of reliability is necessary in order to deal with a problem that arises from sources that might be used very infrequently in the actual world. Imagine  $S$  is the only person with a belief source  $\Phi$  and that  $S$  only uses  $\Phi$  once in her lifetime to form a belief of kind  $K$ . As it turns out, the belief is true. If reliability only requires that the source actually produce mostly true belief then  $\Phi$  would be 100% reliable. But imagine that, *had*  $S$  used  $\Phi$  more often,  $\Phi$  *would* have only produced true beliefs of kind  $K$  45% of the time. Clearly we wouldn't want to characterize  $\Phi$  as reliable. Thus, our notion of reliability is a modal notion.

My explanation of reliability thus far contains an ambiguity that we must consider in order to determine which notion JR is making reference to. Ernest Sosa distinguishes the following two notions of reliability:

**GR** A source  $\Phi$  is generically reliable *iff* it would tend to produce true beliefs.

**SR** A source  $\Phi$  is specifically reliable *iff* it would tend to produce true beliefs on this particular occasion, i.e. in the current circumstances.<sup>57</sup>

My color vision could be generically reliable (i.e. would produce mostly true beliefs in the environments that I'd normally encounter) yet be specifically unreliable since I'm in a room lit by red lights (i.e. color vision wouldn't tend to produce mostly true beliefs in the current circumstances). To which kind of reliability does JR refer?

JR concerns the epistemic role played by *beliefs* about reliability rather than the role of reliability itself. *Beliefs* about reliability ought to play an important role in any theory of epistemic justification. We can determine which kind of reliability JR should be taken to refer by considering uncontroversial roles played by reliability beliefs and asking whether such roles are best played by beliefs about generic or specific reliability.

Reliability beliefs play both positive and negative epistemic roles. Imagine I form the belief that a top secret CIA mission is taking place in Afghanistan on the basis of Paul's testimony. Paul is a complete stranger whom I met five minutes ago. Intuitively my belief isn't justified. Gaining justification for believing Paul is a reliable source regarding CIA missions, however, would thereby provide me with justification for believing Paul's testimony. Alternatively, beliefs about unreliability play a negative role by acting as defeaters. Imagine that I use perception to form the justified belief that there is an elephant in front of me but *subsequently* I gain justification for believing that my perception is unreliable; I'd

---

<sup>57</sup> See Sosa (2009), p. 225.

thereby *lose* my justification for believing there is an elephant in front of me. John Pollock calls this higher-order belief about a source's unreliability an *undercutting defeater*.<sup>58</sup>

Both the positive and negative roles of reliability beliefs are best played by beliefs about *specific* reliability. One reason for thinking this is that beliefs about specific reliability have a justificatory force that overrides that of beliefs about generic reliability. Consider a case where I form the belief that Susanna is walking toward me from a distance in the low light of dusk by relying on visual perception. If I'm justified in believing my visual perception is unreliable in these conditions then visual perception doesn't provide justification for my belief *even if I'm justified in believing my visual perception is generically reliable regarding persons' identities*. Similar remarks apply when I have justification for believing that a source is generically unreliable but specifically reliable; in such a case, the source would provide justification. JR should be interpreted as a principle governing beliefs about *specific* reliability.<sup>59</sup>

### 2.1.2 Contemporary Theories and the Denial of JR

Views commonly classified as externalist clearly deny JR. A variety of externalist theories exist today—process reliabilism, indicator reliabilism, virtue reliabilism, safety theories, sensitivity theories, proper function theories, etc. (see chapter 1)—but process reliabilism is still considered the prototypical externalism. A very crude process reliabilism

---

<sup>58</sup> See Pollock (1974).

<sup>59</sup> One might object that beliefs about generic reliability still play these crucial roles in situations where we're unaware of what the current environmental conditions are (or at least the conditions that are relevant to the reliability of a belief source). Beliefs about generic reliability can play these roles in these kinds of situations only because justification concerning generic reliability or unreliability transfers to beliefs about specific reliability or unreliability when we're unaware of the environmental conditions. A belief that  $\Phi$  is generically reliable amounts to the belief that  $\Phi$  would produce mostly true beliefs relative to *the entire set of environmental conditions we're likely to encounter*. If we lack information about what the current conditions are then this generic belief makes it probable that the source is reliable in the current conditions. Similar remarks apply concerning unreliability.

holds that a belief is justified iff the belief is the output of a reliable process.<sup>60</sup> The *de facto* reliability of color vision is (absent defeaters) sufficient for the justification of the resulting color beliefs; a subject could be completely oblivious to color vision's reliability. Process reliabilism clearly denies JR. Similar points apply *mutatis mutandis* to other non-controversially externalist views. It's important to note, however, that even most internalists deny JR. In fact, every type of internalism mentioned in 1.3 can, and usually does, deny JR.

Alston's indicator reliabilism meets the access requirements of local accessibilism. He requires that a belief's grounds *but not their adequacy* (i.e. that these grounds are reliable indicators of the relevant propositions) are accessible. Eschewing a requirement of access to grounds' adequacy (when adequacy is understood in this reliabilist fashion) is an explicit rejection of JR.

Perhaps more astonishing is the fact that even *global* versions accessibilism, which require access to *all* J-factors, could (and often do) deny JR! Assuming a belief has grounds, the adequacy of these grounds will be a J-factor and global versions of access internalism will require that a subject have access to this adequacy. However, internalists traditionally *do not* analyze adequacy in terms of reliability; this is the upshot of the New Evil Demon thought experiment. In which case, reliability wouldn't be a J-factor and global access internalists could deny JR (and might do so for anti-skeptical reasons).

Mentalist versions of internalism also usually deny JR. Mentalist internalism holds that justification strongly supervenes on a subject's mental states, a view that could conceivably be combined with an acceptance of JR, but most mentalists reject such a

---

<sup>60</sup> This is just the barest sketch of process reliabilism. A fully developed version of the view would need to modify the view so as to distinguish between foundational and inferential justification, explain how defeaters work, address the "generality" problem, etc. For further developments and defenses of reliabilism more generally (including process reliabilism, indicatory reliabilism, and virtue reliabilism) see the sources cited in fn. 33.

requirement.<sup>61</sup> Conservativist epistemologies are extremely popular today and provide a nice example of mentalist internalism that explicitly denies JR. I illustrate the view in the context of perception but it should be clear how the view can be extended to other sources as well. Conservativism's central thesis is that S's merely *having* a perceptual experience as of *P* provides (absent defeaters) S with at least some degree of justification for believing *P*. Hannah's merely undergoing a perceptual experience as of a Blue Jay provides her with *some* degree of justification for believing there is a Blue Jay in front of her.<sup>62</sup> This experience provides this justification even if Hannah lacks any reason for believing perception is reliable; all that matters is that Hannah *in fact* undergoes said experience. Conservativism is internalist in the sense that justification is solely a function of the subject's mental (e.g. perceptual, memorial, etc.) states, but it clearly denies JR.<sup>63</sup>

Even inferential internalism is consistent with a denial of JR. Inferential internalism is defined by its acceptance of IJ2: S has justification for believing *Q* on the basis of her

---

<sup>61</sup> Two exceptions might be Matthias Steup's (2004) internalist reliabilism and Anil Gupta's (2008) non-doxastic coherentism. Steup's view claims that perceptual experiences only provide justification when accompanied by memorial experiences that make a track-record argument available for both perception and memory's reliability. As such, perception and memory do not constitute a counterexample to JR on Steup's internalist reliabilism. It isn't clear but I think Steup is committed to the view that the set of memorial experiences that have propositional contents matching those of a suitable track-record arguments are available can justify even if a subject lacks justification for believing that such a coherent set of experiences is a reliable source of beliefs about perception and memory's reliability. And so I think even Steup's view could be pushed into a rejection of JR regarding the ability of coherence to justify various beliefs.

<sup>62</sup> Difficulties are lurking in the background here regarding the metaphysics of perception. Two relevant issues are: (1) Do perceptual experiences have intentional content that represents the physical world or are they to be treated representations of physical objects merely in the sense of being signs or symptoms of such objects in the same way that dark clouds are a sign of a coming storm? and (2) assuming that perceptual experiences have intentional content, which properties are represented?

<sup>63</sup> Views closely resembling the above sketch of conservativism have been defended by Foley (1983), Huemer (2001), (2006), (2007); McCain (2008), (2012); Pryor (2000); Skene (2013); Tucker (2010). For objections to conservativist epistemologies see Bonjour (2004); DePoe (2011); Fumerton (2007); Hannah (2011); Hasan (2013); Littlejohn (2011); and White (2006). See Tucker (*forthcoming*) for an anthology of papers both defending and critical of conservativist theories of justification.

belief that  $P$  iff (a)  $S$  is justified in believing  $P$  and (b)  $S$  is justified in believing that  $P$  makes probable  $Q$ . Clause (b) may appear to be an explicit acceptance of JR but IJ2 is a principle that applies only to *inferential* justification. An inferential internalist is free to deny JR at the level of foundational justification. In fact, the principal defender of inferential internalism, Richard Fumerton, accepts an acquaintance theory of foundational justification. On his view acquaintance is capable of providing a subject  $S$  with justification even if  $S$  lacks justification for believing that acquaintance is a reliable belief source; if this were required then acquaintance wouldn't be a source of *foundational* justification after all! A reader might thus be tempted to identify basic justification sources and foundational justification sources. This is a mistake. A source of justification might both be inferential and basic where the latter notion is defined in terms of being a counterexample to JR.

Understanding the notion of a basic *inferential* justification source requires a grasp of the distinction between foundational and inferential justification. A belief's justification is foundational when its justification isn't even partially constituted by another justified belief. A belief's justification is inferential when its justification *is* partially constituted by another justified belief. Within the context of reliabilism the distinction can be made in terms of a process' inputs. A process provides inferential justification when it takes beliefs as inputs *and* the process' reliability is conditional upon those inputs being true.<sup>64</sup>

Illustrating the possibility of basic inferential justification sources is easiest by considering inductive reasoning. As I'm currently using the phrase, inductive reasoning refers only to singular and general predictive inferences. In the singular case one infers that the next  $F$  will be  $G$  from one's belief that all observed  $F$ 's have been  $G$ 's. In the general case one infers that *all*  $F$ 's are  $G$ 's. This is a crude formulation of inductive reasoning but it

---

<sup>64</sup> The second clause preserves the intuition that introspection of one's beliefs could provide one with foundational justification. This introspective process would obviously take beliefs as inputs but this process' reliability might not essentially depend on those input beliefs being true.



suffices for our current purposes. Disagreement crops up regarding almost every claim in philosophy, but the claim that this inductive reason will only provide justification for the conclusion if one possesses justification for the premises is uncontroversial. Nevertheless, one can maintain (and many philosophers do maintain) that the inductive reasoning can justify these beliefs for a subject S even if S lacks justification for believing that induction is reliable. If one holds this view of inductive reasoning then induction would be an inferential justification source *and* a basic justification source.<sup>65</sup>

It may seem incredible but basic inferential justification sources are possible even within the framework of inferential internalism! Some notions of probability are analyzed in such a way that they aren't equivalent to reliability. Fumerton himself adopts such a notion of probability—he even suggests that the only chance of avoiding external world skepticism is stripping probability claims of *any* empirical content (reliability claims clearly have empirical content). Therefore, Fumerton's inferential internalism places a requirement on inferential sources that one have justification for believing *P* makes *probable Q* but this

---

<sup>65</sup> What is important for our puzzle in the case of inferential sources is whether the application of an inference can provide justification absent prior justification for believing the inference is *conditionally* reliable. Whether induction is reliable in general is irrelevant since this would be determined partly by applications of enumerative induction to unjustified beliefs. However, since induction is an inferential source of justification one must have justification for believing the premises. Induction only confers justification if one is justified in believing that one is making an inference from *true* premises. Therefore, the important question is whether an induction applied to true premises is reliable, i.e. whether induction is conditionally reliable. Adjusting JR in order to make this distinction between basic foundational sources and basic inferences is a bit tedious and can easily distract from the main point. And so I've relegated this point to a footnote. Notice, however, that perceptual justification might be a basic *inferential* justification source. This is easiest to illustrate within the context of reliabilism (but similar remarks could apply mutatis mutandis within an internalist framework such as mentalism). It isn't obvious that the inputs to our perceptual processes don't include both perceptual experiences *and* beliefs! An expert birdwatcher might form a belief that a certain bird is a Goldfinch via perceptual mechanisms that take as input both the perceptual experience (an experience even a child could have) and the background beliefs that Goldfinches have some feature F that distinguishes them from Canaries, that Goldfinches live in environment E, that she (the bird watcher) is currently in environment E, etc. If this is how this perceptual process works then it's likely that the reliability of the bird watcher's perceptual mechanisms that produce her Goldfinch belief is conditional upon the truth of her background beliefs. Finally, it's possible that *all* of these background beliefs are distinct from the belief that perception is reliable. If all of this were true then the birdwatcher's perceptual processes would be a basic *inferential* justification source.

requirement can be met even if one lacks justification for believing that  $P$  is a *reliable* indicator of  $Q$ .

A vast majority of contemporary internalist and externalist views deny JR. Said denial is supposed to inevitably lead to a problem of easy justification. Cohen's canonical article develops the problem in two forms: a problem from closure and a problem from bootstrapping. Fixing certain difficulties with Cohen's presentation of the problem from closure, however, shows that the two problems aren't distinct.

## 2.2 The Problem of Easy Justification

### 2.2.1 The Problem from Closure

These days perception is a common candidate for a basic justification source—although this was a minority view until quite recently. Assume that perception is indeed a basic justification source. Further assume that I lack any justification for or against perception's reliability; I have no reason to believe that in current circumstances, usually, if I were to have a perception as of  $P$  then it would be the case that  $P$ .<sup>66</sup> Now imagine that I have a perception as of a red table in front of me. Being a basic justification source, my perception provides me with justification for believing that the table in front of me is red (presuming I possess no defeaters). I can also know a priori that a table's being red entails that it isn't a white table illuminated by red lights. I can therefore *gain* justification for believing the consequent of this conditional by inferring it from my justified belief in the antecedent.<sup>67</sup> We can represent the structure of my justification as follows:

---

<sup>66</sup> Some claims about what I perceive, remember, etc. are such that these terms are used in a "factive" sense. Consider the locution "perceive an X" or "perceive that  $P$ ." Each of these locutions is usually used in a factive sense such that one cannot perceive an apple unless there is an apple and one cannot perceive that there is a table in front of one unless there really is. Whenever I refer to perception or memory I mean to be discussing their non-factive senses. The locutions "I have a perception *as of*" and "I have a memory *as of*" naturally suggest this non-factive use.

<sup>67</sup> This assumes a principle of closure regarding justification. Roughly one must assume that: if a subject  $S$  has justification for believing  $P$  and  $S$  knows that  $P$  entails  $Q$  then, if  $S$  has no

## PERCEPTUAL CLOSURE REASONING

1. The table in front of me is red (Justified by my perception as of a red table).
  2. If the table in front of me is red then it isn't white with red lights shining on it (Known a priori).<sup>68</sup>
- 
3. Therefore, the table in front of me isn't white with red lights shining on it.

When my belief in 1 is justified by an apparent perception as of a red table it doesn't seem as if I could *gain* justification for believing this conclusion simply by noticing the entailment described in 2. *Gaining* justification for believing the conclusion *in this way* seems all too easy!

Some epistemologists articulate the problem from closure in terms of a subject gaining justification for believing that the environmental conditions are favorable to perception (or some other source). This may help make the problem especially vivid but it's ultimately misleading. Characterizing the problem in this way suggests that the reasoning provides justification for believing that there aren't red lights shining on the table. But there is no question here: the reasoning does no such thing!

---

independent justification for believing  $\neg Q$ , S has justification for believing  $P$ . That this principle requires S to know the entailment and that S have no independent justification for believing  $\neg Q$  are vital to its plausibility. If S knows the entailment then S's belief that the entailment holds has an epistemic probability of 1 for S. As such, justification won't decrease through the inference. Moreover, if S had independent reasons for rejecting  $Q$  then S's knowledge of the entailment might defeat S's justification for believing  $P$ . Frege's intuitions about his basic law V might have initially provided him with justification for believing basic law V. However, when Frege came to know that basic law V entailed a contradiction (i.e. Russell's paradox) this didn't thereby provide Frege with justification for believing the contradiction. Frege's knowledge of the entailment defeated the justification he had for believing basic law V.

<sup>68</sup> One thing to note about this reconstruction is its explicit reference to my knowledge of the entailment cited in 2. It shouldn't be assumed that my knowledge of the entailment serves as a critical *premise* in my reasoning to the conclusion. In fact, knowing the entailment might not be necessary for justification to transfer from a belief that 1 to a belief that 3. What one says about this will crucially depend on one's take on Lewis Carroll's (1895) infamous regress—issues that will be discussed in detail later in this dissertation. The reason I raise the problem by citing knowing of the entailment is simply because this is clearly *sufficient* (absent any independent reason for rejecting the conclusion—see fn. 67) for justification to transfer from the belief that 1 to the belief that 3. I raise the problem this way for the sake of easy of exposition.

My conclusion of the PERCEPTUAL CLOSURE REASONING is logically equivalent to “either the table in front of me isn’t white or it doesn’t have red lights shining on it.” In our imagined scenario, my justification for believing this disjunction *essentially* depends on my justification for believing the first disjunct. A priori reasoning *cannot* produce justification for believing the second disjunct in isolation. A more perspicuous reconstruction of my supposed structure of justification is:

PERCEPTUAL CLOSURE REASONING\*

1. The table in front of me is red (Justified by my perception as of a red table).
  2. If the table in front of me is red then it isn’t white (Known a priori).
  3. The table in front of me isn’t white (From 1 and 2).
  4. If the table in front of me isn’t white, then it isn’t white *or* it doesn’t have red lights shining on it (Known a priori).
- 
5. Therefore, the table in front of me isn’t white *or* it doesn’t have red lights shining on it, i.e. the table in front of me isn’t a white table illuminated by red lights (From 3 and 4).<sup>69</sup>

Deductive reasoning (as well as legitimate non-deductive reasoning) cannot get you from 1 to the conclusion that there aren’t red lights shining on the table.

There is a trivial sense in which we could get justification for believing that the perceptual conditions aren’t deceptive if by describing the conditions as non-deceptive we just mean that the conditions have not in fact led to an inaccurate perception. But this says nothing about whether the room is lit with red lights, white lights, green lights, etc. However, this does point to a natural explanation of our uneasiness with the original PERCEPTUAL

---

<sup>69</sup> Remember that I am just *displaying* the structure of my supposed justification in a way that clearly illustrates that my justification won’t transfer to the belief that certain perceptual conditions obtain or fail to obtain. I’m not suggesting that the subject goes through this specific reasoning process where each of the a priori entailment relations is explicitly considered and used as a premise. One could legitimately infer 5 directly from 1 (although understanding the conditions under which this would be a legitimate inference is difficult).

CLOSURE REASONING that doesn't rely on the idea that we learn anything about the environmental conditions. Our uneasiness is due to the fact that we're well aware that *if* the table were white and illuminated by red lights then we would have an inaccurate perception that is subjectively indistinguishable from our current perceptual state. Appealing to the scenario of a room lit by red lights isn't essential to our problem. Mentioning the "skeptical hypothesis" is simply a nice heuristic for reminding us that perceptions can be inaccurate. Our uneasiness with PERCEPTUAL CLOSURE REASONING is due to the fact that a subject appears to illegitimately gain justification for believing that her perception is in fact accurate. Sadly, our perception of the subject's reasoning hasn't reflected this very well.

Let's consider a better way of articulating the problematic reasoning. *Ex hypothesis*, my perception provides justification for my belief that there is a red table in front of me even if I lack justification for believing that my perception would, reliably, be accurate in the current conditions (whatever they are). I can introspectively note that I have this perception as of a red table and know a priori that these two claims entail that my perception is in fact accurate on this occasion. At which point I can conclude that my perception is accurate. The supposed structure of my justification can be reconstructed:

PERCEPTUAL CLOSURE REASONING\*\*

1. There is a red table before me (Justified by my perception as of a red table).
  2. I have a perception as of a red table (Known introspectively).<sup>70</sup>
- 
3. Therefore, my perception is accurate on this occasion.

---

<sup>70</sup> I stipulate that the subject *knows* 2 in order to make the presentation simpler. If the subject had less than knowledge level justification for 2 then the probability of the conclusion would be lower than that of the first premise and in certain instance might thereby fail to be justified in believing the conclusion despite having justification for the premises. One might also worry that it is actually quite difficult to figure out what the sources of our beliefs actually are. Nevertheless, it is still odd that in order to find out that a belief source is accurate or reliable we need merely figure out what that source is. Such a claim is still problematic.

My perceptual state provided justification for believing 1 and 2 despite the fact that, by hypothesis, I lacked justification for believing that my perception would reliably be accurate. *If* perception is a basic justification source, it would appear that simply having a perception thereby provides new justification for believing that that very perception is accurate. A particular perception *vouches* for its own accuracy. That a particular perception could tell its subject that that very perception is (contingently) accurate seems incredible! It's analogous to the absurdity of thinking one might have no reason to think that a person's testimony would be accurate but then magically that very person tells me that her testimony is true and I thereby gain justification for believing her testimony is accurate!<sup>71</sup>

Our worry isn't limited to perception. The problem generalizes to any supposed basic source of justification. Consider a case of memory. At time  $t_1$  I lack any justification for or against memory's reliability in the current conditions (whatever they might be). At time  $t_2$  I have a particularly clear and vivid memory as of previously breaking me leg skiing. Now assume that memory is a basic justification source. As such, my memory at  $t_2$  generates justification for believing that I previously broke my leg skiing. I can then introspectively note this memory, recognize that these two claims entail that my memory is accurate, and thereby *gain* justification for believing that my apparent memory is accurate on this occasion.

My supposed justification in the memorial example can be represented thusly:

#### MNEMONIC CLOSURE REASONING

1. I previously broke my leg skiing (Justified by memory).
  2. I have a memory as of breaking my leg skiing (Known introspectively).<sup>72</sup>
- 
3. Therefore, my memory is accurate on this occasion.

---

<sup>71</sup> It's important to note that this problem is specific to cases where the source supposedly providing justification for 1 is identical to the source cited in 2.

<sup>72</sup> See fn. 68.

Yet again, gaining justification for believing the conclusion in this way seems all too easy. If I lack justification for believing that my memory would reliably be accurate in the current conditions then justification cannot be generated for believing that a memory is accurate simply in virtue of having that memory. MNEMONIC CLOSURE REASONING would inevitably lead to the same conclusion in any situation where the relevant memory was in fact inaccurate. Imagine that my older siblings repeatedly told me a false story about my breaking my leg skiing which eventually produced in me an inaccurate memory as of breaking my leg skiing. Even in this situation, the style of reasoning embodied in MNEMONIC CLOSURE REASONING (where the belief that 1 is true is produced by the source cited in 2) would inevitably lead to the same conclusion that the memory is accurate.

Clearly we can abstract away from both perception and memory to show that the problem of easy justification would appear to result from allowing *any* source to be a basic justification source. In the next section I show how this style of reasoning can be extended so as to create the “problem from bootstrapping.”

### 2.2.2 Easy Justification and Epistemic Bootstrapping

When we iterate the closure style reasoning we can arrive at an inductively strong argument for the *modal* conclusion that the relevant source is reliable. The intuition that this form of epistemic bootstrapping is problematic is even stronger than in the case of the closure style reasoning. But if the closure style reasoning is legitimate then the bootstrapping reasoning is as well. They stand or fall together.

In order to help illustrate the problem I want to begin by considering one of the original bootstrapping cases discussed by Jonathan Vogel:

Roxanne drives a car with a well-functioning, reliable gas gauge. She has never looked into the status of the gauge or other like it, she has no information whatsoever on the subject. Rather, Roxanne automatically forms beliefs about the level of gas in the car’s tank simply by consulting the gauge. For example, if the gauge reads “F” she immediately and directly forms the belief that the car’s tank is full. Given that the gauge is reliable, it seems clear that Roxanne’s belief that the car’s tank is full is formed by a reliable process. Now,

Roxanne can also observe what the state of the gauge itself is, if she chooses to. Roxanne notes that the needle reads “F” at the time when she believes, by reading the gauge, that the tank is full. Roxanne conjoins her belief that the gauge reads “F” with her belief that the tank is full, and deduces that the gauge reads accurately on this occasion. We can suppose that Roxanne repeats this strange procedure a good number of times, accumulating beliefs that the gauge reads accurately at various times  $t_i$ . Roxanne goes on to conclude by induction that the gauge is accurate in general, that is, that the gauge is reliable.<sup>73</sup>

Roxanne’s reasoning here is just an iteration of the closure style reasoning discussed in the previous section. Roxanne uses a source  $\Phi$  in forming a belief that  $P$ . She then notes that this very source  $\Phi$  tells her that  $P$ . At which point she infers that  $\Phi$  was accurate on this occasion. Finally, Roxanne iterates this reasoning in order to conclude that  $\Phi$  was accurate on a wide variety of representative occasions and thereby concludes that  $\Phi$  is reliable.

We all recognize that Roxanne’s reasoning is defective. She couldn’t discover that her gas gauge is reliable *in this way*. There is something about the very structure of her reasoning that makes us uneasy. When Roxanne’s beliefs about how full the tank is have as their source the reading on the gas gauge we don’t think these beliefs can serve as a source of new justification for the claim that the gas gauge is reliable. If Roxanne’s beliefs about how full the tank is had another source (perhaps she used a dipstick to arrive at these beliefs) then there would be nothing wrong with the structure of her supposed justification.<sup>74</sup>

---

<sup>73</sup> Vogel (2008), p. 519.

<sup>74</sup> There is a sense in which this same inductive data could give her justification even without using a separate source for forming beliefs about the level of the gas tank. Cars require gas to run and presumably Roxanne knows this much. And so Roxanne might reason as follows: if my gas gauge were unreliable then, if I were to only fill up when the gauge was close to empty, it would be likely that I’d eventually run out of gas despite my gauge indicating otherwise. The best explanation for why I haven’t run out of gas is that my gas gauge really is reliable. The longer Roxanne relies on her gauge to decide when to fill up her tank without running out of gas the greater her justification for believing the gauge is reliable. However, this justification is better represented as an inference to the best explanation than the problematic inductive structure that Vogel is attempting to illustrate. And such distractions could easily be avoided by modifying the case in such a way that cars don’t need gas to run (but such cases get incredibly artificial).



Most epistemologists (externalists included) are likely to cry foul by explaining that no one thinks that reading a gas gauge is in fact a basic justification source. Roxanne must have implicitly been relying on a premise that her gas gauge is reliable in moving from the reading on the gas gauge to the belief that the gas tank is full. And so if she is to have justification for believing gas tank is full she must have *prior* justification for believing the gauge is reliable. Filling in this detail makes the inductive argument suffer from premise circularity. Clearly arguments that are (essentially) premise circular cannot generate *new* justification for their conclusions. Any justification that Roxanne has for the conclusion must have been in place independently of this inductive argument.

While I think we could vindicate the claim that Roxanne's reasoning isn't objectionable *merely* because gas gauges aren't basic justification sources, the issues involved are complicated and controversial. Thus, I think it's best to move to a parallel case involving a source that many believe is a basic justification source: testimony. I'm not advocating the view that testimony is in fact a basic source. I'm appealing to testimony here simply because it's currently a rather popular view that testimony is a basic justification source—although this was again a minority view until recently.<sup>75</sup>

Roxanne's favorite newspaper is The Roxy. Roxanne, however, has never looked into The Roxy's reliability. She has no information whatsoever regarding The Roxy's reliability (including information about the reliability of newspapers in general). One day Roxanne begins to wonder whether The Roxy is a reliable source of information. Roxanne reads " $P_1$ " in The Roxy and directly forms the belief that  $P_1$ . She then perceptually notes that The Roxy reads " $P_1$ " and infers that The Roxy was accurate regarding  $P_1$ . She repeats this process for claims  $P_2, P_3, \dots, P_N$  and concludes by induction that The Roxy is in fact reliable.

---

<sup>75</sup> See Coady (1992), (1994); Goldman (1999); Goldberg (2006), (2008); Reid ([1764] 1983); and Rysiew (2002) for defenses of non-reductionism. See Hume ([1739] 1975); Lipton (1998); Lyons (1997); and Van Cleve (2006) for defenses of reductivism. See the edited volume by Lackey and Sosa (2006) for a nice collection of articles on the epistemology of testimony.

Clearly Roxanne's reasoning cannot generate justification for believing that The Roxy is reliable. There is something about the very structure of this reasoning that rules out the possibility of its generating new justification for its conclusion. And since this is a claim about the general structure of Roxanne's reasoning we should assume that any reasoning structurally isomorphic to Roxanne's can't generate new justification for believing a source is reliable. However, it appears allowing a basic justification source allows this style of bootstrapping to generate justification for believing a source is reliable.

Someone might again reply that there is nothing wrong with the structure of the reasoning but rather with the supposition that testimony isn't actually a basic justification source. They might further reply that *when* we utilize this kind of reasoning with a genuine basic justification source it won't strike us as problematic.

It may seem that I'm beating a dead horse but I want to consider yet another example of a source that many philosophers have considered to be an even clearer example of a basic justification source: induction. Induction is particularly useful because it allows us to compare two attempts to generate justification for believing induction is reliable in a way that clearly illustrates the problematic nature of the bootstrapping reasoning.<sup>76</sup>

A variety of philosophers today have attempted to defend an inductive justification of induction.<sup>77</sup> The idea is that you begin by forming numerous predictions on the basis of induction. For instance, from my (justified) belief that all observed swans have been white I infer that the next swan will be white. Then I observe the next swan and gain justification via

---

<sup>76</sup> This same contrast case is considered by Vogel (2008). Vogel brings the example up to suggest that "rule-circularity" is bad. As he notes, however, not even advocates of an inductive justification would accept the kind of inductive justification considered in the bootstrapping case below. Intuitively, the classic inductive justification of induction isn't as problematic as the bootstrapping version. This difference has to be accounted for and his appeal to rule-circularity cannot do this. Van Cleve (2003) mentions the same contrast between the two kinds of inductive justifications of induction.

<sup>77</sup> Black (1958); Papineau (1992); Sanford (1990); and Van Cleve (1984).

perception for believing that it's in fact white. So induction was accurate on this occasion. I also infer from my (justified) belief that all observed emeralds have been green to the belief that the next observed emerald will be green. I then observe the next emerald and gain justification via perception for believing that it's in fact green.<sup>78</sup> I iterate this reasoning and, according to the advocates of this argument, gain justification for believing induction is reliable. We can represent the supposed structure of justification as:

#### INDUCTIVE JUSTIFICATION OF INDUCTION

1. Induction tells me that  $P_1$ .
2.  $P_1$  (Justified by perception).
3. Induction to  $P_1$  was accurate (From 1 & 2).
4. Induction tells me that  $P_2$ .
5.  $P_2$  (Justified by perception).
6. Induction to  $P_2$  was accurate (From 4 & 5).

[Repeat for  $P_3$ - $P_N$ ]

---

7. Therefore, induction is generally reliable.

Whether this inductive justification of induction is legitimate is controversial. My point here is only that the inductive justification of induction presented here is significantly different than the bootstrapping reasoning that would be possible regarding induction's reliability if induction is a basic justification source.

If induction is a basic justification source then it can justify claims prior to having any justification for believing it's reliable. As such, one needn't actually use perception (testimony, etc.) to verify induction's predictions in order to get an inductive justification of induction. I can use induction to gain justification for believing that the next swan will be

---

<sup>78</sup> Obviously inductive arguments might arise where the next observed F is observed by someone else and my justification for believing that the F is G will be due to testimony but this complication won't be relevant to the point I'm making here.

white. I can then note that an application of the inductive inference pattern predicts that the next swan will be white, at which point I can infer that induction is accurate on this occasion. Finally, by repeating this process I can conclude by induction that induction is reliable. Such bootstrapping on induction can be represented thus:

#### EASY BOOTSTRAPPING ON INDUCTION

1.  $P_1$  (Justified by induction).
2. Induction tells me that  $P_1$ .
3. Induction to  $P_1$  was accurate.
4.  $P_2$  (Justified by induction).
5. Induction tells me that  $P_2$ .
6. Induction to  $P_2$  was accurate.

[Repeat for  $P_3$ - $P_N$ ]

---

7. Therefore, induction is reliable.

Generating justification for believing that induction is reliable in *this* way is too easy and illegitimate. Acquiring justification for believing that induction is reliable requires more demanding cognitive work. Proponents of inductive justifications of induction have been wise not to utilize EASY BOOTSTRAPPING ON INDUCTION. Had they proposed the second argument lacking independent checks it's doubtful that advocates of the inductive justification of induction would have been capable of convincing anybody. Clearly EASY BOOTSTRAPPING ON INDUCTION cannot generate new justification for believing that induction is reliable. But if justification is a basic justification source then it appears that such reasoning would do just that.

Yet again, one could reply that the reason our intuition that EASY BOOTSTRAPPING ON INDUCTION is illegitimate has nothing to do with the supposed structure of justification. Rather, we all implicitly realize that one needs prior justification for

believing induction is reliable in order for it to be justification conferring. There are two reasons to think that we shouldn't explain our uneasiness in this way.

First, appealing to the supposed non-basicality of induction in order to explain our uneasiness with the bootstrapping reasoning fails to explain why such reasoning is intuitively problematic even for those who don't find the original **INDUCTIVE JUSTIFICATION OF INDUCTION** illegitimate. If all of our uneasiness with the bootstrapping reasoning can be explained in terms of the intuition that induction is a non-basic source of justification then this should give rise to the intuition that the original inductive justification of induction is illegitimate as well. I'm not advocating the position that the inductive justification of induction isn't problematic, my point is that the bootstrapping case is intuitively *more* problematic which means that this intuition concerns something particular to the supposed structure of the bootstrapping reasoning.

Second, the same bootstrapping argument would seem illegitimate for any non-deductive inferential source of justification. Imagine I use inference to the best explanation (IBE) to form the belief that *P*, note that this was formed by IBE, infer that IBE was accurate on this occasion, repeat, and finally conclude that IBE is reliable.<sup>79</sup> No matter the non-deductive inferential source, the bootstrapping strikes us as problematic. Thus, if you want to chalk up the intuition to the fact that the sources are not basic (in the sense defined relative to JR) rather than to an intuition that the justificatory structure is problematic, you'll have to say that there are *no* basic non-deductive inferential sources. You'll commit yourself to the view that no non-deductive inferential source can provide one with justification unless one already possesses justification for believing it to be reliable. And this would appear to lead to a far reaching skepticism since, on such a view, non-deductive inferences could only be justification conferring if we can find some foundational justification for believing these

---

<sup>79</sup> Note that it wouldn't make the bootstrapping any better if the conclusion were viewed as being justified by induction on the success rate of IBE or as an IBE of IBE's previous success.

inferences are reliable or we can find some deductive argument from our foundational justification to the reliability of these inferences. How could we have foundational justification for believing that induction, IBE, etc. are reliable? None of the common candidates for foundational sources appear capable of delivering such a result. Can I perceive induction's reliability? Can I remember induction's reliability?<sup>80</sup> Can induction's or IBE's reliability be foundationally justified *a priori*? Reliability is an empirical notion and clearly the reliability of these non-deductive inferences is a contingent matter. As Hume pointed out, induction will only be reliable if the world is in fact regular (which it might not be). Similar remarks apply to IBE's reliability. Nor does it seem that we could provide a deductive argument from foundational beliefs to the reliability of non-deductive inference. It's therefore a better strategy to admit that what we find problematic is the supposed structure of justification involved in the bootstrapping and hope there is a way to avoid allowing this structure while maintain the basicity of at least some non-deductive inference.

Now that it's clear that what is intuitively problematic with the bootstrapping is the structure of the supposed justification, *not* the assumption that a certain source is basic, I want to show how this problem arises even for foundational sources. As such, denying JR appears to entail the legitimacy of the problematic structure of justification illustrated by the bootstrapping reasoning.

Assume that perception is a basic foundational justification source. Suppose I lack any justification for or against the reliability of my perception. Now imagine that I consider the question of whether my perception is reliable regarding the existence of mid-sized ordinary objects (e.g. tables, chairs, cars, computers, etc.). If perception is a basic justification

---

<sup>80</sup> Note that this would require more than remember induction being accurate in the past. Memory may be able to provide you with foundational justification for an inductive argument supporting induction's reliability but it doesn't provide foundational justification for the reliability claim itself. And on the view under consideration an inductive argument won't justify a conclusion unless one has prior justification for believing such inferences are reliable.

source, I can look around the room and use my perceptual experiences to verify their own accuracy via the closure style reasoning, and thereby supply myself with inductive data that generates new justification for believing that my perception is reliable regarding the existence of mid-sized ordinary objects. My supposed structure of justification would be as follows:

#### EASY BOOTSTRAPPING ON PERCEPTION

1.  $P_1$  (Justified by a perception as of  $P_1$ ).
2. I have a perception as of  $P_1$  (Known by introspection).
3. My perception as of  $P_1$  is accurate (From 1 & 2).
4.  $P_2$  (Justified by a perception as of  $P_2$ ).
5. I have a perception as of  $P_2$  (Known by introspection).
6. My perception as of  $P_2$  is accurate (From 4 & 5).

[Repeat for  $P_3$ - $P_N$ ]

---

7. Therefore, my apparent perception is reliable.

Gaining justification for believing that one's perception is reliable in this way appears all too easy and illegitimate. By now it should be clear that this problem generalizes to any supposed basic justification source (e.g. memory).

Admitting basic justification sources seems to be tantamount to accepting the legitimacy of the problematic structures of justification illustrated by the closure and bootstrapping style reasoning. Collectively these two forms of reasoning can be referred to as easy justification reasoning and the problem of legitimating these structures of justification can be referred to as the problem of easy justification.

#### 2.2.3 It isn't a Compliment When Someone Calls You Easy

Despite its intuitive force, the charge that justification is *too easy* doesn't provide an adequate appreciation of our problem. So here I take a few steps towards a more clear articulation of what is problematic about easy justification. My discussion here will take us quite a ways in appreciating the problems surrounding these issues but my discussion of

various proposed solutions in the next chapter also helps provide a fuller understanding of what is problematic about easy justification.

As I've been at pains to point out, our problem concerns the *structure of justification* and one step towards appreciating this problem involves noting that this is distinct from any problem with the *structure of the arguments*. Easy justification reasoning relies on deductively valid and inductively strong arguments. Our problem doesn't concern the relationship between premises and conclusion. *If* one is justified in believing the premises then one is also justified in (or at least has some justification for) believing the conclusion. What is problematic about the closure and bootstrapping reasoning is related to the fact that the first premise (of each iteration in the bootstrapping form) is supposedly justified by the very source—cited in the second premise—whose accuracy/reliability is assessed in the conclusion. Were the first premise justified by a distinct source (remember the example of the dipstick in the gas gauge scenario), the supposed structure of justification wouldn't be problematic.

Our problem doesn't concern the arguments themselves but, rather, the supposed relationship between our justification for believing the premises and our justification for believing the conclusion. But what about this relationship is problematic?

### 2.2.3.1 *Running in Circles?*

It's tempting to diagnose the problem as one of circularity. For example, Richard Fumerton complains:

You cannot *use* perception to justify the reliability of perception! You cannot *use* memory to justify the reliability of memory! You cannot *use* induction to justify the reliability of induction! Such attempts to respond to the skeptic [or attempts to justify the reliability of a belief source] involve blatant, indeed pathetic, circularity.<sup>81</sup>

---

<sup>81</sup> Fumerton (1995), p. 177.



But this is precisely what appears to happen in cases of easy justification. Perception is used to justify the reliability or accuracy of perception. Memory is used to justify the reliability or accuracy of memory. And so on.

Nevertheless, we must be careful with such a charge. Roger White explains:

There are plenty of legitimate lines of reasoning which might broadly be characterized as using a [source] to support the reliability of that [source]. I can read the optometrist's report from my eye exam and learn that my eye-sight is quite good. Doing well in a memory game can suggest that I have a good memory, even though I can't help but use my memory to evaluate my performance.<sup>82</sup>

In cases of easy justification one uses a source to justify that source's accuracy or reliability, but this isn't an adequate explanation of why it's problematic. Easy justification is problematic because of the particular *way* that we use a source to justify that source's accuracy or reliability. White's examples involve circularity in some sense but, even if these examples are problematic, easy justification is clearly *more* problematic.

More contrasts can bolster this point. Compare PERCEPTUAL CLOSURE REASONING\*\* to a case where my belief that there is a red table in front of me is supposedly justified by looking at a tag on the table that reads "Red Mahogany Table." At which point I introspectively note my perception as of a red table and infer that this perception is accurate. Label this TAG REASONING. Such reasoning is still susceptible to a charge of circularity; you must *use* perception to read the tag on the table. Nevertheless, the supposed structure of justification in TAG REASONING is clearly epistemically better than the supposed justification in PERCEPTUAL CLOSURE REASONING\*\*.

This same point is illustrated by a comparison of MNEMONIC CLOSURE REASONING to the following:

APPARENT ACCURATE MEMORY (AAM) REASONING

---

<sup>82</sup> White (2006), p. 530.

1. At time  $t_1$  I seemed to remember that  $P$  (Justified via a memory at  $t_2$  as of *having seemed to remember at  $t_1$  that  $P$* ).
  2.  $P$  (Justified by a memory at  $t_2$  as of  $P$ ).
- 
3. Therefore, my memory at  $t_1$  was accurate.

AAM reasoning is superficially similar to the closure style reasoning involving memory but the differences are important. Easy justification reasoning utilizes a memory as of  $P$  to justify  $P$  and then the subject subsequently notes *that very memory*. AAM REASONING notes a memory as of *previously having a memory as of  $P$* , utilizes a *distinct* memory as of  $P$  to provide justification for  $P$ , and combines these to arrive at the conclusion that the earlier memory was accurate. AAM REASONING could also be iterated to form an inductive argument for memory's reliability resembling bootstrapping reasoning. Perhaps AAM REASONING and its iteration are problematic but the easy justification reasoning is clearly *more* problematic.

Similar remarks apply to the comparison between INDUCTIVE JUSTIFICATION OF INDUCTION and EASY BOOTSTRAPPING ON INDUCTION. Perhaps both are problematic but the structure involved in the bootstrapping is obviously *more* problematic. An account of why easy justification is problematic must distinguish it from White's optometrist example, TAG REASONING, AAM REASONING and its iteration, INDUCTIVE JUSTIFICATION OF INDUCTION, and similar forms of reasoning. This means an appreciation of our problem requires a diagnosis that goes beyond mere circularity.

I maintain that easy justification has two problematic features but I should note that our comparisons have highlighted something that we'll do well to keep in mind (especially when evaluating proposed solutions in the next chapter). We intuitively think that generating justification for believing a source is accurate or reliable is *more* demanding than simply going through the easy justification reasoning. Whether other forms of "circular" reasoning such as those discussed above are demanding enough is an interesting but distinct question.

### 2.2.3.2 Failure Is Not an Option

A key problem with the easy justification reasoning is that we realize that we could go through this same reasoning in a variety of scenarios where the source is inaccurate or unreliable but the reasoning would reach the same conclusion *in all of these scenarios* (i.e. would lead to a false positive). The way we use our belief sources in easy justification reasoning *guarantees* that we reach the favorable conclusion. Our general reasoning procedure is such that it could never detect, what are surely possible, false positives. As Crispin Wright explains:

[S]o long as my sole evidence concerning the truth of a verdict issued by the source is the source's own word for it, so to speak, I won't be in a position to *detect* the occurrence of any counterexamples. Since I know in advance that I won't be [in] a position to detect any counterexamples should they occur,—that I will falsely believe of any counterexample that does occur that it is not one—it seems plainly irrational to regard my “findings” as a confirmation of reliability [or accuracy].<sup>83</sup>

There is *no chance* that the general reasoning procedures might lead you to data that supports the inaccuracy or unreliability of the belief source despite the fact that said inaccuracy or unreliability is clearly possible.<sup>84</sup>

This “impossibility of failure” suggests that such reasoning cannot reveal justification that wasn't in some sense available without the reasoning. To be clear, the problem isn't that one couldn't gain justification for the conclusion by going through the easy justification reasoning. Our realization that such reasoning inevitably leads to the favorable conclusion illustrates that *if* any justification is conferred upon the conclusion by such reasoning then such justification must have been implicit prior to the reasoning itself. The arguments are deductively valid and inductively strong, so if we grant that one has justification for the premises then it seems we must allow that the reasoning confers justification on the

---

<sup>83</sup> Wright (2011), p. 36.

<sup>84</sup> See Titlebaum (2010) and Wright (2011) for similar diagnoses.

conclusion (we'll consider proposed solutions in chapter 3 section 3.5 that seek to avoid this). Now we encounter two problems. First, it seems that such justification must have been implicit prior to the reasoning but this seems inconsistent with a denial of JR. Second, the only thing present prior to the reasoning was the perceptual experience itself and a perceptual experience cannot vouch for its own accuracy or reliability. This leaves us with a puzzle of what could provide implicit justification for the conclusion prior to the reasoning and why the necessity of such implicit justification isn't tantamount to accepting JR.

In order to lend credence to this diagnosis reconsider Vogel's gas gauge example. If Roxanne's only reason for her beliefs about the level of gas in her tank is the reading on the gas gauge and she relies on these beliefs in forming beliefs about the gauge's accuracy or reliability then clearly she couldn't but reach the conclusion that the gauge is accurate or reliable. Compare this to the case where Roxanne's beliefs about the level of the gas in the tank are formed via her use of a dipstick, which are then used to form beliefs about the accuracy or reliability of the gauge. Clearly this latter procedure could fail to lead to the favorable conclusion. There is no guarantee that the gauge and the dipstick won't give conflicting results. Thus, this diagnosis correctly identifies easy justification as having a problematic feature not shared by cases where independent sources are used.

Our diagnosis also identifies a feature of easy justification that separates it from other forms of "circular" reasoning. White's optometrist example, for instance, involves *using* perception to form a belief about perception's reliability but clearly there is a chance that this type of reasoning leads to a negative verdict. It's possible that you look at an optometrist's report and have a perception as of a report that says your perception is unreliable. Similarly, in TAG REASONING, it's possible for your perception of the table to conflict with the perception of how that table is described on the tag. Perhaps Isabella, a double major in philosophy and art, was inspired by her epistemology class to make her next art piece a white table illuminated by red lights (clearly one shouldn't look to epistemology for artistic inspiration). Were a subject to visit Isabella's exhibit he would have a perception as of a red

table but his perception of the tag would be of a tag that reads “White Table Illuminated by Red Lights.” In this scenario the TAG REASONING would lead to the conclusion that one’s perception as of a red table is inaccurate. Similarly, there is a chance of failure in AAM REASONING and its iteration. It’s possible that I have a memory as of *previously having a memory as of P* while simultaneously having a memory as of  $\neg P$ . I might have a memory as of previously seeming to remember parking my car in the east parking lot while simultaneously having a memory as of my car having turned out to be parked in the west parking lot. Finally, the INDUCTIVE JUSTIFICATION OF INDUCTION, as opposed to EASY BOOTSTRAPPING ON INDUCTION, has a possibility of failure as well. It’s entirely possible that induction makes the prediction that the next observed swan will be white but when I actually run into the next swan I have a perception as of a black swan. In all of these cases one couldn’t know in advance that such methods for forming beliefs about the accuracy or reliability of a source would lead to the favorable conclusion.

Such a diagnosis of why easy justification is problematic is closely related to William Alston’s reasons for taking track-record arguments for perception’s reliability to be illegitimate in general—however, I’m explicitly *not* rejecting all track-record arguments. Alston’s penultimate reasons for rejecting track-record arguments are presented in the following passage:

[W]hen we ask whether one or another source of belief is reliable, we are interested in *discriminating* those that can reasonably be trusted from those that cannot. Hence merely showing that *if* a given source is reliable it can be shown by its record to be reliable, does nothing to indicate that the source belongs with the sheep rather than with the goats... I shall disqualify epistemically circular arguments on the grounds that *they do not serve to discriminate between reliable and unreliable doxastic practices* [my emphasis].<sup>85</sup>

---

<sup>85</sup> Alston (1993), p. 17. Whether Alston’s remarks here are consistent with his advocacy of a kind of reliabilism about epistemic justification is an interesting issue but it isn’t my concern here.

An *inability* to discriminate accurate/inaccurate and reliable/unreliable sources is my reason for taking easy justification to be problematic. When attempting to acquire new justification for believing a source to be accurate in the current instance or more generally, we want some indication that would discriminate good and bad cases. Easy justification reasoning doesn't utilize anything that discriminates these cases in *any* interesting sense; the reasoning could be employed in bad cases and would still *inevitably* lead to the same positive conclusion.<sup>86</sup>

Yet again, this distinguishes easy justification from circular reasoning. Some circular reasoning might discriminate in the relevant sense. Reading the optometrist's report, TAG REASONING, AAM REASONING and its iteration, and INDUCTIVE JUSTIFICATION OF INDUCTION are all possible candidates for reasoning that would at least imperfectly discriminate cases where the relevant source is reliable from cases where it is unreliable. Such reasoning *could* produce false positives (e.g. I could place a tag that reads "Red Table" in front of a white table illuminated by red lights). Thus, these forms of reasoning do not perfectly discriminate scenarios where the source is accurate/reliable from scenarios where the source is inaccurate/unreliable. Nevertheless, it's consistent to maintain that such reasoning is *unlikely* to produce a false positive. A request that one's reasoning must discriminate reliable from unreliable sources in order for it to generate *new* justification for believing a source is reliable is only a request for reasoning that increases the likelihood

---

<sup>86</sup> We must be careful with this kind of diagnosis and distinguish *perfect* and *imperfect* discrimination. I've already suggested as much, but the proposed diagnosis doesn't classify easy justification as problematic because it fails to be even *imperfectly* discriminating. Take two cases, one where  $P$  holds and one where  $\neg P$  holds. Now take some item  $X$ . The fact that  $X$  might be present in both cases entails that  $X$  doesn't *perfectly* discriminate  $P$  from  $\neg P$  cases or vice versa. Perfect discrimination is what *bad* skeptical reasoning relies on. This *poor* skeptical reasoning argues: it's possible that  $X$  obtains yet your belief that  $P$  is false; therefore,  $X$  cannot justify  $P$ . I'm not arguing in this way. A principle of the form [ $X$  can *generate* new justification for believing  $P$  iff  $X$  perfectly discriminates  $P$ -cases from  $\neg P$ -cases] is obviously false. However, it's plausible to maintain that  $X$  cannot *generate* justification for believing  $P$  if  $X$  fails to *at least imperfectly* discriminate  $P$  from  $\neg P$  cases. This idea of *imperfect* discrimination is difficult to define. As a crude formulation we might say that  $X$  imperfectly discriminates  $P$  from  $\neg P$  iff  $P$  would be highly unlikely relative to  $X$ 's occurrence. But spelling this idea out more fully is difficult.

that the source is accurate or reliable. Easy justification reasoning fails to discriminate even in this relatively weak sense.

Notice that this diagnosis fits very closely with the desire for justification to be essentially tied to the notion of providing one with assurance from the egocentric perspective that one's belief is true (see chapter 1). Something that generates *new* justification should provide someone taking on an egocentric examination of her beliefs with more assurance than they had previously. The fact that the easy justification reasoning would lead to the conclusion that a source is accurate or reliable in *every* case where this is false suggests that combining the beliefs (when gained in the way imagined in easy justification) in this reasoning cannot provide you with any assurance for the accuracy/reliability of one's source that wasn't already implicitly present when forming those beliefs.<sup>87</sup>

---

<sup>87</sup> One might object that this diagnosis of what we find problematic about the easy justification reasoning is inadequate since cogito-style reasoning can't fail to produce the favorable result but can still provide assurance and justification that wasn't already implicitly present. Vogel (2000) p. 16 briefly considers the fact that his bootstrapping case (e.g. the gas gauge case) couldn't fail to reach the conclusion it does. Whatever reasoning (or procedure if it isn't best described as reasoning) we use in the cogito to form the belief that we exist couldn't fail to reach this conclusion. But clearly such reasoning can provide justification an assurance. Therefore, Vogel suggests, that the putative defect isn't really grounds for thinking that a procedure cannot *generate* justification for a certain belief. In response, I want to suggest that the difference between the easy justification reasoning and the cogito is that the cogito can't fail to lead to a negative conclusion because the resources used in the reasoning are *perfectly* discriminating. Easy justification and the cogito are both concerned with contingent propositions. There are cases where sources are accurate/reliable and where sources fail to be accurate/reliable. There are cases where I exist and where I don't exist. Easy justification is such that it will always reach the positive conclusion even in those situations where one employs the relevant resources in cases where the source is inaccurate or unreliable. It's for this reason that the availability of the presence of the resources used in the easy justification reasoning fail to change the ratio of reliable cases to non-reliable cases in favor of the former and therefore fails to be even *imperfectly* discriminating. Things are different in the case of the cogito. I'm not sure how to characterize the procedure used in the case of the cogito but I agree that such a procedure would always lead to a positive verdict. But, whatever the correct characterization of the procedure, the resources that must be in place in order to use this procedure is such that we could never have such resources in cases where the belief that we exist is false—perhaps the correct characterization of the cogito procedure is something akin to the following: if you're aware of a sensory state, belief state, etc. then form the belief that you exist otherwise don't form this belief. The fact that employing the cogito procedure would always result in a positive verdict would never reach a negative verdict is the product of the fact that the resources used in such a procedure *perfectly* discriminate in the sense that the procedure couldn't be carried out in situations where the positive verdict is false.

### 2.2.3.3 *I'm Sorry but I Have a Prior Commitment*

Ernest Sosa pinpoints another problem involved in easy justification: it appears to generate justification *ex nihilo*. Sosa explains that forming a belief using a source  $\Phi$  manifests a commitment to (or presupposition of)  $\Phi$ 's reliability. In an informative footnote he argues:

Can't we agree that there are dispositions to take visual experience at face value in play when we visually perceive, e.g., that a wall we see is red? Won't there then be a state that is the basis for that disposition, and isn't it plausible to describe the person as being disposed to believe that he sees something red upon having a visual experience as if he does? If so, then won't it be proper to attribute to such a person an implicit "commitment" whose content is something like "if it looks red, then it is red," an implicit assumption, or presupposition, or mindset, to that effect. This would be analogous to the unspoken and perhaps even unrecognized prejudices of someone who nonetheless makes evident in his conduct that he regards members of a certain target group as inferior.<sup>88</sup>

Sosa explains that these commitments are *operative* when we use a belief source, in the same way a person's prejudices are actually operative when the person behaves as if a certain group is inferior. In trusting a belief source I manifest a certain disposition and therefore rely on the categorical state that grounds such a disposition; this categorical state can be plausibly described as having a propositional content corresponding to the content used in ascribing the disposition (e.g. "Reliably, if it looks red, then it is red"). The categorical state seems to be a commitment or presupposition of the source's reliability in the current circumstances.

This brings us to the problem of *ex nihilo* generation of justification. Sosa explains, "The normative status of the subject's belief that he sees a red wall cannot derive from the presence and operation of that commitment (whatever its nature) while the commitment acquires *its* normative status only *posteriorly* to the belief's having its proper status."<sup>89</sup> If this were to happen this would amount to what I, following Sosa, am calling *ex nihilo* generation of justification. The idea is that, in the imagined scenario, a commitment is unjustified but

---

<sup>88</sup> Sosa (2009), p. 223 fn. 14.

<sup>89</sup> Sosa (2009), p. 234.



then a subject S's reliance on that unjustified commitment generates justification for S to believe the content of that very commitment. Surely this is problematic but it appears to be precisely what happens in the easy justification reasoning.

One might be tempted to question the idea that these commitments could be justified or unjustified prior to their being made explicit by the easy justification reasoning. But the implicit nature of a commitment should be irrelevant. Epistemologists too often focus on idealized kinds of intentional states. Epistemologists often attempt to vindicate the intuition that we have justified beliefs about the existence of tables, computers, cars, etc. It's probably more accurate to say that we rarely form any explicit beliefs about the existence of these ordinary objects. If we're phenomenologically honest with ourselves it's more accurate to describe the states we usually form on the basis of our perceptual experiences as expectations rather than explicit beliefs about ordinary objects. Debates between epistemologists regarding the justification of common sense beliefs wouldn't be rendered moot by this phenomenological point. These debates would apply equally to these expectations even though they are subtly different from explicit beliefs. Making subtle distinctions between explicit beliefs, expectations, and implicit commitments is incredibly natural. However, provided that all of these states have truth conditions (as opposed to merely satisfaction conditions like desires), it's natural to think that these intentional states can be evaluated as justified/unjustified or rational/irrational.<sup>90</sup>

## 2.3 The Problem of the Criterion

### 2.3.1 Justification: From Easy to Impossible!

In section 2.2 I developed the worry that denying JR appears to commit one to recognizing various ways of gaining justification as legitimate that are clearly problematic.

---

<sup>90</sup> Fumerton (1995) chapter 2 makes this same point regarding the intentional states about ordinary objects that we form on the basis of our sensory experiences.

Where do we go from here? Perhaps this just shows that we ought to accept JR. Perhaps a belief source only confers justification if one already has justification for believing the source is reliable. Accepting JR, however, appears to have dire skeptical consequences.

JR's skeptical worries arise when we combine it with the following plausible epistemic principle:

**JT** Necessarily, a subject S has justification for believing a source  $\Phi$  is reliable only if S *already* has justification for believing that particular deliverances of  $\Phi$  are true.

Whether a source  $\Phi$  is reliable is a contingent matter.<sup>91</sup> And so it seems straightforward that the only way we could gain justification for believing that  $\Phi$  is reliable is by inferring this from justified beliefs correlating  $\Phi$ 's delivering certain verdicts and those verdicts being true. How else could you gain justification for believing a source is *contingently* reliable?

JT doesn't lead to skepticism by itself. You can satisfy JT by using an independent source  $\Psi$  to gain justification for believing that the verdicts of  $\Phi$  are true. Suppose you wanted to acquire justification for believing that your auditory faculties/experiences are reliable. You might note the verdicts of your auditory faculties and then confirm these claims using visual perception. You could note that you have an auditory experience as of a loud engine in the sky and that this inclines you toward believing that there is a plane nearby. Subsequently your visual experience might provide you with justification for believing that there is indeed a plane nearby and thereby confirm the accuracy of our auditory faculties.

Skeptical worries arise from the *combination* of JR and JT.<sup>92</sup> In the previous example JT was satisfied relative to my auditory faculty (or experiences) partly in virtue of a belief B1 that my auditory faculty was accurate on a particular occasion. In order for B1 to serve as part of my justification for believing my auditory faculty is reliable B1 must be justified. B1 is

---

<sup>91</sup> The a priori creates difficulties for this point but these needn't concern us currently.

<sup>92</sup> My discussion of the skeptical worries lurking in the background of the problem of easy justification is very much influenced by Van Cleve's (2003) discussion of externalism, easy knowledge, and skepticism.

a conjunction of two claims: (i) my auditory faculties are telling me that a plane is nearby, and (ii) there is a plane nearby. This second claim is most important for our purposes. Once we've accepted JT and JR, however, vision cannot provide the required justification for (ii) unless I already have justification for believing my vision is reliable (by JR). At which point JT entails that I must already have justification for believing that particular deliverances of vision are true. How do I get this justification? I cannot move in a circle and appeal to my auditory faculties. Vision, remember, was meant to justify the reliability of my auditory faculty *prior* to the latter being capable of justifying any beliefs. You might appeal to a third source  $\Omega$  in an attempt to vindicate the testimony of vision, but (by JR)  $\Omega$  cannot provide justification for any beliefs unless I already have justification for believing  $\Omega$  is reliable. And so the same reasoning applies to  $\Omega$ . This repeats *ad infinitum*. Sooner or later, the strategy of appealing to an independent source as a means to justifying a source's reliability will be frustrated by the fact that we have a finite number of belief sources. In fact, the regress is even more vicious than this suggests. Even if one had an *infinite* supply of sources the regress would still be vicious and lead to an extreme form of skepticism. If JR and JT are both true then there could never be an initial source that could "get the justificatory juices flowing." JR and JT cannot both be satisfied.

Our skeptical worry here is just the ancient problem of the criterion. Ancient skeptics worried about how, in light of the persistent disagreement across all forms of inquiry, we can reasonably trust that our beliefs get at truth. According to these skeptics, we ought to withhold assent from a claim until we find some feature of the claim that distinguishes it from false claims. Determining whether a claim is true or false requires us to apply some criterion distinguishing true and false appearances.

Unfortunately, disagreement exists about what the correct criterion of truth is and a criterion won't help distinguish true and false appearances if it's a bad criterion. Deciding which truth criterion is correct, however, would require an application of an independent criterion. In *Outlines of Pyrrhonism*, Sextus Empiricus explains the resulting dilemma:

[I]n order to decide the dispute that has arisen about the criterion, we have need of an agreed-upon criterion by means of which we shall decide it; and in order to have an agreed-upon criterion, it is necessary to have decided the dispute about the criterion. Thus, with the reasoning falling into the circularity mode, finding a criterion becomes aporetic; for we do not allow them to adopt a criterion hypothetically, and if they wish to decide about the criterion by means of another criterion we force them into an infinite regress.<sup>93</sup>

This is the reasoning that led us to the conclusion that JR and JT cannot *both* be satisfied.

The problem of easy justification motivates accepting one of the principles, i.e. JR, constituting one horn of this ancient dilemma. Our considerations above motivate accepting the other horn, i.e. JT. We can thereby view the problem of easy justification as motivating one piece of the skeptical worry posed by the problem of the criterion:

“EASY” SKEPTICISM

1. If  $\neg$ JR then easy justification reasoning is legitimate [see 2.2.1 and 2.2.2].
  2. Easy justification reasoning is illegitimate [see 2.2.3].
  3. JR [From 1 & 2].
  4. JT [See the beginning of 2.3.1]
  5. If JR & JT, then it's impossible for any belief source to provide justification for belief  
[The Problem of the Criterion; see the previous three paragraphs].
- 

6. Therefore, it's impossible for any belief source to provide justification for a belief.

This conclusion is a radical form of skepticism. First, the skepticism concerns *all* beliefs rather than a specific subset of beliefs such as those about the external world. Second, the skepticism concerns the very *possibility* of beliefs possessing *any degree of justification!*

This is the puzzle I hope to dissolve. We want a theory of justification that neither makes justification too easy nor impossible. Our discussion up to this point suggests that this isn't merely a difficult tight rope to traverse but, rather, there isn't even a tight rope that one

---

<sup>93</sup> Sextus Empiricus (1996), pp. 128-129.

could attempt to walk. If you deny JR justification appears too easy. If you accept JR justification appears impossible. There appears to be no conceptual space leftover.

### 2.3.2 Easy Responses to “Easy” Skepticism

Some will argue that we can simply ignore “EASY” SKEPTICISM either by appealing to its self-defeating nature or by invoking a kind of Moorean response. While it’s possible that there is a grain of truth in these replies, we’ll see that they are far from offering a philosophically satisfying solution to our puzzle. Despite the name, it’s anything but easy to come up with a satisfying response to “EASY” SKEPTICISM.

Our skeptical argument is self-defeating: if the conclusion is true, we have no justification for believing its premises and therefore the argument provides no justification for believing its conclusion. If the conclusion of the skeptical argument is correct then none of our beliefs are justified, including the belief that none of our beliefs are justified!

Can we ignore “EASY” SKEPTICISM on account of this self-defeating nature? A skeptic finds herself in an odd predicament of asserting a claim that *by her own lights* is completely unjustified. There is something quasi-paradoxical about believing or asserting *P* while simultaneously believing you’re unjustified in believing *P*. Having recognized that “EASY” SKEPTICISM is self-defeating also makes it appear peculiar that anyone would spend as much time as I have offering what appear to be putative *reasons* for believing JR and JT. My discussion looks like an attempt to show that JR and JT are *reasonable* principles so as to argue that believing these principles or any other claim is unreasonable.

There is something odd here, but none of this is sufficient to dismiss “EASY” SKEPTICISM. Even if the skeptic is in a bad (or at least odd) position, pointing this out shouldn’t alleviate any worries for an anti-skeptic. Just because the skeptic is in a bad position by the skeptic’s own lights doesn’t mean the anti-skeptic is in any better a position. We can simply view the skeptical discussion as an attempt to illustrate that our ordinary thinking about justification appears inconsistent. Common sense has it that we are justified

in the majority of our ordinary beliefs. However, our ordinary intuitions about the canons of rationality also push us toward an acceptance of JR and JT, which leads to global skepticism. Our anti-skeptical intuitions and our intuitions about the canons of rationality are inconsistent. Clearly this is a bad position to be in. Simply showing that there is something bad about the skeptic's position does nothing to resolve our own inconsistency.

A satisfying response to our puzzle must remove the inconsistency of our anti-skeptical intuitions with our intuitions concerning the canons of rationality. How might the anti-skeptic do this? Many philosophers will invoke a Moorean response to our skeptical worries. G.E. Moore famously turned skeptical arguments against the very principles of justification used to generate the skeptical conclusion.<sup>94</sup> In order to illustrate the general move, consider three claims:

- (i) I'm justified in believing that *this* is a pencil (uttered while looking directly at a pencil in ideal conditions).
- (ii) JR and JT
- (iii) If JR and JT are true then global skepticism about justification is true.

These three claims are individually compelling. The first strikes us as common sense. The latter two strike us as intuitive when we take into account the considerations of 2.2 and 2.3.1.

Despite being *individually* compelling, these three claims are *jointly* inconsistent. At least one of the claims is false and consistency demands that we give up at least one of them. How might we argue against one (or more) of these claims? Interestingly, we could use any two of these claims to produce a valid argument for the denial of the remaining claim. The two arguments of interest here are, first, the skeptical argument which argues from (ii) and (iii) to the denial of (i) and, second, the anti-skeptical argument that argues from (i) and (iii) to the denial of (ii). Which of these two arguments should we accept? Presumably we should

---

<sup>94</sup> Moore's discussion was actually carried out in terms of knowledge.

accept the argument whose premises have the most pre-theoretic plausibility. Since the skeptical and anti-skeptical arguments both accept the conditional premise our decision of whether to accept the skeptical or anti-skeptical argument boils down to whether (i) or (ii) has more pre-theoretic plausibility.

Moore's basic suggestion was that a claim like (i) is always more plausible than any of the complicated and abstract principles concerning justification or knowledge upon which skeptical arguments rely. Thus, an anti-skeptic can use the anti-skeptical argument to justify her rejection of the skeptic's epistemological principles (i.e. the combination of JR and JT in our case) and thereby remove the inconsistencies in her beliefs. With these insights in hand, are we finally in a position to simply ignore "EASY" SKEPTICISM?<sup>95</sup>

These considerations might provide us with reasons for thinking that "EASY" SKEPTICISM must go wrong *somewhere* and thereby justify a rejection of its conclusion. Perhaps considerations of self-defeat and Mooreanism are sufficient for preventing ourselves from *believing* global skepticism. I'm sympathetic to such a claim but what I want to note is that, even if all of this is right, we still lack a philosophically satisfying response to our puzzle. I don't think these considerations are even necessary for preventing ourselves from believing global skepticism. One of the grains of truth in pragmatic replies to skepticism is that almost no one is ever, practically speaking, convinced by skeptical arguments that target significant sets of our commonsense beliefs. We *naturally* believe that most of our beliefs are rational and it's highly doubtful that clever skeptical arguments (even if they are correct!) will ever override the part of our nature responsible for this anti-skeptical attitude.

But our goal as epistemologists isn't the mere *preservation* of the belief that many of our beliefs are rational. Our goal should be to achieve an understanding of what justification

---

<sup>95</sup> For Moore's writings on skepticism see his essays (1993a), (1993b), (1993c), and (1993d): *The Refutation of Idealism*, *Hume's Theory Examined*, *A Defense of Common Sense*, *Proof of an External World*, and *Certainty* reprinted in Baldwin (1993).

is and *how* we have justification for the beliefs we do (assuming our anti-skeptical attitude is correct).<sup>96</sup> The Moorean considerations can only tell us that the skeptical argument goes wrong *somewhere*.<sup>97</sup> These considerations don't provide us with any kind of understanding of justification itself nor do they provide an understanding of *how* particular beliefs are justified. Is JR false? Is JT false? Are both false? If so, then where did our reasoning in favor of those principles go wrong? Answering these kinds of questions is essential to a satisfying response to our puzzle and the Moorean considerations can't answer these questions.

Achieving a philosophically satisfying solution to our puzzle requires us to identify where "EASY" SKEPTICISM goes wrong. We also need to identify *why* it goes wrong where it does. And, finally, we need to have good reasons by our own lights for making these identifications. With this in mind, we'll look at attempts to provide a more satisfying solution to our puzzle in the next chapter.

---

<sup>96</sup> Notice that I am here implicitly advocating for the primary philosophical importance of metaepistemology over normative epistemology.

<sup>97</sup> Here I'm admitting that metaepistemology and normative epistemology cannot be completely separated even though this might be ideal.



## CHAPTER THREE

## PROPOSED SOLUTIONS TO OUR PUZZLE

Chapter two focused on two principles concerning the relationship between a source's ability to confer justification and a subject's possessing justification for believing that the source is reliable:

**JR** Necessarily, a belief source  $\Phi$  provides S with justification for believing a particular deliverance (Bp) of  $\Phi$  only if S already has justification for believing that  $\Phi$  is reliable.

**JT** Necessarily, a subject S has justification for believing a source  $\Phi$  is reliable only if S already has justification for believing that particular deliverances of  $\Phi$  are true.

JR and JT both seem well motivated.

Our puzzle is that JR and JT can't both be satisfied and the result would be an extreme global skepticism regarding justification. As such, we arrived at the skeptical argument that serves as a foil for various theories of justification:

“EASY” SKEPTICISM

1. If  $\neg$ JR then easy justification reasoning is legitimate [see 2.2.1 and 2.2.2].
  2. Easy justification reasoning is illegitimate [see 2.2.3].
  3. JR [From 1 & 2].
  4. JT [See the beginning of 2.3.1]
  5. If JR & JT, then it's impossible for any belief source to provide justification for belief  
[The Problem of the Criterion; see the previous three paragraphs].
- 
6. Therefore, it's impossible for any belief source to provide justification for a belief.

Solving our puzzle requires giving a philosophically satisfying response to “EASY” SKEPTICISM. Achieving this is anything but easy. A satisfying response requires identifying where the argument goes wrong, why it goes wrong there, and we better have good reasons for making such identifications.

In “EASY” SKEPTICISM premise 3 is entailed by 1 and 2. And I’m at a loss as to how one could deny premise 5. This leaves three possible strategies for developing a philosophically satisfying response to the argument: (i) deny premise 1 and argue that denying JR doesn’t commit one to easy justification; (ii) deny premise 2 and argue that easy justification is, contrary to appearances, unproblematic; or (iii) deny premise 4 and argue that we can have justification for believing a source is reliable *prior* to possessing any justification for believing that any particular verdict given by that source is true. Most of the solutions we discuss are complicated and don’t necessarily fit perfectly into a single strategy to the exclusion of the other two; most involve some mixture of the three solutions.

This chapter discusses proposals that ultimately fail. Such an investigation serves multiple purposes. First, grappling with the proposals illustrates how treacherous our puzzle is. Second, I chose to discuss these particular proposals because, despite their inadequacies, we can glean important lessons from the proposals. Finally, a general conclusion that I hope to draw from this discussion is that an independent investigation of evidence, possession of evidence, and evidential support will be crucial to resolving our puzzle.

### 3.1 Dialectical vs. Epistemological Failure

Peter Markie has argued that easy justification isn’t actually illegitimate (i.e. he denies premise 2 of “EASY” SKEPTICISM).<sup>98</sup> Markie argues that our intuition that there is something problematic with the easy justification reasoning is correct but that the problem amounts to a dialectical rather than an epistemological failure.

---

<sup>98</sup> See Markie (2005). Van Cleve (2003) also takes the tack of arguing that easy justification is legitimate. Essentially, he argues against a variety of proposed solutions that attempt to avoid skepticism without admitting “easy justification”—most of these arguments consist of Van Cleve showing that the proposed solutions don’t actually prevent easy justification. Van Cleve then concludes that since the only remaining options are to either allow easy justification or to accept a global skepticism, easy justification is actually epistemically legitimate. This might be motivated by an adherence to a kind of Mooreanism (refer back chapter 2 section 2.3.2). Such a response isn’t philosophically satisfying since it doesn’t help identify why our reasons for taking the easy justification reasoning to be illegitimate go wrong, whereas Markie’s discussion attempts just this.

Cohen's original presentation of the problem of easy justification focused on the following scenario in an attempt to provide a vivid illustration of the problem:

Suppose my son wants to buy a red table for his room. We go in the store and I say, "That table is red. I'll buy it for you." Having inherited his father's obsessive personality, he worries, "Daddy, what if it's white with red lights shining on it?" I reply, "Don't worry—you see, it looks red, so it's not white but illuminated by red lights... [Or] I reply, "Well—look, the table is red... since it's red, it can't be white with red lights shining on it. See?" ... Surely he should not be satisfied with this response.<sup>99</sup>

Markie's discussion focuses heavily on this imagined social interaction. He suggests that the father's belief that his current apparent perception is accurate does in fact gain justification. We, however, *mistakenly* believe that the father cannot gain justification for his belief in this way because it's obviously a horrible response to his son's concerns.

We can bolster Markie's claims by distinguishing the *property* of justification from our *practice* of justification. A belief is justified when that belief is epistemically likely to be true in some sense yet to be specified. *Justifying* a belief is a social practice that involves *appropriately* citing reasons for one's belief so as to illustrate to an interlocutor that one's belief is likely to be true. This practice is governed by a variety of argumentative and social norms, and it's an open possibility that a person's belief possess justification despite the fact that the person cannot satisfy the relevant argumentative/social norms when trying to justify her belief to an interlocutor. Having good reasons for a belief and being able to *appropriately* cite those reasons in a way that should convince a third party are very different matters—the "should" is important since convincing a third party depends on psychological factors and an interlocutor could be inappropriately stubborn.

Easy justification reasoning is problematic, according to Markie, because it begs the question against a skeptic. Such reasons can't be appropriately cited in that particular

---

<sup>99</sup> Cohen (2002), pp. 314-315. The extended quote bundles together remarks that Cohen makes while discussing the problem in both the context of evidentialism and of reliabilism.

argumentative context. But this is a problem for meeting the norms of argumentation not for meeting the epistemic norms governing the proper ways of forming beliefs.

In Cohen's scenario involving the father and son, the son (i.e. the skeptic) questions whether the father's particular perception is accurate and whether the table is in fact red. Therefore, the father's appeal to the fact that the table is red on the grounds that it appears red is an appeal to the very claim that the son (i.e. the skeptic) is challenging. These reasons cannot be appropriated cited in this argumentative context since doing so begs the question. Citing these reasons *shouldn't* convince the son (i.e. the skeptic) to abandon his withholding assent. None of this, however, shows that the reasons the father cites aren't actually good reasons that provide justification for the father's belief.

Markie is right that the distinction between dialectical and epistemic failure/success is important. And so his discussion provides us with an important lesson: we shouldn't infer that a subject's reasons fail to be justification-conferring from the fact that citing such reasons in an attempt to justify her belief to a skeptic would be question begging.

My discussion in 2.2.3 is surely sufficient for establishing that easy justification is not merely a dialectical problem. In my discussion I purposefully ignored the kind of social interaction that Cohen utilizes and Markie focuses on. My presentation of the problem wasn't in terms of a skeptic challenging the accuracy or reliability of S's belief source where S then utilizes the easy justification reasoning in an attempt to meet the challenge. I merely stated that a subject in fact arrived at her belief that her source is accurate or reliable in the way illustrated by the easy justification reasoning. I then pushed the intuition that one could not generate *new* justification for a source's accuracy or reliability in the way suggested. Such reasoning is guaranteed to reach the favorable conclusions even if employed in situations where the conclusion is false. Moreover, the reasoning relies on an operative commitment to a source's reliability and, therefore, *if* it generated *new* justification for believing the content of this commitment then it would generate justification ex nihilo.

I hope that this is enough to convince the reader that what is problematic about easy justification is more than the fact that citing the relevant reasons would be question begging in a certain dialectical context. However, contrasting easy justification reasoning with other reasoning that is intuitively *legitimate* but that nevertheless cites the relevant reasons in a dialectical context that would be question begging provides further support. If easy justification reasoning strikes us as problematic in comparison to other forms of reasoning that beg the question then there must be something more to the problematic nature of easy justification than the fact that citing such reasons would be question begging against a skeptic.

We've already considered such contrast cases. In chapter 2 we compared the easy bootstrapping on induction with the traditional inductive justification of induction. There we noted that even if the traditional inductive justification of induction is problematic the easy bootstrapping on induction is clearly *more* problematic. It's clear, however, that presenting either argument to an inductive skeptic would beg the question. And so what strikes us as problematic in cases of easy justification must go beyond its mere question begging-ness. Similar remarks apply to the contrast between easy justification with memory and the AAM style reasoning (where one combines an apparent memory of having had an apparent memory as of  $P$  with an apparent memory as of  $P$ ). Each style of reasoning relies on memory and would therefore beg the question against the skeptic about memory. Nevertheless, the easy justification reasoning strikes us as *more* problematic and so we cannot explain our uneasiness with easy justification purely in terms of its being a dialectical failure. Easy justification is epistemically problematic.

### 3.2 A Reidian Inspired Rejection of JT

Thomas Reid is another source of inspiration for an attempt to resolve our puzzle.<sup>100</sup> Reid famously argued that the belief that our belief sources (at least the natural ones)<sup>101</sup> are reliable is what he calls a “first principle”:

Another first principle is, that the natural faculties [belief sources], by which we distinguish truth from error, are not fallacious.<sup>102</sup>

First principles are propositions that are foundationally justified due to their self-evidence.

Self-evident propositions are those that are, in some sense, evidence for their own truth. The idea here is that the truth of a self-evident proposition is made immediately evident to a subject *S* in virtue of *S*'s understanding the relevant proposition. Paradigm examples include propositions such as that  $2+2=4$  and that *all bachelors are unmarried*. I know these propositions to be true because my mere understanding of the relevant propositions makes their truth evident to me. Such propositions are evidently true to me without requiring any kind of empirical investigation, e.g. the evident truth of the proposition that *all bachelors are unmarried* doesn't require anything like performing a survey of bachelors.

That a proposition is self-evident doesn't entail that its truth is apparent to everyone; a toddler might fail to understand the proposition that  $2+2=4$  and therefore fail to see its truth. Moreover, I might only have a partial understanding of various complex mathematical propositions such that their truth isn't immediately obvious to me even though *were* I to have

---

<sup>100</sup> I want to emphasize that the following discussion pertains to a response *inspired* by Reid's writings. I'm not committing myself to the claim that the views discussed below are a correct interpretation of Reid's views.

<sup>101</sup> I don't know if there is a good gloss for what Reid means in referring to the natural faculties or belief sources. I take it that the contrast is meant to be between sources like perception, memory, induction, testimony, etc. and sources like gas gauges, rain clouds, litmus paper, etc. The former being the so-called natural sources and the latter being the non-natural sources. Such a claim becomes more difficult, however, since Reid distinguishes between non-acquired and acquired perception. Thankfully, the discussion that follows doesn't turn on any of these difficult issues.

<sup>102</sup> Reid ([1785] 2002), p. 480. I've adapted Reid's discussion of knowledge to a discussion of justification.

a fuller understanding of the propositions their truth would be immediately evident to me. The upshot is that Reid holds that we have *foundational* a priori justification for believing that our (natural) belief sources are reliable.

In line with our discussion of the problem of easy justification (2.2.3.3) Reid explains:

If any truth can be said to be prior to all others in the order of nature, this [that our sources are reliable] seems to have the best claim; because in every instance of assent, whether upon intuitive, demonstrative, or probable evidence, the truth of our faculties is taken for granted, and is, as it were, one of the premises on which our assent is grounded.<sup>103</sup>

Reid appears to explicitly recognize Sosa's point that the forming a belief using a source  $\Phi$  manifests a commitment to  $\Phi$ 's reliability. Reid also recognizes that we wouldn't form a belief via  $\Phi$  if we lacked this commitment. He explains that the truth of our faculties' reliability is *prior* to all others, and context suggests that "prior" here indicates epistemic priority. Therefore, Reid appears to give an explicit endorsement of JR. In order for a source to justify a belief we must already justifiably presuppose its reliability. Thankfully, according to Reid, the general belief that our sources are reliable is justified a priori and foundationally. This general belief thereby provides justification for believing of each particular source (e.g. perception, memory, induction, testimony, etc.) that it's reliable. Reid is taking on the herculean task of denying JT, i.e. denying premise 4 of our skeptical argument.

An initial hurdle for the Reidian response must overcome is what we can label the "mysteriousness objection." The proposition that a belief source is reliable is both general and contingent. Van Cleve claims to understand how general necessary propositions such as that *all triangles have interior angles that add up to 180°* and specific contingent propositions such as that *I exist* can be self-evident. However, Van Cleve complains: "[A]re there any

---

<sup>103</sup> Reid ([1785] 2002), p. 481.

immediately evident propositions that are both general and contingent? This is a question to which I find it hard to say yes.”<sup>104</sup>

I take it that the proposition that *I exist* is foundationally and infallibly justified but I also think it is misleading to describe the proposition as self-evident. More importantly, however, it appears as if Van Cleve and the Reidian are simply begging the question against one another. Van Cleve claims that he finds it hard to think that there are general contingent propositions that can be self-evident. The Reidian provides an example of such a proposition and claims that it is self-evident. How are they not just butting heads?

The main force of Van Cleve’s comments can be reconstructed into a more compelling objection. Intuitively, general contingent propositions about a source’s reliability aren’t the kind of proposition that can be justified merely in virtue of our understanding the proposition. Insofar as the Reidian solution depends on denying this it departs from our intuitions concerning the canons of rationality. A philosophically satisfying defense of the Reidian solution must therefore develop and defend an account of a priori justification as it applies to paradigmatic cases (including paradigmatic beliefs that cannot be justified a priori) and go on to show that this account can be extended to the justification of general contingent propositions such as that *my belief sources are reliable*.

Traditionally, a priori justification has been understood as involving a direct grasp of conceptual relations necessitating the truth of a proposition. This direct grasp is to be understood as a *factive* state: you can’t directly grasp conceptual relations involved in the proposition *P* that necessitate *P*’s truth unless there really are such relations necessitating *P*’s truth. Consider one of the prime candidates for a proposition that can be justified a priori: that *all bachelors are unmarried*. When I understand and entertain this proposition I have a direct grasp of the concept UNMARRIED’s relation of being a proper-part of the concept

---

<sup>104</sup> Van Cleve (2003), p. 51.



BACHELOR. This proper parthood relation necessitates the truth that *all bachelors are unmarried*. I have a direct grasp of the very necessity of this proposition's truth when I understand and entertain it properly and thereby have a priori justification for my belief.<sup>105</sup>

Much more needs to be said in order to defend this traditional view—specifically, what is it to “directly grasp” various conceptual relations?—but it's clear that a Reidian solution to our puzzle must reject the traditional conception of the a priori. We cannot have a direct grasp of a relationship between our concepts BELIEF SOURCE and RELIABLE that necessitates the truth of the proposition that *my belief sources are reliable*. This proposition isn't necessary and so there is no such relationship for us to grasp! We're well aware that our belief sources (e.g. perception, memory, induction, etc.) can get things wrong. Indeed, we're well aware that it's possible for our belief sources to get things wrong most of the time!

One way to defend the Reidian solution is to abandon a traditional view of the a priori in favor of what I'll call an experiential theory of a priori justification. These theories maintain that a priori justification is provided by a paradigmatic non-sensory phenomenal experience distinct from the entertaining of a proposition. Given that this non-sensory experience is by hypothesis logically distinct from the entertaining of a specific kind of proposition, there isn't any in principle mysteriousness to the idea that this experience could be prompted by our entertaining general contingent propositions such as that *my belief sources are reliable*. If this kind of account of a priori justification is correct then the Reidian view can successfully deflect the mysteriousness objection.

---

<sup>105</sup> This captures the notion of *foundational* a priori justification. There are also cases where my understanding of a proposition is fuzzy and I cannot grasp the conceptual relations that necessitate its truth directly. A good example would be a very complicated mathematical theorem. Label this theorem *T*. Despite lacking a clear enough understanding of *T*, my fuzzy grasp might be clear enough to grasp the conceptual relations that necessitate the truth that *T\** entails *T*. If I also have a clear enough grasp of *T\** such that I have a direct grasp of the conceptual relations that necessitate the truth of *T\** I would thereby have *inferential* a priori justification for believing *T*.

Why would one adopt an experiential theory of a priori justification? One way might proceed from the assumption that certain beliefs are justified a priori but others are not, comparing the two cases so as to find a difference in why we hold the two beliefs, and concluding that such a difference explains why we have a priori justification in the former but not the latter case. In carrying out such a method you might start with the assumption that you possess a priori justification for your belief that  $2+2=4$  but not for your belief that *it's raining outside*. When we introspect what do we find that is present in the former case, absent in the latter case, and that might plausibly contribute to my justification for the former? When you carry out this method you might come to the conclusion that the relevant difference is a kind of experience or feeling as of obviousness upon entertaining the former proposition. When I entertain that  $2+2=4$  the claim just feels obvious but when I entertain that *it's raining outside* it doesn't feel this way.

Alternatively, an advocate of a traditional conception of a priori justification holds that the difference involves my direct grasp of the necessity of the former proposition which might also explain my feeling of obviousness. An advocate of the experiential theory will likely be quick to argue that it's the separate feeling of obviousness that is responsible for my belief: (i) if I didn't have the experience of obviousness I wouldn't believe  $2+2=4$  even if I were to directly grasp its necessity, and (ii) I would still believe  $2+2=4$  if it were to prompt this feeling of obviousness absent a direct grasp of the proposition's necessity. I don't endorse this experientialist line but I do want to suggest how someone might be led to this view of a priori justification.<sup>106</sup>

---

<sup>106</sup> Reid proposed a distinct but intimately related account concerning *foundational* a priori justification. Reid suggested that when I consider the denial of a first principle, a faculty that he calls commonsense causes us to have an experience of absurdity or "an emotion of ridicule", and it's this experience that supposedly provides us with justification for accepting the first principle. See Reid ([1785] 2002). See Bergmann (2004) and (2008) for some discussion of Reid's appeal to the emotion of ridicule. The difference between the view that a priori justification arises from a feeling of obviousness when entertaining a proposition and the view that it arises from a feeling of ridicule when entertaining the denial of a proposition are subtle but irrelevant to my current purposes. For

An experientialist must be careful in her characterization of the relevant experience as one of “obviousness.” Such a feeling is meant to do technical work and is being introduced via ostension. We need to be careful in recognizing exactly what we’re ostending. We shouldn’t take the relevant feeling of “obviousness” to apply too widely. When I entertain the proposition that *it’s raining here and now* while undergoing a visual and tactile experience as of rain falling on me there is certainly a sense in which this proposition strikes me as obvious, but no one holds that I thereby possess a priori justification for believing this proposition. The experience as of obviousness when entertaining the proposition *it’s raining here and now* in these conditions is different in kind from the feeling of obviousness when I entertain the proposition  $2+2=4$ . Adequately characterizing the difference between the two feelings of “obviousness” is difficult but important since it’s only the latter kind of feeling that could be used in an experiential account of a priori justification.

So what exactly is the Reidian solution to our puzzle? If we allow that belief sources such as perception, memory, testimony, induction, etc. can be justification conferring absent prior justification for believing these sources are reliable then we allow easy justification. And so the Reidian accepts that these sources are only justification conferring when we have prior justification for their reliability. However, we were wrong in the last chapter when we suggested that the fact that the reliability of these sources is a contingent matter entails that the reliability of these sources could only be justified empirically. Once we’ve rejected the traditional theory of a priori justification in favor of the experientialist theory we can see that it’s possible for the contingent reliability of these sources to be justified a priori. Thus, JT is false and we avoid skeptical impalement.

Now we must discuss the critical issue for evaluating the Reidian inspired solution. Everything I’ve said about this proposal crucially depends on an experience as of

---

ease of presentation I focus on the view that utilizes the feeling of obviousness as a stand in for all experiential views of a priori justification.

obviousness involved in a priori justification being capable of providing justification for believing that perception, memory, induction, etc. are reliable. Such an experience constitutes the *grounds* or *basis* for these beliefs. These beliefs about reliability are themselves delivered via a belief source, *the a priori*, which can be assessed in regards to its accuracy and reliability. What should we say about the relationship between our justification for believing the verdicts of the a priori and our justification for believing that the a priori is reliable? Should we accept JR, *or* should we hold that the a priori is the sole basic justification source that provides a counterexample to JR?

Our Reidian inspired solution should *not* accept JR in its full generality as this would lead back to global skepticism. If we accept JR in full generality then the a priori will not confer justification unless we have prior justification for believing that the a priori is reliable. From where would our justification for believing the a priori is reliable derive? We can't appeal to perception, memory, induction, etc. in order to justify the a priori's reliability. There doesn't seem to be any independent source to which we could appeal! You might suggest that we can have a priori justification for believing that the a priori is reliable—a claim to which the Reidian inspired solution will eventually be pushed—but clearly the a priori cannot confer justification on the belief that the a priori is reliable *prior* to its being justification conferring!

It seems as if the Reidian solution ought to deny JR as applied to the a priori. Unfortunately, this brings back the problem of easy justification. If the a priori is a basic justification source then we can justify the accuracy/reliability of the a priori as follows:

#### EASY BOOTSTRAPPING ON (EXPERIENTIAL) A PRIORI

1.  $P_i$  [Justified by an a priori experience as of  $P_i$ ].
2. I have an a priori experience as of  $P_i$  [Known on the basis of introspection].<sup>107</sup>

---

<sup>107</sup> We're supposing a rejection of JR's applicability to the a priori and so it can render introspection justification-conferring by justifying the reliability of introspection absent justification for believing the a priori is reliable.

3. My a priori experience as of  $P_1$  is accurate [From 1 & 2].
  4.  $P_2$  [Justified by an a priori experience as of  $P_2$ ].
  5. I have an a priori experience as of  $P_2$  [Known on the basis of introspection].
  6. My a priori experience as of  $P_2$  is accurate [From 4 & 5].
- [Repeat for  $P_3$ - $P_N$ ]
- 

7. Therefore, the a priori is reliable.<sup>108</sup>

The structure of the easy justification reasoning is just as worrisome in the case of the experiential a priori. Easy bootstrapping still exhibits the problematic features of no failure and ex nihilo generation of justification.

Forming beliefs via an a priori experience (i.e. forming the beliefs  $P_1$ - $P_N$ ) manifest a commitment to the reliability of such a priori experiences. Rejecting JR's application to the a priori entails that this commitment could be unjustified yet its operation would provide justification for  $P_1$ - $P_N$ . These beliefs then confer *new* justification for the belief that the a priori is reliable, a belief with the same propositional content as the previously unjustified commitment. Thus, the problem of generating justification ex nihilo arises even when bootstrapping on the experiential a priori. The bootstrapping is also guaranteed from the outset to reach the favorable conclusion regarding the accuracy and reliability of the a priori. Such reasoning shouldn't be capable of generating *new* justification for the conclusion.

One complication arises. You might think the a priori is *necessarily* accurate and reliable and *couldn't* be deployed in a scenario where the conclusion is false. The a priori would perfectly discriminate between cases where the conclusion is false from those where it is true. This objection, however, turns on a traditional understanding of the a priori where it involves a factive state of grasping a proposition's necessary truth and is therefore infallible.

---

<sup>108</sup> Fn. 107's points about introspection's justification conferring ability on this Reidian proposal apply *mutatis mutandis* to induction.

We must remember that the Reidian inspired solution is committed to a rejection of this traditional picture of a priori justification.

Once we've conceptualized a priori justification in terms of an experiential theory it's clear that the a priori could produce inaccurate beliefs some or even all of the time! This is precisely the point Descartes is making when he appeals to the evil genius scenario in an attempt to throw doubt on  $2+2=4$ . If my only reason for believing  $2+2=4$  is the feeling of obviousness then there are grounds for doubt. An evil genius could bring it about that I have a feeling of obviousness only when considering propositions that are in fact false. One needn't even appeal to such exotic possibilities to illustrate this point. People have supposedly had this feeling of obviousness when considering Euclid's fifth postulate and thereby believed this postulate. Non-Euclidian geometries appear to illustrate the possibility that the postulate is false. Scientific evidence even suggests that space is *in fact* Non-Euclidian. Similarly, people have previously had a feeling of obviousness when entertaining the comprehension axiom of naïve set theory and thereby believed naïve set theory. As we know, naïve set theory is actually inconsistent and the culprit is the axiom of comprehension. It's clear that *if* an experiential theory of a priori justification is correct then we could go through the easy bootstrapping reasoning in cases where the a priori is inaccurate and unreliable. Such reasoning would lead to the favorable conclusion in *all* of the scenarios where the a priori is inaccurate and unreliable and so suffers from the problem of no failure.

If the Reidian solution is to work it needs to identify something distinctive about the a priori that allows the Reidian inspired solution to deny JR without being susceptible to the problem of easy justification. There must be something special about the a priori such that this reasoning doesn't generate *new* justification for the conclusion. In other words, the Reidian solution must also deny premise 1 of "EASY" SKEPTICISM. Yet again, a tack of this sort can be gleaned from Reid's writings.

Notice that Reid's first principle mentioned above refers to the reliability of *all* natural faculties. This principle vouches for the reliability of perception, memory,

introspection, testimony, induction... *and the a priori itself!* The principle contains an element of self-reference. Keith Lehrer explains, “The principle vouches for itself. It loops around back on itself.”<sup>109</sup> Moreover, the a priori vouches for itself during every instance that it delivers a belief. When I form the belief that  $2+2=4$  via the relevant a priori experiential episode, I manifest a commitment to a priori experiences being reliable and in this very instance the a priori confers justification on this very commitment as well. This explains one of Reid’s most cryptic remarks concerning his “most fundamental” first principle:

How then come we to be assured of this fundamental truth on which all others rest? Perhaps evidence, as in many other respects it resembles light, so in this also, that as light, which is the discoverer of all visible objects, discovers itself at the same time: so evidence, which is the voucher for all truth, vouches for itself at the same time.<sup>110</sup>

Every time the a priori delivers a belief such as that  $2+2=4$ , that *perception is reliable*, etc. it provides correlative justification for believing that the a priori is reliable. The “looping” nature of the a priori entails that when forming the beliefs  $P_1-P_N$  in EASY BOOTSTRAPPING ON (EXPERIENTIAL) A PRIORI we must have correlative and independent justification for believing the conclusion that the a priori is reliable. Therefore, such reasoning doesn’t provide any *new* justification for the conclusion. The reasoning suffers from what has been labeled transmission failure. If we have justification for the premises then we also have justification for the conclusion. What is important is that, given the “looping” nature of the a priori, this justification doesn’t transmit from the premises to the conclusion.

The Reidian solution must reject JR since even the a priori cannot confer justification on the proposition that *the a priori is reliable* prior to the a priori being justification conferring. As such, a Reidian inspired solution must adopt a modified version of JR:

---

<sup>109</sup> Lehrer (1989), p. 43.

<sup>110</sup> Reid ([1785] 2002), p. 481.

**JR\*** Necessarily, a *non-looping* belief source  $\Phi$  (i.e. a belief source that doesn't always immediately vouch for its own reliability) provides S with justification for believing a particular deliverance (Bp) of  $\Phi$  only if S already has justification for believing that  $\Phi$  is reliable.

We've seen that the looping nature of the a priori stops the easy justification reasoning from generating *new* justification, and therefore it's the looping nature that allows one to deny JR and accept JR\* so as to avoid easy justification.

On the Reidian inspired solution the belief that the a priori is reliable is justified in a way that doesn't fit into our mold of easy justification reasoning; it isn't a product of "reasoning" at all. It's nevertheless useful to consider if the way in which justification is conferred on this belief shares the problematic features of easy justification reasoning.

Acquiring justification for believing the a priori is reliable in the way imagined by the Reidian solution exhibits a kind of circularity (even if it isn't circular "reasoning"), but I've argued that circularity isn't the correct way to characterize the problematic nature of easy justification. We need to consider whether allowing this foundational justification for *the a priori is reliable* on the basis of a feeling of obviousness suffers from issues of no failure and whether it generates justification ex nihilo.

Due to the self-referential character of Reid's "most fundamental" first principle, generating justification for *the a priori is reliable* in this way avoids the problem of generating justification ex nihilo. My feeling of obviousness while entertaining the proposition *the a priori is reliable* is what provides me with justification for believing that the a priori is reliable. In forming this belief I manifest a commitment to experiences as of obviousness being a reliable basis for belief. Justification ex nihilo would occur if the operation of this unjustified commitment could result in a justified belief which would then provide justification for the previously unjustified commitment. This doesn't occur in the imagined procedure. In the self-referential case, the propositional content of the commitment and that of the resulting belief are identical. Therefore, the experience as of obviousness simultaneously provides justification for both the commitment and the resulting belief.



It also appears that the procedure used in generating justification for the reliability of the a priori on the Reidian proposal avoids the problem of no failure. The procedure is best characterized as follows: (i) entertain the proposition *the a priori is reliable*, (ii) if you have an experience as of obviousness then form the belief that the a priori is reliable; (iii) otherwise, don't form the belief that the a priori is reliable. One could follow this procedure and *fail* to arrive at the favorable conclusion about the reliability of the a priori. Entertaining *the a priori is reliable* might not induce a feeling of obviousness.

Whether this procedure really lacks the problematic feature of no failure, however, is more complicated. As I've characterized a looping source it's a source that vouches for its own reliability. But does a looping source vouch for its own reliability *necessarily* or *contingently*? Reids' analogy to light suggests that it's part of the very *nature* of the a priori that it "discovers" its own reliability. If this is correct then the problem of no failure reoccurs since there are clearly scenarios where feelings of obviousness wouldn't be reliable. If the a priori necessarily vouches for its own reliability then in each of the scenarios the general procedure would inevitably lead to the false conclusion that the a priori is reliable.

I also want to suggest that this proposed Reidian solution that appeals to the looping nature of the a priori is ad hoc and cheats. Such a view must admit that a *mere* a priori experience as of  $2+2=4$  or as of *perception being reliable* wouldn't provide justification for believing the relevant propositions. We've seen that this would allow the easy justification reasoning to produce illegitimate gains in justification. In an attempt to avoid this untoward consequence, the proposed solution admits that a *mere* a priori experience as of  $2+2=4$  won't provide justification. It's only an a priori experience as of  $2+2=4$  accompanied by an a priori experience as of *the a priori is reliable* that provides justification for believing  $2+2=4$ . Such a move is ad hoc and appears utterly magical. Consider an analogy with testimony. Imagine that, having realized the problem of easy justification, I admit that Abe's testifying that there are 8 black marbles in a certain urn cannot *by itself* provide me with any justification for believing that there are in fact 8 black marbles in the urn. Surely I should not go on to claim

that Abe's testifying "My testimony (including this one) is reliable and there are 8 black marbles in this urn" can generate justification for believing that there are in fact 8 black marbles in a certain urn.

The previous criticism essentially boils down to the truth of the following conditional: if a source  $\Phi$  cannot by itself provide justification for believing its *object-level* deliverances  $P, Q$ , etc. then  $\Phi$  cannot by itself provide justification for believing the *meta-level* deliverance that  $\Phi$  is reliable. Examples such as the testimony case support such a principle. But the Reidian solution must deny this.

A related worry is what I call the "direction objection." We worry about the easy justification reasoning because, intuitively, acquiring justification for the reliability of a certain source is *more* cognitively demanding than it would be if easy justification reasoning were legitimate. When we attempt to acquire justification for believing that our color vision is reliable we actually attempt to learn how our perceptual mechanisms work, check our color vision against the testimony of others, get doctors to make sure the mechanisms are working properly, etc. It's only when we have inductive reasoning involving separate faculties which provide independent checks on a source's outputs that we think we can *gain* justification for believing that our color vision is reliable. The problem with the Reidian solution is that it goes in the opposite direction; it makes acquiring this justification *less* demanding. I can just think really hard about the propositions *perception is reliable, memory is reliable, the a priori is reliable*, etc. and thereby acquire justification for these beliefs.

My final two objections to the Reidian solution are directly related to its appeal to an experiential theory of the a priori. First, experiential theories of the a priori are clearly inadequate. Second, the experiential theory might vindicate the *possibility* of justification but it would ultimately lead to a *de facto* global skepticism.

Let's first consider why the kind of experiential theory of a priori justification that the Reidian solution requires must be incorrect. To this end it's interesting to note that we often characterize the notion of *foundational* a priori justification, which we're attempting to

give an account of, in terms of the proposition's being "self-evident". The idea is that there is something about the proposition itself that allows it to be evidence for its own truth. An explanation of a priori justification must in some way appeal to features of a proposition itself that makes it a *candidate* for possessing foundational a priori justification. Experiential theories are woefully inadequate in this respect. According to experiential theories, the only requirement for a proposition's being a candidate for possessing foundational a priori justification is that it be possible that entertaining the proposition produces a distinct state involving a specific kind of experiential episode. The proposition isn't evidence for itself in any sense. It's this distinct experiential state that is evidence for the proposition's truth. Alas, *any* proposition *could* possess foundational a priori justification. But this result is clearly unacceptable and gives us reason to reject the account.

Consider the proposition *it's raining here and now*. This is a paradigmatic example of a proposition for which we don't possess a priori justification. More importantly, this is a paradigmatic example of a proposition for which we *couldn't* possess a priori justification. Allowing for even the possibility that we possess a priori justification for this proposition is reason to reject the view.<sup>111</sup> Given that experiential theories hold that the experiential state which confers a priori justification and the entertaining of a proposition are distinct states, these theories are committed to the possibility that I possess a priori justification for *it's raining here and now*. There is nothing incoherent in my entertaining *it's raining here and now* prompting in me this separate feeling of obviousness. Thus, the experiential account of a priori justification that the Reidian inspired solution requires is false.

---

<sup>111</sup> Some might claim that if God exists then his knowledge of such a claim would be a priori. I leave this complication to the side since I'm not sure how to conceive of knowledge in the case of a being with the characteristics that God is claimed to have. Though I imagine that any proposed account of God's knowledge of such truths would construe this as more akin to introspective (where God's knowledge of the contingent fact rests on his knowledge of his willing such a state-of-affairs) rather than a priori knowledge. However, if God exists, my intuitions concerning his knowledge are so unclear that I am simply at a loss as to what to say about epistemic concepts applied to such a being. My intuitions about epistemic concepts applied to a God-like being get even more unclear once the doctrine of divine simplicity is adopted.

You might attempt to save the experiential theory by tacking on a requirement that the proposition entertained have certain intrinsic features in addition to its prompting a feeling of obviousness. The problem is that, as far as I can tell, any intrinsic features of the proposition one might appeal to would rule out not only the possibility that *it's raining here and now* possesses a priori justification but also the possibility that *my belief sources are reliable* possesses a priori justification. In which case, the Reidian solution wouldn't be available.

Finally, even if the (unmodified) experiential theory of the a priori is correct, it would provide a rather hollow victory against our skeptical threat. The Reidian inspired solution might parry the threat of skepticism posed by "EASY" SKEPTICISM and vindicate the *possibility* of justification—assuming it can avoid the problems discussed so far. I worry that such a view will *as a matter of contingent fact* lead to global skepticism regarding justification. Global skepticism will only be *non-actual* if we actually have this feeling of obviousness when considering the proposition that our belief sources are reliable. Here we need to take care to be phenomenologically honest with ourselves. I want to suggest that we simply *don't* have the phenomenal experiences of obviousness that would, according to this proposed solution, confer a priori justification on the claim that our belief sources are reliable. Thus, this response still leads to a contingent but *actual* global skepticism.

Why think we lack the relevant experience? Isn't it clear that it's obvious to us that perception is reliable? Yes, *in some sense*. It's important to remember, however, our distinction between two different kinds of feelings of obviousness. There is one sense in which we have a feeling of obviousness when entertaining  $2+2=4$ , *there are no round squares*,  $P \vee \neg P$ , etc. We have a feeling of obviousness in another sense when entertaining *it's raining here and now* while undergoing the visual, tactile, and auditory experience as of its raining. When we say that it feels obvious that perception is reliable, is this feeling of obviousness comparable to the former or latter experience? In order for the Reidian inspired solution to avoid global skepticism it must be comparable to the former. But clearly the sense in which it's obvious that perception is reliable is the latter sense.

The best candidate for distinguishing the two kinds of experience of obviousness is by construing the former as a perception of the proposition being necessary.  $2+2=4$  and *there are no round squares* strike us as obviously true because we take it that they couldn't possibly be false. I mean to construe this in a way that is consistent with the experiential theory's insistence that this state is distinct from the state of entertaining a proposition and its commitment to the fallibility of the a priori. We can construe this as an intentional state of *representing* a proposition  $P$  as being necessary. One could be in this state even if  $P$  isn't actually necessary and even if  $P$  is actually false. Entertaining certain propositions may prompt in some people a *mistaken* representation of the proposition's being necessarily true.<sup>112</sup> However, clearly it doesn't strike us as obvious *in this sense* that our belief sources are reliable. We think it's entirely possible that our belief sources aren't reliable. This is illustrated by the fact that we take Descartes' dream and evil genius scenarios to be intelligible. Thus, the Reidian inspired view might vindicate the *possibility* of justification since entertaining the proposition *my belief sources are reliable* could prompt an experiential representation of this proposition being necessarily true. The view leads to *actual* skepticism, however, since entertaining the relevant proposition *doesn't* prompt this experiential representation. I conclude that even if the Reidian view can deal with the other difficulties it could only provide a rather hollow victory against "EASY" SKEPTICISM. We should therefore hold out hope for a better resolution of our puzzle.

### 3.3 The Wright Way of Rejecting JT

Crispin Wright's theory of epistemic entitlement provides another possible solution that denies JT. Wright proposes that we can have justification for believing a belief source is reliable absent any prior justification for believing of any particular verdict of that source that it's true. Unlike the Reidian strategy, however, this justification isn't conferred via any belief

---

<sup>112</sup> For a defense of this kind of view see Bonjour (1998).

source.<sup>113</sup> Our justification for presupposing that perception, memory, induction, the a priori, etc. are reliable is a kind of default status that doesn't rest on any kind of basis, ground, evidence, or on features of the process producing such presuppositions. These presuppositions have a kind of non-evidential justification that Wright refers to as "entitlement." You might think of this "entitlement" as a kind of a priori justification but it's best to keep the two notions separate.

The entitlement view avoids allowing easy justification for the same reasons as the Reidian inspired view. As such, it initially appears to avoid worries about generating justification *ex nihilo* and of no failure. The beliefs that our most fundamental sources such as perception, memory, introspection, the a priori, etc. are reliable will possess this default non-evidential justification, and this must be in place *prior* to one's possessing justification for the premise beliefs involved in the easy justification. No *new* justification is produced by the reasoning. Yet again, the reasoning suffers from transmission failure.

Most of the difficulties for the Reidian approach were due to the way it construed the source of justification for these reliability beliefs. Wright's view holds that there is no belief source that confers justification on these reliability beliefs. The justification for such beliefs isn't *produced* at all. We simply have it. It's in this way that Wright's view is the natural next step for someone attracted to the Reidian inspired view.

Even if "entitlement" is a default status and it's "just there", Wright owes an account of *why* these propositions possess this status. There must be some story since a proposition like *I'm drinking a cold beer* doesn't possess such entitlement. Toward this end Wright explains:

[T]he presuppositions to which one is entitled all articulate conditions which, in the course of a particular cognitive project, one may rationally trust in, or take for granted, precisely *without any requirement*

---

<sup>113</sup> Perhaps these beliefs about the reliability of sources themselves *have* a belief source in the sense described in 2.1.1. The point here is that these sources do not confer justification on these beliefs about reliability. As such there is no question of whether the sources of the beliefs about reliability can or cannot confer justification prior to having justification for believing they are reliable.

*of cognitive work.* Trust is, in the nature of the case, something one does without evidence; if you have evidence that a source is trustworthy, you don't need to trust. There are a variety of ways of arguing that trust may nevertheless be a rational stance, or frame of mind. Those that I myself find most impressive flow from the observation that *all* enquiry involves so far untested presuppositions, some specific to the particular enquiry at hand, others generic and recurrent; and that the attempt to improve one's epistemic position in this respect is doomed to failure, either because counterparts of the original presuppositions recur or because they themselves recur as presuppositions of the second investigation (as for example when we find ourselves reasoning in ways which presuppose Modus Ponens in attempting to justify that very rule).<sup>114</sup>

Wright isn't clear on the point in this passage but I take it that it is the presuppositions of this latter kind—those that recur as presuppositions of the second investigation—that have the status of “entitlement.” The idea appears to be that there are propositions that are presuppositions of *any* kind of rational enquiry and that, in any enquiry, “we *cannot but* take such things for granted.”<sup>115</sup> The very idea of justification only makes sense within a domain where such presuppositions are made, and it's for this reason that we possess this non-evidential justification for these presuppositions.<sup>116</sup>

Before arguing that this view still won't solve our puzzle, I want to argue that the account offered for the justification of these presuppositions is clearly unacceptable even when construed as this lesser kind of *non-evidential* justification. Wright argues that trusting these presuppositions may be a rational stance despite our lacking evidence for them. We're concerned with *epistemic* rationality, a concept that is intimately tied to likelihood of truth unlike the concept of *practical* rationality. Does the supposed fact that in all inquiry we must presuppose perception's reliability make perception's reliability at all likely? I don't see how it could. If anything this only grounds the *practical* rationality of making these presuppositions.

---

<sup>114</sup> Wright (2011), pp. 33-34.

<sup>115</sup> Wright (2004), p. 189.

<sup>116</sup> This idea is inspired by Wittgenstein's *On Certainty* (1969).

I can't help but presuppose that perception is reliable so *practically* I ought to continue making this presupposition. But this doesn't make the presupposition's truth at all probable.

More importantly, Wright's reasoning could only lead to the conclusion that it's (practically) rational to make presuppositions *in general*. Such reasoning can't account for the practical rationality of *specific* presuppositions. Even if the justification is *non-evidential*, presumably we have this justification for *some* presuppositions but not others. We think we possess justification for our commitment to perception's reliability. If a person, however, were systematically and naturally disposed to move from a perceptual experience as of *P* to the belief that there is an evil genius with the desire of producing in her a deceptive perceptual experience as of *P* this would reflect her presupposition that perception is deceptive. But we don't think this subject's commitment is justified. However, her presupposition would allow the subject to undertake an enquiry into what desires the evil genius has. And, perhaps, any attempt by the subject to justify this commitment would rely on this very commitment. Thus, Wright's reasoning can't even establish that *specific* commitments are (practically) rational.

Wright's example of Modus Ponens (MP) provides further illustration. He suggests that we possess "entitlement" for presupposing that reasoning in accord with MP is reliable because we inevitably "find ourselves reasoning in ways which presuppose MP in attempting to justify that very rule." Tim and Lydia McGrew (2007) nicely illustrate how this happens. You might suggest that you can give an argument from truth tables to the claim that MP is necessarily truth-preserving and, therefore, conditionally reliable. We can argue as follows:

1. Every line of the truth table that assigns truth to both *P* and ( $P \rightarrow Q$ ) also assigns truth to *Q*.
  2. But if this is so, then whenever *P* and ( $P \rightarrow Q$ ) are true, *Q* must also be true.
- 
3. Therefore, whenever *P* and ( $P \rightarrow Q$ ) are true, *Q* must also be true.<sup>117</sup>

---

<sup>117</sup> See Tim and Lydia McGrew's (2007), p. 128.



Our attempt to justify MP must itself presuppose the reliability of MP since this reasoning is itself an instance MP. That our attempts to justify the reliability of MP will always be frustrated in this way is Wright's reason for claiming that our proposition of MP's reliability must have this default "entitlement". The problem is that this same reasoning would lead to the conclusion that a subject who always and naturally reasons in accord with affirming the consequent has "entitlement" for her presupposition of affirming the consequent's reliability. But clearly one doesn't have "entitlement" for this presupposition. This means that, at the very least, Wright's account of "entitlement" is seriously incomplete. Thus, I conclude that not only does Wright's proposed account fail as an account of *epistemic* rather than *practical* rationality but it also fails to give a complete enough account that will result in *specific* presuppositions being rational and not others.

Despite Wright's failure to provide an account that is developed fully enough to assess whether it really would vindicate a rejection of JT, let's consider whether a fully developed account of this sort *could* provide a solution to our puzzle. Wright's "entitlement" view makes this supposed justification for believing our sources are reliable even *easier* than going through the easy justification reasoning. It doesn't seem to vindicate our intuition that acquiring this justification should be *more* cognitively demanding. As such, it might be thought to fall prey to the direction objection mentioned earlier. However, it might not be problematic since Wright admits that this is a *lesser* kind of justification, a kind of *non-evidential* justification which doesn't represent any kind of cognitive achievement. The justification we arrive at via perceptual, memorial, inductive means, however, is *evidential* justification and represents a genuine cognitive achievement.

Distinguishing between evidential and non-evidential justification in this way leads to the worry that Wright's view will actually fall prey to the problem of easy justification as well. Consider the case of perception. On Wright's account, my perception as of a red table provides me with *evidential* justification for believing that there is a red table in front of me *because* I have *non-evidential* justification (or "entitlement") for believing perception is reliable.

Using introspection I gain *evidential* justification for believing that I have a perception as of a red table. These two beliefs entail that my perception is accurate on this occasion. If I've now *gained* evidential justification for these two beliefs, don't they now provide *evidential* justification for the conclusion that my perception is accurate on this occasion? A belief which acts as (minimal) evidence for the belief that my perception is reliable? Wright's view therefore gives rise to the problem of easy justification regarding *evidential* justification. Evidential justification for the reliability of a source can be generated *ex nihilo* via the easy bootstrapping reasoning. Moreover, evidential justification seems to be gained by "no failure" procedure. Wright's appeal to "entitlement" doesn't solve our puzzle.

Wright might claim that evidentially justified propositions which rely on a non-evidentially justified presupposition can't act as evidence for that presupposition. If this is correct, then easy justification reasoning wouldn't produce *evidential* justification for its conclusion. With this restriction regarding the conditions under which something can act as evidence for something else, Wright's view prevents easy justification. Thus, Wright would seem to have a genuine solution to our puzzle.

This response fails for two reasons. First, it reintroduces the direction objection. The idea that we have this default "entitlement" for the presupposition that perception, memory, etc. are reliable is only satisfying if it represents a rather minimal status and achieving a more desirable status requires more *cognitive* work. Now we're told that this is the best status these presuppositions can have; doing more cognitive work produces nothing. But hasn't our epistemic situation regarding the reliability of our faculties *in fact* improved? Haven't our scientific investigations of perceptual and memorial mechanisms provided evidence that these are more reliable than we thought in certain circumstances and less reliable in others? Second, it would seem that this admission just rules out the possibility of easy justification reasoning generating justification on account of its circularity. I've already argued at length that the problem isn't circularity, and the same contrast cases show that this move in defense of entitlement is illegitimate. Wright is committed to saying that in White's optometrist case

a subject must already have *non-evidential* justification for believing her perception is reliable and so the optometrist report can't provide evidential justification for this reliability. Similar remarks apply to TAG REASONING, AAM REASONING and its iteration, and INDUCTIVE JUSTIFICATION OF INDUCTION. Wright is committed to saying that these types of reasoning are as equally problematic as the easy justification reasoning. But the easy justification reasoning strikes us as intuitively *more* problematic, and so Wright doesn't provide an acceptable solution to our puzzle.

### 3.4 Sosa, Levels of Justification, and Coherence

Ernest Sosa's *Reflective Knowledge* contains one of the best and most interesting attempts to provide a philosophically satisfying solution to the problem of easy justification while avoiding the skeptical worry posed by the problem of the criterion. Sosa, like most authors, focuses his discussion at the level of knowledge (although he does fall into discussion of justification as well). I'll attempt to apply his remarks to a discussion of justification. However, while I think the view I *find* in Sosa is incredibly interesting, I'm not entirely sure that it's the correct interpretation of Sosa's oftentimes obfuscating remarks.

Sosa has famously distinguished between two kinds of knowledge: animal and reflective. Sosa explicates these concepts within the context of a kind of virtue reliabilism—although, I see no reason that one couldn't apply many of his remarks about the distinction to other theories. Animal knowledge has what he calls an AAA structure. Animal knowledge is belief that is accurate, adroit, and apt. Belief aims at truth and it's accurate insofar as it achieves this aim. Clearly, however, accuracy (i.e. true belief) is insufficient for knowledge. We can arrive at true beliefs in any number of epistemically inappropriate ways, e.g. random guesses, crystal ball gazing, wishful thinking, etc. *How* we arrive at our belief is also important to whether a belief qualifies as knowledge. It's at this point that adroitness comes in. A belief is adroit when it arises from a stable character trait (i.e. disposition) that “would in

appropriately normal conditions ensure (or make highly likely)” that the belief is accurate.<sup>118</sup> In other words, a belief is *adroit* when it’s the result of an epistemic competence—where “epistemic competence” is understood as this stable character trait that would reliably produce true beliefs. Appealing to these stable character traits makes this a virtue account. That the virtuous and non-virtuous character traits are distinguished by their reliability makes this a reliabilist account; hence, the label “virtue reliabilism.” Finally, a belief is *apt* iff it’s accurate *because* *adroit*. This last feature rules out Gettier type cases. In Gettier cases a subject’s belief is both accurate and *adroit* but the belief’s accuracy isn’t *due to* its *adroitness*. Usually its accuracy is the result of a lucky environmental factor rather than the subject’s epistemic competence. Since *aptness* entails accuracy and *adroitness*, we can equate Sosa’s notion of animal knowledge with *apt* belief.

What’s most important to notice is that *adroitness* is the clear analogue of epistemic justification. Sosa is sometimes reluctant to actually give this the label “justification” and sometimes uses the label “epistemic competence” since this is meant to be a rather minimal achievement. A belief will be *animally* justified when it’s the result of an epistemic virtue (understood in Sosa’s reliabilist way).

Sosa is understandably unsatisfied with the idea that this is all human justification or knowledge amounts to. Animal justification and knowledge are modest achievements that only amount to proper attunement with one’s environment. Sosa’s account of animal justification and knowledge is a kind of “thermometer” account. Proper attunement with one’s environment can be had by unreflective animals or even (insofar as we could describe them as having beliefs) thermometers and calculators! Surely reflective human animals are capable of a kind of justification and knowledge more deserving of the name.

---

<sup>118</sup> Sosa (2007), p. 29.

It's for this reason that Sosa introduces the notion of *reflective* knowledge. What is distinctive about much human knowledge is that unlike thermometers, calculators, or unsophisticated animals, we can gain an epistemic perspective on our own belief source and epistemic competencies. I can endorse the reliability of perception, memory, induction, etc. I can also reject as unreliable crystal ball gazing, random guessing, wishful thinking, etc. These meta-level endorsements and rejections of reliability can themselves be aptly formed. In such a scenario I'd have animal knowledge that *P* and that *P* is reliably formed. Sosa describes this as an apt belief aptly noted. When I have this kind of perspective on my epistemic faculties I achieve what Sosa calls *reflective knowledge*. Similarly, we can develop a notion of an adroit belief adroitly noted, or reflective justification.

Bifurcating justification and knowledge in this way initially seems to make our problem more difficult. Doesn't the problem now arise at both levels? If my perception is an epistemic competence then I can use it to form animally justified perceptual beliefs *P*, *Q*, *R*, etc. I can then use the bootstrapping reasoning to form an inductive argument for the conclusion that my perception is reliable and thereby gain *animal* justification for this latter belief. This is the problem at the animal level. But now that I've achieved animal justification for believing that my perception is reliable, don't I possess an epistemic perspective on perception that raises my beliefs that *P*, *Q*, and *R* to the level of *reflective justification*? It seems that the reflectively *unjustified* belief *P*, *Q*, and *R* thereby generate (via the easy bootstrapping reasoning) *reflective* justification for themselves. This reflective justification is created ex nihilo. Moreover, wouldn't this reflective justification for *P*, *Q*, and *R* transfer (via the easy bootstrapping reasoning) to the conclusion that perception is reliable?

Sosa adopts different strategies for resolving the problem at the two levels but I take it that the solutions are to be integrated in such a way as to avoid the difficulties when the solutions are taken separately. At the animal level Sosa adopts a view similar to those discussed in 3.2 and 3.3 and (implicitly) rejects JT. Consider the relevant principles modified to apply solely to animal justification:

**JR-A** Necessarily, a subject  $S$  has *animal* justification for believing a particular deliverance  $B_p$  of a belief source  $\Phi$  only if  $S$  already has *animal* justification in believing that  $\Phi$  is reliable.

**JT-A** Necessarily, a subject  $S$  has *animal* justification for believing that a source  $\Phi$  is reliable, only if  $S$  already has *animal* justification for believing that particular deliverances of  $\Phi$  are true.

Sosa himself argues that we manifest a commitment to a source's reliability when we use that source to form a belief, and he agrees that the resulting beliefs couldn't be justified if the commitment wasn't itself *already* justified. He therefore accepts JR-A. Since we must already possess animal justification before going through the easy justification reasoning, no *new* animal justification is transferred from the premises to the conclusion.

As I've stressed, however, in order to be philosophically satisfying we need an account of the *animal* justification that these commitments enjoy independently of any animal justification for believing any of the particular verdicts of these commitments. Like Wright, Sosa rejects the Reidian inspired view that this justification derives from these commitments having a non-sensory phenomenological experience as their basis. Sosa also rejects Wright's Wittgensteinian inspired account. We've seen that such an account is better suited to offer a kind of *pragmatic* justification for the relevant commitments (although we've seen that it isn't developed enough to provide even this for *specific* commitments). This kind of account isn't well suited for pairing with Sosa's reliabilist leanings that indicate a desire for an intimate connection between justification and truth. Thus, Sosa offers us a purely reliabilist explanation of the animal justification that these commitments enjoy.

Sosa's penultimate explanation of why the relevant commitments have animal level justification is found in the following remarks:

What gives it *epistemic* standing, moreover, is in essential part its animal reliability in enabling the harvest of needful information... This just means that the epistemic standing of taking experience at face value, absent special reason for suspicion, derives from its serving us reliably well in the harvest of information proper to a well-functioning human organism.

Our trust in our animal competences is thus a source of epistemic standing for the belief thus acquired simply because those competences themselves, those animal faculties, have a proper

epistemic standing of their own, which they derive from being part of the animal endowment of an epistemically well-functioning human being.<sup>119</sup>

Most of us have a commitment to taking apparent perception at face value. *We* have a disposition to form the belief that *P* when we have an apparent perception as of *P*—it’s helpful to note that characterizing dispositions requires utilizing *ceteris paribus* clauses since a disposition can be present yet overridden in particular cases. However, as I pointed out in discussing Wright’s “entitlement” view, a person, Richard, could have a commitment to taking perceptions as indicative of an evil genius’ intention to produce such a hallucinatory experience in him. Our commitment has a positive epistemic standing such that animal justification is conferred upon beliefs that manifest such a commitment. Richard’s commitment lacks a positive epistemic standing and therefore his beliefs manifesting his commitment lack animal justification (absent other factors). One of the difficulties with Wright’s “entitlement” account was that it failed to explain the differing epistemic statuses of these commitments. Sosa’s view avoids this difficulty with his reliabilist account. *Our* commitment to taking perception at face value is animally justified because it “serves us reliably well in the harvest of information proper to a well-functioning human organism”, whereas Richard’s commitment isn’t animally justified since it does no such thing.

Despite avoiding the pitfalls of both Wright’s and the Reidian inspired rejection of JT, Sosa’s account strikes me as odd. Sosa admits that our commitment to taking perception at face value can be described as having a propositional content akin to “Reliably, if I have a visual experience as of seeing *P*, I tend to see *P*”. Similarly, Richard’s commitment can be described as having a propositional content akin to “Reliably, if I have a visual experience as of seeing *P*, there tends to be an evil demon with the intention/desire to produce in me a hallucinatory experience as of *P*”. Sosa’s account explains that the former commitment has animal justification that the latter lacks simply because, as a matter of fact, it reliably serves

---

<sup>119</sup> Sosa (2009), p. 237-238.

us well in gaining important information. But this is just to say the relevant commitments are animally justified when they're true! Why would truth be necessary and sufficient for the animal justification of these commitments when truth is *unnecessary* and *insufficient* for animal justification regarding other propositional attitudes?!

Sosa's account of the animal justification for our commitments to perception's, memory's, induction's and other sources' reliability is unbelievable. Such commitments could be produced in a number of inappropriate ways. If I had reliabilist tendencies, a much more natural and attractive account regarding the positive epistemic status of these commitments would be in terms of the reliability of the environmental processes (not seated in the subject) that produced the commitments. Many of our belief forming commitments/dispositions, such as taking experience at face value, are due to evolutionary forces and social factors (e.g. I imagine that the training involved in learning various concepts is responsible for a large number of our commitments/dispositions related to forming beliefs involving those concepts). Thus, we could plausibly maintain that these commitments have positive epistemic status in virtue of being produced by evolutionary and training processes that reliably produce commitments, dispositions, or belief forming character traits that serve organisms well in the harvest of information. As such, we get a kind of *double* reliabilism: a belief will possess animal justification when it's the product of a stable character trait that reliably produces true beliefs and where this stable character trait was *itself* produced by environmental and social processes that reliably produce belief forming character traits that are reliable. I think such an account is far from implausible so I'll focus my criticisms of Sosa's view elsewhere.

It should be clear how this view resolves our puzzle at the animal level. Like the Reidian and "entitlement" views, no *new* justification for the conclusion is generated using the easy justification reasoning. On Sosa's view, easy justification reasoning merely *uncovers* the implicit commitment to a source's reliability by raising it to the level of conscious belief, but the animal justification for this commitment was in place prior to any of this reasoning.



Sosa's solution, like Wright's, is only plausible when we construe animal justification as a minimal achievement. If animal justification wasn't such a modest achievement that can be had by thermometers and unreflective animals then this solution would succumb to the DIRECTION OBJECTION. Achieving the full-fledged justification to which human animals aspire for the belief that a source is accurate or reliable ought to be *more* demanding than going through the easy justification reasoning. When the reasoning involved in White's optometrist case, TAG, AAM and its iteration, and even INDUCTIVE JUSTIFICATION OF INDUCTION are available to us our beliefs are in a much better position than when we only have available the easy justification reasoning. Sosa has yet to vindicate this intuition. In conclusion, Sosa's solution at the animal level is only appealing given his admission that there is a higher level of justification to which we can aspire. Our main concern should therefore be with how Sosa attempts to avoid a problem of easy *reflective* justification.

It's only in the last few pages of *Reflective Knowledge* that Sosa tackles the problem at the reflective level. Here he suggests that the problem is to be resolved by adopting a "web-like" rather than "pipeline" model of reflective justification. Sosa is rejecting a foundationalist in favor of a coherentist model of *reflective* justification. Animal justification is reliabilist in nature, and provides our beliefs with an epistemic normative status which both recognizes the foundational role that experience plays (albeit contingently) and that connects our beliefs to the environment.<sup>120</sup> Moreover, animal justification has a clear *flow* or *direction*. If a subject S's belief, or commitment, that *P* generates animal justification for believing that *Q* then *Q* doesn't generate animal justification for S's believing *P*. It's in this way that animal justification is *transmitted* and has a clear direction. Reflective justification, however, is

---

<sup>120</sup> Thus, with his notion of animal justification in place, Sosa's appeal to a coherentist model of reflective justification doesn't raise the oft-cited worry that coherentism leads to a "frictionless spinning in a void". This phrase comes from McDowell (1994). This complaint hints at the popular "isolation" objection to coherentist theories of justification.

symmetrical, does *not* transmit, and has no clear *direction* or *flow*. Reflective justification is *constituted* by mutual support or a kind of coherence that holds between one's beliefs.

For Sosa, however, reflective justification doesn't seem to consist in the *mere* coherence of one's belief system. Reflective justification consists in a special kind of coherence, a kind of *cross-level coherence* which arises from an epistemic endorsement of one's belief sources as reliable. Consider a set of beliefs that includes the object-level beliefs *P*, *Q*, and *R* delivered via perception and the meta-level belief that one's perception is reliable. There is a kind of mutual support holding between these object and meta-level beliefs. That perception is reliable provides reason to believe that the perceptually delivered beliefs *P*, *Q*, and *R* are true. Similarly, that these perceptually delivered beliefs are true provides reason for believing that perception is reliable. We have a kind of symbiotic relationship between the object and meta-level beliefs. One has reflective justification to the degree that this cross-level coherence obtains. Reflective justification, however, does *not* transfer from the object-level beliefs to the meta-level belief or vice versa. Reflective justification for *both* levels supervenes on the degree to which these beliefs mutually support on another.

The idea behind Sosa's solution at the reflective level is that our reasoning (the easy justification reasoning included) is never the *source* of reflective justification. Reflective justification is never generated by or transferred from the premises to the conclusion of the reasoning. Reflective justification is, in some sense, produced by the reasoning but it's an emergent phenomenon and comes all at once instead of in pieces—although it can gradually increase in virtue of a subject incorporating new beliefs into her belief system which increase the degree of cross-level coherence. Neither the reflective justification of the premises or of the conclusion is prior to the other. It's this resulting cross-level coherence that generates reflective justification. Easy justification reasoning produces reflective justification only in

the sense that it produces beliefs with this cross-level coherence which in turn generates (perhaps even *constitutes*) reflective justification.<sup>121</sup>

Sosa's appeal to cross-level coherence in his explication of reflective justification is incredibly plausible when considering the complex justificatory interplay that often occurs between object and meta-level beliefs, especially in scientific contexts where the interplay is rather conspicuous. We design experiments and use perception to form a set of beliefs {B}. We then use {B} to construct a theory T about how perception works which informs our beliefs about perception's reliability in various conditions. We then use T (and its information about perception's reliability/unreliability) to regulate our perceptual beliefs and form a revised set of object-level perceptual beliefs {B\*}. We then use {B\*} to revise our theory T and arrive at theory T\* concerning the reliability of perception. T\* can then go on to, yet again, help us revise our perceptual beliefs (i.e. help us correct various errors that arose due to perception's unreliability in various circumstances). And so the cycle continues.<sup>122</sup> This process relies on particular instances of perceptual beliefs in forming beliefs about perception's reliability in those very instances. I agree, however, with Sosa's general intuition that it'd be dumbfounding to discover that there isn't a way to understand this interplay between object and meta-level beliefs that *in some way* helps us achieve a better epistemic position regarding both levels. Sosa's account of reflective justification provides a natural explanation of how this process can do just this. This process would cause us to revise our beliefs in various ways that, as a matter of contingent fact, would create more

---

<sup>121</sup> *Doxastic* reflective justification arises when the mutually supporting beliefs causally sustain one another in a symmetric way such that the cross-level coherence thereby forms a basis of the subject's beliefs. This mutual basing contributes to how the subject maintains their belief system. If cross-level *incoherence* arises, this will pressure the subject to revise her beliefs so as to rid the system of such incoherence. Note, however, that I take it that on Sosa's view this sensitivity to cross-level coherence/incoherence might take place entirely at the sub-personal level. As such, this sensitivity doesn't necessarily require that the subject *notice* the cross-level coherence or incoherence in order for it to constitute a partial basis for her beliefs.

<sup>122</sup> This idea expressed here is that of Rawls' (1971) notion of reflective equilibrium.

cross-level coherence by getting rid of inconsistencies between the object and meta-level beliefs. And this cross-level coherence seems to be of epistemic value.

But doesn't one inevitably arrive at said cross-level coherence when I engage in the easy justification reasoning? At this point, Sosa takes on the strategy of arguing that arriving at reflective justification, as opposed to animal justification, via the easy justification reasoning isn't problematic. He explains:

How can that [i.e. the scientific reasoning above] be acceptable, if it is acceptable *neither* to indulge in *simple* bootstrapping *nor* to reason from our perceptual belief directly to the conclusion that we are well situated for such perception? Answer: It must be recognized that, by parity of reasoning, the mutual support even in these latter cases might add *something* of epistemic value. Coherence through mutual support seems a matter of degree, and even the minimal degree involved in simple bootstrapping is not worthless. Nor does it seem worthless even when it turns out that both the particular perceptual belief *and* the commitment are false. Mutually supportive comprehensive coherence is always worth something, even if its value is vanishingly small when it remains this simple, especially when the web is detached from the surrounding world because it is false through and through.<sup>123</sup>

When raising the problem of easy justification we assumed that the subject *S* arrived at her conscious belief that a belief source  $\Phi$  is accurate/reliable by using the easy justification reasoning. *S* therefore lacked the cross-level coherence distinctive of reflective justification prior to the reasoning. After going through the reasoning, *S* does have a set of beliefs that exhibit this cross-level coherence and *simultaneously* generates reflective justification for both the premises and the conclusion of the easy justification reasoning. Isn't this to gain something of epistemic value, even if it's extremely minimal in this particular case?

It's important to note that *this* achievement of reflective justification is minimal when this cross-level coherence is arrived at via the easy justification reasoning. This "simple bootstrapping" only provides *some* reflective justification; in fact, the cross-level coherence might be so small that even though the beliefs have some degree of reflective justification

---

<sup>123</sup> Sosa (2009), p. 242.

they haven't reached the threshold for being reflectively justified *simpliciter*. Moreover, the degree of reflective justification achieved by such reasoning will pale in comparison to that involved in the incredibly comprehensive cross-level coherence achieved via the methods of cognitive science, or even of a sophisticated reasoner who uses a variety of sources to check each other's reliability. Sosa's solution avoids the DIRECTION PROBLEM. Despite the fact that easy justification reasoning brings about minimal reflective justification, the degree of reflective justification (i.e. comprehensive cross-level coherence) that humans both strive for and are capable of achieving requires much more cognitively demanding work.

Should allow a subject to arrive at even this *miniscule* amount of reflective justification via the easy justification reasoning? If we have reason give a negative answer, we have reason to deny Sosa's solution at the level of reflective justification.

We can't object to the gain in reflective justification on the grounds of an ex nihilo generation of justification. Sosa's coherence model of reflective justification solves this difficulty. When a subject goes through the easy justification reasoning and arrives at reflective justification, this justification isn't *transferred* from object-level perceptual (or memorial, etc.) beliefs to the meta-level conclusion or vice versa. Reflective justification has its source in the resulting cross-level coherence, which isn't a belief but a *feature of a set of beliefs*. This miniscule reflective justification is in no way constituted by an operative but previously reflectively unjustified propositional attitude. Nevertheless, Sosa must give up *both* JR and JT at the level of reflective justification. Sosa's coherentism, however, allows him to adopt modified principles of *correlative* reflective justification:

**JR-R** Necessarily, a subject *S* has *reflective* justification for believing a particular deliverance *B<sub>p</sub>* of a belief source  $\Phi$  only if *S* has *reflective* justification in believing that  $\Phi$  is reliable.

**JT-R** Necessarily, a subject *S* has *reflective* justification for believing that a source  $\Phi$  is reliable, only if *S* has *reflective* justification for believing that particular deliverances of  $\Phi$  are true.

These principles regarding the nature of reflective justification attempt to avoid both easy justification and skepticism by dropping any requirement of *epistemic priority*.

Sosa's coherentist solution at the reflective level resolves the problem of generating justification *ex nihilo*. However, even the supposedly *miniscule* reflective justification produced by the easy justification reasoning is illicit. Sosa's argument from parity fails to establish that the cross-level coherence achieved by the easy justification reasoning is of even miniscule epistemic value. Cross-level coherence is only of epistemic value *under certain conditions*. When cross-level coherence is produced via easy justification reasoning, contrary to Sosa's remarks above, it's epistemically worthless! It's epistemically worthless because the reasoning by which the subject arrives at the cross-level coherence *couldn't have failed to produce said cross-level coherence*. Sosa's solution falls prey to the no failure problem. In fact this is a general worry for any coherence theory of justification whether it appeals to coherence in general or more specifically to cross-level coherence. Coherence can be arrived at in incredibly easy ways that make it of no epistemic value at all.

Let's first illustrate the general worry for run of the mill coherence theories. Imagine that I form the unrelated beliefs that  $P$  and that  $Q$ . According to these traditional coherence theories,  $P$  and  $Q$  are unjustified for me since they are unrelated and haven't been incorporated into a coherent set of beliefs. But what would it even be for a set of beliefs to be coherent? On a supposedly very strict understanding, the coherentist might *hope* for a set of beliefs such that each individual belief is *entailed* by the remaining beliefs. But as Richard Fumerton points out, by simply using the rules of elementary logic I can too easily construct a coherent set of beliefs by using the following procedure: when I believe  $P$  and  $Q$ , form the beliefs that  $(P \text{ or } \neg Q)$  and that  $(Q \text{ or } \neg P)$ .  $P$  entails the first disjunction.  $Q$  entails the second disjunction.  $Q$  and the first disjunction jointly entail  $P$ . Finally,  $P$  and the second disjunction jointly entail  $Q$ . Thus, via this procedure, I arrive at a *minimally* coherent set of beliefs and

these beliefs therefore gain a *miniscule* amount of justification.<sup>124</sup> But this gain seems illicit. Coherence gained *in this way* seems to be epistemically worthless. Why? Going back to the problem of no failure, the procedure used to arrive at these beliefs could *not* have failed to produce a coherent set of beliefs (in the previous sense). As such, it seems absurd that the coherence has any justificatory value when gained this way.<sup>125</sup>

Admitting even a miniscule amount of reflective justification as a result of easy justification reasoning is inadmissible. Easy justification reasoning is guaranteed to produce a set of beliefs with this *minimal* cross-level coherence. If this reasoning *couldn't* bring about cross-level *incoherence* then “achieving” cross-level coherence *in this way* shouldn't provide any new justification even if it's miniscule. We can cement this criticism of Sosa's proposal by pointing out that he is committed to saying that the cross-level coherence achieved by mnemonic bootstrapping provides as much reflective justification as the iterated AAM reasoning since both involve the same *beliefs* and therefore the same degree of cross-level coherence. The important difference between these two forms of reasoning involves the way in which the subject arrives at these *identical* beliefs, a difference that can't be captured in terms of the beliefs' coherence.

What if Sosa were to defend his solution by arguing that only cross-level coherence that isn't produce by easy justification reasoning provides reflective justification? This would allow that the optometrist reasoning, the AAM reasoning and its iteration, the inductive justification of induction, and the sophisticated reasoning involved in the cognitive sciences to provide reflective justification (the last probably providing a greater *degree* than the former forms of reasoning) whereas the easy justification reasoning would not. But this just notes the problem rather than solves it. This is exactly the distinction that I'm arguing Sosa's view

---

<sup>124</sup> See Fumerton (2006b), p. 46.

<sup>125</sup> As Fumerton (2006b) notes, I see no reason why this won't apply equally well if one gives an account of coherence in terms of probabilistic relations.

of reflective justification can't vindicate. To simply add the proviso that the cross-level coherence achieved by easy justification reasoning doesn't count is *ad hoc* and completely unsatisfying. We need an account of justification that vindicates our intuition that the easy justification reasoning cannot produce new justification, not a mere stipulation. Similarly, we don't want a mere stipulation that justification can't be generated *ex nihilo*. Sosa's account of justification vindicates this last intuition but Sosa doesn't offer an account of justification that vindicates the intuition that reasoning which has the feature of "no failure" can't produce new justification. Modifying his view in the suggested way gives us stipulations where we want vindications. I conclude that Sosa's proposal doesn't provide us with a philosophically satisfying solution to our puzzle.

### 3.5 Restricting Evidential Support

The past three sections focused on attempts to resolve our puzzle that utilize specific accounts of justification. In this last section I want to consider some attempts to resolve the problem by utilizing specific epistemic principles considered independently of any specific theory of justification.

#### 3.5.1 Vogel on Rule-Circularity

Jonathan Vogel is one of the principal players in the debate regarding the problem from bootstrapping. Vogel argues that this bootstrapping problem is a problem for reliabilism and argues that internalist evidentialism can avoid the problem by adopting the following principle:

##### NO RULE CIRCULARITY (NRC)

A belief that an epistemic rule R is reliable cannot be justified by the application of R. That is, neither the conclusion itself nor any belief which supports the conclusion may be justified in virtue of the application of R.<sup>126</sup>

---

<sup>126</sup> Vogel (2008), p. 531.



Vogel is wrong that this proposed solution depends on any specific theory of justification. If NRC blocks the problem of easy justification, I see no reason that would prevent a reliabilist from adopting the solution. A reliabilist can simply give a reliabilist account of legitimate epistemic rules and then adopt NRC. I'm at a loss as to why Vogel claims otherwise.

NRC is shown to be too strong by our previous contrasts of circular reasoning with easy justification reasoning. Consider White's optometrist reasoning for perception's reliability. Assume the following epistemic rule: if S has a perception as of *P* then (absent defeaters) S has at least some degree of justification for believing *P*. My perception of the optometrist's report will provide me with some degree of justification for believing that it says that my perception is reliable. Nevertheless, by NRC, this belief about the optometrist's testimony couldn't (under any circumstance!) provide me with any justification for believing that my perception is reliable since this belief would itself rely on perception. The problem is that this reasoning is clearly better than the easy justification reasoning and, assuming that my belief about the optometrist's report does possess some degree of justification, *should* be able to provide me with at least some *minimal* degree of justification for this conclusion. Similar remarks apply to the iterated AAM reasoning. It would seem that the justification for the relevant beliefs would rely upon the rule: if you have a memory as of *P* then (absent defeaters) you have justification for believing *P*. NRC entails that these beliefs couldn't provide *any* support for believing that my memory is reliable. However, this reasoning is again *better* than the easy bootstrapping reasoning and should provide me with at least a minimal degree of justification for believing that my memory is reliable.

NRC is also too weak. It doesn't prevent the closure style reasoning. As I was at pains to point out, when I go through the closure reasoning and conclude that a certain epistemic rule or belief source was accurate on a specific occasion, I'm not reaching any kind of modal conclusion about that rule or source's reliability. NRC is therefore not applicable to

such reasoning. I could even use the bootstrapping reasoning to reach the justified (non-modal) conclusion that an epistemic rule or belief source has *always* been accurate.<sup>127</sup> Vogel gives the impression that he is okay with such results and might even encourage them. I hope, however, that my discussion so far has shown that the reasoning in closure cases is problematic as well. It's also incredibly counterintuitive to think that following an epistemic rule or using a belief source could provide justification for believing *all* of its deliverances have been true without thereby providing *any* degree of justification for believing that the rule or source is reliable.<sup>128</sup>

Jonathan Weisberg has pointed out a probabilistic version of easy justification reasoning that also illustrates the inadequacy of NRC. Consider again the case of Roxanne and her favorite newspaper The Roxy. Imagine that Roxanne reads the sentence “*Q*” in The Roxy and forms the justified belief that *Q*. Roxanne then forms the justified belief that The Roxy reads “*Q*” on the basis of her perception. Finally, Roxanne infers that The Roxy was accurate on this occasion. At this point Roxanne can ever so slightly raise her confidence that The Roxy's next sentence will be true. Surely such reasoning is as illegitimate as our original easy justification reasoning and suffers from the same defects. But Roxanne's conclusion makes no claim about the reliability of The Roxy.

Perhaps there is more room for Vogel to use NRC to respond to the previous two worries. NRC claims that an application of an epistemic rule *R* cannot justify any belief that even supports the reliability of *R*. Doesn't the claim that *R* was accurate on some occasion

---

<sup>127</sup> This will require that the justification of each premise is high enough that the probability of the conclusion doesn't drop below the threshold for justification simpliciter since probabilities less than 1 will decrease over conjunction (e.g. consider the lottery case). Of course, one could also modify the reasoning so that it concludes that the epistemic rule or belief source was accurate on most occasions so as to avoid this worry of diminishing probabilities. Moreover, even if the probability *diminishes*, it should be obvious that beliefs in the easy bootstrapping reasoning should be able to provide *any degree* of justification for either of these conclusions (even if the degree provided doesn't suffice for meeting the threshold for justification simpliciter).

<sup>128</sup> See White (2006) for a development of this point.

constitute data that *minimally* supports the reliability of R? It's tempting to give a positive answer. It may not be enough to justify the claim that R is reliable but surely this data provides *some* support for such a claim. In which case, NRC has the result that we cannot have justification for the conclusion of the closure reasoning. The problem is that NRC now has the result that I can be justified in believing *P* and justified in believing that I have a perception as of *P* but lack justification for believing that the perception was accurate on that occasion. But surely these two beliefs *entail* such a conclusion! NRC now pushes us dangerously close to denying closure. Of course the general point that one might have justification for believing *P* and for believing *Q* without having justification for believing *P* & *Q* isn't troubling once we realize that this might be due to the fact that the probability of *P* and of *Q* are each right at the threshold for justification. NRC, however, is committed to a failure of transmitting justification of the conjuncts to the conjunction even in cases where this is inconsistent with probability theory. Surely this is a bad result.

### 3.5.2 Weisberg's Probabilistic Approach

Jonathan Weisberg has proposed an alternative principle to block the problematic reasoning. Weisberg focuses on the fact that probabilistic support is intransitive. In other words, it's possible that *P* makes probable *Q* and that *Q* makes probable *R* but *P* does not make probable *R*. It's along these lines that Weisberg wishes to rule out the bootstrapping reasoning and proposes the following principle:

#### NO FEEDBACK (NF)

If (i) *B1-Bn* are inferred from *A1-Am*, and (ii) *C* is inferred from *B1-Bn* (and possibly some of *A1-Am*) by an argument whose justificatory power depends on making *C* at least *x* probable, and (iii) *A1-Am* do not make *C* at least *x* probable without the help of *B1-Bn*, then the argument for *C* is defeated.<sup>129</sup>

In order to get a better grip on Weisberg's solution we need only note that probabilistic support is intransitive. Various cases where each step in an inference makes the relevant

---

<sup>129</sup> Weisberg (2010), p. 533-534.

premise likely but where there are so many steps that the probability slowly diminishes are common knowledge. But there are other cases as well. Let's assume for the sake of argument that more than 50% of car owners in the U.S. own a Ford. Now consider my sister Dana who lives in the U.S. and owns a Kia. That Dana lives in the U.S. and owns a Kia makes it probable (in fact it entails) that Dana is a U.S. car owner. That Dana is a U.S. car owner makes it probable (by assumption) that she owns a Ford. That Dana own a Kia does not make it probable that she owns a Ford. In fact, if most car owners only own one car then this makes it probable that Dana *does not* own a Ford.

Weisberg's suggestion is that something similar happens with easy justification reasoning. My justification for believing that there is a red table in front of me ultimately rests on this being made probable by my perceptual experience as of a red table. This belief is then conjoined with the claim that I have such a perceptual experience so as to produce support for the reliability of my perception. However, Weisberg correctly notes that my perceptual experience as of *P* doesn't make it probable that my perception is accurate or reliable all by itself. And so NF kicks in and gives the result that the premises made probable by my perceptual experiences in the easy justification reasoning don't provide justification for the relevant conclusions.

The main difficulty for NF has been noted by both Cohen (2011) and Neta (2013). Such a proposal appears to commit one to appalling denials of closure. A perceptual experience might provide me with justification for believing that there is a red table in front of me and I might know that this entails that there isn't a white table illuminated by red lights shining on it in front of me but fail to have justification for this latter claim. My perceptual experience taken by itself doesn't make probable that there isn't a white table illuminated by red light shining on it in front of me. Moreover, as Cohen (2011) notes, I could be justified in believing the conjunction *the card looks red and is red* but fail to have any

justification for believing *the card's apparent color matches its actual color*.<sup>130</sup> Surely this is an unpalatable consequence that we should hope to avoid.

Weisberg considers two responses to this problem. The first possibility he considers is restricting NF's applicability to cases that use only inductive reasoning.<sup>131</sup> This runs into all of the same problems as NRC. The closure reasoning strikes us as problematic but the restricted NF doesn't prevent such reasoning from providing new justification. Moreover, if one has justification for believing that one's perception was accurate on some occasion this provides at least *some degree* of justification for believing that perception is reliable. The restricted version of NF, however, would be forced to deny this. Even worse is again the fact that one could have justification for believing that *all* of one's perceptual experiences have been accurate and that this wouldn't provide any degree of prima facie justification for believing that perception is reliable. Surely we ought to reject the restricted NF.

A second possibility is that when the premises of the deductive inference are probable the conclusion is made probable on independent grounds. The question then becomes what these independent grounds for believing the deductive consequence are. Such a view must hold that whenever my perceptual experience as of *P* provides justification for believing *P* I have independent grounds for believing that my perceptual experience is accurate. What are these independent grounds? And how could one ensure that they are always present whenever a perceptual experience as of *P* provides me with justification for *P*? More importantly, if these independent grounds make it likely that my perceptual experience is accurate then wouldn't such grounds make it (at least to some degree) likely that my perception is reliable? In which case, one ought to give the same treatment to the

---

<sup>130</sup> Cohen (2011), p. 149.

<sup>131</sup> The cases we considered are inductive-deductive cases. Inductive reasoning (here I am referring simply to defeasible reasoning) leads to the beliefs that there is a table in front of one and then one uses deductive reasoning to infer that there isn't a red table illuminated by red lights.

bootstrapping cases. Leading us to the view that when one has justification for believing the premises of the bootstrapping reasoning one has justification for the conclusion but where the argument suffers from transmission failure. The justification for the conclusion rests on independent grounds making the conclusion probable. But now the question concerns what the independent grounds are for believing the conclusion of the bootstrapping reasoning? More importantly, how could this be consistent with a rejection of JR?

### 3.6 Lessons Learned

We've considered a number of responses to our puzzle and found them wanting. Our discussion shows how difficult it is to find a satisfactory solution. Despite their failure, the views we've considered provide lessons that we would do well to keep in mind.

Markie taught us an important distinction between dialectical and epistemological failure. Eventually I develop an epistemology such that, necessarily, whenever a source provides us with justification we're aware of evidence for its reliability. When a skeptic challenges us to *cite* our reasons for thinking we're aware of such evidence we might ultimately end up begging the question. As Markie is right to stress, however, we shouldn't infer an epistemological failure on our part from this dialectical failure.

Reid, Wright, and Sosa's (at the animal level) attempt to avoid the problem by allowing that our beliefs about the reliability of certain sources have a kind of a priori or *quasi-a priori* status. These views avoid easy justification by accepting JR and avoid skepticism by denying JT. One difficulty is coming up with a half-way plausible account of a kind of a priori justification we could have for the contingent claim that our sources are reliable. Another difficulty is that these views seem to make gaining justification for believing our sources are reliable even easier. Such a response needs to make sense of the increased assurance we get when our sources have independent checks on their outputs. Nevertheless, if we could provide a plausible account of the a priori status of the reliability of a source

while also making sense of how we could gain further justification empirically this would provide us with a satisfying solution to our puzzle.

Sosa's proposed solution at the reflective level is also important. Sosa advocated a coherentist structure of reflective justification where reflective justification for both particular deliverances of a source and its reliability arose simultaneously in virtue of their presence in a coherent system. An appeal to a coherentist structure of justification can't resolve our puzzle, but the notion of coherence might play other important epistemic roles. I'll suggest at the end of chapter eight that, once we've seen the proper solution to the problem of easy justification, our uneasiness with certain *instances* of easy justification might be due to the fact that we simply don't think certain experiences make various propositions probable unless accompanied by other experiences with which they converge or "cohere".

Finally, both Vogel and Weisberg's proposals are right to focus on principles of evidential support. Their proposals don't solve our problem but do help to more clearly highlight the problematic. Combining a perceptually delivered belief that P with an introspective belief that one has a perceptual experience as of P doesn't provide new evidence for perception's accuracy or reliability. Any evidence must have been implicitly present when forming the initial perceptual belief. But how can we accept that, necessarily, evidence for a source's reliability is implicitly present prior to the reasoning without accepting JR? Moreover, it seems the only evidence available prior to the reasoning is a perceptual experience as of P. Here Weisberg rightly points out that a perceptual experience doesn't make its own accuracy or reliability probable. If we want to resolve these issues we need to take a closer look at what evidence is, what kind of evidence might support the reliability of a source (and, therefore, the nature evidential support), and how a person comes to have evidence.

## CHAPTER FOUR

### FACTS AND EVIDENCE

We've focused on a dilemma for theories of justification. On the first horn, if we allow sources to generate justification absent prior justification for believing the source is reliable then we appear to allow illegitimate gains in justification. On the second horn, if we require prior justification for believing a source is reliable in order for that source to confer justification then we appear to be led to a radical skepticism. A resolution emerges from a theory of justification that requires a subject be aware of the evidential connection between her evidence and her belief (avoiding easy justification), but where the required awareness is weaker than justified belief (avoiding radical skepticism). Chapters 4-7 develop, explain, and defend the kind of theory just sketched. Chapter 8 applies this theory to our dilemma.

In chapter 1 I introduced the evidentialist/non-evidentialist debate that revolves around the evidentialist maxim or EM:

**EM** The epistemic justification of any doxastic attitude holds in virtue of the subject's possession of evidence and what this evidence supports.

EM is strikingly plausible. However, as I explained, EM is relatively uninformative. A full evidentialist theory of justification requires detailed accounts of evidence, evidence possession, and evidential support. More worrisome, however, was that a bare evidentialism doesn't rule out any substantial theory of justification. For instance, the reliabilist and virtue theories of justification with which evidentialism is often contrasted can accommodate EM by simply building in reliability and responsible inquiry conditions into their accounts of evidence, evidence possession, or evidential support.

Nevertheless, EM can play an important role in epistemological theorizing. First, EM pre-theoretically distinguishes epistemic justification from other species of justification such as pragmatic, moral, economic, etc. Second, we can treat EM as a helpful heuristic. When considering hypothetical cases of justification we can ask: What is acting as evidence? What



explains why these items are included in the subject's evidence? How does this evidence support this belief? This is a top-down strategy. Alternatively, we can adopt a bottom-up strategy. We can investigate the evidential concepts independently of our considerations of justification and *subsequently* apply these results (via EM) to our theory of justification.<sup>132</sup> I utilize both strategies in developing my theory of justification.

EM suggests three evidential notions that we must investigate: evidence, evidence possession (i.e. a subject S's *having* evidence), and evidential support. This chapter focuses on the ontology and metaphysics of evidence. Chapter 5 explores the relation between awareness, evidence possession, and justification. Chapter 6 will develop and defend the idea that, in the case of *justificatory* evidence, the evidential support relation is *sui generis* and *unanalyzable*. Together these chapters investigate the evidential notions in order to provide constraints that a theory of justification must respect. These insights, however, need to be combined into a full theory of justification and this is the topic of Chapter 7.

#### 4.1 A Fact Ontology of Evidence

Our notion of evidence is such that it's possible for something to *be* evidence despite the fact that no one *has* this evidence. There might be evidence that Joe Smith was Jack the Ripper that no one has yet discovered or even that no one will *ever* discover. What kind of thing is this evidence that is out there waiting to be discovered?

---

<sup>132</sup> When we do this we might find that applying our investigation of the evidential concepts to our theory of justification (by filling them into EM) leads to various counterintuitive results. If so, there are two possible reactions: (i) revise one's account of the evidential concepts or (ii) abandon/modify the evidentialist maxim. I imagine that in many cases there won't be a principled reason to prefer one reaction over the other. In the end the choice will often hinge on personal preference or various pragmatic considerations. I think the reason that EM has often been *thought* to be inconsistent with certain theories of justification is just that there are certain intuitive judgments concerning what kinds of things might be part of a subject's evidence or items upon which the subject might base her beliefs (i.e. her putative evidence) such that if we were to hold onto these intuitive judgments about evidence possession then certain theories (such as process reliabilism) amount to a rejection of EM. However, the reason that the *bare* EM isn't inconsistent with these theories is that these theories have the option of rejecting these intuitive judgments about the nature of evidence and evidence possession.

Something X can correctly be characterized as evidence independently of any subject's bearing a *having* relation to X. But it's important to realize that there is at least one sense in which there is no non-relational answer to the question of whether something *is* evidence. BEING EVIDENCE isn't an intrinsic property. Things are always EVIDENCE FOR or EVIDENCE AGAINST some claim. Is the bloody knife evidence? If this question isn't relativized to some claim (either explicitly or implicitly) then it's ill-formed. In order to answer the question we need to know whether the person is asking whether the bloody knife is evidence relative to the claim that Penelope killed Ruth or to some other claim. X is evidence iff it stands in an evidential relation to some proposition. The property BEING EVIDENCE just is the property BEING (SOME KIND OF) AN INDICATION OF A PROPOSITION'S TRUTH (OR FALSITY).

In ordinary contexts we identify a bloody knife, a collection of fossils, a paper trail, an X-ray, a strand of hair, a note on the refrigerator, etc. as evidence. But notice that we don't identify these objects as evidence *simpliciter*. Context always makes clear that these attributions of evidential status are relativized to certain propositions. The bloody knife is (supposedly) evidence because it's (supposedly) an indication of various claims such as that Penelope killed Ruth. The fossils are (supposedly) evidence because they're (supposedly) an indication of various claims such as that the earth has existed for more than 6,000 years. We classify these objects as evidence because we take them to stand in the relevant kind of relation to any number of propositions.

Philosophical reflection might lead us to reject the claim that physical objects like bloody knives, fossils, etc. stand in the kinds of relations constitutive of evidence. In turn, this would force us to revise our beliefs about the ontology of evidence. But this wouldn't show that the ordinary and philosophical discussions of evidence are concerned with two distinct *concepts*—an impression you might get from some philosophers. Such a finding would only show disagreement concerning the *extension* of the *same* concept. Our concept of

evidence in both philosophical and ordinary discourse is the concept of something's being an indication of a proposition's truth or falsity.

In fact, philosophers have given good reasons to abandon the view that physical objects are evidence. It's convenient but misleading to characterize bloody knives, fossils, etc. as evidence. And this is something we discover by attending to our *ordinary* concept of evidence. Our ordinary concept holds that standing in evidential relations is constitutive of something's being evidence. Physical objects, however, cannot stand in the relevant relations. To see this, consider a scenario discussed by Ram Neta:

A particular knife is bloody. The blood on the knife is—unlike most human blood—of the same unusual chemical composition as the blood of the murder victim. The knife was lying in the kitchen, which is right next to the living room, and that is where the murder victim was found. The knife had the defendant's fingerprints, and nobody else's fingerprints, all over the handle. Furthermore, there is no evidence pointing to the guilt of anyone other than the defendant.<sup>133</sup>

The evidence in this scenario, whatever ontological category it belongs to, offers a unique degree of evidential support for the claim that the defendant is guilty.<sup>134</sup> Perhaps it makes this claim 90% likely. Perhaps it makes this claim exactly 75% likely. But the evidence in this scenario cannot make it *both* exactly 90% and exactly 75% likely.<sup>135</sup>

---

<sup>133</sup> Neta (2008), p. 98. Neta's is concerned with showing that the evidence that a subject has cannot be constituted by objects since these wouldn't constrain rational belief in the right ways. But this leaves open the possibility that these objects might be evidence even though it isn't the kind that a subject could have. My discussion of this case parallels Neta but I intend to show that Neta's reasoning applies equally well to evidence that isn't possessed by anyone.

<sup>134</sup> This is a claim about the specific set of evidence suggested by the description of the scenario; clearly, the support could be increased or decreased if the description identified a larger (or smaller) set of evidence.

<sup>135</sup> This isn't to say that we know the precise degree of support provided by the evidence in this scenario. Our capacity for discriminating evidential support might only be good enough to provide us with knowledge that the evidence stands in a determinable relation of making it more probable than not that the defendant is guilty. Nor am I claiming that the mere existence of this evidence determines a unique rational attitude that people ought to take towards the claim that the defendant is guilty. But any indeterminacy here seems to be epistemic rather than ontological.

Evidence can't consist of a set of physical objects for the simple reason that no set of physical objects determines a unique evidential relation. You might identify the evidence in the preceding scenario with the following set of physical objects: {the knife, the corpse, the fingerprints, the blood, the kitchen, the living room}. But Neta asks the reader to consider an alternative scenario where the same objects exist in a different arrangement where "the defendant's fingerprints are all over the kitchen,... the victim's corpse was found in the living room with blood all over it, and the knife was found free of both blood and fingerprints."<sup>136</sup> The evidence in this new scenario provides less support for the claim that the defendant is guilty than the evidence in the original scenario. The difference in support must be explained by a difference in evidence. But the same set of physical objects is present in each scenario. Therefore, evidence must not consist of a set of physical objects.

A difference in evidential support in these scenarios isn't due to a difference in what physical objects are present but, rather, *how* the objects are related. The relations between these objects (e.g. the knife, the blood, and fingerprints) need to be incorporated into the relations of evidential relationships. We can do this by identifying evidence with facts. A fact is a non-linguistic complex that is composed of an entity or entities instantiating various properties and relations. It isn't the knife, the fingerprints, etc. that constitute evidence that the defendant is guilty in the original scenario. The evidence is the fact *that* the defendant's fingerprints are on the knife. Facts gather their constituents (including objects and relations) into a unified whole and can thereby stand in a unique truth-indicating relation to a proposition. This leads us to accept an ontology of evidence such that the kinds of things that can be evidence for or against a claim are facts or sets of facts.

---

<sup>136</sup> Neta (2008), p. 99.

#### 4.1.1 Advantages of a Fact Ontology of Evidence

Elucidating the nature of the evidential relation and the conditions for a subject's *having* some piece of evidence will be more crucial to my project than merely characterizing the ontology of evidence. So I mention a few advantages to this fact ontology of evidence but I do not develop these in excruciating detail.

First, the fact ontology explains why epistemic rationality requires that a subject "respect" her evidence where this amounts to not forming beliefs that one has reason to believe to be inconsistent with one's evidence. In other words, an epistemically rational person attempts to avoid beliefs that are inconsistent with her evidence. If evidence consists of facts, there is a clear explanation of this requirement. One aspect of the epistemic goal is the avoidance of false beliefs. Assuming my fact ontology, it's clear that this goal is best fulfilled if we don't form beliefs inconsistent with our evidence. Since any evidence we possess will be a fact, any belief that is inconsistent with our evidence is bound to be false.

Second, this fact ontology provides a way to clean up our ordinary discourse concerning evidence. Earlier I explained that we ordinarily characterize various objects as evidence. But we shouldn't generalize too quickly from a small sample of ordinary discourse. We clearly attribute the status of evidence to much more than just physical objects. We commonly characterize testimony, actions/behaviors, and mental states as evidence. Bob's testimony is evidence that Lucy fled a crime scene. Sara's grabbing her chest and making a face of agonizing pain is evidence that she is having a heart attack. My perceptual experience is evidence that there is a computer in front of me.<sup>137</sup> We ordinarily apply our concept of

---

<sup>137</sup> J.L. Austin gives an ordinary language challenge to the idea that the relationship between perception and our justification for perceptual beliefs is to be construed as an evidential relationship. He explains, "[T]h situation in which I would properly be said to have *evidence* for the statement that some animal is a pig is that, for example, in which the beast itself is not actually on view, but I can see plenty of pig-like marks on the ground outside its retreat. If I find a few buckets of pig-food, that's a bit more evidence, and the noises and the smell may provide better evidence still. But if the animal then emerges and stands there plainly in view, there is no longer any question of collecting evidence; its coming into view doesn't provide me with more *evidence* that it's a pig, I can now just *see* that it is, the question is settled" (1963), p. 115. While I think Austin is capturing an insight of

evidence to entities that fall under a wide variety of ontological categories. Considerations of parsimony should lead us to prefer, *ceteris paribus*, a more unified ontology. Facts offer a way to capture these ordinary claims about evidence in a way that provides theoretical unity at little to no cost. All of these kinds of entities are *parts* of our evidence in the sense that they are constituents of the facts that make up our evidence.<sup>138</sup>

Finally, the fact ontology of evidence provides a unifying ontology for different kinds of evidence. Evidence is a genus among which we might distinguish the following: justificatory evidence, scientific evidence, legal evidence, historical evidence, etc. It's plausible that facts provide an ontology suitable for any of these categories. Evidence of any kind consists of facts that (in some way) indicate a proposition's truth or falsity.

We might categorize different kinds of evidence along any number of different properties: by the kind of subject that possesses (or could possess) the evidence, what conditions must be met in order for the subject to *have* that evidence, when a subject could properly *appeal* to that evidence, or even by the way in which the fact indicates a proposition's truth or falsity. For instance, legal evidence seems to consist only of facts that

---

ordinary language I don't think this shows that we don't think this challenges the idea that in our ordinary everyday thought we think of perception as providing *evidence* for the belief that a certain animal is a pig. What Austin shows here is simply that it would be conversationally inappropriate to make the claim that one's perception provides one with evidence for the claim that an animal is a pig unless that claim were somehow challenged. However, when the belief that it is a pig is challenged it is natural to cite the experience itself as evidence. See Conee and Feldman (2008) for a development of this kind of response to Austin.

<sup>138</sup> Actually, if you prefer, you can actually construe the physical objects, testimony, actions, etc. as themselves being evidence by taking the evidential relation to be a tertiary rather than a binary relation. The idea would be that the evidential relation is a relation that holds between an entity, a fact, and a proposition. Evidence, on this view, would be any entity X such that X is a constituent of a fact F which is some sort of indication of some proposition p. However, it's simpler to construe the evidential relation as a two-place relation that holds between facts and propositions. After all, it's the facts that bear on the truth or falsity of the propositions. Moreover, on this alternative proposal, the same entity might simultaneously be evidence for p and evidence against p in virtue of being the constituent of two different fact F and F\*, where F is indicative of p's truth and F\* is indicative of p's falsity. And our concept of evidence doesn't seem to be one with some piece of evidence can simultaneously be evidence for and evidence against the same claim.

one can appeal to in defending a claim  $P$  within a legal context—where this is governed by rules produced by a certain legal system.<sup>139</sup> Scientific evidence, on the other hand, might be distinguished along various lines. Perhaps scientific evidence should be understood as facts that one can appeal to in a scientific context (yet again, where the norms for appealing to these facts are constructed by a social entity). Or, alternatively, perhaps scientific evidence is the evidence possessed by (or that *could* be possessed by) a scientific community as a whole rather than any particular individual. However, one of the most natural ways of categorizing scientific evidence is in terms of *how* it indicates a proposition's truth or falsity: scientific evidence consists of facts which are indicative of a proposition's truth or falsity in virtue of the existence of a reliable correlation between F-like facts and the truth or falsity of  $P$ -like propositions (as opposed to logical or quasi-logical relations between F and  $P$ ).

I'm not going to attempt to answer the difficult questions of how properly to categorize different kinds of evidence. Nor am I going to attempt to answer questions concerning how these categories are interrelated. My point here is simply that the fact ontology provides a plausible account that *unifies* the ontology of all categories of evidence. As such, the fact ontology describes a similarity across all categories of evidence that explains why each qualifies as a category of the general genus EVIDENCE.

#### 4.1.2 Challenging the Fact Ontology of Evidence

I've focused on developing and motivating a fact ontology that applies to the *general* category of evidence. At this point I want to consider some challenges to this ontology. My discussion of the issues surrounding these objections will contain important insights for defending the theory of justification I eventually develop.

---

<sup>139</sup> It's for this reason that a juror might have legal evidence that supports  $p$  but justificatory (or even scientific evidence) that supports  $\neg p$  since much or all of the evidence the juror has that provides them with justification for believing  $\neg p$  has been deemed inadmissible (i.e. the judge asked them to ignore this evidence since the prosecution failed to meet certain legal requirements for introducing that evidence).

Various philosophers have either explicitly argued for or provided resources from which we might extract arguments for an ontology of evidence that identifies it with propositions rather than facts. Principal defenders of this “propositionalism” include Trent Dougherty, Ram Neta, and Timothy Williamson.<sup>140</sup> Before diving into these arguments, however, I need to make clear how my proposed fact ontology differs from the propositions ontology. As such, I must put forward what are sure to be controversial metaphysical views.

Facts are distinct from propositions in general and even from *true* propositions in particular. Facts are non-linguistic complexes constituted by an entity or entities instantiating various properties or relations. Propositions, on the other hand, are types of entities which accurately or inaccurately, represent the world as being one way or another. Every individual proposition can be identified as a *type* of subjective mental entity (i.e. a *type* of thought). These propositions are true when they correspond to or “picture” a fact and they are false when they fail to correspond with or “picture” a fact. In other words, a proposition is true iff it represents the world as being a way that the world actually is.

Consider the fact that there is a computer in front of me and the true proposition that there is a computer in front of me. The fact is the state of the world that *makes it true* that there is a computer in front of me. The true proposition, however, is a mental representation made true by the fact. Alternatively, propositions are *expressed* by our use of ‘that’ clauses and, when these propositions are true, these ‘that’ clauses (typically) *refer to* the corresponding fact.<sup>141</sup>

---

<sup>140</sup>See Dougherty (2011), Neta (2008), and Williamson (2000).

<sup>141</sup> I say that ‘that’-clauses only *typically* refer to the corresponding fact because, as Frege ([1892] 1997) pointed out, in certain contexts a ‘that’-clause refers to a proposition. For instance, in a context of belief such as “Brad believes that Sally is taller than Suzy” the ‘that’-clause *that Sally is taller than Suzy* refers to the proposition usually expressed by this sentence rather than to any fact.



The metaphysical view I've sketched is far from uncontroversial, but defending it would take us too far afield. Having distinguished facts and propositions (including *true* propositions), let us turn to the challenges to my fact ontology of evidence.

#### 4.1.2.1 *Intuitive Resistance and the Notion of False Evidence*

One source of intuitive resistance to my fact ontology might come from the idea that there can be *false* evidence. If evidence can be *false* then it would have to be identified with the kind of thing that has a truth-value. But propositions are the kinds of things that have truth-values not facts. Therefore, evidence must consist of propositions rather than facts.

What motivation is there to think evidence can be false? The best motivation can be explained using an example from Aaron Rizzieri:

I believe that nobody can enter my office (O for now) because I believe that I have just locked the door (LD for now). Let us stipulate that I have inferred (O) from (LD). I pushed the lock in and gave it a quick twist to the left, which usually does the trick; however, my lock is damaged and does not work. Hence, (LD) is false.<sup>142</sup>

Imagine that you're in my office with me. While in conversation I express my belief that nobody can enter my office. You then ask what reason or evidence there is for believing this to be true. I sincerely respond by expressing my belief in (LD): "I just locked the door." All of this looks like it fits well with the fact ontology. When asked for evidence that nobody can enter my office I express my belief that I just locked the door in an attempt to cite evidence. It appears that I'm attempting to refer to a putative fact as evidence for my belief. It turns out, however, that in this scenario no such fact exists since the lock is damaged and I didn't actually lock my door. But we wouldn't say that I lacked supporting evidence for my belief. As Rizzieri notes, "[i]t is very plausible that (LD) does partially constitute my evidence for

---

<sup>142</sup> Rizzieri (2011), p. 236.

(O). After all, I'm justified in believing (LD), (LD) supports (O), and an explicit inference from (LD) is my most immediate basis or ground for (O)."<sup>143</sup>

This intuitive criticism of the fact ontology of evidence focuses on its applicability to *justificatory* evidence. Clearly a false proposition *P* can play an important role in providing a subject with justification for believing some further proposition *Q*, which thereby gives the impression that the false proposition is part of the subject's evidence for *Q*.

Not everything that plays a role in providing a subject with justification for a certain belief plays an *evidential* role. For instance, the basing relation plays an important role in providing a subject with doxastic justification but the basing relation does not play this role in virtue of being part of the subject's evidence for her belief. In the case considered by Rizzieri I admit that the belief that nobody can enter my office is supported by and based on my evidence. However, it seems natural to identify the evidence for this claim not with the *belief* that I have just locked my door but rather with the evidence I possess for this false belief. My evidence in this case might consist of the facts that I seem to remember locking the door and that I seem to remember my memory being accurate in the past.

There are at least two reasons for identifying these facts as the subject's evidence instead of the false proposition that I locked my door. The first form of support comes from our linguistic practice of retracting our claims about evidence and citing further putative facts when confronted with a belief's falsehood. Second, the evidential probability of my belief that no one can get in the office cannot exceed the probability of this proposition relative to my evidence for the false proposition that I locked the door. The belief that I locked my door may play a role in justifying my belief that no one can get in my office, but the probabilistic support offered by my false belief shouldn't be construed as part of the ultimate supervenience base that determines the evidential probability of my belief.

---

<sup>143</sup> Rizzieri (2011), p. 236.

First the linguistic support. Imagine again that you and I are in my office, I express my belief that no one can get in, and you ask what evidence I have for believing this. I sincerely assert, “I just locked the door.” However, you happen to know that the lock is damaged and that twisting the lock doesn’t lock the door. Once you inform me that the door isn’t locked it strikes us as incredibly unnatural to respond, “Nevertheless, that I just locked the door was part of my evidence when you asked me.” The natural response is to retreat to further claims that purportedly refer to what both you and I accept as statements of matters of fact. Once you inform me that the lock is broken I’d naturally retreat and cite the fact that I seemed to remember locking the door as evidence I had for believing that no one could get in my office: “But it still seemed to me that I remembered locking the door and this was still a good reason to believe that the door was locked and that nobody could get in my office.”

When you explain to me that I didn’t actually lock the door I wouldn’t continue citing this putative fact as my evidence. We both recognize that the putative fact *would* be a reason to believe that nobody could get in my office. When I find out that the putative fact doesn’t exist and retreat to facts about what I seem to remember I’m explaining that I had reasons R to believe that (what we both recognize would be) evidence E for believing that no one can get in my office existed. Having reasons R to believe that E exists thereby gave me reasons and justification for believing that no one could get in my office despite the fact that the putative evidence E didn’t actually exist (and, therefore, I didn’t have evidence E). This idea that evidence E\* that evidence E for *P* exists is, under the right conditions, itself evidence for *P* (i.e. the original evidence E acts as indirect evidence for *P*) will be important for my understanding of the nature of inferential justification.<sup>144</sup>

---

<sup>144</sup> Here I’m appealing to the evidence for evidence is evidence thesis. However, I’m not advocating the idea that *necessarily* evidence for evidence is evidence thesis. I’m only appealing to the fact that under the right conditions this relation will hold. Specifying what those particular conditions are, however, is an incredibly complicated matter. For more on the evidence for evidence is evidence principle see Conee (2010); Christenson (2010); Feldman (2005), (2009); and Fitelson and Feldman (2012).

Second, there are probabilistic reasons for holding that the false proposition isn't playing an evidential role in justifying my belief. The evidential probability of my belief that no one can get into my office cannot exceed the support provided solely by the evidence that lends support to my false belief that I locked the door. Let's make two simplifying assumptions: (i) *that I have locked the door* actually entails *that no one can get in my office* and (ii) the fact that I seem to remember locking the door makes the claim *that I have locked my door* .8 probable. The evidential probability of a proposition *P* for a subject *S* is determined by how probable *P* is *relative to the evidence S has*. Should we include the false proposition that I locked the door as part of my evidence? The false proposition makes my belief that no one can get in my office certain. But surely my belief that no one can get in my office doesn't have an evidential probability of 1 for me! We must take account of the fact that the belief from which this is inferred is only .8 probable for me. And probability claims are relational. That my belief that I locked the door is .8 probable for me just means that this belief is made .8 probable by my evidence. This, in turn, suggests that my evidence only includes the facts that make it probable for me that I locked the door. My belief that no one can get in my office will only be .8 probable for me and this is fully determined by the fact that I seem to remember locking the door. There is no need to include the false proposition itself as part of my evidence. If we included the proposition *that I locked the door* as part of my evidence we would either: (a) illegitimately boost the evidential probability of my belief that no one can get into my office to 1, or (b) we would take into consideration the evidential probability of my belief that I locked the door so that including the false proposition in my evidence set wouldn't make the evidential probability of my belief any higher than the support it receives from the fact about my memory alone. It therefore appears that the false proposition isn't playing a specifically *evidential* role. The false proposition plays a justifying role but not in virtue of providing evidential support for the relevant belief<sup>145</sup>.

---

<sup>145</sup> The false proposition *that I locked the door* is of course playing *some* justificatory role and I

I've merely stopped one objection to the fact ontology. Nevertheless, I think the previous considerations can be modified to provide positive support for the fact ontology. First, when challenged by an interlocutor to find evidence for a certain claim we're always attempting to put forward certain statements that both we and our interlocutor accept as factual statements. Second, in defending the idea that the false proposition *that I locked the door* is irrelevant to determining the evidential probability of my belief we appealed to the fact that any appeal to a proposition must make explicit its evidential probability as well. The evidential probability conferred on  $Q$  by  $P$  always depends upon the evidential probability of  $P$  which is a relative notion. The evidential probability of  $P$ , *whether it's certain or merely probable*, is always relative to our evidence. In order to account for the evidential probability of a proposition  $Q$  we will ultimately be led to some kind of item  $X$  that can confer probability in a way that needn't take account of  $X$ 's evidential probability. Facts, i.e. the *instantiation* of certain properties and relations, are perfect candidates. By appealing to facts we ensure that evidential probability is rationally constrained by input from the world itself and thereby ensure an objective truth-connection.

#### 4.1.2.2 Neta: Same Facts, Different Justification

Ram Neta 2008 has also attempted to challenge the fact ontology of evidence. I've reconstructed his main argument as follows:

##### SAME FACTS DIFFERENT JUSTIFICATION (SFDJ)

1. A subject's total evidence fully determines how rationally confident she can be that a certain proposition is true or false.

---

have yet to explain this role. In chapter 7 I suggest that our awareness of a proposition  $P$ , for which we have evidence  $E$ , can play an essential role in inferential justification for  $Q$  not because  $P$  is part of our evidence for  $Q$  but rather because it allows us to gain an awareness of the connection between our evidence  $E$  and  $Q$  thereby allowing  $E$  to provide justification for  $Q$ .

2. *If a subject's total evidence consists solely of facts then there are two situations such that my evidence set is identical where the appropriate rational confidence in a proposition differs.*

---

3. Therefore, a subject's total evidence doesn't consist of facts.

Neta's first premise makes clear that this challenge is only a challenge to a fact ontology of *justificatory* evidence. Nevertheless, if one could establish that one kind of evidence must consist of propositions rather than facts then considerations of parsimony should lead us to apply this propositions ontology to all categories of evidence. Fortunately, the fact ontology withstands Neta's argument.

Neta takes his first premise to be a desideratum that any adequate theory of evidence must accommodate. However, we'll see that there are two important ways to understand the notion of a "subject's total evidence." On the first understanding Neta is right that premise one is a desideratum for an adequate theory of evidence but, on such an understanding, Neta fails to establish premise 2. On a second understanding of "a subject's total evidence" premise two is true but premise 1 ought to be rejected.

Because Neta takes premise 1 to be a desideratum, let's turn our attention to his defense of premise 2. In order to get the argument off the ground Neta needs to presuppose more than the fact ontology, he needs a working theory of which facts are part of a subject's evidence. Neta must assume a theory of evidence possession in order to even begin to make a case for premise 2. And so Neta supplies the toy theory that a subject S's evidence consists of the facts with which S is acquainted.<sup>146</sup>

---

<sup>146</sup> Neta wants to put aside worries about how to understand the acquaintance relation. I take it Neta assumes that the difficulty will arise no matter how you explain the nature of acquaintance. Of course, an advocate of the fact ontology of evidence might eschew the idea of acquaintance all together. In which case, Neta's criticism wouldn't threaten the view. Nevertheless, the proposal that a subject's evidence consists of the facts with which they are acquainted is closely related to the theory I'll be developing and defending in chapters 6 and 7, and so it is incredibly important for me to consider this objection to the fact ontology.

With this in mind, Neta challenges the fact ontology via the murder scenario discussed earlier. In the scenario a knife is found covered in blood “of the same unusual chemical composition as the blood of the murder victim” (2008, p. 98). Neta then attempts to construct a scenario where two subjects, *S* and *S'*, are acquainted with the same fact but differ in the rational degree of belief that they can place in the truth of the proposition that the defendant is guilty. In this vein Neta constructs the following scenario:

#### MURDER SCENE SCENARIO

- *Proposition P* =<sub>df</sub> The defendant killed the murder victim
- *Fact F* =<sub>df</sub> A knife with blood on it was found near the murder victim
- *Fact F'* =<sub>df</sub> A knife with chemical composition *C* was found near the murder victim
- *Assumption 1* =<sub>df</sub> Fact *F* makes probable *P*
- *Assumption 2* =<sub>df</sub> Fact *F* is identical to Fact *F'*
- *Assumption 3* =<sub>df</sub> *S* and *S'* are each acquainted with fact *F'*
- *Assumption 4* =<sub>df</sub> *S* but not *S'* has reason to believe the identity statement of assumption 2

According to assumption 3 *S* and *S'* are both acquainted with *F'* which makes probable the defendant's guilt. *S'*, however, cannot place the same degree of rational confidence in *P* since *S'* has no reason to accept that *F* and *F'* are identical. Neta therefore takes himself to have established premise two of SFDJ since the facts with which a subject is acquainted fail to fully determine how much confidence she can rationally place in the defendant's guilt.

Many acquaintance theorists think a subject can only be acquainted with her own sensory experiences and not with physical facts like *F*. These philosophers might be tempted to think the difficulty only threatens those with a naïve theory of acquaintance. But if we admit at least the epistemic possibility of the mind-brain identity theory we can construct an introspective analogue of Neta's objection:

#### BRAIN STATE SCENARIO

- *Proposition P* =<sub>df</sub> I am in brain state *X*
- *Fact F* =<sub>df</sub> I am in brain state *X*
- *Fact F'* =<sub>df</sub> I am in pain
- *Assumption 1* =<sub>df</sub> Fact *F* makes probable *P*
- *Assumption 2* =<sub>df</sub> Fact *F* is identical to Fact *F'*
- *Assumption 3* =<sub>df</sub> *S* and *S'* are each acquainted with fact *F'*
- *Assumption 4* =<sub>df</sub> *S* but not *S'* has reason to believe the identity statement of assumption 2 (perhaps *S* is a really good neuroscientist).

Surely we must admit that S' might be acquainted with the fact that he is in pain but not be able to rationally believe *P* since he has no reason to accept the identity statement. And surely S's acquaintance with this same fact might thereby allow him to rationally believe *P* since he does have good reason to accept the identity statement. And so we have an introspective analogue of Neta's case for premise two of SFDJ.

One move you might make in response to Neta's concern is to simply refuse to identify the two facts that Neta attempts to identify. And one might do this for reasons independent of the physicalism/dualism debate. At least some philosophers who accept the correspondence theory of truth (as I have) might be tempted to posit a distinct fact for every logically distinct true proposition. Such a view is incredibly controversial and I regard this move as a last resort since.<sup>147</sup> I will go as far as I can in responding to these and other worries without appealing to such bloated ontology of truth-makers.

The first problem with Neta's argument is that even if a subject's evidence were to consist of the facts with which they are acquainted the advocate of the fact ontology can insist that the intelligibility of the scenario requires us to imagine S and S' being acquainted with different facts. According to the fourth assumption in both scenarios, one of the subjects S has good reason to believe the identity statement and the other subject S' doesn't. Thus, S must have evidence supporting the identity statement that S' lacks. The acquaintance theorist should insist that this means that the *total set of facts* with which S and S' are acquainted are non-identical even though they are both acquainted with F. Presumably S will have reason to believe the identity statement because S is acquainted with various memory experiences as of reading about the identity statement in textbooks or with memory experiences of having done research that supports the identity statement. Either way, making

---

<sup>147</sup> I myself am suspicious of certain kinds of facts (such as the thought that in addition to the facts <A> and <B> there is an additional conjunctive fact <A & B> that we need to posit as the truth-maker for the proposition *A and B*



sense of the idea that S has reason to accept the identity statement while S' lacks said reason requires imagining that S has evidence that S' lacks. We can insist that we understand this difference in terms of the facts with which the two subjects are acquainted. I therefore conclude that Neta has failed to give any reason to accept premise 2 of the SFDJ argument and that he hasn't posed any insuperable challenge to the fact ontology.

Moreover, it's important to realize that the propositions ontology of evidence is subject to a similar worry. Imagine two subjects who *know* the axioms of set theory. One of the subjects, Fred, is a math dilettante and can only perform very simple proofs from these axioms. Wilma, on the other hand, is an expert mathematician and can perform incredibly complex proofs well beyond Fred's abilities. If both Fred and Wilma *know* the axioms of set theory and can use these in at least some proofs, then it seems as if the propositional theory of evidence ought to think the axioms are to be included as parts of both Fred and Wilma's evidence set. But surely there will be some complicated mathematical theorem that is outside of Fred's but inside of Wilma's ability to prove. In this kind of case Fred couldn't rationally have the same degree of confidence in T as Wilma. Thus, it isn't clear that a propositional ontology of evidence is any better suited than a fact ontology for avoiding Neta's worry. I'm offering a kind of partners in guilt response.

Finally, the "partners in guilt" response suggests that the source of Neta's worry is something other than the ontology of evidence. I want to suggest that the problem lies in the way Neta is implicitly thinking about a subject's total evidence. One way to understand the notion of a subject's total evidence is as a single and unique total "*body of evidence*, for use in the assessment of hypotheses."<sup>148</sup> Neta appears to think of a subject's total evidence in this way. If we think of a subject's evidence as consisting of the set of propositions or facts (depending on one's preferred ontology) such that each of these could individually be used

---

<sup>148</sup> Williamson (2000), p. 187.

to assess some hypothesis or other then it seems that Neta's case and my own case of pertaining to the axioms of set theory illustrate the possibility of two subjects will sharing the same total evidence *in this sense* but where the rational confidence they can place in a proposition *P* differs. Premise 2 of SFDJ comes out true on this understanding of a subject's total evidence whether we adopt the fact or propositional ontology. But this just shows that premise 1 is false when we understand a subject's total evidence this way.

It's natural to understand a subject *S*'s evidence (i.e. the evidence *S has*) in relation to *S*'s ability to *use* that evidence in the assessment of hypotheses.<sup>149</sup> As I've pointed out this might lead one to include in *S*'s total evidence anything that *S* can *use* in the assessment of *some* hypothesis. Notice however that an ability to use some evidence might be relative to a particular proposition. *S* might have the ability to *use* some evidence *E* in the assessment of an hypothesis *H* but not in the assessment of another hypothesis *H'* even though the *E* is probabilistically relevant to both. This suggests a notion of total evidence relativized to particular propositions. Consider Fred again. The axioms of set theory are part of his total evidence relative to very simple arithmetical theorems. However, relative to more difficult theorems like *T*, these axioms aren't to be included in his total evidence. When considering a theorem like *T* Fred isn't able to take the axioms of set theory into account when assessing the truth of *T*. As such these axioms play no role in determining the degree of confidence that Fred can rationally place in *T*.

---

<sup>149</sup> This gloss on the nature of evidence possession creeps up in a number of philosopher's discussions of evidence possession. Consider the following examples: "[a] natural idea is that *S* has a *body of evidence, for use in the assessment of hypotheses*" (Williamson, 2000); "evidence must be accessible to the knower in at least the following sense: **it must be available for use in reasoning and inference** (Greco, 2011); "[f]or any bit of evidence that you have, **it is possible for you to use that evidence in the rational regulation of your degrees of confidence**, or of your attitudes more generally (Neta, 2008); "[i]t's a necessary condition of *S*'s being rationally required to proportion her confidence across hypotheses in proportion to the support that those hypotheses receive from [the evidence] **that *S* has the ability to do this**" (Neta, 2008); "[o]ne useful way to identify what we take to be a person's evidence at a time is to say that it is the information or data **the person has to go on in forming beliefs**" (Conee and Feldman, 2008). The bold portions mark out *my* emphasis.

This shows that the notion of a subject's total evidence where premise 1 comes out true is the relativized notion. But now it's clear that Neta's scenario doesn't support premise 2 of SFDJ. F' might be part of S's total evidence relative to the claim that the defendant is guilty since S is acquainted with F' in a way that allows him to *use* this fact in assessing this claim. Nevertheless, F' might be part of S's total evidence relative to the claim that a liquid of chemical composition C is on the knife but not relative to the claim that the defendant is guilty since his acquaintance with the fact (along with background knowledge) allows him to use this fact in assessing the truth of the former but not the later proposition.

I want to stress that I'm not rejecting the requirement of total evidence but merely suggesting a new way of understanding it. There is a distinction to be drawn between *parts* of S's evidence relative to *P* and S's *total* evidence relative to *P*. Fred's evidence relative to the theorem T might include the fact that Wilma provided testimony that T is true. Such evidence supports T. Nevertheless, if Fred's evidence relative to T also includes testimony from six other expert mathematicians that Wilma's proof is flawed, Fred is rationally required to proportion his belief to the degree of support offered by this total evidence relative to T rather than Wilma's testimony alone.

All of this falls out of the idea of evidence possession that stresses a connection between a subject S *having* some evidence and S's being able to *use* that evidence in assessing a certain proposition *P*. I haven't yet given an account of evidence possession that will vindicate the ideas that: (i) a subject might be able to use a fact F in assessing a proposition *P* but not the proposition *Q* and, thereby, (ii) a subject's total evidence is to be relativized to a proposition. I think Neta is right that mere acquaintance with a fact F that is probabilistically relevant to a proposition *P* is insufficient for possessing evidence *in a way that allows F to play a*

*justificatory role*. What more is required is developed further in chapters 5 and 7. But what I've said here is sufficient for responding to Neta's objection to the fact ontology.<sup>150</sup>

#### 4.1.2.3 Williamson: Evidence & Confirmation Relations

The previous objections target the fact ontology's application to justificatory evidence. As such, my replies partly depend on features of justification that will only be vindicated by the work in chapters 5-7 where I develop and defend a fully formed theory of justification. There is, however, one remaining objection that targets the fact ontology of evidence that I must respond to in order to even get my theory of justification off the ground. And I take this to be the most important objection to the fact ontology of evidence.

Timothy Williamson has developed an influential type of argument for the claim that evidence consists of propositions. He begins by noting that appeals to evidence enter our reasoning in virtue of this evidence standing in explanatory relations, probabilistic relations, and relations of logical inconsistency. Williamson then argues that only propositions can enter into these relations, and thereby concludes that evidence must consist of propositions (or a set of propositions). The general argumentative structure can be represented thusly:

1. Evidence is the kind of thing that can enter into relations of kind K.
2. Only propositions can enter into relations of kind K.

---

3. Therefore, evidence consists only of propositions.

---

<sup>150</sup> As a preview: In chapter 5 I argue that possessing evidence in a way relevant to justification requires that a subject S not only be aware of E but also that S appreciates the evidential connection between E and the proposition P. Applying this to the initial view that a subject's evidence consists of facts with which she is acquainted (that Neta attacks), we would arrive at the full view of chapter 7 that her evidence relative to P consists of those facts F that she is acquainted with and such that she is acquainted with the connection between F and P (or acquainted with a chain of probability relation that mirrors this connection). This easily deals with the case of Fred and Wilma. Fred doesn't have the axioms of set theory as evidence relative to the theorem T because Fred doesn't grasp the connection between these axioms and T. Wilma, however, does grasp this connection and so these axioms are included in her evidence relative to T. Similar remarks apply *mutatis mutandis* to the other cases.

Substituting each type of confirmation relation (mentioned above) for K we get three different arguments for the conclusion that evidence must consist only of propositions.

I take premise 1 of each resulting argument to be uncontroversial. My discussion of evidence suggests that evidence isn't merely the kind of thing that stands in these confirmation relations but that this is *constitutive* of being evidence. Thus, I must reject premise 2: facts enter into the relevant relations. Why does Williamson hold that only propositions (i.e. possible contents of beliefs) can enter into these relations?

Evidence must be the kind of thing that can be explained by various hypotheses. But hypotheses explain *why* such-and-such is the case. Why-explanations “can be put in the form of ‘--- because ...’, which is ungrammatical unless declarative sentences, complements for ‘that’, fill both blanks.”<sup>151</sup> What hypotheses explain, i.e. the evidence, is *that P*. It makes no sense to simply explain an event, sensation, or object. An explanation is always an explanation *that* the event occurred, *that* I have some sensation, or *that* some object exists. Williamson therefore concludes that evidence must consist of propositions.

Evidence is also the kind of thing that can make hypotheses probable. However, Williamson explains that “what gives probability must also receive it.”<sup>152</sup> For instance, consider the fact that we can compare the probability of two competing hypotheses H and H\* relative to a common body of evidence E using the following ratio:

$$\frac{\Pr(H) \times \Pr(E | H)}{\Pr(H^*) \times \Pr(E | H^*)}$$

Such a comparison depends on our contrasting the *probability of the evidence* conditional upon the truth of the competing hypotheses. Similarly, Bayes theorem expresses the probability conditional upon some evidence as follows:

---

<sup>151</sup> Williamson (2000), p. 195.

<sup>152</sup> Williamson (2000), p. 196.

$$\Pr(H|E) = \frac{\Pr(E|H) \times \Pr(H)}{\Pr(E)}$$

Making sense of Bayes' theorem appears to require that we can attribute a probability to the evidence. If evidence is the kind of thing that can make various hypotheses probable then evidence itself would appear to be the kind of thing that can be made probable. But propositions are the bearers of truth-value and so only propositions can be probable.

Evidence is also the kind of thing that we can use to rule out various hypotheses on the grounds that these hypotheses are *inconsistent* with our evidence. The most natural way to understand this inconsistency is in terms of our evidence *entailing* the denial of the relevant hypothesis. Entailment is a logical relation defined over propositions.  $P$  entails  $Q$  when it's absolutely impossible for  $P$  to be true while  $Q$  is false. The bearers of truth-value, however, are propositions not facts.

A first line for defending the fact ontology is to point out that our expressions of various propositions are attempts to cite putative facts as putative evidence. Moreover, true propositions used in explanatory and probabilistic reasoning successfully refer to actual evidence that exists (and which, in the right conditions, we actually possess). When we engage in explanatory reasoning we use a hypothesis to explain *that* entity X exists, *that* event X occurred, or *that* entity X has property  $\Phi$ , etc. However, it's plausible to hold that these 'that' clauses are being used to refer to the facts that make the expressed proposition true. Similarly when we talk about the probability of some evidence given the truth of some hypothesis we can think of this as the probability of the existence of this evidence (i.e. the probability of this fact's existence relative to the truth of the relevant hypothesis).

Let's first consider the idea that what we explain in explanatory reasoning are facts and not propositions. This idea can be well motivated by considering a context where explanatory reasoning is especially prevalent: scientific contexts. Scientists offer (possible) explanations for a variety of phenomena. For instance, scientists often consider a variety of explanations for the following:

- 1) Why the tides rise and fall.
- 2) Why fossils were found in certain sedimentary layers.
- 3) Why the planets in our solar system have certain orbits.

Are these scientists attempting to explain propositions (i.e. possible contents of belief) about tides, fossils, and planets? An affirmative answer is odd. Intuitively, the scientists are offering possible explanations of the *fact* that the tides rise and fall, the *fact* that such-and-such fossils were found at location L, and the *fact* that our planets have such-and-such orbits.

This unreflective view of explanation within scientific context holds up upon reflection. If we're explaining anything about a proposition when we explain why the tides rise and fall, we're explaining the proposition's *truth*. In a possible world where tides don't rise and fall there can be no explanation of the rise and fall of the tides. At most, in the possible world under consideration, certain hypotheses concerning gravitational forces are such that they *would* explain the rise and fall of tides in a world where certain other properties hold. But an explanation of *P*'s truth would seem parasitic upon explaining the existence of the fact <P> that makes *P* true (I will use a capital letter in brackets to represent a fact or group of facts that together act as *a* truth-maker for the proposition expressed by the sentence placed between the brackets and I'll use a capital letter in italics to represent a proposition).

Even though facts are the ultimate target of explanation, we ought to admit three kinds of explanatory relations. First, facts stand in explanatory relations in the primary sense. Scientists often offer *possible* explanations for various facts about the world. However, it seems that only a *true* hypothesis could *actually* explain the existence of the explanandum fact. So in the primary sense, the *fact* that certain laws pertaining to gravitational attraction obtain along with various *facts* about the mass, location, and orbits of the earth and moon stand in an explanatory relation to the *fact* that the earth's tides rise and fall.

We can also make sense of the idea of competing hypotheses (propositions that *purport* to pick out some fact about the world) being *possible* explanations of a certain fact. For

instance, prior to the discovery of Neptune the hypothesis that there was another planet (i.e. Neptune) was one among many *possible* explanations for the perturbations in the other planets orbits. As it turns out the proposition that Neptune exists is the *actual* explanation. However, we ought to insist that there is an explanatory relation that holds between the competing hypotheses and the fact that the other planets had certain perturbations in their orbits that accounts for these being *possible* explanations. The best way to understand this is that a proposition  $P$ 's standing in an explanatory relation to a fact  $F$  just means that *if*  $P$  were true then there would be a fact  $F^*$  that made  $P$  true and  $F^*$  would explain the existence of  $F$ .

Similarly, we can make sense of the idea that propositions stand in a kind of explanatory relations to one another even though the entity that *actually* gets explained or *actually* explains is always a fact. For instance, that  $P$  stands in an explanatory relation to  $Q$  just means that *if*  $P$  and  $Q$  were both true then there would be: (i) a fact  $\langle P \rangle$  that makes  $P$  true, (ii) a fact  $\langle Q \rangle$  that makes  $Q$  true, and (iii)  $\langle P \rangle$  would explain the existence of  $\langle Q \rangle$ . This relation can hold independent of the truth-value of both  $P$  and  $Q$ .

It's important to notice that when the propositions that stand in an explanatory relation are true, these relations will *necessarily* be isomorphic as pictured in figure 1.



Figure 1: Isomorphic Explanatory Relations

The fact that these isomorphic relations *necessarily* hold when the relevant propositions are true follows from the fact that the proposition-fact and proposition-proposition pairs of



explanatory relations can be characterized in a way that is parasitic upon our understanding of fact-fact explanatory relations.<sup>153</sup>

We can make similar moves in order to avoid Williamson's arguments from probability and inconsistency relations. Usually we think of making probable and inconsistency relations as holding between propositions. But consider what it means to say that  $Q$  is inconsistent with  $P$ . This means that  $P$  entails  $\neg Q$  which, in turn, means that it's absolutely impossible for  $Q$  to be true while  $P$  is true. The truth of either proposition requires a fact to make it true. So we can define a kind of inconsistency relation between fact-proposition pairs. A proposition  $Q$  will be inconsistent with a fact  $\langle P \rangle$  when  $\langle P \rangle$  makes true a proposition  $P$  which entails  $\neg Q$ . It can *never* be the case that the fact  $\langle P \rangle$  exists and  $Q$  is true. This relation between that fact  $\langle P \rangle$  and the proposition  $Q$  is an internal relation that necessarily holds given the existence and intrinsic natures of  $\langle P \rangle$  and  $Q$ .<sup>154</sup> A proposition can be inconsistent with our evidence even if we accept the fact ontology.

This same move works in the case of probabilistic relations. A set of evidence, i.e. a set of facts, can make a proposition  $P$  probable *and* the existence of this evidence (i.e. the existence of this set of facts) can be made probable by  $P$ . There is nothing unnatural about this. After becoming aware (perhaps *knowing*) that some event  $E$  occurred or that fact  $F$  obtains we might be certain of  $E$ 's occurrence or  $F$ 's existence and still truly claim that the

---

<sup>153</sup> Explanatory relations are actually much more complicated than is suggested in this main text. These complications are due to the fact that most of our discussions of explanations actually only pick out a salient part of an explanation where this salience is determined by context. Such issues, however, are tangential to my main points.

<sup>154</sup> A characterization of this inconsistency that proceeds in terms of a fact  $\langle P \rangle$  making true  $P$  which entails  $\neg Q$  is somewhat misleading. The relation isn't a tertiary relation but a binary relation between the fact  $\langle P \rangle$  and  $Q$ . However, explaining the relation by means of the intermediary proposition that is made true by the fact and entails  $\neg Q$  is helpful for explaining and thinking about the relations. Moreover, this way of characterizing the relations is helpful for seeing that this inconsistency relation would necessarily mirror inconsistency relations between propositions.

occurrence of this event or the existence of this fact was incredibly unlikely relative to the truth of some other propositions.

Nevertheless, we must explain the relation between probability relations holding between proposition-proposition, fact-proposition, and fact-fact pairs. This is quite easy on a relative frequency interpretation of probability. According to the relative frequency model,  $P$  makes probable  $Q$  iff there is a reliable correlation (perhaps construed as a counterfactual frequency) between the truth of  $P$ -like propositions and the truth of  $Q$ -like propositions. Since this relation must be characterized in terms of the *truth* of kinds of propositions and facts are what make propositions true, we can extract a notion of probability relations that hold between fact-proposition and fact-fact pairs. A fact  $\langle P \rangle$  makes probable (in a relative frequency sense)  $Q$  when there is a reliable correlation between  $\langle P \rangle$ -like facts and the truth of  $Q$ -like propositions. A fact  $\langle P \rangle$  will make probable (in a relative frequency sense) the fact  $\langle Q \rangle$  when  $\langle P \rangle$ -like facts are reliable correlated with  $\langle Q \rangle$ -like facts. Provided that we individuate types of propositions and facts in the same way, when the relevant propositions are true the three relations will be isomorphic to one another as depicted in figure 2.




Figure 2: Isomorphic Reliability Relations

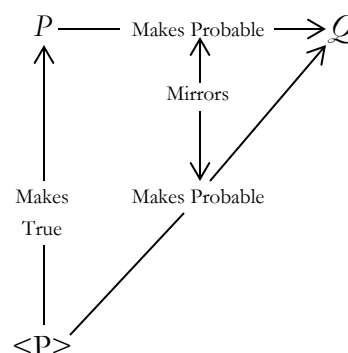
It's quite easy to make sense of probability relations holding between fact-fact, fact-proposition, and proposition-proposition pairs when probability is understood along the

lines of reliable co-variation. I see no reason not to extend this to the case of logical theories of probability that construe probability as a kind of partial entailment. When a proposition  $P$  entails  $Q$  it's absolutely impossible for  $P$  to be true while  $Q$  is false. This relation is an internal relation that necessarily holds between  $P$  and  $Q$  given the existence and intrinsic characters of  $P$  and  $Q$ . Moreover, this is a relation that can be known to hold *a priori*. Keynes (1921) suggested that probability be modeled on this notion of entailment. On such a view, the relation of one proposition  $P$  making probable  $Q$  is an internal relation that can be known to hold a priori, and where entailment is the *upper limit* of making probable. Since this relation is an internal relation its existence must be grounded solely by the existence and intrinsic nature of its relata rather than external relational facts such as facts about relative frequencies. An entailment (the upper limit of this logical making probable relation) between  $P$  and  $Q$  isn't constituted by the fact that the truth of *P-like* propositions is invariably correlated with the truth of *Q-like* propositions. Even if *P-like* and *Q-like* propositions are always true together,  $P$  might not entail  $Q$  since this might be an accidental correlation. Moreover, entailment can't even be identified in terms of invariable co-variation across all possible worlds. Only the actual world exists and we must find truth-makers for possibility claims in the one and only actual world. The truth-maker for the claim that  $P$  entails  $Q$  isn't that the truth of *P-like* propositions is always correlated with the truth of *Q-like* propositions across all possible worlds. If anything, the existence of this entailment relation *in the actual world* is the truth-maker for this claim about invariable co-variation across possible worlds. And since entailment is simply construed as the upper limit of this making probable relation similar remarks apply to the making probable relation.

The idea behind the partial entailment theory of probability can be extended in such a way to posit a *sui generis* making probable relation that holds between fact-proposition and fact-fact pairs. The relevant making probable relations will internal relations. A fact  $\langle P \rangle$  will

make probable a proposition  $Q$  iff  $\langle P \rangle$  makes true a proposition  $P$  that makes probable  $Q$ .<sup>155</sup> Such a relation will be an internal relation. The correspondence/truth-making relation between  $\langle P \rangle$  and  $P$  is an internal relation and the making probable relation between  $P$  and  $Q$  is an internal relation. As such, there will be an internal relation of making probable between  $\langle P \rangle$  and  $Q$  isomorphic to the relation between  $P$  and  $Q$  as depicted in figure 3.

PROPOSITIONS:



FACTS:

Figure 3: Isomorphic “Logical” Probability Relations

Of course, there is an important difference between the making probable relation at the level of propositions and the making probable relation between fact-proposition pairs. The former can be known to hold *a priori*. Nevertheless, even though a making probable relation holding between fact-proposition pairs is an internal relation one couldn’t know that the relation exists *a priori*. The reason for this is that the existence of the relation depends on the

---

<sup>155</sup> If one allows that the same fact can make true more than one proposition then this might have to be restricted in various ways. Let’s assume that the fact  $\langle \text{Sam is in pain} \rangle$  makes true the proposition *Sam is in pain* and the proposition *someone is in pain*. The latter proposition increases the probability of the proposition *Siri is in pain* but the former does not. So what should we say about the fact  $\langle \text{Sam is in pain} \rangle$ ? It makes true a proposition that increases the probability that *Siri is in pain* and makes true a proposition that fails to increase the probability that *Siri is in pain*. Does the fact increase or fail to increase the probability that *Siri is in pain*? Presumably we want to claim that the fact  $\langle \text{Sam is in pain} \rangle$  does not increase the probability that *Siri is in pain*. But then this constitutes a counterexample to my claim that  $\langle P \rangle$  makes probable  $Q$  iff  $\langle P \rangle$  makes true a proposition  $P$  that makes probable  $Q$ . What we need is some way to pick out the relevant generality of proposition made true by the fact. Perhaps the best way is to claim that  $\langle P \rangle$  makes probable  $Q$  iff the most specific proposition  $P$  made true by  $\langle P \rangle$  makes probable  $Q$ . Alternatively, one might become tempted to posit a distinct truth-maker for each logically distinct true proposition.

existence of the relata, and the existence of the facts that stand in evidential relations may not be knowable *a priori*. What is knowable *a priori* is only that *if* some fact were to exist it would make probable some proposition's truth. I can know *a priori* that *if* the fact that George is in pain *were* to exist then this fact *would* make it probable that someone is in pain *a priori*. I can't know *a priori* that the fact that George is in pain does make it probable that someone is in pain simply because I can't know *a priori* that George is in pain.

Our discussion provides a metaphysic of explanatory, inconsistency, and probability relations that allows us to reject Williamson's 2<sup>nd</sup> premise for the conclusion that evidence must consist of positions. Facts can stand in the relevant confirmation relations.<sup>156</sup> This metaphysics is surely to be controversial but we'll see in chapter 7 that it has great theoretical benefits in terms of developing a theory of justification that respects the motivations of classical foundationalism while avoiding many of its costs.

#### 4.2 Factoring Accounts of Justifying Evidence

In this chapter I've focused on developing and defending a characterization the *general* category of evidence. Specifically I've been focused on defending a fact ontology of

---

<sup>156</sup> Victor DeFate (2007) has suggested that one could bolster Williamson's argument for the propositional character of evidence by stressing his appeal to the roles that evidence plays in a subject's reasoning. Williamson's contention isn't just that evidence is the kind of thing that can enter into explanatory, probabilistic, and inconsistency relations but also that it's the kind of thing that we can use in these kinds of reasoning. As such, evidence must be the kind of thing that we can take into account in our reasoning. DeFate explains how this might be used to bolter Wiliamson's argument: "[I]f one is to reason with one's evidence, either probabilistically, deductively, or explanatorily, the evidence must be the sort of thing that one can *grasp* or *understand*, namely, a proposition. (It makes little sense to grasp an event, although we can grasp *that* an event took place). So, while there may be theories of probability or explanation whereby events are implicated, when we turn to explanatory, probabilistic or deductive *reasoning* with the evidence, we are arguably dealing only with what is propositional" DeFate (2007). It isn't clear however that this favors the proposition over the fact ontology of evidence. We can *grasp*, *understand*, and *think about* facts or propositions. In fact it is initially plausible that I can take into account various facts about my experiences in my reasoning without actually entertaining a proposition that represents the fact that I have these experiences. I discuss this idea further in chapter 5. If I'm right about this then we ought to construe facts as evidence in order to capture the sense in which these experiential facts can enter into our reasoning.

evidence. Before moving on to the next chapter I want to briefly consider an objection to the very methodology of which my defense of the fact ontology is a part.

For my purposes in this dissertation a characterization of *justificatory* evidence is most important. In this chapter I've focused on characterizing evidence (including justificatory evidence) by focusing on the ontology of evidence independently of anyone's possessing it. This method presupposes the idea that we can factor out the conditions that make something justificatory evidence or an epistemic reason for belief and conditions under which some person *has* that evidence or reason. Mark Schroeder (2008) refers to this as a factoring account of a subject's evidence. Such an approach treats *having* evidence much like *having* an ice cream cone. The ice cream cone qualifies as an ice cream cone independently of my *having* it. However, one might object that a subject's *having* evidence should be thought of as being more analogous to the notions of having a father or a golf partner—these examples come from Schroeder. A person cannot *be* a father independently of someone's *having* that father. Someone can't *be* a golf partner without someone's *having* that person as a golf partner. In these examples we can't *factor* the notion of a person *having* something of kind K into the components of some entity X belonging to kind K *and* some person's being a having relation to the thing of kind K. In these examples, someone's being the *having* relation is constitutive of what it is to be a thing of kind K.

If there are reasons to think that *having* evidence is like *having* a father or golf partner—if there are reasons to be suspicious of the idea that epistemic reasons for belief can exist independently of anyone having those reasons—then the method I'm employing for developing an evidentialist theory of justification would be inappropriate.

Why might you think epistemic reasons fail to exist independently of a subject's having those reasons? You might reason as follows: justifying evidence or epistemic *reasons for belief* existing independently of someone *possessing* that evidence or those reasons is nonsense since evidence that no one *has* is irrelevant to justification. Conee and Feldman seem to make just this complaint:

Evidentialism holds that a person's doxastic justification<sup>157</sup> is a function of the evidence that the person has. One might think that the way to understand this is first to define, or characterize, what evidence *is*, and then to explain what it is for a person to *have* a particular bit of evidence. To see why this is true, consider, for example, your current perceptual experience. Is this experience evidence? The best answer seems to be that it is evidence for you, but it is not evidence for the rest of us. It is part of your evidence for the proposition that you are reading an essay on epistemology. But your experience itself is not evidence for the rest of us.<sup>158</sup>

In this passage, Conee and Feldman presuppose an ontology of evidence that identifies evidence with mental states, but their reasoning can easily be translated into my fact ontology. Is the fact that you have a perceptual experience with a particular character C to be characterized as evidence for the claim that you're reading an essay in epistemology? Conee and Feldman suggest that this fact only provides *you* with justification for believing the relevant claim. As such, the experiential fact constitutes a reason for you to believe that you're reading an essay in epistemology but it doesn't constitute a reason for me to believe this proposition. Or, in other words, this fact is evidence for the claim that you're reading an essay in epistemology relative to you but not relative to me.

This moves to fast. Considering an analogy with reasons for action helps illustrate why. Consider Isabelle and Rosalie. Both are incredibly thirsty and there is a glass in front of them with a transparent liquid that looks, smells, feels, and tastes just like water. The liquid, however, is actually a very cleverly disguised poison. Isabelle has no idea that the liquid is a poison and (with good reason) believes that it's simply a glass filled with water. Rosalie, on the other hand, *knows* that the liquid is poison. There is a sense in which Isabelle doesn't *have* any reason not to drink from the glass since she can't take the fact that it's filled with poison

---

<sup>157</sup> The reader should note here that Conee and Feldman are using the phrase "doxastic justification" in this passage differently than I and most other epistemologists use the phrase. What they refer to as "doxastic justification" is supposed to pick out a person's having propositional justification for a belief that he or she actually has. Conee and Feldman use the label "well-founded belief" to describe what I and others refer to as doxastic justification.

<sup>158</sup> Conee and Feldman (2008), p. 88.

into account when deciding how to act. Rosalie, however, does *have* a reason not to drink from the glass since her knowledge allows her to take account of this fact in her decision process. But none of this threatens the idea that the fact that the glass is filled with poison constitutes a reason for both women not to drink from the glass. If Isabella were to ask you whether there was any reason for her not to drink from the glass we wouldn't tell her "no" simply on the grounds that she failed to know the liquid is poison. The reason exists for both women not to drink from the glass, the difference between Isabella and Rosalie is simply that the latter *has* this reason and the other doesn't.

Similarly, in the case considered by Conee and Feldman, the experiential fact failing to provide me with justification for believing that you're reading an essay in epistemology doesn't show that it isn't a reason for me to so believe. It only shows that this experiential fact is a reason which I fail to possess for it is only the *possession* of evidence that provides justification. I can't take the epistemic reason for me to believe *P* into account when assessing the truth of *P*, but this doesn't show that it's not an epistemic reason for me to believe *P* nonetheless.

#### 4.3 Conclusion

This concludes my discussion of the ontology or metaphysics of evidence. Two theses developed here are most important for my development of a theory of justification. First, there is my fact ontology of evidence. Evidence consists of a set of facts that is some kind of an indication of the truth or falsity of a proposition. Second, there is an isomorphism between the evidential relations that hold between sets of facts and propositions and various confirmation relations that hold between propositions made true by the set of facts and other propositions.



## CHAPTER FIVE

## HAVING EVIDENCE, AWARENESS, AND JUSTIFICATION

I've adopted the methodology of starting with the evidentialist maxim (EM) as a framework for developing a theory of justification. Developing this into a substantive theory requires giving accounts of evidence, evidence possession, and evidential support. This is a nice starting point for epistemological theorizing even if developing an account of the relevant evidential concepts were to eventually lead one to reject EM.

Chapter 4 discussed the metaphysics of evidence. Evidence consists of a fact or a set of facts which is an indication of a proposition's truth or falsity. This applies to all categories of evidence: scientific evidence, legal evidence, historical evidence, justificatory evidence, etc. What a theory of justification requires is an understanding of *justificatory* evidence specifically. There is some truth-indicating relation R such that a set of facts F standing in R to P is constitutive of F's being the kind of thing that can provide subjects with justification for believing P. For the moment let's stipulate that justificatory evidence for a proposition P, i.e. an epistemic reason for believing P, consists of a set of facts that stands in the truth-indicating relation R to P while leaving an investigation of R for the next chapter.

Our discussion only offers a characterization of a set of facts being justificatory evidence *independent of anyone having said evidence*. If a subject doesn't *have* evidence E, however, E indicates nothing *to her*. E might be a reason for S to believe P but if S doesn't *have* that reason it doesn't provide S with any justification. Even assuming conclusive evidence for believing that Joe Smith was Jack the Ripper exists, if this isn't evidence I possess then it indicates nothing *to me* about Jack the Ripper's identity. Therefore, it initially appears that only evidence I have can provide me with justification for believing anything. And so we must now turn to a consideration of evidence possession and its relation to justification.

### 5.1 Evidence Possession and Awareness

A natural way to explain the notion of a subject's evidence is as those facts that the subject can *use* to assess the truth or falsity of some claim. Such a characterization appears trivial at first blush and this is reason to think I've accurately characterized the concept.<sup>159</sup> Our gloss nonetheless illustrates that our concept of having evidence is intimately tied to the thought that a subject has certain abilities relative to that evidence thereby providing direction for developing a substantive account of evidence possession. We need to explain what these abilities are and the nature of the cognitive relation grounding these abilities.

Any apt characterization of evidence possession must (minimally) account for S's ability to use this evidence in reasoning concerning a proposition's truth. S *has* evidence E for or against *P* only if S can use E in her reasoning about *P*. In other words, E is something S can take into account when forming a belief concerning *P* in the sense that S can form beliefs about *P on the basis of E*.<sup>160</sup>

I mean for "reasoning" and "basis" to be taken more liberally than some philosophers. There is a sense in which we can use some of our experiences in an attempt to determine the truth of a proposition. In determining whether it rained last night I can take

---

<sup>159</sup> As pointed out in fn. 149, this gloss on the nature of evidence possession creeps up in a number of philosopher's discussions of evidence possession. Consider the following examples: "[a] natural idea is that S has a *body of evidence, for use in the assessment of hypotheses*" (Williamson, 2000); "evidence must be accessible to the knower in at least the following sense: **it must be available for use in reasoning and inference**" (Greco, 2011); "[f]or any bit of evidence that you have, **it is possible for you to use that evidence in the rational regulation of your degrees of confidence**, or of your attitudes more generally" (Neta, 2008); "[i]t's a necessary condition of S's being rationally required to proportion her confidence across hypotheses in proportion to the support that those hypotheses receive from [the evidence] **that S has the ability to do this**" (Neta, 2008); "[o]ne useful way to identify what we take to be a person's evidence at a time is to say that it is the information or data **the person has to go on in forming beliefs**" (Conee and Feldman, 2008). The bold portions mark out *my* emphasis.

<sup>160</sup> This aspect of my view on evidence possession and the resulting awareness requirement has been greatly influenced by the discussion of "grounds" by William Alston (1988) and Jack Lyons (2009) chapter 2. However, unlike Lyons, I think all justification is evidential justification. Unlike both Lyons and Alston, I hold that awareness of (a non-conceptual awareness) the adequacy of evidence is also necessary for justification.

into account my current perceptual experience as of the roads looking wet without actually forming any occurrent belief that I have such an experience. Even if I use this experience in conjunction with various background beliefs (e.g. that wet roads are usually caused by rainfall) in assessing whether it rained last night, I also utilize the perceptual experience itself. The fact that I have a perceptual experience as of the roads looking wet can be at least part of my reason(s), motivation, or basis for believing that it rained last night.

A subject *S* will possess some evidence *E* for *P* only if *S* is in a position to base a belief that *P* on *E*. There are two things to note here: (i) this only specifies a necessary condition on evidence possession and (ii) I'm not advocating that the contraposition of this conditional is true since surely a subject might base a belief on something that she takes to be but isn't actually evidence for *P*. In order to get some requirements on evidence possession we need to get a grasp on this notion of "being in a position" to base one's belief on *E*. Most accounts of the basing relation include a causal component and so a natural thought is that *having* evidence *E* for *P* is to be understood solely in terms of causal sensitivity to that evidence. Causal sensitivity, however, is surely inadequate.

For a dramatic illustration that evidence possession cannot be explained solely in terms of causal sensitivity consider the case of Descartes' evil demon. Imagine that the details are as follows: the all-powerful evil demon's desire to produce in you a hallucinatory experience as of a red table directly causes you to have said experience, which then causes you to form the belief that there is a red table in front of you. Your belief about the external world would be causally sensitive to the demon's desire but it strains credulity to think that the fact that the demon desires to produce in you a hallucinatory experience as of a red table is something you're in a position to use as a basis for a belief about whether there is a red table in front of you. This fact about the demon's desire isn't something that you're in a position to take into account in forming your belief.

A more instructive point, however, is to consider that even in ordinary perception my beliefs are clearly causally sensitive to a number of factors, not all of which are things

that I can take into account in determining the truth of a proposition. We might, for instance, attempt to isolate the visual processes that begins with light hitting the retinas and eventually leads to the production of a belief about the color of the table in front of me. The resulting belief would be causally sensitive to each step in this process but it's a conceptual stretch to hold that I base my belief on each step. Surely I don't base my belief on the fact that such-and-such image is projected onto my retina. Nor do I base my belief on the fact that such-and-such neural activity is taking place in the V1 area of the visual cortex. These aren't facts that I (*the person or subject*) can take into account in forming beliefs about the table's color. At best these are facts my sub-personal mechanisms can take into account (though such a description is best construed as metaphor). Being causally sensitive to a fact F is insufficient for being in a position to take F into account in assessing the truth or falsity of a proposition, and causal sensitivity is thereby insufficient for a fact F's being evidence that the subject has (i.e. possesses).

So what is it that I (the person or subject) am in a position to take into account when forming a belief about the table and its color? If anything, I base my belief that the table is red on the fact that I have a certain visual experience with a certain character. What singles out this particular causal factor as a candidate for the basis of my belief that the table is red? Intuitively the best answer is that this is the only part of the process of which I'm aware. My awareness accounts for this fact being a potential basis for my belief, i.e. something that I'm in a position to take into account. Awareness is necessary for a fact being part of my evidence. When my awareness of this fact (perhaps along with other conditions being met) causes my belief, the fact thereby acts as an *actual* basis of my belief.<sup>161</sup>

---

<sup>161</sup> Notice that I'm hesitant to say that my awareness of a fact F causing me to believe *P* is sufficient for F's constituting one of my bases or reasons for believing *P*. Basing might require other conditions to be in place as well. For instance, in a case of self-deception I might be aware of a desire that *P* and this might thereby cause me to believe *P*. But it's odd to say that I *base* my belief on my desire that *P* or that I *base* my belief on self-deception. It strikes us as natural to describe this as a case where the relevant belief is the result of my desire/self-deception and it strikes us that this is importantly different than basing my belief on these things. Lyons (2009) makes this point regarding

Our discussion here motivates the idea that the following awareness requirement captures a significant part of our concept of evidence possession:

**AR** S *has* evidence E only if S is in some way aware of E

Conjoining the evidentialist maxim EM with an acceptance of AR produces a substantive theory of justification that can be contrasted with what I refer to as “radical externalism.”

Certain theories of justification allow that a belief-forming process such that every step in the process leading to the belief occurs at the sub-personal level (i.e. these steps make no difference to the subject’s conscious life) can confer justification. Tracking theories, information-theoretic theories, process reliabilism, and proper-functionalism have traditionally been developed so as to allow for this possibility.<sup>162</sup> Such views are exemplars of externalism and I characterize them as “radically” externalist in order to distinguish them from views that I discuss later that appear to be more moderately externalist theories.

Radical externalists can accept either EM or AR but they cannot accept the conjunction. Radical externalism is clearly in conflict with EM+AR since this combination entails that a subject’s awareness of supporting evidence is necessary for justification. When I consider AR I find it hard to construe it as anything less than an analytic truth. Characterizing a sub-personal input to a belief-forming process as *being* evidence is harmless (my characterization of evidence in the previous chapter allows as much) but characterizing a sub-personal input as evidence that a subject/person *has* seems to radically alter our concept of having evidence. Nevertheless, it’s best to take AR as simply stipulating how I’m using the

---

self-deception. What more is involved in the basing relation? It’s hard to say. My point here is only that possessing evidence in any relevant sense involves an awareness of that evidence.

<sup>162</sup> Information-theoretic and tracking theories have traditionally been used in attempt to provide analyses of knowledge rather than justification. See Dretske (1981) and Nozick (1981). Nevertheless, one could easily modify such views into an analysis of justification. Despite focusing on resurrecting Nozick’s (1981) tracking theory of knowledge, many of the insights from Roush (2006) could easily be translated so as to give an account of justified belief where justified is understood as a binary notion or (with a bit more work) used in an attempt to develop a very interesting tracking theory of justification where this is understood as a gradated notion.

phrase “evidence possession.” At this point the crucial epistemological question becomes: should we accept the evidentialist idea that having supporting evidence is necessary for justification once we’ve adopted AR?

## 5.2 Justification and Awareness: Part I

### 5.2.1 Appealing to the Cases

Why think awareness of supporting evidence is necessary for justification? The best motivation is illustrated by well-known and widely discussed counterexamples to various externalist theories of justification:

THE BRAIN LESION PATIENT: Suppose *K* suffers from a serious abnormality—a brain lesion, let’s say. This lesion wreaks havoc with *K*’s noetic structure, causing him to believe a variety of propositions, most of which are wildly false. It also causes him to believe, however, that he is suffering from a brain lesion. *K* has no evidence at all that he is abnormal in this way... but surely *K* does *not* know [or have justification for believing] that he is suffering from a brain lesion. He has no evidence of any kind—sensory, memory, introspective, whatever—that he has such a lesion; his holding this belief is, from a cognitive point of view, no more than a lucky (or unlucky) accident.<sup>163</sup>

TRUETEMP: Suppose a person, whom we shall name Mr. Truetemp, undergoes brain surgery by an experimental surgeon who invents a small device which is both a very accurate thermometer and a computational device capable of generating thoughts. The device, call it a tempucomp, is implanted in Truetemp’s head so ... as to transmit information about the temperature to the computational system of his brain. This device, in turn, sends a message to his brain causing him to think of the temperature recorded by the external sensor... All told, this is a reliable belief-forming process. Now imagine, finally, that he has no idea that the tempucomp has been inserted in his brain... [and] never checks a thermometer to determine whether these thoughts about the temperature are correct. He accepts them unreflectively.<sup>164</sup>

NORMAN THE CLAIRVOYANT: Norman, under certain conditions which usually obtain, is a completely reliable clairvoyant

---

<sup>163</sup> Plantinga (1993), p. 195.

<sup>164</sup> Lehrer (1990), pp. 163-164.

with respect to certain kinds of subject matter. He possesses no evidence or reasons for or against the general possibility of such a cognitive power or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power under circumstances in which it is completely reliable.<sup>165</sup>

These cases are naturally interpreted as involving processes of the kind mentioned in the previous section where each step up to the belief takes place sub-personally.<sup>166</sup> The brain lesion patient, Mr. Truetemp, and Norman all lack any awareness of supporting evidence for their respective beliefs. As I defined it, radical externalism is any view that allows a subject's belief be rational in these kinds of cases.<sup>167</sup> The difficulty is that when cases where a subject lacks any awareness of supporting evidence are explicitly laid before us it's almost irresistible to describe the relevant beliefs as obviously irrational!

Considering these kinds of cases strongly suggests that these radical forms of externalism are extremely implausible. The most natural and best explanation for why these subjects lack justification is that they lack any awareness of supporting evidence or reasons for holding that their beliefs are true. Such an explanation applies equally to each of the three scenarios. Therefore, we ought to reject radical versions of externalism that would allow a source to confer justification absent any awareness of supporting evidence.

---

<sup>165</sup> Bonjour (1985), p. 41.

<sup>166</sup> Notice that, in presenting the case of the brain lesion and of Norman, both Plantinga and Bonjour claim that the subjects fail to *have* evidence for the relevant belief. This suggests that these authors implicitly accept AR, which provides further evidence of how natural a requirement AR is.

<sup>167</sup> It's possible that a belief produced by a process that takes place entirely sub-personally could be rational in virtue of the subject's being aware of a positive track-record for beliefs produced in this way. Awareness of such a track-record would, however, constitute an awareness of supporting evidence that her belief is true. The important point is that radical externalism allows that processes such as these can produce justification even absent awareness of this track-record evidence.

However, other explanations are available. And before further developing the case for an awareness requirement on justification we need to consider how a radical externalist can attempt to accommodate the intuition that these individuals lack justification.

### 5.2.2 Radical Externalist Attempts at Accommodation

Certain varieties of radical externalism are well situated for explaining the brain lesion patient's lack of justification without invoking an awareness requirement on justification. Both process reliabilism and proper functionalism give a verdict that accords with our intuitions in this case and provide plausible explanations of our intuition that a negative epistemic evaluation of the belief is appropriate.

Consider first the reliabilist explanation. It's impossible that a brain lesion causes a false belief that you have a brain lesion. A theory that made justification a function of a belief's being caused by its truth-maker or even the belief's tracking the truth would thereby be threatened by the brain lesion case. Nonetheless, the description explicitly states that the lesion causes the patient to form many other beliefs that turn out to be predominantly false. Again, on a crude formulation, process reliabilism is a radically externalist view according to which a belief is justified iff it was produced by a generally reliable process. If the process reliabilist can provide a principled way of "typing" the process that leads from the brain lesion to the patient's true belief so that the false beliefs caused by the brain lesion get categorized as instances of this type of process then it can explain the lack of justification in the brain lesion case without abandoning radical externalism; the irrationality of the brain lesion patient's belief is due to its being produced by an unreliable process.<sup>168</sup>

---

<sup>168</sup> This is a big "if". For one, the generality problem is a monster of a problem for reliabilism and it isn't clear that one can give a principled way of typing any belief-forming process. More worrisome, however, is that one of the best attempts to solve the generality problem is Alston's (1995) and (2005) psychological solution that types the process as the input-output function such that the relevant input includes all of the properties that play a causal role in outputting the belief. In the brain lesion case it's stipulated that the brain lesion causes the belief "I have a brain lesion." And so Alston's solution would appear to make the process type so specific that it would in fact be



Proper functionalism is another form of radical externalism that can explain the brain lesion patient's lack of justification. The lack of justification, according to proper functionalism, is because the patient's belief results from a malfunction or abnormality of the subject's cognitive system; intuitively, the patient's cognitive system isn't supposed function the way it does in the scenario. On a proper functionalist account, a belief is rational iff (i) the belief is produced by a faculty that was designed for the purpose of forming true beliefs and avoiding false beliefs, (ii) the faculty is reliable relative to the environment for which it was designed, and (iii) the faculty is functioning properly (i.e. in the way it was designed to function).<sup>169</sup> The brain lesion patient's belief wasn't produced by a properly functioning faculty but, rather, by a cognitive malfunction. Our picking up on this when we read the description thereby explains our inclination to judge that the patient's belief is unjustified.

We need only modify the brain lesion case slightly in order to show that both the process reliabilist and proper functionalist explanations fail to identify the fundamental epistemic defect. Suppose the lesion produces a large number of beliefs that just pop into the patient's mind spontaneously and, miraculously, this lesion reliably produces true beliefs. Now imagine that the very first belief produced by the brain lesion, and which seems to pop spontaneously into his mind, is the belief that he has a brain lesion. Such a modification does nothing to quell the intuition that his belief is completely irrational. And we can push the example even further. Imagine that (unbeknownst to anyone) God herself directly causes the brain lesion for the purpose of producing these true beliefs. There is still a strong intuition that the first belief that seems as if it just pops into his head is highly irrational. Therefore,

---

perfectly reliable since the other false beliefs would be products of a different input-output function. In which case, process reliabilism lacks a solution to the brain lesion case.

<sup>169</sup> See Bergmann (2004), (2006), and (2008) and Plantinga (1993) for defenses of proper functionalism. Plantinga's discussion is framed around what he labels "warrant" rather than justification which is his technical name for that which turns true belief into knowledge. Plantinga gives a rather penetrating discussion of the etymology of "justification" and its connection to a kind of epistemic deontology that I think he is right to reject.

process reliabilism and proper functionalism cannot explain the full extent of the epistemic defect with the patient's belief that he has a brain lesion.

The failure of the process reliabilist and proper functionalist explanations is as easily illustrated by moving to the case of Mr. Truetemp. Truetemp's tempucomp is a reliable device designed by a scientist for the purpose of producing true beliefs about the temperature. Moreover, the tempucomp is working precisely as the scientist designed it. Truetemp's beliefs about the temperature are products of a process that meets the process reliabilist and proper functionalist requirements for justification. Nevertheless, it's almost irresistible to describe the resulting beliefs as entirely irrational. Each of these views fails to explain why Truetemp lacks justification for his beliefs about the temperature.

Jack Lyons has recently provided an interesting attempt to modify radical externalism in an attempt to explain why we're inclined to describe the subject's beliefs in these supposed counterexamples cases as irrational. Lyons' response relies on a distinction between what he refers to as basic and non-basic beliefs. His distinction, however, denotes something different than beliefs produced by basic and non-basic justification sources as these terms were understood in chapters 2 and 3. Moreover, Lyons' notion of a basic belief is to be distinguished from that of a foundationally justified belief.<sup>170</sup> On the distinction that

---

<sup>170</sup> As I mentioned, Lyons' notion of a basic belief is distinct from the notion of a belief that is produced by basic belief source discussed in connection with the problem of easy justification that in chapters 2 and 3. Lyons' notion of basicity is probably best construed as a kind of *psychological* as opposed to *epistemic* basicity. However, he does appear to suggest that the basicity of a belief has implications regarding *modal* claims about a beliefs epistemic status. Basicity implies (on Lyons' distinction) that a belief *can* possess foundational prima facie justification. What's the modal status of such a claim? Presumably this should be construed as a kind of psychological possibility. But now it appears that basicity is *merely* a psychological notion. Meeting the psychological criteria for basicity is a necessary condition for a belief's being foundationally justified. Surely, however, it's conceivable that *any* belief meets this psychological criteria and the reliability condition. And so it's conceivable that any belief be foundationally justified on Lyons' view. Thus I'm at a loss as to why Lyons presents his view as a theory of basic and non-basic beliefs conjoined with an additional theory of how each of these might enjoy justification. Why not just make a distinction between foundationally and inferentially justified beliefs and add the relevant psychological criteria into one's account of foundational justification?

Lyons is trying to make there can be unjustified basic beliefs. And he proposes that process reliability is necessary and sufficient for (prima facie) justification in the case of basic but not of non-basic beliefs. He intends to thereby avoid the previous counterexamples by arguing that the relevant beliefs are non-basic and, therefore, his version of process reliabilism isn't committed to giving a positive epistemic evaluation of the subject's beliefs.

What is Lyons' all important basic/non-basic belief distinction? He holds that three conditions are constitutive of a belief's being basic but, for our purposes, the most important condition is that the belief must be the result of a psychological system that "has developed as the result of the interplay of learning and innate constraints."<sup>171</sup> Lyons puts the burden of avoiding the counterexamples on this condition and so, for the ease of presentation, I ignore the remaining conditions for basicity in Lyons' sense. Lyons holds that, in all of the supposed counterexamples our intuition that the subject lacks justification for his belief is the result of our imagining the belief-forming process having an unusual etiology (and thereby failing to meet the previously mentioned condition for being a basic belief).

Such a strategy appears to deal with the cases of the brain lesion patient (in its original and modified forms) and Truetemp. Neither of these beliefs are the output of a process that has developed as the result of the interplay between learning and innate constraints. Truetemp's beliefs, for instance, are due to a process that takes place in a tempucomp whose development is due to the design of a scientist and whose presence is due to its being implanted by this same scientist.<sup>172</sup>

---

<sup>171</sup> Lyons (2009), p. 136. Lyons intends for this condition to be read very broadly so that a system which develops *solely* as a result of innate constraints meets this condition. Lyons' remaining conditions for basicity are that that none of the inputs the subsystems of the process are under the voluntary control of the larger organism and that the system is "inferentially opaque."

<sup>172</sup> It isn't clear to me that Lyons' view actually deals with the modified brain lesion patient case where the brain lesion happens to reliably produce true beliefs. I do tend to think of the belief forming process produced by the brain lesion as having an unusual etiology. But what if we were to make explicit that the development of the brain lesion and the way it produces beliefs was solely the result of genetic factors? Surely we ought to construe a subject's genetics as innate constraints. So isn't the belief forming process associated with the brain lesion one that developed as the result of

In the end the adequacy of Lyons' proposal all hinges on what he can say about the case of Norman. At least initially, however, it appears that in this case the view flounders. Norman's case is the most important and clearest illustration of what is wrong with radical externalism because it's stripped down to the bare essentials, forcing the reader to consider a case where everything appears fine with the subject's belief *except* for the fact that the subject lacks awareness of any supporting evidence for her belief. Bonjour's case of Norman makes no mention of the clairvoyance leading to many false beliefs as does the original brain lesion case; Bonjour makes it clear that Norman's clairvoyance is "completely reliable." Moreover, important for current purpose, Bonjour's presentation of Norman makes no mention of an unusual etiology of the belief source as does the case of Truetemp. It therefore appears that Lyons' proposal is incapable of responding to the clairvoyance challenge. And it's instructive to see why Lyons' protests to the contrary fall flat.

Why does Lyons suggest that his proposal can deal with the case of Norman? He appeals to the fact that the case is under-described—as all thought experiments inevitably are—and argues that Bonjour's description *implicitly* suggests to us that Norman's clairvoyance has an unusual etiology. Lyons maintains that an unusual etiology is the best explanation of why Norman lacks any evidence relevant to the general possibility of clairvoyance or that he possesses it. When we read Bonjour's description we therefore imagine Norman's clairvoyance as being "the result of a recent encounter with radioactive waste, a neurological prank, or the like."<sup>173</sup> And this is just to imagine that Norman's belief fails to meet the conditions for being a basic belief. Thus, on Lyons' view, the reliability of Norman's clairvoyance doesn't thereby confer justification on the resulting belief. According to Lyons, our imagination filling in this detail of an unusual etiology, as opposed to

---

innate constraints? In which case the belief would be a basic belief and the reliability of this process would be sufficient for the beliefs *prima facie* justification. Lyons' view would still be susceptible to the counterexample. I'm unsure what Lyons might say in response to such a case.

<sup>173</sup> Lyons (2009), p. 118.

Norman's lack of awareness of supporting evidence, explains our intuition that Norman's belief that the president is in NYC is irrational.

#### *5.2.2.1 Norman, Nyrmoon, and Normina*

Lyons is well aware that the advocate of an awareness requirement on justification will quickly modify the description of Norman so as to make explicit that Norman's clairvoyance lacks any unusual etiology and thereby produces basic beliefs (in Lyons' sense). Here, however, Lyons suggests that such a modification alters our intuitions regarding the belief's epistemic status. He supports this contention by offering the following example of such a modified case:

NYRMOON THE CLAIRVOYANT: Nyrmoon is a member of an alien species for whom clairvoyance is a normal cognitive capacity, which develops in much the same way as vision does in humans. Members of Nyrmoon's species have specialized internal organs that are receptive to the highly attenuated energy signals from distant events; as an infant, all was a "blooming buzzing confusion" for Nyrmoon, until, like everyone else, he learned to attend selectively, recognize various objects, and filter out coherent distant events. Nyrmoon, however, is so extremely unreflective that he has no beliefs (a fortiori, no justified beliefs) about the reliability of his clairvoyance. One day he forms, as a result of clairvoyance, the belief that his house is on fire (which it is).<sup>174</sup>

Nyrmoon's belief is produced by a reliable process that produces beliefs that meet the psychological criteria for basicity. Lyons version of reliabilism is thereby committed to giving a positive epistemic evaluation of Nyrmoon's belief. He claims, however, that this more detailed description of a clairvoyant belief elicits an intuition that this positive epistemic evaluation of Nyrmoon's belief is appropriate. And I too admit that, when I read the description of Nyrmoon, I'm inclined to think his belief is justified.<sup>175</sup>

---

<sup>174</sup> Lyons (2009), p. 119.

<sup>175</sup> As with all thought experiments the details here are underspecified. If you don't share the intuition that Nyrmoon's belief is justified then perhaps your imagination fills in the details differently than my own. If you lack the relevant intuition, this is all the better for my contention that awareness of supporting evidence is necessary for justification and you might just skip ahead to section 5.2.3. However, I want to provide an explanation of how Lyons' description suggests filling

Wait! I've admitted that the description of Nyrmoon elicits an intuitive positive epistemic evaluation of his belief. Isn't this tantamount to admitting that awareness of supporting evidence isn't necessary for justification? I'll argue that the description only elicits an intuitive positive evaluation because the metaphorical language actually causes us to imagine that Nyrmoon is aware of supporting evidence. It's only because this metaphorical language suggests that such awareness is present that the description of Nyrmoon elicits an intuitive positive evaluation of his belief's rationality.

The language used in the description of Nyrmoon's case implicitly suggests that Nyrmoon's clairvoyance utilizes some sort of experiential input and we intuitively think of experiential states as being evidence for various claims (just think of the common phrase "the evidence of the senses"). Thus, despite no explicit mention of any awareness of supporting evidence the description of the case suggests as much to our imagination. How exactly does the language suggest that Nyrmoon's clairvoyance utilizes experiential inputs?

First, consider Lyons' comparison of the clairvoyance faculty with our faculty of vision. Our thought about perception is intimately tied up with our thought about perceptual experiences and such a connection is stronger in the case of vision than in our thought about any other modality. And so the comparison to vision thereby suggests that Nyrmoon's clairvoyance utilizes experiential inputs.

Second, consider two key passages in the description of Nyrmoon: (a) "as an infant, all was a "blooming buzzing confusion" for Nyrmoon"; and (b) "he *learned to attend selectively, recognize* various objects, and filter out coherent distant events" (emphasis added). These

---

in the details so that we imagine Nyrmoon as being aware of supporting evidence and that it is this that thereby elicits an intuitive positive evaluation. This will show that Lyons' case fails to challenge my proposed awareness requirement on justification. In fact my discussion of *why* the case elicits a positive evaluation actually provides positive support for my awareness requirement. Finally, our discussion will reveal different kinds of supporting evidence one might have for the outputs of a belief source that will be important for our discussion at the end of chapter 8 regarding the problem of the problem of easy justification in relation to perceptual evidence.

passages also give the impression that Nyrmoon's clairvoyance works by producing experiential states. How could a description of the world being a "blooming buzzing confusion for Nyrmoon" be anything other than a description of his experiential states? Moreover, the natural interpretation of a learning process delivering Nyrmoon from a blooming buzzing confusion is that this learning organizes crazy unintelligible experiential states into intelligible experiential states with more determinate content.<sup>176</sup>

So even though the description of Nyrmoon makes no explicit mention of him being aware of an experiential state that we would intuitively construe as evidence, the language used in the description implicitly suggests such an interpretation. This implicit suggestion of an experiential state of which Nyrmoon is aware and our intuition that experiences usually constitute evidence explains the intuition that Nyrmoon has justification for his belief.<sup>177</sup>

However, even if we explicitly imagine that Nyrmoon's clairvoyance faculty works independently of any experiential inputs, the description suggests that Nyrmoon would be aware of further supporting evidence that would make his perspective on his belief very

---

<sup>176</sup> Lyons will no doubt give operational definitions of "blooming buzzing confusion", "learning", and "selective attention", etc. But, as is so often the case when cognitive science gives operational definitions for terms such as "pain", "belief", etc. that are commonly associated with phenomenal states with which we are intimately acquainted, it is hard to read these terms without thinking about the associated phenomenal states which are no part of the operational definition. And so my explanation of the intuition that Nyrmoon has justification in terms of the language bringing us to imagine Nyrmoon being aware of a distinctive clairvoyant experience that might act as evidence for his belief still stands.

<sup>177</sup> Clairvoyance is a faculty foreign to humans and we usually associate distinct kinds of faculties with distinct (perhaps incommensurable) kinds of experiences (e.g. sight with visual experiences, hearing with auditory experiences, touch with tactile experience, etc.). We might have no idea what the kind of experience associated with Nyrmoon's clairvoyance would be like; we might be in the same position regarding Nyrmoon's clairvoyant experiences as we're in regarding what it's like to be a bat (see Nagel 1974). Lyons' description of nevertheless suggests that Nyrmoon's clairvoyance produces *some* experience, and it's natural to think of the experiences associated with specific faculties as constituting evidence for their output beliefs. Here I'm relying on the intuition perceptual experience can serve as evidence in isolation to explain the intuition that Nyrmoon has justification and this support remains even if we eventually decide that we need to abandon our intuitive judgment and hold that experiences only constitute evidence when accompanied by a wide range of other experiential states (or even background beliefs) that cohere with it.

different than Norman's. Consider again the idea that Nyrmoon "*learned* to attend selectively, recognize various objects, and filter out coherent distant events." If Nyrmoon's clairvoyance develops via this kind of learning then he would intuitively have justification for his belief but only because such learning would bring with it an awareness of supporting evidence for the relevant belief. Considering the analogy with human vision will help see why.

How did I learn to attend selectively to various features via vision, to *visually* recognize various objects/properties, etc.? An adequate psychological explanation of this learning process is going to be incredibly complicated but surely something like the following is *part* of the story. My learning to attend selectively and recognize redness was facilitated by various members in the community (mostly my parents) pointing to paradigmatic red objects in normal conditions and uttering the word "red", repeating this procedure, and providing positive and negative feedback if I uttered the correct color word when an object was presented to me. By the time that this learning process brings it about that I've acquired the requisite color concepts, that I've learned to selectively attend to objects' colors (rather than shapes), and that I've learned to recognize the color of an object, I will have amassed a large body of memorial and testimonial evidence that arriving at beliefs with this kind of content (i.e. beliefs about the colors of objects) in this visual way is reliable. Forming beliefs about the colors of objects in this visual way will *feel familiar* to me, I'll be aware of reasons for thinking that such beliefs are reliable, and all of this will make it likely from my own perspective that a belief with this content and formed in this way is true. The learning process involved in my training in identifying colors ensures that my perspective on color beliefs formed via vision is incredibly different than Norman's perspective on his clairvoyant belief.

The description of Nyrmoon suggests a similar story. The development of clairvoyance is compared to the development of vision and the same kinds of abilities are stipulated to be learned. Thus, we imagine that a similar kind of process takes place in the case of Nyrmoon's clairvoyance. Other members of Nyrmoon's species provide testimony



along with positive and negative reinforcement that helps to develop Nyrmoon's clairvoyance. Moreover, we imagine that he receives similar training regarding other perceptual modalities where these will eventually begin to converge on delivering beliefs with similar and cohering contents. Over time this learning process will provide him with coherent testimonial and memorial evidence about what kinds of spontaneous beliefs are the result of clairvoyance and which he ought to trust and which one he ought not trust. This thereby makes Nyrmoon aware of evidence supporting the reliability of beliefs that strike him in this clairvoyant way (including his belief that his house is on fire) and, therefore, supports the truth of these beliefs.

It's important to note that it's stipulated that Nyrmoon lacks any meta-beliefs about the reliability of his clairvoyance, which threatens the thesis that having a meta-belief about a source's reliability is necessary for justification. But we should be thankful that our intuitions accord with the possibility of justification absent any such meta-belief. Imposing this meta-belief requirement raises worries of over-intellectualizing justification. More importantly, imposing a meta-belief requirement of this sort is precisely what generated the skeptical horn of the puzzle from chapters 2 and 3. Nevertheless, an intuitive positive evaluation of his belief poses no threat to the claim that an awareness of supporting evidence is necessary for justification. What's important about the Nyrmoon case is that even if we imagine that he fails to form the meta-belief we nevertheless imagine that he is aware of evidence that supports the reliability of his clairvoyance. Imagining such awareness is what elicits the intuition that Nyrmoon has justification for his belief.

At this point it's clear that Lyons' description of Nyrmoon is much more than the Norman case with an added stipulation ruling out an unusual etiology of clairvoyance. When we read the Norman case we automatically imagine that his belief just pops into his head randomly as if it came from thin air and as if it was nothing more than a mere guess. We therefore hold that he lacks any justification for his belief. Alternatively, the metaphorical language in the description of Nyrmoon suggests that the Clairvoyance utilizes experiential

inputs. Moreover, when we imagine the learning process, we imagine it as providing Nyrmoon with an awareness of evidence for the existence of clairvoyance, his possession of such a faculty, and even its reliability. And so we imagine that his belief that his house is on fire would strike him incredibly different phenomenologically than the way the belief that the President is in NYC strikes Norman. Forming beliefs in this way would feel *familiar* to Nyrmoon due to the training involved in his learning how to use his clairvoyance appropriately. We imagine his clairvoyant beliefs striking him with the same familiarity that accompanies our memory beliefs. This familiarity would allow Nyrmoon to identify his belief as a clairvoyant rather than a tactile belief or even a random guess, and his training would provide him with evidence that such beliefs are reliable. Lyons' description of Nyrmoon suggests a case that is a far cry from the original case of Norman where we were supposed to imagine Norman as lacking "evidence or reasons for or against the general possibility of such a cognitive power or for or against the thesis that he possesses it." The metaphorical language used in Lyons' description does all the work in priming us for a positive epistemic evaluation of Nyrmoon's belief.

*After* presenting his cases Lyons claims that we could further specify (without any modification of the description) that Norman and Nyrmoon are doxastically and *phenomenologically* identical and that this wouldn't change our intuition that Norman lacks justification but Nyrmoon does. This stipulation, however, drastically changes my intuitions about Nyrmoon's justification! Consider the following adaptation of Norman that makes explicit that his clairvoyance has the right kind of etiology but that avoids the metaphorical language used in Lyons' description of Nyrmoon:

**NORMAN\***: Norman, under certain conditions which usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. Norman was born with this faculty and its nature is due solely to innate features of Norman's noetic structure (features shaped by evolutionary forces). Norman's clairvoyance can deliver beliefs with highly determinate content without any training or learning. As it turns out, however, the kind of energy to which Norman's clairvoyance organs are sensitive is incredibly rare. No one with this organ has yet been exposed to the kind of energy that would activate the clairvoyant process. He possesses no evidence or reasons for or against the general possibility of such a cognitive power or for or against the thesis that he possesses it. Norman goes 25 years

without forming any clairvoyant beliefs. One day Norman comes into contact with the relevant kind of energy and, as a result of his reliable clairvoyance, the true belief that the President is in New York City pops into his head as if from nowhere.

Speaking autobiographically, this is how I imagined the case when I initially read Bonjour's presentation, and so whenever I speak of Norman the Clairvoyant from here on out I'll be referring to Norman\*. Perhaps I haven't read enough science fiction to think that the best explanation of Norman's clairvoyance would be his falling into a pit of radioactive waste. Nevertheless, when described this way, our intuition still delivers the verdict that Norman's belief is highly irrational.

One might begin to worry that ordinary unreflective humans lack any awareness of supporting evidence for their perceptually delivered beliefs. Aren't we in a very similar situation regarding our perceptual beliefs as Norman is regarding his clairvoyant beliefs? And so Lyons considers one last case that we ought to consider:

NORMINA: Normina is an otherwise normal human, with normal, reliable perceptual systems, but she is quite unconcerned with anything other than what's immediately in front of her and so extremely unreflective that she has no metabeliefs about the reliability of her perceptual faculties. One day, she forms the (visual) belief that there's a chair in front of her.<sup>178</sup>

Surely, given the caveat that Normina is an "otherwise normal human", it's intuitive to grant that Normina's perceptual belief is justified.

It's important to note again that this intuitive verdict fails to challenge the idea that awareness of supporting evidence is necessary for justification. At best it challenges the necessity of a meta-belief about reliability. But this raises the question of whether in ordinary cases of perception we're aware of any better evidence for our beliefs than Norman is for his belief that the President is in NYC. We need to find some *intuitive* difference in the evidence for their beliefs of which Norman and Normina are aware if we're going to account for our differing intuitions. This, however, isn't very difficult.

---

<sup>178</sup> Lyons (2008), pp. 129-130.

Intuitively we think that of normal cases of perception as cases where subjects are aware of two kinds of evidence for their beliefs. First, Normina's conscious awareness of her perceptual experience intuitively constitutes an awareness of evidence for believing that there is a chair in front of her. Admittedly, there is a lot going on in Normina's experience including visual, tactile, proprioceptive, etc. experiential states and sorting out which features are evidence and why they constitute evidence for believing that there is a chair in front of her is incredibly difficult. Nevertheless, what is important is only our intuition that such experiences constitute evidence that there is a chair in front of her since this is what explains our intuition that she has justification for her belief.<sup>179</sup>

Second, even if we were to add that Norman's clairvoyance included an experiential input that might be construed as evidence, ordinary humans such as Normina would have additional evidence for the reliability of their perception that provides them with additional justification for these beliefs. Normina, for instance, will have had the perceptual training (mentioned previously) that would provide her with testimonial and memorial evidence that beliefs formed in this perceptual way are reliable. In turn, this would support the truth of her perceptual belief that there is a chair in front of her. As Conee and Feldman explain:

Even if there are such visual experiences that could serve as a basis for a clairvoyant's beliefs, still there is a relevant epistemological difference between beliefs based on normal perceptual experience and the clairvoyant's belief in Bonjour's example. We have collateral evidence to the effect that when we have perceptual experience of certain kinds, external conditions of the corresponding kinds normally obtain...This includes evidence from the coherence of these beliefs with beliefs arising from other perceptual sources, and it also includes testimonial evidence. This latter point is easily overlooked. One reason that the belief that one sees a book fits even

---

<sup>179</sup> Lyons (2008) chapter 3 poses a kind of Sellarsian challenge to the idea that experiences can act as evidence for our beliefs since they don't have propositional content and so can't stand in the right kind of confirmation relations in order to act as evidence. I have already defended against such an attack in chapter 4 by arguing that facts can stand in evidential relations and, therefore, experiential facts can act as evidence just as well as *propositions* about our experiences. Of course, whether the evidential support provided by our experiential states is *sufficient* support for justifying our perceptual beliefs is a more difficult matter.

a child's evidence when she has a perceptual experience of seeing a book is that children are taught, when they have the normal sort of visual experiences, that they are seeing a physical object of the relevant kind. This testimony, typically from people whom the child has reason to trust, provides evidence for the child. And of course testimony from others during adult life also gives evidence for the veridicality of normal visual experience. On the other hand, as Bonjour describes his example, the clairvoyant has no confirmation at all of his clairvoyant beliefs.<sup>180</sup>

I want to stress that I'm relying on our *intuition* that the experiential state, the coherence amongst outputs of independent perceptual sources, and the memories as of testimony constitutive of the learning process constitute evidence for our perceptual beliefs of which we're ordinarily aware. The existence—not the accuracy—of this intuition accounts for the difference in our intuitive judgments of the cases. As such, these cases really do support the idea that our *concept* of justification requires that a subject be aware of supporting evidence.

Lyons responds to such a challenge by claiming that even if we imagine that Normina suffers from a memory deficit that prevents her from recalling any of this track-record evidence for the reliability of perception we're still intuitively judge her perceptual belief to be justified. I'll admit that my intuitions regarding how to epistemically evaluate a subject's perceptual beliefs with such a memory deficit get hazy. Nevertheless, even *if* one has the intuition that Normina would still have justification, this wouldn't threaten the support for an awareness requirement. We must remember that we think of ordinary cases of perception as being intimately tied to perceptual experiences and we intuitively judge these experiences to constitute evidence for various claims about the external world. Moreover, we imagine that within the specious present Normina is aware of various deliverances from different perceptual modalities converging in on the same object and properties. She has visual experiences as of a chair. These experiences cohere with her tactile experiences as of a chair. And these experiences change ever so slightly as her proprioceptive experiences change (e.g. changes due to tilts of the head, stepping closer to the chair, moving

---

<sup>180</sup> Conee and Feldman (2004), p. 98.

ones hand across the back of the chair, etc.) in just the way one would expect if there were a three-dimensional object of a chair-like size and shape. Even assuming such a memory deficit, intuitively these all constitute supporting evidence of which Normina is aware for her beliefs that there is a chair in front of her.

A case more analogous to Norman would be beliefs produced via blindsight. Certain studies have suggested that subjects with damage to their primary visual cortex lack visual experiences of items presented in certain parts of their visual field. Amazingly, however, when prompted in various ways they can identify such things as the shape and orientation of objects in these “blind” areas of the visual field with an amazing accuracy.<sup>181</sup> So it has been suggested that the damage to the primary visual cortex in these patients prevents the visual system from producing a conscious visual experience but that visual information is still being processed and is available in the production of beliefs about various stimuli in a way that has a surprising degree of accuracy. When prompted for certain information, however, such patients find themselves inclined to give a specific answer but it initially feels like nothing more than a guess.

Imagine a hypothetical case where blindsight affects the entire visual field. This hypothetical subject has no visual experiences whatsoever but visual information obtained via light hitting the retinas is still processed by the visual system in a way that can produce beliefs about the sizes, shapes, and orientation of various physical objects. Imagine a case where a subject S with blindsight is like Norman and lacks any awareness of evidence for or against the general possibility of blindsight or for or against the claim that she has this ability. Also imagine that no training was required for the development of S’s blindsight and that it’s generally reliable. Consider the very first belief delivered via blindsight. Perhaps, it’s the belief that there is a red rectangular object off to S’s left. S has no visual experiences and has

---

<sup>181</sup> See Weiskrantz (1990) and (1997) for detailed discussions of the phenomenon of blindsight.

no previous familiarity with forming beliefs in this way, she has no idea where this belief came from, and her belief seems to her to pop out of thin air. The case of Normina isn't analogous to Norman's case but this hypothetical blindsight case is. Intuitively, however, the blindsighter's belief is irrational. She might eventually gain collateral evidence in the form of an independently confirmed track-record that would justify her in believing the deliverances of blindsight but surely the beliefs prior to any such collateral evidence are irrational.

The support for an awareness requirement on justification provided by comparing how we think about ordinary cases of perception, blindsight, and the clairvoyance thought experiment derives from the fact that our intuitions about justification line up with our intuitions about which objects of the subjects' awareness constitute evidence for their beliefs. It's important to note that this support remains intact even if philosophical reflection on skeptical hypotheses were to force us to retract our intuitive judgments about what kinds of things constitute evidence for ordinary perceptual beliefs.<sup>182</sup>

### 5.2.3 The Subject's Perspective Objection (SPO)

The dialectic between me and the radical externalist can continue indefinitely. Radical externalists can continue to add conditions for foundational *prima facie* justification (short of accepting the awareness requirement) in an attempt to fend off the counterexamples. And I can continue to modify the cases in order to illustrate that these additional conditions continue to be conceptually insufficient for justification. Instead of continuing down this path of endless rejoinders, however, let's turn to a consideration of how our cases can be

---

<sup>182</sup> Later I argue against the idea that a single visual experience as of *P* is evidence for *P*. Our most important evidence for ordinary perceptual beliefs is the fact that sets of different perceptual experiences home in on the same object and properties. I usually seem to be able to see what I seem to be able to touch, I usually seem to be able to touch what I seem to be able to see, and the way my visual and tactile experiences change in relation to proprioceptive changes is usually what one would expect given that our ordinary perceptual beliefs are true. Though, as I've stressed, such a point is tangential to my goals here. We'll come back to a discussion of what our *actual* as opposed to *intuitive* evidence is for our ordinary perceptual beliefs towards the end of chapter 8.

used to provide positive support (as opposed to the purely negative support provided by the counterexample method) for an awareness requirement on justification.

When we consider Norman's belief that the President is in NYC we're justified in following his lead and believing that the President is in NYC since *we're* aware of the fact that his belief was produced by a reliable process. That Norman's belief was produced by a reliable process confers justification on our belief since we can "see" the likelihood of this belief's truth. Herein lays the problem. The likely truth of the propositional content of Norman's belief can be seen from a third-person or God's-eye perspective but not from Norman's own perspective.

*We* can (in some sense) "see" what Norman's belief has going for it. From *our* perspective, the proposition that the President is in NYC is likely to be true. Norman's own perspective on the proposition, on the other hand, is entirely different. Norman lacks any awareness of anything that would favor the claim that the President is in NYC rather than that the President is in Chicago, Mulan, or Paris. When we think about what the situation would be like from Norman's perspective, when we so to speak "put ourselves in Norman's shoes"), we realize that it would be as if the belief popped into his head from thin air. There is nothing in Norman's own perspective on the proposition that distinguishes the belief from a mere guess or arbitrary conviction.

It's Norman's 1<sup>st</sup>-person perspective on the relevant proposition's truth that explains why we find Norman's belief to be epistemically defective. From his perspective the belief's truth is nothing more than a cosmic accident and he sees no epistemic reason to prefer this belief to any other belief about the President's location. The only way the truth of his belief wouldn't be a cosmic accident from his own perspective is if he were aware of something counting in favor of the proposition's truth thereby distinguishing it (*to him!*) from a mere guess. Any view that doesn't require awareness of supporting evidence therefore allows a subject's belief to be justified even when that belief appears to the subject as no better than a



guess or hunch. Michael Bergmann—who ultimately rejects its legitimacy—nicely captures this motivation for an awareness requirement on justification:

THE SUBJECT’S PERSPECTIVE OBJECTION (SPO): If the subject isn’t aware of what the belief has going for it, then she isn’t aware of how its status is any different from a stray hunch or arbitrary conviction. From that we may conclude that from her perspective it is an accident that her belief is true. And that implies that it isn’t justified.<sup>183</sup>

When we begin the project of reflecting on our beliefs in an attempt to determine which are true and which are false we’re searching for some sort of assurance. It isn’t enough that some objective factors outside of our ken connect our beliefs to the truth. We want to get at the truth in a way that this isn’t mere happenstance from our own perspective. To borrow an analogy from Sextus Empiricus, if we get at the truth but in a way that is a mere coincidence from our own perspective then our situation is akin to someone who strikes gold while searching in the dark:

Let us imagine that some people are looking for gold in a dark room full of treasures... [N]one of them will be persuaded that he has hit upon gold even if he *has* in fact hit upon it. In the same way, the crowd of philosophers has come into the world, as in a vast house, in search of truth. But it is reasonable that the man who grasps the truth should doubt whether he has been successful.<sup>184</sup>

Here we have a unified diagnosis of what is epistemically defective in our considered cases. The brain lesion patient lacks any awareness of anything that indicates *to him* that he has a brain lesion. From his perspective the belief that he has a brain lesion is just as likely as its negation. Truetemp again lacks any awareness of anything that indicates to him what the precise temperature is. From his perspective the beliefs that it is 95°, 99°, 102°, etc. are equally likely to be true; nothing from his own perspective suggests that he ought to prefer any of these beliefs over the others. Similar remarks apply to Norman and the blindsighter.

---

<sup>183</sup> Bergmann (2006), p. 12.

<sup>184</sup> Sextus Empiricus (1935), pp. 480-481.

They are, so to speak, in the dark regarding the likelihood of the relevant beliefs' truth. Nothing indicates *to them* that the truth of their beliefs would be any better than grasping gold in the dark (though surely getting true beliefs and grasping gold *in any way* are desirable... the point is that when we begin to reflect on our beliefs we desire more so to get at truth in a way that isn't merely a cosmic accident from our own perspective we want some sort of assurance).

### 5.3 Justification and Awareness: Part II

The SPO provides the main motivation for holding that awareness of supporting evidence is necessary for justification. However, I now want to suggest that the SPO motivates more than this *generic* awareness requirement. The SPO also motivates a rejection of less radical forms of externalism and even weak forms of internalism. In order to explain why, I need first to distinguish three levels of evidence possession.

#### 5.3.1 Three Forms of Evidence Possession

S's being aware of some fact F which has the property of being evidence for *P* is necessary for S's having supporting evidence for *P*. Moreover, it seems that any kind of conscious awareness of F is sufficient for S's having supporting evidence for *P* *in some sense*. But we need to distinguish this minimal form of evidence possession from more demanding notions. The possibility for the distinctions I have in mind is due to the fact that one might be aware of some entity X without being aware of or appreciating all of X's properties and relations. Of specific interest here are distinctions that arise when we consider how S's awareness of evidence E might also involve a relation between S and E's property of evidentially supporting *P*. And there are three kinds of awareness or evidence possession that we should distinguish in this regard.

Our first two cases fall under what I'll refer to as "unconnected" awareness and evidence possession. In these kinds of cases a subject S is aware of a fact that *is* supporting evidence for *P* but S fails to appreciate, or show some kind of relevant "sensitivity" to (I'm

not using this in the technical epistemological sense), this relational property.<sup>185</sup> And cases of unconnected evidence can be further broken down into cases where the subject either lacks or possesses the ability to appreciate the evidential connection without further training or development. We can refer to the subject *S*'s evidence in the former case as her unconnected evidence simpliciter and as her potentially connected evidence in the latter case. Finally, there are cases where a subject is aware of a fact *F* that evidentially supports *P* such that this awareness grounds an *actual* appreciation of the evidential relation. I will say that such a subject has “actually” connected awareness and possession of supporting evidence *E* for *P*.

Consider the following illustrations of unconnected evidence simpliciter. Imagine a child with limited mathematical abilities who knows a mathematical proposition *P* which entails *Q*. Given the child's current cognitive development, however, the entailment is too difficult for the child to grasp. Such a child is aware of *P*'s truth but neither appreciates nor has the ability to currently appreciate the evidential connection between *P*'s truth and *Q*.<sup>186</sup> Or, assuming you know the axioms of set theory, you're aware of the truth of the axioms of set theory and these entail any number of incredibly complicated mathematical truths. But it isn't difficult to convince yourself that some of these entailments, absent further training or some expert testimony that such a relation obtains, are beyond your current mathematical ability to “see”. In these cases one has unconnected evidence simpliciter supporting a proposition *P* (i.e. you're aware of supporting evidence but lack even an ability to appreciate that support).

---

<sup>185</sup> In these kinds of cases I might refer to *E* as part of *S*'s unconnected evidence but the reader should keep in mind that “unconnected” here modifies *S*'s relation to the evidence not the nature of the evidence itself.

<sup>186</sup> Of course, the child might have the ability to appreciate this connection in some looser sense of ability since *after* much mathematical training and cognitive development she could come to appreciate this entailment.

Cases of potentially connected and actually connected evidence can both be helpfully illustrated by considering cases where a subject has a kind of “eureka!” moment that we’re all familiar with. For instance, there is often a point in detective novels where the detective has gathered all the information they need to solve the case but they’ve failed to actually put those pieces together. After mulling over the evidence for a while the detective has that “eureka!” moment where it all comes together and she can finally just “see” the connection between some set of facts *F* and the claim that so-and-so committed such-and-such crime. Prior to that moment *F* was part of the detective’s potentially connected evidence and then at the “eureka!” moment *F* became part of the detective’s actually connected evidence.

Of course, there are more mundane intuitive examples of actually connected evidence. My awareness of my current experiences seems to ground an appreciation of an evidential connection between the existence of these experiential states and the claim that there is currently a computer in front of me. However, it’s controversial whether the appreciation of an apparent evidential connection in this case is a grasp of a *real* evidential connection. A less controversial example is the following: in cases where you’re aware of the fact that you’re in pain you also appreciate the epistemic support this provides for the claim that someone is in pain. Here you have actually connected evidence for the proposition that someone is in pain.

At this point I’m merely stipulating different ways in which a subject might possess supporting evidence for her belief. In the previous section I used the SPO to argue for the generic requirement that *some* awareness or possession of supporting evidence is necessary for justification. But disambiguating these different forms of evidence possession raises the question of whether the SPO motivates adding a generic awareness requirement or whether it motivates a stronger requirement.

Consider three possible versions of the requirement that a subject is aware of or possesses supporting evidence for her belief: a generic, moderate, and strong evidence possession requirement:

**GEP** A subject *S* has justification for believing *P* only if *S* has unconnected, potentially connected, or actually connected evidence *E* that supports *P*.

**MEP** A subject *S* has justification for believing *P* only if *S* has potentially connected or actually connected evidence *E* that supports *P*.

**SEP** A subject *S* has justification for believing *P* only if *S* has actually connected evidence *E* that supports *P*.

The strong evidence possession requirement (SEP) entails the moderate requirement (MEP) but not vice versa. And the moderate requirement, in turn, entails the generic requirement (GEP) but not vice versa. In 5.3.2 I argue that the problematic kind of accidentality utilized by the SPO can only be avoided if we adopt SEP. Thus, our discussion of cases such as Truetemp and Norman along with the SPO motivate accepting the stronger requirement.

### 5.3.2 Motivating a Strong Evidence Possession Requirement

One might initially respond to the SPO by adopting the generic evidence possession requirement (GEP) while rejecting the stronger MEP and SEP. On such a view, a subject's having either unconnected, potentially connected, or actually connected awareness or possession of supporting evidence is necessary for *prima facie* justification but any of the three forms of evidence possession is sufficient.

A theory that accepts GEP but rejects MEP and SEP can come in both externalist and internalist forms depending on how they understand both the nature of the subject's awareness of evidence and the nature of epistemic support.

Let's begin by considering examples of externalist theories that accept GEP but reject MEP and SEP. Let's assume that the awareness of supporting evidence is to be understood as conscious awareness of, or acquaintance with, the relevant fact. If such a view holds that the support relation is an external relation where the nature and existence of the relata are insufficient for the existence of the relation then the view is what I'll label a "moderate" externalism. The natural proposal for understanding epistemic support in an externalist fashion is to understand it as some sort of reliability connection. But the view that accepts GEP is more demanding than the radical versions of externalism discussed earlier

since justification requires that the subject be aware of something that supports her belief where this support is understood in terms of a reliability connection.

In a sense, this externalist view requires that there be evidence that shows up for the subject. Such a view places more demands on the 1<sup>st</sup>-person perspective (in line with our discussion from chapter 1) and thereby appears to be a less radical form of externalism. Nevertheless, there are factors outside of the 1<sup>st</sup>-person perspective that help determine facts about justification. Given that the evidential support relation is construed as an external relation it's possible that two subjects *S* and *S\**, in worlds *W1* and *W2* respectively, might be aware of the same facts *F* but *S* has justification for believing *P* even though *S\** lacks such justification; an evidential support relation might hold between *F* and *P* in *W1* but fail to hold in *W2*. Examples of views similar to this kind of moderate externalism include Alston's (1988) indicator reliabilism and Comesaña's (2010) "evidentialist reliabilism."

Why might someone be attracted to the idea that the externalist version of adopting GEP is sufficient for avoiding the truth of a subject's belief being a mere accident from *S*'s own perspective? Consider again the contrast between Normina's justified perceptual belief that there is a chair in front of her and Norman's unjustified belief that the President is in NYC. When we read Normina's case we imagine that her belief wouldn't be a cosmic accident from her perspective. What makes the difference? One salient difference in the way we imagine Normina's and Norman's perspectives is that we imagine that Normina is aware of an experiential state that can serve as a ground of her belief whereas Norman lacks awareness of any such potential ground. We imagine that Normina has, for instance, the kind of visual experience that we would all describe as an experience as of a chair. And we intuitively take it that a visual experience as of a chair is reliably associated with the existence of such a chair. Thus, we might conclude that what removes accidentality from Normina's perceptive is her awareness of something that *is* evidence (in this externalist sense) for the relevant proposition.

Such an explanation, however, gets all of its intuitive power by appealing to our familiarity with perceptual experiences. In normal cases of perception such as Normina we intuitively take ourselves to have some appreciation of an evidential connection between our experiential states and our perceptual beliefs. In fact, we're so familiar with these cases it's hard to imagine what it would be like if we didn't have an apparent appreciation of this connection. And in this regard it's helpful to consider a modification of the Norman case that introduces an experiential state where a stipulated connection between that state and some external state-of-affairs is more foreign to us.

Imagine that we add to the case of Norman that his clairvoyance works in a way that utilizes a variety of experiential states. When the President is in NYC and Norman is asked where the President is, his clairvoyance organ happens to be around the rare kind of energy it needs to operate and thereby produces a quasi-visual experience with the following characteristics: (a) there are three side-by-side octadecagons; (b) the farthest left octadecagon in his visual field is a pale green, the center octadecagon is a bright yellow, and the farthest right octadecagon is a deep purple; (c) these features are super-imposed over his perceptual experience in such a way that they don't appear to be representations of external objects in his surroundings (in the same way that an afterimage or phosphenes are experienced). Assume that the causal laws are set up in such a way that clairvoyance would reliably tend to produce this kind of experience only when the president is in NYC—and similar remarks about the reliability of the other clairvoyant experiences apply as well. But let's further imagine that Norman in no way appreciates this reliable connection between experiences of this sort and claims about the President's location. Does Norman's awareness of this experience thereby make the truth of his belief non-accidental from his own perspective? It's possible that his experiential state might distinguish the *existence* of his belief from cases of mere hunches or guesses. Guesses and hunches aren't usually associated with these odd phenomenal experiences. He may even take this experience to be the cause of his belief and thereby take the *existence* of the belief to be non-accidental. Nevertheless, if he fails to

appreciate the connection between this experience and the *truth* of the claim that the President is in NYC, his awareness of this experience does nothing to make the truth of his belief any less accidental. Norman lacks any appreciation of how this experience bears on the claim that the President is in NYC as opposed the claims that he is in Mulan, Chicago, or Paris. Norman's awareness of what might very well *be* evidence fails to distinguish, from his own perspective, the *status* of his belief's truth from a mere hunch or arbitrary conviction

One might worry that my example relies on an experiential state that we don't intuitively think would be reliably correlated with the truth of the President's being in NYC. You might suggest that the intuition changes if we imagine that the clairvoyance produces a kind of faint image of the President in NYC similar to a visual experience that is reliably associated with the truth of this proposition. But avoiding this familiar case is what we need to do if we are to properly assess the view's plausibility. When we imagine what it would be like to have experiential states that are imagined to be very similar to our own visual experiences we easily fall into the trap of bringing into our imagination what we take to be an awareness of a connection between the experience and some external world proposition. Moving to an unfamiliar experience that is merely stipulated to have the relevant reliability connection helps us imagine what it would truly be like from the subject's own perspective if he failed to have any appreciation or sensitivity to the connection. When he moves from an awareness of the experiential state to the belief that the President is in NYC it would be a mere cosmic coincidence from his own perspective that such a belief is true.

Adding this experience to Norman's perspective without an appreciation of the relevant connection fails to remove the relevant kind of accidentality. Norman's perspective on the truth of his belief is no better than the person who knows the axioms of set theory but fails to "see" any relevance to the truth or falsity of Goldbach's conjecture. And this latter case works just as well to illustrate that the moderate externalist's acceptance of GEP is insufficient for forestalling the SPO—and we'll see helps to show that the weak internalist's acceptance of GEP is also insufficient.



Now let's move to a consideration of internalist views that accept GEP but reject the stronger MEP and SEP. Let's assume that epistemic support isn't to be construed as an external relation such as reliability but rather an internal relation such that the nature and existence of a set of facts and a proposition necessitates the existence of a determinate evidential relation holding between these two relata. And now we add to this an adoption of GEP such that justification for believing *P* requires awareness of evidence that stands in this necessary evidential support relation to *P*. Such a view would appear to be a version of internalism since facts within the 1<sup>st</sup>-person perspective fully determine facts about justification. However, I'll refer to this view as "weak" internalism since it holds that justification can be conferred by the subject's unconnected evidence. A more demanding version of internalism holds that justification is fully determined solely by connected evidence, i.e. an internalism that accepts SEP.

There are various possible weak internalist views that resemble (though might be a bit different from) prominent views in contemporary epistemology. For instance, one might adopt a version of phenomenal conservatism that qualifies as a version of weak internalism.

Phenomenal conservatism (PC) is characterized by its acceptance of the principle that, necessarily, if it seems to *S* that *P* then *S* has some degree of prima facie justification for believing *P*. Most philosophers defending PC have suggested that these "seeming states" are propositional attitudes that represent their content as actualized (unlike desires, fears, or imaginings) but that are nonetheless distinct from beliefs.<sup>187</sup> The distinction between seemings and beliefs is often made by appealing to known perceptual illusions. When I look at the Muller-Lyer lines it seems to me that the lines are unequal in length even after I've become privy to the illusion and no longer actually believe that the lines are unequal.<sup>188</sup>

---

<sup>187</sup> See especially Cullison (2010) and Huemer (2001) and (2007).

<sup>188</sup> I'm actually doubtful that one needs to posit a sui generis propositional attitude in order to make sense of the way in which the lines still seem unequal in length. My own view is that this use of seems merely indicates a felt disposition to believe the relevant proposition. For the sake of

It's plausible to construe PC as holding that the fact that it seems to S that *P* as evidence supporting *P* for S. And such an evidential relation according to PC is a necessary relation and is independent of the existence of any reliability connection between the fact that it seems to you that *P* and *P*'s actually being true. Finally, you might amend PC so that only seemings of which one is aware can provide justification—it's tempting to think that seemings (like pains) are by their very nature states of which we're consciously aware and so this might be trivially satisfied by all versions of PC. If such a view then rejects any requirement that the subject appreciate the fact that it seems to one that *P* epistemically supports *P*'s truth, it will qualify as a version of weak internalism. Similar remarks might apply mutatis mutandis to Conee and Feldman's version of mentalism.

An advocate of weak internalism as I've conceived the view might claim that the criticisms of moderate externalism fail to show that awareness of something that *is* supporting evidence for *P* fails to make *P*'s truth non-accidental from a subject's perspective but merely show that the externalist has the wrong theory of epistemic support—we'll see in the next chapter that I do think the externalist has the wrong theory of epistemic support. The weak internalist will insist that our intuition that adding an awareness of the octadecagon experience to Norman's clairvoyance doesn't remove the relevant accidentality simply because such an experiential state clearly *isn't* evidence for believing that the President is in NYC. When we add to Norman's perspective awareness of something that might intuitively be construed as evidentially supporting his belief there is no need to add an appreciation of this evidential support. For instance, if we add that it seems to Norman that the President is in New York and that this seeming is accompanied by qualia distinctive of clairvoyance then perhaps our intuitions concerning whether the truth of Norman's belief is an accident from his own perspective (and therefore unjustified) become hazier.

---

argument I admit the existence of "seemings" as a propositional attitude distinct from beliefs and felt dispositions to believe.

Consider, however, the case of entailment. Surely if  $P$  entails  $Q$  then an awareness of  $P$ 's truth constitutes an awareness of supporting evidence for  $Q$  even on internalist-friendly construals of evidential support. However, there are surely cases where one might know that  $P$  is true but fail to have any degree of justification for believing an entailed proposition  $Q$ . The entailment might be so complicated that one simply cannot “see” or appreciate that connection. Again, the fact that I know the axioms of set theory doesn't thereby make the truth of any belief entailed by these axioms non-accidental from my own perspective. Even if my knowledge of the axioms of set theory cause me to believe Goldbach's conjecture (assuming that this turns out to be true) I simply cannot convince myself that my knowledge of these axioms thereby provides me with justification for this belief. My lack of appreciation of, or sensitivity to, the relevant entailment relation makes the truth of my belief no less surprising than the truth of a random guess or hunch. Since I lack an appreciation of the entailment holding between the axioms of set theory and a complicated theorem  $T$  I cannot actually take the axioms into account when forming a belief whether  $T$ . From my own perspective I'm still forming beliefs in a way that is analogous to grasping for gold in the dark. And so the SPO threatens even this weak internalist view that accepts GEP.<sup>189</sup>

Notice that our remarks also suggest that a requirement of a subject's having potentially connected evidence for  $P$  is insufficient for avoiding the SPO. If a subject has the ability to currently appreciate the connection between  $E$  and  $P$  but doesn't *actually* appreciate the connection then her awareness of  $E$  doesn't make her perspective on  $P$ 's truth any different than the person with unconnected evidence simpliciter. Potential appreciation of the connection merely means that the subject's awareness of  $E$  provides her with the *ability*

---

<sup>189</sup> One might think that this only applies at the level of inferential justification but I see no reason to think the same insight doesn't apply at the level of supposed foundational justification as well. Experiential facts of which I'm aware might entail a proposition  $Q$  where this entailment is too complex for me to appreciate. Surely here we still want to admit that  $Q$ 's truth is accidental from my perspective and I lack justification for believing it.

to improve her perspective on *P*'s truth. Until she manifests this ability, however, *P*'s truth is still an accident from her own perspective. Such a person is analogous to the subject searching for gold in the dark while holding a flashlight but who hasn't yet turned it on. Even if she stumbles upon gold she shouldn't be persuaded that she has hit upon it until she actually flips on the flashlight and thereby illuminates the gold in her hand.

We might admit that someone with potentially connected possession of supporting evidence has a kind of dispositional justification for believing *P*. Such justification is of a degenerate sort but in some sense it puts *S* in a better epistemic position than the subject who *lacks* the ability to connect her evidence to *P*. Though it's important that this dispositional justification derives its epistemic worth from the ideal case where one has *actually* connected supporting evidence to which one is disposed.<sup>190</sup>

Therefore, the commonly used counterexamples to externalism and the associated SPO motivate a strong requirement on evidence possession (SEP). Justification requires that a subject is aware of supporting evidence and that she appreciates this support. I mentioned earlier that Michael Bergmann, who gave our formulation of the SPO, ultimately rejects its legitimacy. He presents a dilemma even for those who accept my stronger awareness requirement on justification. The dilemma targets the nature of a subject's appreciation of the evidential connection. He argues that either this appreciation is a conceptual awareness of the fact that *E* is relevant to the truth of *P* or it's something like a non-conceptual awareness. If it's a conceptual awareness then we're led to a vicious regress that results in a radical global skepticism. If, however, our appreciation of the relevance of our evidence to *P*'s truth is merely non-conceptual then the view is still vulnerable to the SPO. At this point I merely want to mention the dilemma and mention that I think a non-conceptual awareness in the form of direct acquaintance with an evidential connection is sufficient to forestall the

---

<sup>190</sup> See Brett Coppenger's (2012) dissertation for an in-depth study of degenerate justification and its relation to ideal justification.

SPO. When we have direct acquaintance with such an evidential connection we have everything constitutive of a belief's likely truth directly before our mind and it doesn't seem we could ask for more in making a belief's truth non-accidental from a subject's own perspective. We will revisit Bergmann's dilemma in chapter 7 section 7.1.6.

### 5.3.3 An Outline of a Theory of Prima Facie Justification

Our discussion of the SPO thereby leads us to an acceptance of the strong evidence possession requirement (SEP): justification for believing *P* requires that S be aware of supporting evidence *E* for *P* and that S appreciate this evidential connection. When a subject meets this requirement she "sees" something the belief has going for it. Her perspective on this belief's truth is very different than her perspective on the truth of a mere guess or arbitrary hunch.

Of course, this requirement does not ensure that the subject doesn't see further reasons for rejecting the belief and so we will eventually need to consider the nature of epistemic defeat. However, I do want to suggest that when we combine our discussion in this chapter with the fact ontology of evidence we arrive at a general outline of a theory of *prima facie* justification that will be helpful in our investigation of the epistemic support relation in the next chapter.

**PF** A subject *S* has prima facie justification for believing *P* iff (i) *S* is aware of a set of facts *F*, (ii) *F* epistemically supports *P*, and (iii) *S* appreciates the evidential connection between *E* and *P*.

PF rules out various theories of justification. It rules out radical externalist theories of justification such as process reliabilism, safety theories, sensitivity theories, proper functionalism, etc. *as traditionally construed*. It also rules out more moderate versions of externalism such as Alston's (1988) indicator reliabilism and Comesaña's (2010) evidentialist reliabilism.<sup>191</sup> And, finally, it even rules out various theories that many would consider

---

<sup>191</sup> As well as Goldman's (2011) attempt to bring together a synthesis of (mentalist) evidentialism and reliabilism.

internalist theories such as Conee and Feldman's (2001) mentalism and the usual developments of phenomenal conservatism. Though views like the "seeming evidentialism" that Conee (2004) considers have a strong affinity with phenomenal conservatism but might very well be consistent with PF.<sup>192</sup>

However, I want to be clear that I don't take an acceptance of PF to commit one to internalism. I think versions of externalism can (and should!) accommodate the insights contained in PF and our consideration of the SPO. PF is a general outline of a theory of prima facie justification and whether a development of it qualifies as internalist or externalist depends on how one develops the critical notions of awareness, epistemic support, and the appreciation of epistemic support.

On Goldman's traditional version of process reliabilism a subject needn't have any kind of awareness of the input to the process nor any appreciation of that process' reliability in order for that process to confer justification on a belief. Such a view has been found wanting. This, however, doesn't mean that process reliabilism doesn't give an account of the epistemic support relation. The difficulty for combining PF with a process reliabilist account of epistemic support is in explaining how conditions (i) and (iii) can be satisfied, but I see no reason to think that this is impossible.

Giving an account of the required awareness of facts that support *P* (i.e. PF's first condition) is actually fairly easy. The facts that are inputs to the process must be psychological facts that occur at the personal as opposed to the sub-personal level. Giving an account of the appreciation of the evidential connection (i.e. the reliability of the process)

---

<sup>192</sup> Chris Tucker (2012) incorporates a requirement of appreciating the evidential connection in the case of inferential justification into a theory closely related to phenomenal conservatism. In the introduction to Tucker (*forthcoming*) he indicates in a footnote that he plans to incorporate a similar requirement into foundational justification and he mentions this specifically in relation to developing a version of phenomenal conservatism that can meet the SPO. Needless to say I'm incredibly interested to see a development of a view along those lines and I think it would lead to a much more plausible theory in the spirit of phenomenal conservatism.

appears to be the more difficult task. But my discussion of the SPO suggests that even an externalist should attempt to make sense of this notion. Doing so would help externalists deal with the supposed counterexamples of Mr. Truetemp and Norman.

It might be nice to at least gesture at some of the paths an externalist might go down in an attempt to explain the notion of a subject's appreciation of an evidential connection. One possibility is just to appeal to a reliably formed belief that the connection holds.<sup>193</sup> This view, however, would suffer from a worry of over-intellectualizing justification. Thus, one might turn to a second externalist friendly proposal: a subject *S* appreciates an evidential connection between *F* and *P* when *S* has a reliably formed seeming that the process producing *P* is reliable (perhaps this should be construed as having a kind of *de re* content). A third possibility is to just appeal to counterfactuals of the form: if *F* didn't support *P* (i.e. the process going from *F* to *P* wasn't reliable) then *S* wouldn't form the belief that *P* on the basis of *F*. For instance, Greco considers the following case:

[C]onsider the case of the Careless Math Student. Suppose that *S* is taking a math test and adopts a correct algorithm for solving a problem. But suppose that *S* has no understanding that the algorithm is the correct one to use for this problem. Rather, *S* chooses it on a whim, but could just as well have chosen one that is incorrect. By hypothesis, the algorithm is the right one, and so using it to solve the problem constitutes a reliable process. It seems wrong to say that *S* thereby knows [has justification for believing] the answer to the problem, however.<sup>194</sup>

The problem here is that the Careless Math Student doesn't seem to be sensitive to the evidential connection. It was just lucky happenstance that the process he used was reliable.

---

<sup>193</sup> This wouldn't amount to having a justified belief. And notice that this higher-order belief wouldn't be acting as a subject's evidence for the 1<sup>st</sup>-order belief that *P*, but rather it would be a condition for the subject's *having* the fact *F* that is input to the reliable belief-forming process as evidence for the 1<sup>st</sup>-order belief.

<sup>194</sup> Greco (2010), p. 149. I think Greco's case is a bit misleading, however, since we could construe the relevant belief-forming process as including the process used to pick the relevant algorithm and this process might very well be unreliable. Nevertheless, I think Greco is one externalist who has been especially receptive to the idea that one needs to be in some sense sensitive to evidential connections (though Greco would reject that one needs to always be aware of evidence).

Had the algorithm been unreliable the student still would have used it.<sup>195</sup> What would make it the case that a subject's having a belief-forming process that is reliable isn't mere luck? An externalist might appeal to evolution: the subject has this belief-forming process *because* such a process was *selected for* its reliability.<sup>196</sup> <sup>197</sup>

Here I'm only *gesturing* at ways the externalist might go in developing a theory of justification that is consistent with PF. I think an externalist has a lot more work to do in developing the notion of an appreciation of evidential connections in a way that is consistent with these connections being understood in terms of an appropriate reliability connection. At this point I'm not actually too worried with finding the best externalist account of this notion of appreciation because *however* it's developed (even if it helps avoid the

---

<sup>195</sup> This is a bit tricky since the correct algorithm will *necessarily* be reliable and so we have a necessarily false antecedent. How are we to assess subjunctive conditionals with necessarily false antecedents? Honestly, I don't have any very promising proposals. I think the intuitive point, however, still stands.

<sup>196</sup> I take it that this sort of move is precisely what makes proper functionalism (PF) seem plausible and to be an improvement over Goldman's traditional process reliabilism. The difference between the view gestured at here and traditional PF, is that PF holds that the fact that when F is input to a proper functioning process it outputs belief that P is constitutive of the epistemic support relation itself. This is what makes F and epistemic reason for believing P. On the view developed here it is being an input to a reliable process that determines epistemic reasons. The design plan is just what makes a subject sensitive to this process' reliability and is thereby one of the conditions for a subject's *having* this epistemic reason (i.e. the view I've sketched doesn't make the design plan part of what it is for a fact to *be* an epistemic reason for believing P).

<sup>197</sup> One problem with this view is that it doesn't appear to be *any* help in responding to Norman and Truetemp type cases since this kind of sensitivity to the reliability of a process can be present without making *any* difference to the subject's conscious life. This might suggest another externalist friendly view of sensitivity to evidential connections. I confine my description to a footnote because, while I think such a view might actually be the best version of externalism, it strikes me as the one that would be most difficult to develop. You might think that there are certain ways of forming beliefs that just *feel* right, where this doesn't amount to any kind of representation of these processes as being reliable. The idea here derives from the fact that the manifestation of *know how* actually has a phenomenal feel. Consider the golfer who knows how to hit a ball 300 yards. Certain swings just *feel* right. You can take a practice swing and it feel like a good swing. You can also take a practice swing and it feel like a bad swing. Similarly, certain forms of reasoning just *feel* right. It feels wrong to move from a blurry perceptual experience to a belief that a certain object has some determinate shape. It feels wrong to move from one observation of a white swan to the belief that all swans are white. In this way the appreciation of evidential support might amount to knowing how to reason that was produced by a reliable process and where this know how *shows up* in the 1<sup>st</sup>-person perspective.



counterexamples of Truetemp and Norman) it will be of no help in responding to the problem with externalist accounts of the support relation itself that I develop in the next chapter. I just want to point out that our discussion doesn't commit one to internalism, and I think the most plausible forms of externalism will be those that attempt to develop an externalist version of PF since these will be in the best position in accommodating the SPO and the associated thought experiments.

## CHAPTER SIX

## THE IRRUDICIBILITY OF EPISTEMIC SUPPORT

In chapter 4 I defended a characterization of the genus “evidence”. Evidence is a set of facts that is indicative of the truth or falsity of some proposition. Standing in a truth-indicating relation is *constitutive* of *a*'s having the property of being evidence for *P*. Chapter 4 also hinted that different species of truth-indicating relations can distinguish different kinds of evidence. When an agent (e.g. a person, community, or species) represents a fact *F* as being indicative of *P*'s truth we can refer to *F* as subjective evidence for *P*. When *F* is an indication of *P*'s truth in virtue of *P*'s being part of the best explanation of *F* we can refer to *F* as explanatory evidence for *P* (this is a bit misleading since the evidence is actually the explanandum). When *F-type* facts are reliably correlated with the truth of *P-type* propositions we can refer to *F* as scientific evidence for *P*.

I have no qualms with the view that a set of facts is *a kind* of evidence for *P* whenever it stands in some kind or other of truth-directed relation to *P*. Considering evidence in terms of one kind of truth-indicating relations might be useful in one context and considering evidence in terms of another kind of truth-indicating relation might be useful in another context. Our concern, however, is with epistemic justification. In chapter 5 I argued that justification requires possessing supporting evidence *E* for *P* in a way that includes an awareness of *E* and an appreciation of the evidential connection between *E* and *P*. But what kind of truth-indicating relation allows facts to play this justificatory role? There is a relation *R* such that a relational fact of the form  ${}_F R_P$  is a truth-maker for the following claim: if *S* were to have evidence *F* (in the sense outlined in chapter 5), *S* would possess supporting evidence that thereby provides her with *prima facie* justification for believing *P*.

My attempt to home in on the relation constitutive of what we can simply refer to as “justificatory” evidence appeals to the notion of a subject's *having* that evidence. Homing in on the relation of interest in this way is compatible with my insistence that a fact can have

the property of being evidence (including justificatory evidence) for  $P$  even if no one has that evidence.  $F$  is justificatory evidence for  $P$  iff  $F$  and  $P$  stand in a relation  $R$  that makes the previous subjunctive conditional true. In order to simplify the discussion that follows let's refer to the relation that makes this subjunctive conditional true as the epistemic support relation.<sup>198</sup> Standing in the evidential support relation is *sufficient* for  $F$ 's being justificatory evidence for  $P$ ; no one need actually bear the having relation to  $F$ .

Pointing to the epistemic support relation in this way shows that intuitions about justification can be used in assessing proposed accounts of said relation. Consider what we might label the “Jerry Command Theory of Epistemic Support” of JCT: a fact  $F$  epistemically supports  $P$  iff Jerry issues a command to believe that  $F$ -type facts are reliably correlated with the truth of  $P$ -type propositions.<sup>199</sup> Imagine that I'm aware of the fact that I desire that my lottery ticket is a winner and that I'm aware that Jerry has issued a command to believe that facts about my desiring  $P$  are reliably correlated with the truth of  $P$ . My awareness of these facts doesn't provide me with justification for believing that my lottery ticket is a winner and JCT therefore fails as an account of the relation constitutive of

---

<sup>198</sup> The “epistemic” modifier is meant to indicate that we're interested in the support relation that is constitutive of justificatory evidence. Clearly any kind of evidence for  $P$  will support  $P$ 's truth in some sense but this support relation isn't necessarily essentially tied to epistemic notions such as justification, rationality, knowledge, etc. Also notice that a fact standing in the epistemic support relation isn't to be confused with its being a J-factor. The epistemic support relation is a relation that allows a fact to act as a specific kind of J-factor: the subject's evidence for the proposition. Some J-factors help justify a belief that  $P$  by acting as the subject's evidence for  $P$  or by “so to speak” bearing on  $P$ 's truth. Other J-factors help justify a belief by playing a non-evidential role, i.e. its role in the subject's justification isn't due to its bearing on the truth of the relevant proposition. Clearly the fact  $F$  that I'm aware of some fact  $F^*$  which epistemically supports  $P$  is a condition that contributes to my having justification for believing  $P$ , but it is  $F^*$  not my awareness of it that bears on the truth of  $P$ ... my awareness of  $F^*$  is what we can think of as an enabling condition that allows  $F^*$ 's being on the truth of  $P$  to contribute to my justification.

<sup>199</sup> Even JCT is compatible with justificatory evidence for  $P$  existing even if no one—including Jerry himself—*has* that evidence. Imagine that Jerry issues a command to believe that  $F$ -type facts are reliably correlated with the truth of  $P$ -type propositions. Now imagine that the fact  $F^*$  exists and is an  $F$ -type fact. According to JCT this is sufficient for  $F^*$ 's epistemically supporting  $P$ . Nevertheless, it's possible that no one, including Jerry, is *any* way aware of  $F^*$ . In such a situation there exists justificatory evidence for  $P$  that no one has.

justificatory evidence, i.e. it fails as an account of the epistemic support relation.<sup>200</sup> Clearly arguing against more plausible accounts of epistemic support will be more difficult than JCT, but we'll see that the general argumentative strategy remains the same.

Many philosophers will find my view of epistemic support wanting. I hold that epistemic support is a *sui generis* and *unanalyzable* relation. A fact F's epistemically supporting P is a fundamental relation that is an irreducible part of reality. As we'll see, however, this doesn't mean we can't say anything informative about the relation. How do I plan to defend this conclusion? First, I will provide some examples of fact-proposition pairs that we intuitively take to either stand in or to fail to stand in an epistemic support relation. Second, I argue against proposed reductive accounts of this relation that ultimately attempt to identify justificatory evidence with subjective or scientific evidence. Finally, I explain the non-reductivist view and defend it from a variety of objections.

### 6.1 Examples of Epistemic Support

Most uncontroversial examples of epistemic support are cases where a fact F *conclusively* supports P. Consider the following pairs:

(E) <I'm in pain>

(P) Someone is in pain

(E) <Abe is in my office>

(P) Either Abe is in my office or Brian is in my office

(E) <George is a bachelor>

(P) George is an unmarried male

(E) <Sally's shirt is red>

(P) Sally's shirt is colored

---

<sup>200</sup> We can also argue against accounts of epistemic support relations by appealing to cases where a subject *has* justification for believing P but where a theory of epistemic support would entail that a subject fails to be in possession of anything that stands in the epistemic support relation. For instance, if I were aware of the fact that I was in pain and aware of this fact's entailing that someone is in pain surely this would provide me with justification for believing that someone is in pain even if Jerry hadn't issued a command to believe that facts about an individual's being in pain were reliably correlated with the existential proposition that someone is in pain.

In each of these examples we intuitively take the putative fact *E* to provide epistemic support for believing *P*. We take it that *E* and *P* stand in the relation *R* such that: if *S* is aware of *E* and appreciates *R* then *S* has some prima facie justification for believing *P*. Of course, this isn't to give an account of what kind of relation *R* is. My point here is merely that these pairs elicit in us an intuition that *E* epistemically supports *P*, and such an intuition gives us a starting point for theorizing about this relation.

Not all facts that we take to provide epistemic support are thought to provide *conclusive* support. The extent to which a fact is indicative of *P*'s truth is a matter of degree. Our concept of the epistemic support relation is of a relation that can come in various strengths. Such minimal information about our concept already shows that we cannot *identify* our concept of epistemic support with our concept of entailment—though our concept of entailment might be a concept of a particularly strong instance of epistemic support (more on this later).<sup>201</sup>

Uncontroversial examples of fact-proposition pairs where the relevant fact provides conclusive epistemic support are easy to come by. Sadly, any purported example of non-conclusive support is bound to be controversial.<sup>202</sup> But thankfully I only need to rely on certain examples eliciting an initial intuitive judgment that some fact provides non-conclusive epistemic support for a proposition. I want to understand what the content of

---

<sup>201</sup> Allowing non-conclusive epistemic support is also important if one hopes to forstall a far reaching (though not global) skepticism. If all epistemic support were conclusive support then only propositions that either cite our evidence, entailed by that evidence, or perhaps whose truth-makers are nomologically necessitated by the existence of our evidence could be justified. If epistemic support is always conclusive support then we could never have a justified false belief (since evidence consists of facts). Surely this is a result that we should, *ceteris paribus*, aim to avoid.

<sup>202</sup> Though, notice that what we need here are *particular* examples of non-conclusive support not generalized *types* of fact-proposition pairs such that the relevant type of fact always provides non-conclusive support for the relevant type of proposition. Such particular examples are often easier to identify than general rules. Consider an analogy. A student may be able to clearly see that the premises of a particular *instance* of modus ponens provide good support for the conclusion without seeing the validity of modus ponens when considered as a general rule. This is why it is often so helpful to beginning logic students to actually insert particular propositions for the *P*'s and *Q*'s.

this intuitive judgment is. My purpose here isn't to vindicate the accuracy of that judgment. And so I take it that the fact-proposition pairs that elicit such a judgment most easily are examples taken from what we often consider to be the most fundamental belief sources: perception, memory, and induction (here I mean enumerative induction).

Most of us take our experiences at face value. I have a perceptual experience as of a computer in front of me and thereby form the belief that there is a computer in front of me.<sup>203</sup> I have a memory experience as of having previously broken my leg skiing and thereby form the belief that I have previously broken my leg skiing. More importantly, we unreflectively take such processes to be justification conferring. We appear to implicitly accept that epistemic support relations exist between the following:

- (E) <It perceptually appears to me as if there is a computer in front of me>
- (P) There is a computer in front of me
  
- (E) <It memorially appears to me as if I previously broke my leg skiing>
- (P) I have previously broken my leg skiing

If our belief in *P* were challenged, we would unreflectively cite *E* as the fact of which we're aware that provide epistemic support for our belief.<sup>204</sup>

---

<sup>203</sup> Thorny issues concerning the contents of perception are lurking in the background here. In ordinary discourse we talk about having perceptions as of computers, Barack Obama, etc. Do perceptual experiences actually represent high-level properties like something being a computer or are the contents of perception limited to more basic properties like color and shape? I can't resolve the issues here. I talk as if the contents of perception include these high-level properties purely for the sake of readability, and nothing in my argument hinges on this. For more on the debate regarding the admissible contents of perception see Brogaard (2013), Hawley & MacPherson (2011), Lyons (2005), and Siegel (2011).

<sup>204</sup> Some readers might be hesitant to identify these single experiential states as supporting evidence for the corresponding proposition. There is much more going on in our conscious life than just a single perceptual experience when we look at a computer in front of us and you might think that all of the fine-grained details included in our experiences is essential to our experience being able to support the corresponding claims. If one is attracted to this idea then one can simply replace the first fact cited in these pairs with the fact picked out by the expression by "I'm experiencing thusly" said while when looking at a computer in ordinary circumstances and replace the second fact in these pairs with the fact picked out by "I'm experiencing thusly" said while in an ordinary instance of remembering previously breaking one's leg. These demonstratives will pick out all of these fine-grained details making up the whole of one's experiential life in these situations.

We also intuitively take facts about past correlations to epistemically support certain claims about the future. Intuitively, the fact that all previously observed swans have been white provides epistemic support for the claim that the next swan will be white. In fact, we intuitively take it that this fact supports the stronger claim that all swans are white. And as the observed sample of the swan population becomes larger the intuitive epistemic support for these claims increases. We would cite these putative facts about past correlations were the rationality of our convictions about the future to be questioned. We implicitly appear to judge that epistemic support relations exist between the following:

(E) <All previously observed swans have been white>  
 (P) The next observed swan will be white

(E) <All previously observed swans have been white>  
 (P) All swans are white

Intuitively, *E* and *P* stand in a relation *R* such that: if a subject *S* is aware of *E* and appreciates *R* then *S* has some prima facie justification for believing *P*.<sup>205</sup>

In putting forward these examples I'm merely attempting to elicit the natural and intuitive judgment that *E* epistemically supports *P* to some degree. I'm not committing myself to the claim that these intuitive judgments are correct.<sup>206</sup> Our goal is to determine what the content of such an intuitive judgment is. We can investigate what this content is even if the intuitive judgment turns out to be false; we need only consider what kind of information would make us feel as if we had to "correct" our judgment (i.e. what kind of information is such that our awareness of that information would and would not dispose us retract our intuitive judgments).

---

<sup>205</sup> As in the perceptual and memorial cases, one might think that other facts are essential for the existence of an epistemic support relation. For instance, one might think facts about past observations only epistemically support claims about the future when taken in conjunction with facts about the observations being random, representative, etc. The reader is free to insert whatever facts intuitively strike them as necessary for the existence of the support relation. The key here is merely to identify fact-proposition pairs that we intuitively judge to stand in an epistemic support relation so as to provide us with data for theorizing about this relation.

<sup>206</sup> See the previous two footnotes.

It's also helpful to lay down some examples of cases where we intuitively judge that a fact does not epistemically support a certain proposition. However, we ought to distinguish between two ways a fact can fail to epistemically support  $P$ . Consider the fact that Eugene has a desire that God exists. Label this fact  $F^D$ .  $F^D$  neither supports believing *that no one desires that God exists* nor believing *that God exists*. But the way in which  $F^D$  fails to support the former is very different than the way it fails to support the latter.  $F^D$  fails to support the belief that no one desires that God exists because it actually supports this proposition's negation.  $F^D$  is evidence *against* the claim that no one desire that God exists. On the other hand,  $F^D$  fails to support the belief that God exists simply because this desire has no bearing on either the truth of this proposition or its negation.  $F^D$  is what we might call an epistemic *non-reason* for believing that God exists. Further examples of epistemic non-reasons are:

- (E) <I asked the magic 8-ball whether Max will get the promotion over me, shook the 8-ball, and it read "Most Definitely!">
- (P) Max will not get the promotion over me
  
- (E) <Electrons have a negative charge>
- (P) The number of dinosaurs that have existed in the past is even

In each of these cases there is a strong intuitive judgment that the putative facts neither support believing  $P$  nor believing  $\neg P$ . Such facts are epistemically irrelevant to the truth or falsity of the target propositions; such facts are epistemically inert relative to these propositions.

## 6.2 Epistemic Support is Not Subjective

### 6.2.1 The Subjectivist View(s)

One might be tempted toward the view that epistemic support is wholly due to subjective attitudes that a subject takes towards a fact. Such a view makes epistemic support a subject relative notion. Some fact  $F$  might epistemically support believing  $P$  for me since I have the relevant attitude toward  $F$  but  $F$  might not epistemically support believing  $P$  for you since you lack this attitude.



On the most straightforward subjectivism, F epistemically supports *P for S* iff S believes that F makes probable *P*. What is the content of this belief? If “F makes probable *P*” is analytically equivalent to “F epistemically supports *P*” then the account is viciously circular and leads to a conceptual regress. Thus, the best way to understand the subjectivist view is to construe the relevant belief as either the belief that: (i) F is a reliable indicator of *P* or (ii) F stands in a logical probability relation to *P*. And the subjectivist will be adamant that neither of these relations constituting the intentional content of the subject’s belief is to be identified with the epistemic support relation. The actual existence of the believed connection is neither necessary nor sufficient for F’s supporting *P*. What’s important is the belief. Call this “doxastic subjectivism.”<sup>207</sup>

Doxastic subjectivism is too strong and too weak. It’s too strong because it over-intellectualizes epistemic support. Subjects might be aware of certain facts F in a way that provides them with justification for believing *P* even if these subjects lack the conceptual resources to form the connecting beliefs. A young child’s awareness of the fact that she is in pain is an awareness of a fact that epistemically supports *for her* the proposition that she is in pain. It’s also tempting to think some children might be aware of their having an experience as of a red table in a way that provides them with justification for believing that there is a red

---

<sup>207</sup> Combining doxastic subjectivism with chapter 5’s discussion of *having* evidence leads to the following view of prima facie justification: if S is aware of F and appreciates that she believes that F makes probable *P*, S has prima facie justification for believing *P*. In chapter 7 I’ll explain that the appreciation of the evidential connection is best construed as an awareness of that connection. Regress threatens, however, if the awareness is construed as a higher-order conceptual awareness of the evidential connection. Thus, the required awareness of an evidential connection must ultimately be a non-conceptual form of awareness. Doxastic subjectivism has a natural way of understanding this and preventing the regress. Epistemic support amounts to a subject’s believing that a certain probabilistic relation obtains. How might I be aware of such a belief? Well, it seems I can have a conscious belief that *P* without forming a higher-order state representing myself as believing *P*. And consciousness is a form of awareness. Such a view, however, must reject high-order thought theories of consciousness such as those defended by Armstrong (1968); Lycan (1987), (2001); and Rosenthal (2012).

table in front of them. But it's doubtful that a child has the conceptual resources to form beliefs about probabilities.

Doxastic subjectivism is also too weak. A subject's belief that *F* makes probable *P* might be wildly irrational. If *S*'s belief about a probabilistic connect is a wild guess or the result of wishful thinking, it's hard to convince oneself that this belief thereby makes it the case that *F* epistemically supports *P* for *S*. Awareness of *F* and of an irrational guess that *F* makes probable *P* doesn't thereby provide prima facie epistemic justification for believing *P*.

A subjectivist about epistemic support is better off appealing to a non-doxastic attitude in her account of epistemic support. Two kinds of mental states that might be helpful for developing a non-doxastic subjectivism immediately come to mind.

First, there are the so-called "seemings" that have piqued the interest of many epistemologists.<sup>208</sup> Sometimes it just seems to us that *P*; we just find ourselves pulled or compelled towards accepting a certain proposition. We're all familiar with what it feels like for a proposition to just *seem* true. And certain philosophers have argued that these "seeming states" are propositional attitudes that represent their content as actualized (unlike desires, fears, or imaginings) but that are nonetheless distinct from beliefs. The distinction between seemings and beliefs is often made by appealing to known perceptual illusions. When I look at the Muller-Lyer lines it seems to me that the lines are unequal in length even after I've become privy to the illusion and no longer actually believe that the lines are unequal.<sup>209</sup> And

---

<sup>208</sup> See Bergmann (*forthcoming*); Brogaard (2013); Cullison (2010); Huemer (2001), (2006), (2007); Tolhurst (1998); and Tucker (2010) and (*forthcoming*).

<sup>209</sup> I'm actually doubtful that one needs to posit a sui generis propositional attitude in order to make sense of the way in which the lines still seem unequal in length. My own view is that this use of seems merely indicates a felt disposition to believe the relevant proposition. For the sake of argument I admit the existence of "seemings" as a propositional attitude distinct from beliefs and felt dispositions to believe.

so one might hold that a fact  $F$  epistemically supports believing  $P$  for  $S$  iff it seems to  $S$  that  $F$  makes probable  $P$ .<sup>210</sup>

Second, a non-doxastic subjectivism might appeal to even less intellectualized states such as a subject's disposition to reason in certain ways in certain conditions.  $F$  epistemically supports  $P$  for  $S$  iff, in conditions  $C$ ,  $S$  has a disposition to form a belief (or at least increase her confidence) that  $P$  on the basis of  $F$  or on the basis of the belief that  $F$  exists.<sup>211</sup>

Non-doxastic subjectivism avoids the worry of over-intellectualizing epistemic support since  $S$  might have the relevant non-doxastic states even though  $S$  is incapable of forming the belief that  $F$  makes probable  $P$ . However, it isn't clear that it avoids the worry of being too weak. Couldn't these seemings be irrational? Defenders of the view will be quick to point out that it at least seems odd to evaluate seemings or dispositions as justified/unjustified or rational/irrational since these aren't under our voluntary control.

The first response to this defense is that the fact that a mental state  $M$  isn't under our voluntary control doesn't entail that it isn't epistemically evaluable. There is a huge debate about whether, and to what extent, even our beliefs are under our voluntary control. Can you simply decide to *right now* form the belief that you're not currently reading a dissertation in epistemology? No. You could get yourself to form the relevant belief by taking a break and cracking open a cold beer (having put up with such a long dissertation I might

---

<sup>210</sup> Earl Conee toys with a view he labels "seeming evidentialism" very similar to that proposed here in Conee and Feldman (2004). However, Conee and Feldman (2008) appears to reject this view in favor of the explanationist view to be discussed in 6.4.4.

<sup>211</sup> It's important to note that this kind of "rule-following" approach to epistemic reasons must appeal to dispositions to reason in certain ways *under very specific conditions*. Without this limitation on which dispositions are relevant, the view would entail that *every* belief of a subject has some prima facie justification. If  $F$  causes me to believe that  $P$  surely I manifest *some* disposition to move from  $F$  to  $P$ . So how ought we specify the relevant conditions? One possibility is that epistemic reasons will be determined by a subject's disposition to reason in certain ways *after a process of ideal reflection* (see Foley 1987 for a doxastic subjectivism that makes this kind of move). Another option is to claim that epistemic reasons are determined by a subject's reasoning dispositions while they are reasoning conscientiously (see Greco (1993) and (1999) for a theory that employs such a notion—though, in these pieces, Greco combines such subjectivist criteria with objectivist reliabilist criteria as well).

recommend you do just that at this point!). But this isn't the kind of voluntariness that seems to be of interest. Of course, explaining the sense of voluntary control that is in dispute when we discuss the voluntariness of belief is incredibly difficult and I won't attempt as much here.<sup>212</sup> The important point is that even if we reached the conclusion that our beliefs are by and large involuntary (I'm inclined to think so), we would still think they can be assessed as justified/unjustified or rational/irrational. So why can't we similarly assess seemings or dispositions to reason in certain ways?

Moreover, there is still a worry that non-doxastic subjectivism is too weak since, even if we reserve the terms 'justified' and 'unjustified' for the evaluation of beliefs, these non-doxastic states can be acquired in deviant ways. It's surely possible for a seeming that *F* makes probable *P* to be the result of wishful thinking. And if this were the case we wouldn't think that the seeming makes it true that *F* epistemically supports *P* for the relevant subject. Or consider Lucy. Lucy witnessed her mother's murder at a young age. The murder happened on a Tuesday and the perpetrator was wearing a red shirt. Seeing this traumatic event caused Lucy to form an unjustified belief that the fact that someone wore a red shirt on a Tuesday makes it probable that he or she has or will commit a murder. Lucy eventually realizes that this belief is completely unjustified and gives it up. However, due to its emotional significance, the belief had been so firmly held that Lucy still has a seeming that the probabilistic connection holds and she remains disposed to move from the belief that a person is wearing a red shirt on a Tuesday to the belief that the person has or will commit murder. It stretches credulity to think a seeming/disposition acquired in this way makes it

---

<sup>212</sup> For more on the doxastic voluntarism debate see Alston (1989); Audi (2001), (2008); Bennett (1990); Feldman (2001); Ginet (2001); Nottelman (2006); Montmarquet (1986), (2008); and Steup (2000), (2008), (2011).

true that the relevant fact epistemically supports for Lucy the claim that some person has or will commit murder.<sup>213</sup>

And this leads to what I take to be the most plausible version of a subjectivist theory of epistemic support. One might respond to the worry of non-doxastic states acquired in deviant ways by appealing to a distinction between learned and unlearned seemings/dispositions. The seeming that certain perceptual experiences make probable that there is an object of a certain shape in our hand is plausibly something that is part of our natural constitution. Other seemings about probabilistic connections, however, must be learned. It seems to some people that litmus paper's turning pink when put in a liquid makes it probable that the liquid is acidic. This seeming is the result of learning rather than our natural constitution. We acquire such a seeming when we perform certain experiments or when we encounter certain kinds of testimony. Similar remarks apply to dispositions to reason in certain ways. How can a non-doxastic subjectivist make use of this distinction? You might suggest that the unlearned non-doxastic states necessarily ground epistemic reason relations but that when a non-doxastic state is acquired it must have been acquired in an appropriate way. F epistemically supports *P* for S iff either (a) S has an unlearned seeming that F makes probable *P* or (b) S has a learned seeming that F makes probable *P* that results from S's having sufficient epistemic support for believing the content of this seeming.

---

<sup>213</sup> Another type of case where we wouldn't want S's seeming that E makes probable *P* to thereby make it that E evidentially supports *P* for S are cases where S has good reason to believe the content of this seeming is false. Consider the Monty Hall puzzle. Initially it seems to us that the information makes it probable that you're equally likely to win the prize whether you switch doors or not. Such a seeming persists (at least in me) even after someone has explained why this isn't so; we're still disposed to reason from the information presented in the puzzle to the conclusion that the strategies are equally likely. Both versions of non-doxastic subjectivism can preserve the intuition that a subject would be *ultima facie* unjustified in believing that the strategies are equally likely to win. Such views, however, are committed to the claim that awareness of this information might make such a subject *prima facie* justified in believing the strategies are equally likely to win *even after the subject has become aware of the illusion*. But intuitively it isn't just that the information in the puzzle provides epistemic support for believing that the strategies are equally likely to win that is then *defeated* but, rather, such information doesn't provide even *prima facie* epistemic support for the relevant belief. Though, I admit that the intuitions in this kind of case aren't entirely clear.

Seemings acquired by wishful thinking, random guesses, or irrational practices like Lucy's fail to ground epistemic reason relations because they fail to satisfy condition (a) on account of their being learned and they fail to satisfy (b) since they aren't learned in an epistemically appropriate manner.

### 6.2.2 Subjectivism as Epistemological Anarchy

Subjectivism leads to a kind of epistemological anarchy. The fundamental worry stems from the fact that certain patterns of reasoning strike us as obviously epistemically absurd. However, subjectivism can give no principled reason why these absurd forms of reasoning couldn't constitute primary forms of epistemic support for a person.

In order to make this point I need to distinguish primary and secondary epistemic support relations. The fact  $F$  stands in a primary epistemic support relation to  $P$  when  $F$ 's supporting  $P$  doesn't depend on our having justification for believing that  $F$  is reliably correlated with  $P$ . According to non-doxastic subjectivism these would be the epistemic support relations determined by the *unlearned* non-doxastic states. So, for instance, we might think the fact that I have certain perceptual experiences epistemically supports the claim that I have a round-hard object in my hand and we might take it that this is a primary epistemic support relation since this relation holds independently of our having justification for believing in a reliable connection between the two. But the epistemic support relation between the fact that litmus paper turned pink and the proposition that a certain liquid is acidic isn't like this. The fact that the litmus paper turned pink epistemically supports the claim that a certain liquid is acidic but this is only because of my background information (testimony from my 5<sup>th</sup> grade teacher and various science textbooks) that justifies me in believing that there exists a reliable connection between the existence of this fact and the truth of the relevant proposition.

Certain facts and propositions strike us as so incredibly unrelated that it appears absurd to admit even the *possibility* that a primary epistemic support relation exists between

them. An example is the fact that a magic 8-ball read “P” and the belief that  $\neg P$ . Awareness of the former fact *could only* provide a subject with justification for  $\neg P$  if the subject had evidence that provided epistemic support for the claim that the fact that a magic 8-ball reads “P” is reliably correlated with the truth of  $\neg P$ . Similarly, awareness of the fact that electrons have a negative charge *could only* provide a subject with justification for a belief about whether the number of dinosaurs that existed in the past was odd or even if the subject was aware of evidence that supported the claim that there is an appropriate reliability connection that holds between these. In other words, while these fact-proposition pairs *could* stand in secondary epistemic support relations, they *could not* stand in primary epistemic support relations.

No version of subjectivism can vindicate this intuition.<sup>214</sup> The move to unlearned seemings or dispositions prevents these facts from *actually* epistemically supporting the relevant propositions. We lack unlearned seemings that readings on magic 8-balls make probable the negation of those readings. Nevertheless, it is *possible* to imagine a creature with the relevant unlearned non-doxastic state. There is nothing incoherent about a situation where a creature has the unlearned seeming connecting magic 8-ball readings with the truth of the negation of that reading. There is nothing incoherent in a creature having an unlearned seeming that the length of a line on someone’s palm makes probable a claim about how long he or she will live. And we can certainly imagine a creature with an unlearned seeming that a person’s smiling at them makes it probable that a group of people are conspiring to kill her. Such creatures strike us as crazy! Not just psychologically crazy, but *epistemically* crazy! Surely we’re right to criticize the resulting beliefs.<sup>215</sup> The problem is that subjectivism appears to make our criticism of such beliefs inappropriate.

---

<sup>214</sup> The fundamental worry I present for externalist accounts of epistemic support is distinct but we’ll see later that this subjectivist worry also afflicts externalist accounts.

<sup>215</sup> A similar point applies to creatures that *lack* certain unlearned dispositions or seemings. Consider again some of the examples of *conclusive* epistemic reasons for beliefs. That I am in pain is

This also leads to a worry of arbitrariness which: (a) makes the connection between epistemic support and truth utterly mysterious, and (b) makes it look as if having epistemic support or justification for one's belief lacks any real value. If subjectivism is correct then creatures which have what appear to us to be monstrous unlearned seemings/dispositions often have justification for the beliefs they form on these bases. What facts support what propositions is a result of how a creature happens to be built. Epistemic support appears arbitrary. Anything can epistemically support believing anything else. The fact that I like chocolate could epistemically support the claim that Napoleon wore a feathered hat, danced like a chicken, and then chopped off his own arm. We need only imagine a creature with the relevant unlearned non-doxastic state. But if any fact could support any proposition and epistemic support relations are just the result of an arbitrary fact about our psychology then why should we even care about epistemic justification? What is the value of believing  $P$  in virtue of my awareness of facts that support  $P$  and my appreciation of this support? Once I realize that the same fact could just as easily have supported  $\neg P$  why think that having such justification makes my belief any more likely to be true? You might insist that the belief will be epistemically likely to be true but if this is just code for the claim that it's epistemically justified then this is utterly trivial. Moreover, once you realize that epistemic support is

---

an epistemic reason for believing that someone is in pain. That Travis is in my office is a reason to believe that either Travis is in my office or Penelope is in my office. Isn't it possible that a subject *lacks* an unlearned disposition to infer the relevant proposition from these facts, or that lacks the related probabilistic seeming? But even if I were to lack the relevant subjective mental state, it's incredibly implausible to deny that this fact is still an epistemic reason for believing that someone is in pain *even for me*. Of course, one might say that if one were aware of whatever this relation is one would thereby have the relevant disposition to form the belief that someone is in pain. And this might be reason to think the person isn't aware of the relation and so doesn't have justification for this belief given what I said about evidence possession in chapter 5. Nevertheless, the intuition that this fact is still an epistemic reason for believing that someone is in pain (even if I don't *have* that reason) remains strong. In these particular examples one might argue that any subject who lacks the relevant disposition or seeming fails to understand the proposition for which the fact is an epistemic reason. And so a subjectivist might attempt to amend the view so as to relativize unlearned dispositions and seemings to situations where the subject has actually achieved an understanding of the target proposition. However, I see no reason to think that the same point wouldn't apply to our examples of non-conclusive epistemic reasons and the appeal to a subject's failure of understanding as an explanation of the lack of a disposition or seeming is much less attractive here.



utterly arbitrary it isn't clear that justification for one's beliefs really provides you with the kind of assurance we seek for thinking our beliefs are true. When I reflect on my beliefs in an attempt to determine which are justified and which aren't, I want some reason to think my beliefs are *actually* likely to be true in a way that the beliefs of these subjects with these monstrous seemings/dispositions aren't. If my beliefs are just as arbitrary as theirs then it isn't clear why I should think that my ways of forming beliefs are any better at getting to the truth than people with unlearned seemings about reliable connections between the readings on magic 8-balls and certain propositions being true.

Subjectivism leads to epistemic anarchy and this doesn't capture how we actually think about the epistemic support relation. In the end these criticisms boil down to an appeal to intuition but the intuition is incredibly strong and convincing. When we reflect on our concept of epistemic support, it strikes us as absurd that when we make a claim about what epistemically supports what we're merely describing our own psychology (e.g. how we happen to be disposed to reason or how things seem to us). Our concept of epistemic support is such that there must be something independent of us that at the very least plays some essential role in determining epistemic support relations.

### 6.2.3 A Blameless Appeal

One way a subjectivist might respond to the previous kind of worry is by appealing to a deontological notion of justification. Such an idea has been advocated by a variety of epistemologists. The central idea is that descriptions of a person as either having or lacking justification for believing a proposition is fundamentally an evaluation. When we evaluate a person's belief as "justified" we use this as a term for praise and we use "unjustified" as a term for blame.

A deontological notion of epistemic justification might provide a way to defend against the previous criticism of the subjectivist theories of epistemic support. Assume that either: (i) S believes that F makes probable P; (ii) it seems to S that F makes probable P; or

(iii) S is disposed to reason from F to *P* when reasoning conscientiously.<sup>216</sup> If someone really does have one of these subjective attitudes then surely we wouldn't *blame* the person for forming the belief that *P* on the basis of F. In forming the belief that *P* on the basis of F they are doing precisely what strikes them as the best way of forming true beliefs and avoiding false beliefs in the current instance. We might still blame someone for these beliefs if he or she *acquired* the relevant attitude in a blameworthy way, e.g. wishful thinking. And this leads one to move to the more sophisticated subjectivist views that distinguish unlearned and learned seemings or dispositions. Surely we can't blame someone for believing *P* on the basis of F when the relevant seeming or disposition is unlearned and a product of her natural constitution. Even if this seeming or disposition strikes us as monstrous, as it does in the case with someone who is naturally disposed to reason from a magic 8-balls reading "P" to the belief that  $\neg P$  while reasoning conscientiously, we still ought not *blame* the person for forming beliefs in this way.

An appeal to blamelessness in defending the subjectivist theory rests on confusing an evaluation of a belief with an evaluation of the believer. Such a distinction is familiar from ethics. A negative moral evaluation of an action doesn't necessarily entail a negative moral evaluation of the person. A person might *non-culpably* perform the wrong action. For instance, it's a contested issue but various features about one's environment such as one's upbringing or culture might excuse one from being blamed for performing a morally wrong action. Similarly, giving a negative *epistemic* evaluation of a belief is to claim that the belief is in some relevant sense not likely to be true but this needn't entail any negative evaluation of the believer. A lack of natural intelligence might be the cause of someone's forming unjustified beliefs. Or, perhaps, a person suffers from various psychological conditions (e.g. paranoia) that are out of her control and lead her to form completely irrational beliefs. Or,

---

<sup>216</sup> The notion of reasoning conscientiously comes from Greco (1993) and (1999).

consider the common example of the person who is raised in a community that adopts epistemically pernicious doxastic practices (e.g. forming beliefs on the basis of entrails, crystal balls, etc.) who then goes on to form beliefs in similar ways. In all of these cases we are likely to evaluate the beliefs as irrational and unjustified while also admitting that the person's self is not to be *blamed* for holding these irrational beliefs. The person's stupidity, psychological conditions, or cultural upbringing excuse the irrationality.

These examples illustrate how justification and these deontological concepts come apart. Thus, the fact that we wouldn't blame people with these monstrous dispositions, seemings, or beliefs doesn't help in the defense against the previous section's objection to subjectivist theories of epistemic support.

### 6.3 Against Externalist Accounts of Epistemic Support

Epistemic support is an objective and mind-independent relation. When we stress the support relation's *objective* truth-conduciveness it becomes almost irresistible to appeal to the kinds of relations that externalists are keen to stress: reliability, safety, sensitivity, proper function, etc. Some philosophers might even try to *define* the internalism/externalism controversy as a debate about whether justification is objectively truth-conducive. We'll see that this is unfair to internalist-friendly accounts of epistemic support.

Reliabilism is the paradigmatic externalist theory of justification, and so it provides a nice illustration of the kind of account of epistemic support to which most externalist theories of justification are committed. Seeing that the reliabilist account fails is instructive because any externalist view will share the problematic feature of understanding epistemic support as a contingent and external relation. My points therefore apply *mutatis mutandis* to any attempt to appeal to proper function, safety, sensitivity, etc. in an account of epistemic support.<sup>217</sup>

---

<sup>217</sup> See chapter 1 for characterizations of these notions.

Goldman's traditional process reliabilism is *crudely* the view that a belief B is justified iff B is the result of a process that would tend to produce true beliefs.<sup>218</sup> Consider again the process of taking experience at face value. We intuitively judge that the psychological facts that are inputs to this causal process (e.g. the perceptual and memorial experiences) epistemically support the output perceptual and memorial beliefs.<sup>219</sup> According to process reliabilism, these psychological facts confer *prima facie* justification on the relevant beliefs because this process is in fact reliable. And so we can conclude that these psychological facts epistemically support the resulting beliefs in virtue of being inputs to the reliable process that outputs the relevant beliefs.

Not all reliabilist theories appeal to the reliability of the process itself in the account of justification. Depending on how one *types* processes it might be possible for the same fact F to be input to two different processes  $\Phi$  and  $\Psi$  that both output a belief that *P* but such that  $\Phi$  is reliable and  $\Psi$  is unreliable. If you were to hold this view you might think that what really matters for justification isn't the reliability of the process but rather that the fact F upon which one bases her belief that *P* is a reliable indicator of *P*'s truth. Something like this seems to be suggested by Alston (1988). In this case epistemic support wouldn't be defined in terms of the reliability of the process but rather in terms of a reliable indicator relationship that is a feature of the fact-proposition pair.

---

<sup>218</sup> I want to stress again that this is a *very* crude formulation of the view. Reliabilists will go on to develop their view in a way that distinguishes foundational and inferential justification, distinguishes propositional and doxastic justification, explains the nature of defeaters, etc. See Goldman (1979) and (1986) and Lyons (2009) for more fully developed versions of process reliabilism.

<sup>219</sup> To be clear, when I refer to perceptual and memorial belief I mean for the former to refer to the beliefs about the external world (e.g. about table and chairs) based upon perceptual experience and the latter to refer to the beliefs about the past based upon the memorial experiences (if there are such things). I do *not* use the phrases "perceptual belief" and "memorial belief" to refer to beliefs about perceptual and memorial experiences.

Despite the differences between these two versions of reliabilism all that matters for my purposes is that each theory attempts to understand epistemic support in terms of there being *some kind* of appropriate reliability connection between F and P. We can refer to this as the reliabilist support thesis (RST):

**RST** A fact (or set of facts) F epistemically supports (or is evidence for) P iff there is an appropriate reliability connection between F and P's truth.

Goldman's traditional version of process reliabilism takes it that a subject needn't have any kind of awareness of F nor any appreciation of a process' reliability in order for the process to confer *prima facie* justification on believing P. Alston's version of the reliable indicator theory is a bit more demanding in holding that the fact F which epistemically supports P must be accessible to S in order for it to confer *prima facie* justification for believing P but S needn't have any access to the epistemic support relation itself.

In the previous chapter I argued against both this radical and the more moderate version of externalism about justification. An investigation of Truetemp and Norman led us to the subject's perspective objections (SPO). In order to meet the SPO we must require for *prima facie* justification that a subject is both aware of supporting evidence and that she appreciate the evidential connection. However, I also explained how these general requirements are compatible with externalist theories of justification since the awareness of evidence and the appreciation of the evidential connection might be understood in externalist-friendly ways (refer back to 5.3.3). Even though the SPO gives us reason to reject the radical and moderate externalist theories of justification, it doesn't necessarily give us reason to reject the view of epistemic support that these theories suggest (i.e. RST).<sup>220</sup> And I take RST and its ilk to be the most plausible competitor to the non-reductive theory that I'll defend later.

---

<sup>220</sup> Comesaña (2010) and Goldman (2011) suggest supplementing a kind of mentalist evidentialism with a reliabilist account of evidential support.

In the next sections I argue against reliabilist accounts of epistemic support by appealing to new evil demon cases. The general argument is as follows: (i) world *W* is a world where we intuitively judge that *S*'s awareness of a fact *F* can give *S* prima facie justification for believing *P*; (ii) our judgment in (i) amounts to the intuition that *F* epistemically supports *P* in world *W*; (iii) *W* is a world where the relevant reliability connection between *F* and *P* fails to hold; (iii) therefore, our concept of the epistemic support relation isn't essentially tied to our concept of the existence of a reliability relation. I develop the argument against reliabilist theories but similar remarks apply *mutatis mutandis* to any other externalist account of epistemic support such as safety, sensitivity, etc.

### 6.3.1 The New Evil Demon Objection

I have perceptual experiences (visual, tactile, etc.) as of a computer in front of me and thereby form the belief that there is a computer in front of me. Label this belief B1. I have a memory experience as of previously breaking my leg skiing and thereby believe that I previously broke my leg skiing. Label this belief B2. We intuitively judge that this is *prima facie* justification-conferring.<sup>221</sup> The so called “new evil demon objection” can be used to show that the content of this judgment cannot be understood in reliabilist (or any other externalist) terms.

Consider two kinds of possible worlds. ORDINARY WORLDS are worlds where we causally interact with physical objects and these interactions give rise to experiences that veridically represent the physical world. DEMON WORLDS are worlds where we're victims of a Cartesian demon and taking experience at face value systematically leads to false beliefs about a physical world. In a demon world our experiences as of physical objects are shared

---

<sup>221</sup> I'll drop the “prima facie” qualifier in what follows—sometimes I use it just as a reminder to the reader—but this is what I have in mind unless I indicate otherwise.

hallucinatory experiences (much like the movie *The Matrix*) produced telepathically by the machinations of an evil demon with the desire to deceive us.

It appears intelligible to imagine that the actual world is a demon world. Surely, however, my beliefs that take experience at face value (e.g. B1 and B2) would be as rational whether or not, unbeknownst to me, the actual world turned out to be a demon world. In his seminal presentation of the problem Stewart Cohen explains:

[W]e are not here supposing that we know that the demon hypothesis is true. Certainly if we were to know that our cognitive processes are unreliable then the beliefs they generate would not be justified. What we want to suppose is the mere truth of the demon hypothesis. Now part of what the hypothesis entails is that our experience is just as it would be if our cognitive processes were reliable.<sup>222</sup>

Beliefs that take experience at face value would be false in a demon world but it's patently absurd to deny that these beliefs would just as rational as belief formed on the same basis in an ordinary world. I'd have exactly the same experiences producing my beliefs and from my own perspective everything would appear just as it would in a world where perception, memory, etc. are reliable. Whatever verdict applies regarding the degree of support that these experiences provide B1 and B2 should apply whether the actual world is an ordinary or a demon world. Assuming that we're unknowingly in a demon world doesn't affect our judgment that forming simple perceptual and memorial beliefs on the basis of experiential facts in the way that we do confers some justification on these beliefs.

Similar points apply to the case of enumerative induction. We intuitively accept that the inference from the fact that all past observed instances of  $\Phi$  have been associated with an instance of  $\Psi$  to the claim that the next (or even all) instances of  $\Phi$  will be associated with an instance of  $\Psi$ . Assume that the external world exists and that perception and memory are reliable; we're not in a demon world. But consider another distinction between two kinds of worlds. REGULAR WORLDS are worlds where events exhibit a kind of lawful

---

<sup>222</sup> Cohen (1984), p. 281.

regularity. IRREGULAR WORLDS are worlds where there is no pair of events such that the occurrence of one event is lawfully *or even reliably* correlated with the occurrence of another but—here is the most important detail—*by pure random happenstance* it turns out that humans exist during a time where there is every *appearance* of lawful regularity. Assume that unbeknownst to us the actual world is an irregular world. Surely our past observations of pencils falling to the ground when we release them in mid-air would make it just as rational to believe that the pencil will fall the next time I let go of it as it does in regular worlds.

It isn't difficult to see that these examples pose a serious threat to externalist theories of justification. It's a purely contingent matter of fact that a process or character trait that takes an experience (or past observation) as its input and outputs a belief that *P* is a reliable, truth-tracking, apt, or a proper functioning process. In demon worlds B1 and B2 are just as rational as in ordinary worlds. In irregular worlds my belief that the next pencil I let go of in mid-air will fall is as rational as it is in regular worlds. But such beliefs are *not* reliably formed in these worlds.<sup>223</sup> Therefore, the reliabilist account of epistemic support (RST) is false.

### 6.3.2 A Blameless Response

In my discussion of subjectivism I distinguished blamelessness and justification. The former is to be understood as an evaluation of the person and the latter as an evaluation of the belief. Moreover, I claimed that a subjectivist would be confusing these two notions if

---

<sup>223</sup> One might attempt to avoid the result that my inductive beliefs are unreliably formed in the irregular worlds by appealing to the notion of a centered world as the relevant environment for determining reliability. Comesaña (2002) makes such a claim in his defense of indexical reliabilism. This, however, doesn't help. First, one might just modify the example to consider the person who has lived 20 years in a world with every appearance of regularity but such that in the next 3 minutes this apparent regularity is going to hell. Surely it still seems that the facts about past observations still support claims about future regularities for that person. Second, even if you were to appeal to centered worlds the past observations and beliefs about the future wouldn't have an appropriate reliability connection. Remember that reliability will have to be characterized *counterfactually*. By hypothesis irregular worlds are worlds where lawful connections fail to hold and this means that in nearby worlds the existence of past observations do not tend to be correlated with the truth of the relevant propositions about the future.



she were to appeal to a blamelessness in defending against the epistemological anarchy objection. Some externalists will be quick to charge that the new evil demon objection rest on this same confusion.

If I were in a demon world we surely wouldn't blame me for holding beliefs that take experience at face value. If I were in an irregular world you wouldn't blame me for forming beliefs via enumerative induction (assuming I'm not blameworthy for holding the premise beliefs). It would still be true in these worlds that I did everything I could to get at the truth; I was just incredibly unlucky to be in worlds that were so unfavorable to forming true beliefs. But the externalist will be quick to claim that we should separate blamelessness and justification. The latter, but not the former, requires that a belief be objectively likely to be true. However, it's easy to *mistakenly* interpret the claim that a subject's belief is unjustified as attributing blame to the subject for holding that belief, especially given the etymology of "justification."<sup>224</sup> Therefore, the externalist might claim that our intuition that the demon victim's beliefs are as justified as our own arises from the fact that we realize she would be equally blameless and that we mistakenly associate evaluations of the justification of a belief with an evaluation of a person's being blameless in holding that belief.<sup>225</sup>

Distinguishing blamelessness and justification is incredibly important. Nevertheless, I doubt that appeal to such a distinction can provide a satisfactory response to the new evil demon objection. Consider two *demon victims* S and S\*. S has the usual experiences that we have when looking at a computer. S thereby forms the belief that there is a computer in front of her. S\* forms the belief that *P* on the basis of a perceptual experience as of a magic

---

<sup>224</sup> Plantinga (1993) has an especially penetrating discussion of this.

<sup>225</sup> While he does allow that a subject who's beliefs are unreliably formed relative to the world she inhabits might have justification since they are reliable relative to the kind of environment for which her processes were designed, Bergmann (2006) is eventually pushed to endorse this kind of blamelessness error theory with respect to our evaluation of a subject who has always been a victim of an evil demon. Goldman's (1988) attempt to respond to the new evil demon problem also seems to be a development of this tactic.

8-ball reading “*P*”. *S\** wouldn’t usually reason in this way but about five seconds ago the demon directly causes *S\** to believe that this form of reasoning is a reliable method for forming truth beliefs and this grounds the relevant disposition in *S\** to reason in this way. Surely both *S* and *S\** are *blameless* in forming their respective beliefs. Nevertheless, there is still a strong intuition that there is something epistemically *better* about *S*’s belief than *S\**’s. Given the details of this case, however, this difference can’t be accounted for in terms of blamelessness. Therefore, our intuition concerning a demon victim’s beliefs that take experience at face value is an intuition about justification and the objective epistemic support that the subject has for her beliefs... it isn’t essentially connected to our judgment concerning blamelessness.<sup>226</sup>

### 6.3.3 Attempts at Accommodation

What’s most interesting about the new evil demon objection is that even many externalists admit to having the relevant intuition. The amount of ink that has been spilled trying to modify externalist theories in response to the problem is a testament to how compelling the thought is. Most externalists won’t just bite the bullet, and many will reject the sort of error theory suggested in 6.3.2. And this leads many externalists to an attempt to develop what I call *accommodation* responses because they attempt to re-construe the externalist’s notion of support in a way that is consistent with the verdict that the relevant beliefs of the demon victim really are justified!

---

<sup>226</sup> There is still this idea that justification requires that the belief be *in some sense* objectively likely to be true. The externalist has a clear notion of objective truth-conduciveness since this can essentially be understood in terms of the frequency in which a process produces a true belief (although an externalist will be quick to appeal to counterfactual frequencies). Such a frequency is a fact of the mind-independent world and holds independently of anyone’s beliefs or attitudes about such a frequency. There is an intuition that a demon victim’s beliefs that take experience at face value are justified but in order to defend this claim and make the objection to externalism stick there is still an uncharged burden of explaining a sense in which these beliefs can be objectively likely to be true despite the fact that these frequencies fail to obtain in demon worlds.

Reliabilism is naturally interpreted as an appeal to reliability in the world in which the belief was formed. Thus, the reliability relevant to the justification for inhabitants of ordinary and of demon worlds is assessed relative to *different* worlds. This is what allows for the same set of facts to stand in different epistemic support relations in different worlds and thereby leads to the counterintuitive results about justification.

Such an interpretation, however, isn't forced on the reliabilist. Reliabilists can privilege reliability in some world(s) rather than others. Making this move guarantees that the reliability which determines justification is relativized to the *same* world(s) for both the inhabitants of ordinary and demon worlds. Therefore, if two subjects use the same process, there can't be any divergence in the reliability that determines facts about justification. This is a strategy externalists have adopted for *accommodating* the intuitions brought out by considering demon worlds. But such a response must identify which worlds to privilege in determining reliability.<sup>227</sup>

### 6.3.3.1 Actual World Reliabilism

One of the most interesting uses of this strategy is "actual world reliabilism."<sup>228</sup> Comesaña (2002) and Sosa (1993), (2001), and (2009) offer the most notable defenses, but Goldman (2008) and (2011) has also shown sympathy for the proposal. This view privileges reliability in the actual world:

**ACT** A subject S's belief B is justified iff B was formed by a process that is *actually* reliable.

---

<sup>227</sup> The first use of this strategy was Goldman's attempt to rigidify the rightness of epistemic rules to reliability in "normal worlds" or worlds that are pretty much as we believe this world to be—see Goldman (1986). My discussion will ignore Goldman's normal worlds response since it has been discussed extensively in the literature and the problems are well known. Goldman himself gave up on this response almost immediately after it was published. See Goldman (1988), Fumerton (1995), and Majors and Sawyer (2005) for discussions of the difficulties that this normal worlds response encounters.

<sup>228</sup> Comesaña (2002) refers to this as "indexical reliabilism."

Privileging actual world reliability in this way shouldn't be construed as an *ad hoc* modification. If this really does deal with the demon world counterexamples then this is good enough motivation. Moreover, there might be good motivation for the view independent avoiding this counterexample. One reason we value justified beliefs is because we take justification to be intimately tied to truth. And, presumably, we care about forming justified rather than unjustified beliefs because this is (positively) tied to getting true beliefs in the *actual* world.<sup>229</sup>

I've been clear that I think even externalists would benefit from adding a condition of awareness of supporting evidence and of appreciating the evidential connection to her theory of justification. I take it that the SPO shows that ACT requires this further supplementation. Nevertheless, more must be said in order to show that the nature of epistemic support upon which ACT relies is incorrect. If ACT succumbs to the new evil demon problem *this* would show that epistemic support isn't to be understood as actual world reliability. So let's consider why advocates think ACT avoids this problem.

ACT's appeal to *actual* reliability creates a kind of ambiguity in the view. "Actual" is an indexical and thereby leads to an ambiguity when discussing counterfactual worlds. Sometimes "actual" refers to the counterfactual world under consideration (the world that *would* have been actual) and sometimes refers rigidly to the world *I* (the speaker) occupy.<sup>230</sup> Comesaña's (2002) development of the view incorporates this fact into the theory of justification by appealing to the framework of two-dimensional semantics so that our

---

<sup>229</sup> This is a complicated issue but I wanted to provide some reason for privileging the actual world. See Cohen (1984), Conee (1992), and Fumerton (2011) for good discussions of the truth-connection.

<sup>230</sup> See Lewis (1970) for a defense of the claim that "actual" is an indexical and that it can be used in two senses when considering counterfactual worlds. Lewis gives the sentence "If Max ate less, he would actually enjoy himself more" as an example where "actual" refers to the subject's world and the sentence "If Max ate less, he would be thinner than he actually is" as an example where the actual world is fixed and "actual" refers to the speaker's world.

concept JUSTIFICATION inherits the indexical nature of the concept ACTUAL (of which it's partially composed). We can thereby use our concept of justification to express two distinct propositions: a “diagonal” proposition which entails reliability in the world of the subject to which we're attributing justification or a “horizontal” proposition which entails reliability in the actual world where “actual” refers rigidly to the world of the speaker making the attribution.<sup>231</sup>

You can think of the diagonal proposition as attributing the property of being subject-world reliable and the horizontal proposition as attributing the property of being speaker-world reliable. An attribution using this indexical notion of justification, according to ACT, can be used to ascribe two distinct properties to a subject's belief:

**JUSTIFICATION-1:** a subject S's belief B is justified-1 iff B is subject-world reliable (i.e. reliable in S's world)

**JUSTIFICATION-2:** a subject S's belief B is justified-2 iff B is *rigidly* speaker-world reliable (i.e. reliable in the world where the attribution is made).

ACT attempts to avoid the new evil demon objection by claiming that an intuitive attribution of justification to demon victims comes out as correct when interpreted as an attribution of justification-2.

Why would an attribution of justification-2 be consistent with the new evil demon intuition? Consider the following counterfactual: if the actual world is a demon world then my beliefs that take experience at face value are *prima facie* justified. One way of understanding the new evil demon objection is just that this conditional strikes us as obviously true but appears to come out false on traditional externalist theories. Does the conditional come out true if we modify the theory so as to appeal to *actual* reliability? Sosa (2009) give a firm affirmative answer. He reasons as follows: (i) either this is a material or

---

<sup>231</sup> The distinction between diagonal and horizontal propositions comes from Stalnaker's (1978) development of 2-D semantics.

subjunctive conditional; (ii) if it's a material conditional then it's trivially true since the antecedent is false; (iii) if it's a subjunctive conditional then we evaluate it by considering the closest possible world where the antecedent is true and if, in that world, the process is reliable relative to the actual world then the conditional is true; (iv) in the demon world it's still true that the process is reliable relative to the actual world (i.e. the world *we* the subjects who are making the attribution, *not the demon victims themselves*, occupy); (v) therefore, the conditional is true.<sup>232</sup>

Sosa is right that there is one sense of the counterfactual where “actual world” refers to the world *we*, the people making the attribution of justification, occupy and not to the possible world under consideration (i.e. the demon world). This is precisely the reading we get if we're attributing justification-2 to a demon victim's beliefs. Nevertheless, when reading Sosa's discussion the reader cannot help but think that Sosa is simply refusing to consider another interpretation of this counterfactual. Comesaña makes room for this by allowing that justification sometimes attributes justification-1 and at other times attributes justification-2. Given this admission, however, an internalist is likely to be worried about the fact that the counterfactual comes out false if interpreted as an attribution of justification-1 to a demon victim's beliefs. I take it that this concern is legitimate but it isn't clear how to turn this into an objection. Surely there is *some sense* in which beliefs that take experience at face value in demon worlds are as epistemically good as those that do so in ordinary worlds. ACT accommodates this by appealing to the notion of justification-2. But there is also clearly *some sense* in which the beliefs formed in this way in an ordinary world are better than those formed in this way in the demon world. Thus, it isn't clear that the ACT's commitment to the claim that taking experience at face value confers justification-1 in ordinary but not demon worlds is problematic. The important question to consider is

---

<sup>232</sup> Sosa (2009), pp. 35-40.

whether the sense in which beliefs that take experience at face value in ordinary worlds are better than those in the demon world is an importantly *epistemic* sense.

I want to raise two initial worries for ACT. The first is a technical worry. The second is a meta-linguistic worry that raises suspicion that something fishy is going on with ACT's apparent solution and thereby sets the stage for raising the more fundamental problem that I raise in section 6.3.4.

Accommodating the new evil demon intuitions with justification-2 essentially relies on the possibility of rigidly designating possible worlds. However, there are reasons to be suspicious that this is actually possible. Consider two claims: (a) I'm a philosopher and (b) I'm actually a philosopher. Intuitively we all think that (a) and (b) express the same proposition. But if I can rigidly designate the actual world then (a) will be a contingent truth while (b) will be a necessary truth. Surely something has gone awry since there doesn't appear to be a sense in which the claim that I'm actually a philosopher would express anything other than a contingent truth. Similarly, we should take it that a process' being actually reliable must remain a contingent truth. A process like taking experience at face value could be reliable and it could be unreliable in *this* world. At the very least I would like advocates of ACT to say more about the mechanisms that rigidly designate entire worlds as I'm skeptical that this is actually coherent. Despite my skepticism, however, I'm willing to let this point slide into the background and rest my case on a more fundamental problem to which my next worry hints (see 6.3.4 for further development).

More worrisome for ACT is a meta-linguistic problem. Susan is in an ordinary world. Sally is Susan's mental counterpart in a demon world. Each has an experience as of red ball in front of her and, on this basis, forms the belief that there is a red ball in front of her. Being the reflective persons they are, they each reflect for a moment and declare, "My belief that there is a red ball in front of me is justified." Given that Susan and Sally are mental counterparts who have the exact same experiences and beliefs it strikes us that either both assertions are true or both assertions are false. In fact it looks like I'm just stating the new

evil demon problem. ACT, unfortunately, entails that Susan's utterance is true but Sally's is false—whether we take them to be asserting the diagonal or horizontal proposition!<sup>233</sup> All of this raises the suspicion that somehow the formal apparatus of 2-D semantics is only providing a superficial solution by preventing us from raising the concern we're trying to raise. In section 6.3.4 I vindicate this suspicion by illustrating how to raise the concern that makes the appeal to the ambiguity of “actual” and the 2-D semantics irrelevant to addressing the new evil demon problem.

### 6.3.3.2 Home World Reliabilism

Further possibilities for fixing justification relevant reliability are available that avoid some of the worries with an appeal to actual world reliability. Brad Majors and Sarah Sawyer (henceforth, M&S) have advocated privileging reliability in what they call “home worlds.” In articulating the view I will appeal to content externalism since this is the framework within which M&S develop the view. However, we'll see that the adoption of content externalism is inessential to M&S's proposed solution to the new evil demon problem (it's actually supposed to be doing anti-skeptical work for M&S).

Largely due to the work of Putnam and Burge, many philosophers now believe that the intentional contents of mental states such as beliefs cannot be individuated without essential reference to the kind of environment a subject inhabits.<sup>234</sup> If sensations have intentional contents, as many philosophers working on perception now believe, these too could only be individuated with essential reference to the subject's environment.

---

<sup>233</sup> This is related to Goldman's (1993) point that there is no indication in ordinary practice that the folk's ordinary notion of justification is relative. Our intuition is that demon victims' beliefs that take experience at face value are justified *simpliciter* not merely relative to some world or other. Goldman's point, however, was directed at Sosa's older (1988) and (1991) view where justification was a relative notion and no world (not even the actual world) was privileged.

<sup>234</sup> See Putnam (1975) and Burge (1979) for canonical articles defending content externalism.



I mentioned that content externalists often reject the idea that a subject who has *always* been the victim of an evil demon could be our mental counterpart (assuming we're in an ordinary world). The demon environment doesn't have the right kind of properties to individuate intentional contents in a subject that represent physical objects like tables, chairs, or even hands. However, the new evil demon problem arises in cases of recent *transportation* to a demon world.

M&S propose a novel way for dealing with these transport cases. They propose a view that they label Home World Reliabilism or HWR:

**HWR** Justification consists in reliability in the subject's home world.<sup>235</sup>

Yet again, the corresponding thesis about epistemic support will result from replacing "reliability" in clause (ii) of RST with "home world reliability." A subject S's home world is "a set of environments—a set of configurations of content determining properties together with relations to the subject"—that would determine mental content of the sort S has. Imagine that a subject S possesses a mental life (beliefs, desires, sensations, etc.) that has the intentional contents  $\{I_1, I_2, \dots, I_n\}$ . S's home world consists of the set of environments such that, had S been placed in these environments, they would have determined a mental life with  $\{I_1, I_2, \dots, I_n\}$ . To say that a subject's belief is home world reliable is to say that the belief was formed in a way that would tend to produce true beliefs when considered across this set of possible worlds.

HWR deals with transport cases because a belief can be unreliably formed relative to a subject's *current* environment yet reliably formed relative to her home world. When I'm transported from an ordinary to a demon world my home world is the set of ordinary worlds. Taking experience at face value is a reliable process relative to ordinary worlds. Therefore, even if I were transported to a demon world yesterday my beliefs that take

---

<sup>235</sup> Majors and Sawyer (2005), p. 272.

experience at face value would still be *prima facie* justified according to HWR. Mental counterparts necessarily have the same home world and therefore the justification relevant reliability is necessarily equivalent for mental counterparts that use the same belief-forming processes. This also shows that the appeal to content externalism is inessential to HWR's solution to the new evil demon problem since the result is guaranteed by the definition of home worlds rather than any account of how content is actually determined.

HWR, in certain respects, is a nice improvement over ACT. HWR's solution doesn't essentially depend on rigidly designating possible worlds. We can pick out home worlds via a description of them as worlds with the relevant content determining properties. HWR also avoids the meta-linguistic problem to which ACT was susceptible. However, the view isn't without its fair share of problems.

Advocates of HWR will face difficult decisions when considering "slow-switching" type cases. There are strong reasons to think a subject who switches from some environment E to an environment E\* with different content determining properties might both *gain* new concepts due to E\* while *retaining* concept attained in E. If I were transported to twin earth, after a while it seems I should be able to form "water" thoughts that refer to XYZ. But consider my apparent memory of drinking a gallon of 'water' after a soccer game which occurred *before* I was unknowingly transported from earth to twin earth. It's at least plausible that my memory would give rise to a thought about the H<sub>2</sub>O that I encountered at earth. So, in such a situation, I'd appear to have the XYZ water concept and the H<sub>2</sub>O water concept.<sup>236</sup> Similar remarks apply *mutatis mutandis* to a person transported from an ordinary to a demon world. This, however, entails that in these slow-switching scenarios there is no single kind of environment that has properties that could by themselves individuate the

---

<sup>236</sup> See Tye (1998) for a discussion of this kind of example. Tye actually suggests that the content of your memory is determined by your present environment rather than that the conceptual content of one's previous environment might be retained via memory.

intentional contents of the subject. How should we understand a subject's home world in this situation? I'll leave this task to the advocates of HWR. Answering this question will be difficult but my hunch is that nothing fatal will arise for HWR along the way.

A more worrisome problem is that a subject's home world actually consists of an *infinite set* of environments. A home world isn't to be construed as the world that actually individuated the contents of a subject's intentional states but as a set of environments that, had S been placed in the environment, the environment *would* have given rise to intentional states with these contents. But now consider a process  $\Phi$ . In an infinite set of environments  $\Phi$  will produce an infinite number of true beliefs and an infinite number of false beliefs. So how are we to assess  $\Phi$ 's home world reliability? We need to compare the ratio of true to false beliefs but, given the infinite number of both true and false beliefs, we can only do this if we have some ordering of the true and false beliefs. On some orderings  $\Phi$  will be classified as reliable but on others  $\Phi$  will be classified as unreliable. The problem is that there doesn't seem to be any principled way of privileging one ordering over another.<sup>237</sup>

#### 6.3.4 The Really New Evil Demon Problem

The claim that these modifications to the externalist's understanding of epistemic support resolve the new evil demon problem has some initial plausibility, but I've tried to raise some worries along the way to make the reader suspicious of these moves. Now I'll show that the appeal to the technical apparatuses of 2-D semantics and home worlds only results in a superficial solution; when we've properly understood the argumentative strategy that an internalist appeals to with the new evil demon problem we can reformulate the

---

<sup>237</sup> Richard Fumerton (2011) pushes this worry against Henderson and Horgan's (2007) transglobal reliabilism. One might think this problem applies just as easily to reliabilism more generally since it must appeal to counterfactuals in characterizing reliability. However, the crude reliabilist can avoid the problem by finding a truth-maker for the counterfactual claim within the actual world. The propensity interpretation of probability that Goldman (1986) appeals to seems particularly well suited for such a task.

problem in a way that neither of these views can solve. Our worry for externalist accounts derives from what we're inclined to say about specific cases, suggesting that the problem is fundamentally conceptual. We need to reformulate the problem so as to stress this conceptual nature.

Consider your total current experiential state and your current belief that you're reading a dissertation on epistemology. You're inclined to think that your current experiential state provides you with some prima facie justification for the relevant belief. Now the philosopher objecting to externalist accounts of epistemic support asks you to imagine that your conscious state is exactly as it is but, unbeknownst to you, *you* (not your mental counterpart!) are the victim of the Cartesian demon. It still strikes you that your belief is justified even if this turned out to be true in *this* world. And so your judgment that your current experiential state provides some prima facie justification isn't essentially tied to your concept of this belief having been reliably formed. This first-person nature of the problem is obfuscated rather than illuminated when we move to a consideration of mental counterparts and certain counterfactual ways of formulating the problem.

Instead of considering the beliefs of our mental counterparts in demon worlds, I suggest raising the concern by appealing to when we would intuitively retain or retract our attributions of justification to our (future) past selves. The idea is as follows: right now, at time  $t_1$ , we intuitively judge that belief B is justified. Instead of asking if B would still be justified if we were in a demon world I suggest consider the question of whether we would retract this claim if we were to become convinced at a later time  $t_2$  that we are in a demon world. Such a thought experiment says nothing about whether our conviction at  $t_2$  is true or false. We needn't say anything about whether the world is an ordinary or demon world in order to imagine the scenario I'm considering. But it appears that we wouldn't and shouldn't retract our attributions of justification *to our past selves* upon becoming convinced that such a

possibility obtains. When we raise the concern in this way it becomes clear why the externalist modifications considered earlier only provide a superficial solution.<sup>238</sup>

Attributions of justification are always indexed to a time. I might attribute justification to a belief *B* at  $t_1$  and then claim that *B* is unjustified at  $t_2$ . But the claim at  $t_2$  doesn't entail that I am retracting my claim that *B* was justified at  $t_1$ . I might have had perfectly good justification for *B* at the earlier time and then lost this justification at  $t_2$  upon encountering defeating evidence. If, at  $t_1$ , I form a belief that *P* on the basis of an article in what I know to be a reliable newspaper this belief is justified even if the same belief fails to be justified at  $t_2$  due to my encountering a reliable source that provides testimony that this is one of the rare occasions when this newspaper gets things wrong. In such a situation I wouldn't and should retract the claim that my belief that *P* was justified *at the earlier time*. I want to suggest that the same thing is true when someone becomes convinced that a certain source is unreliable.

Consider a toy example about temperature. At  $t_1$  I look at a thermometer that reads "76°" and thereby believe that it's 76°. I'm a rather reflective person and I wonder whether

---

<sup>238</sup> An analogy will help illustrate the strategy I'm recommending. Imagine that a theist makes the intuitive judgment that genocide is wrong. Now imagine that you are set the task of convincing the theist of the claim (C) that this notion of wrongness isn't constituted by God's commanding us not to commit genocide. There are two ways of going about this that are very similar but importantly different. You might ask the theist to imagine a world where God doesn't exist and ask if genocide would still be wrong. Even most theists will be strongly inclined to say that it would. And this is a reason for the theist to accept (C). However, one problem with this strategy is that the theist might respond that God necessarily exists and that it's metaphysically impossible for there to be a world where God doesn't exist and perhaps even that it's metaphysically impossible for there to be a world where God doesn't command us not to commit genocide—notice that this looks a lot like the content externalists claim about a mental counterpart who has always inhabited a demon world. The theist might then respond to your request with befuddlement since they don't have any idea how to evaluate counterfactuals with necessarily false antecedents. For this reason I take it that there is a better way to make this same point to the theist in support of (C). An alternative is to ask the theist whether they would be inclined to retract their claim that genocide is wrong if, at some point in the future, they were to become convinced that God doesn't exist. An inclination to give a negative answer provides that same kind of reason for believing (C). And I imagine that this way of pushing the point would actually be much more effective than the method considered in the previous paragraph. One needn't admit the possibility that God fails to exist in order to imagine the possibility that one becomes *convinced* that God doesn't exist.

this belief is justified. After reflecting on the matter, I judge that the belief is in fact justified given my current evidence. I then engage in a further act of reflection and ask myself what I'd do if I were to later become convinced that unbeknownst to me the thermometer was and had always been unreliable. Would I, and should I, *retract* my claim that the belief formed on this basis was justified *at t<sub>1</sub>*? No. It's natural to say that in such a situation I'd *no longer* be justified in trusting the thermometer. But this implies that coming to believe that the thermometer is and has always been unreliable wouldn't affect my judgment that I was justified in trusting the thermometer *at the earlier time*. We can extend the example further. Becoming convinced at *t<sub>2</sub>* that *all* thermometers are and had always been unreliable wouldn't, *by itself*, make me feel pressure to retract my attribution of justified belief *relative to t<sub>1</sub>*. This suggests that the concept of justification applied to my belief at *t<sub>1</sub>* isn't essentially tied to my concept \*PRODUCED BY A RELIABLE PROCESS\*.

Forming beliefs on the basis of a thermometer's reading isn't a foundational source of justification. And this might give the externalist some wiggle room. The problem, however, arises just as easily in the case of supposed basic justification sources like perception, memory, and introspection.<sup>239</sup>

Consider three cases: (i) my believing, B1, that there is a computer in front of me on the basis of my perceptual experiences as of a computer; (ii) my believing, B2, that I previously broke my leg skiing on the basis of my vivid memory experience as of breaking my leg skiing; and (iii) my believing, B3, that I'm having a stinging pain sensation that seems to be located in my leg on the basis of my introspective experience as of such. Prior to entering a philosophy class we all think that these are paradigmatic cases of experiences conferring *prima facie* justification on a person's beliefs. Moreover, these are all examples of

---

<sup>239</sup> And the same worry would apply to the case of testimony if this is a foundational source of justification. This is worth mentioning given the rising popularity of non-reductionism in the epistemology of testimony.

sources that are often supposed to be basic justification sources. So let's apply my reformulation of the new evil demon problem to these cases.

What would happen if we became convinced (*correctly or incorrectly*) at some future time,  $t_2$ , that we have always been in a demon world where beliefs that take perceptual experiences at face value are and have always been unreliably formed. Would we feel any pressure to *retract* the previous claim that *my past self* was justified in believing B1 at  $t_1$ ? No. Having become convinced that the actual world is a demon world we would naturally make the claim that I was *no longer* justified in trusting my perceptual experiences. This "no longer" is incredibly natural and implies that we would, and should, hold on to the claim that the belief was justified *at the earlier time* even though justified relative to the later time  $t_2$ . That we don't think consistency pressures us to retract our claim upon becoming convinced that the demon scenario obtains shows that reliability isn't conceptually necessary for justification.

The same point can be made without appealing to possibilities as exotic as the evil demon. Consider my memory belief B2. I can imagine a neurologist's testimony convincing us (*correctly or incorrectly*) at some future time,  $t_2$ , that I have always had a hitherto undetectable cognitive malfunction which creates mistaken but entirely coherent memories.<sup>240</sup> If we were, in this way, to become convinced that such a possibility obtains would we then feel pressure to retract the earlier judgment that my experiences provided prima facie justification for believing B2 at  $t_1$ ? No. Again we would judge at  $t_2$  that I was *no longer* justified in trusting memory, which thereby suggests that we would hold onto the claim that my belief was justified relative to the earlier time. This is true in spite of the fact that our

---

<sup>240</sup> A very similar case involving memory is considered by Goldman in his (1979) article when he considers how to incorporate the notion of a defeater into his process reliabilism. I'm suggesting here that it's not only intuitive for us as 3<sup>rd</sup>-person observers to view this as a case where justification was present but then defeated. It is also appropriate for the person who receives this testimony about her memory to conceptualize from her own 1<sup>st</sup>-person perspective that this as a case where her justification is *defeated* rather than absent all along. And, as we'll see, this is inconsistent with both actual and home world reliabilism.

conviction at  $t_2$  would amount to the belief that the way in which our memory belief was formed at the earlier time was (and had always been) unreliable. We would construe the situation as one where the neurologist's testimony provided me with a *defeater* rather than a case where I lacked justification all along.<sup>241</sup>

Finally, I can imagine becoming convinced at a future time,  $t_2$ , perhaps on the basis of various empirical studies, that introspection itself is generally unreliable and prone to distortion.<sup>242</sup> However, if I were to become convinced of this, it wouldn't seem to put any pressure on me to *retract* my intuitive attribution of justification to B3 at  $t_1$ . Upon becoming convinced by these empirical studies I'd naturally move to the claim that I was previously justified in trusting introspection, but that I was *no longer* justified in such trust.

ACT and HWR proposed solutions to the new evil demon problem falter when it's reformulated in this way. Consider again actual world reliabilism. In all of the imagined future situations where I become convinced that the relevant skeptical scenario obtains I

---

<sup>241</sup> There are empirical studies that suggest that our apparent memory is much more unreliable and prone to distortion than we might have thought. See Schacter (1999) for a good overview concerning the different ways that memory goes wrong. Once we're exposed to these empirical findings we may no longer be justified in trusting our apparent memories in a variety of situations, but we naturally think that the judgments we made on the basis of these apparent memories *prior to receiving the defeating evidence* were justified at the earlier time despite judging them to be unreliably formed. Of course, the reliabilist could reply that we nonetheless recognize that, despite being unreliable in certain respects, at a more general level memory is reliable. And so our recognition of the earlier beliefs as justified is tied to the judgment that these were formed in a more general way that is reliable. Such a reply actually takes advantage of the "generality problem." It's for this reason that in the main text I consider the possibility of gaining a neurologist's testimony that one's memory is (and has always been) unreliable *in general*. Yet again, the reliabilist might attempt to move to a more general characterization of the process in order to vindicate the idea that we still believe at the later time that the more general level is reliable. For instance, an externalist might say that at the later time we would recognize that the memory beliefs were formed by the general process of taking experience at face value (where this included taking perceptual and memorial experiences at face value) and that this more general process is reliable even if the more specific process of memory is unreliable. Moving to such a general level, however, looks *ad hoc* and simply makes the generality objection more worrisome. Moreover, we do intuitively separate memory and perception as being two distinct kinds of processes. We don't naturally think we ought to assess a process' reliability at a level more general than this.

<sup>242</sup> See Schwitzgebel (2000), (2008), and (2011) for the kinds of empirical considerations that might cause someone to doubt the reliability of introspection in this way.



would then believe that my *earlier* beliefs B1, B2, and B3 had been formed by a process that is unreliable in my past self's world (i.e. it's subject-world unreliable) and my current's self's world (i.e. it's speaker-world unreliable).<sup>243</sup> According to ACT, if my conviction at  $t_2$  is correct then there is *no sense* in which I could correctly characterize my past self as having justification for his beliefs at the earlier time  $t_1$ . If ACT captured our ordinary notion of justification, I should feel pressure to retract my claim that these beliefs were justified at  $t_1$ . But, intuitively, I would not and should not feel pressure to retract these earlier attributions of justification. Therefore, our ordinary notion of justification isn't essentially tied to the notion of actual reliability as suggested by ACT.

Applying this tack to HWR is trickier since believing you're in a demon world doesn't amount to believing that beliefs that take experience at face value are unreliable *relative to worlds that would have determined a mental life with the intentional content of your current mental life*. However, we could still use Descartes' evil demon hypothesis to press the concern against HWR. It's consistent for a person to become convinced that a scenario obtains even if that scenario is actually metaphysically impossible. As M&S admit, "[t]he way content is actually determined is one thing; the way in which we take it to be determined... is another."<sup>244</sup> So one could become convinced (whether correctly or incorrectly) that demon worlds could still individuate intentional contents that refer to physical objects and convinced that *relative to the set containing ordinary and demon worlds* beliefs that take experience at face value are not reliable. Nevertheless, I think it will be easier to apply the tack to HWR by considering the alternative scenario where I become convinced by testimony that my memory is unreliable.

---

<sup>243</sup> Note that Comesaña makes clear that "actual" should really be taken to refer to a centered world, i.e. a triple of world, time, and place (2005, pp. 259-290).

<sup>244</sup> Majors and Sawyer (2005), p. 274.

Assume that my thought contents have been individuated in a world where my perceptual states and thoughts are causally sensitive to mind-independent physical objects like computers, skies, mountains, lakes, canoes, etc. I currently believe that I was at the library writing my dissertation yesterday, that I previously broke my leg skiing, and that I drove a jet ski on Lake Mille Lacs when I was a teenager. All of these thoughts refer to our good old everyday physical objects since my thought contents have been individuated in an environment where I causally interact with these physical objects. And I've formed each of these memory beliefs on the basis of what I would describe as some of the most vivid and coherent memory experiences that I have. These beliefs seem to be paradigmatic examples of currently justified memory beliefs. Right now, at  $t_1$ , we're inclined to judge that my vivid and coherent memory experiences confer *prima facie* justification on each of these beliefs. What would we do if, at some *future* time, we were to become convinced by a neurologist's testimony that I suffer (and have always suffered) from a hitherto undetectable cognitive malfunction that creates mistaken but entirely coherent memories? We'd be convinced that my forming beliefs on the basis of my memory experiences are, have always been, and always would be unreliable in most worlds where I could form these thoughts. Nevertheless, it is still natural to suggest that I am only *no longer* justified in trusting my memory. Becoming convinced that this possibility obtains (which is equivalent to becoming convinced that these beliefs are home world unreliable) wouldn't and shouldn't make us feel pressure to retract the claim that these belief *were justified* at  $t_1$ .

That we don't think we would feel pressured, for the sake of consistency, to retract our *previous* attributions of justification upon becoming convinced that these scenarios obtain intuitively shows that our ordinary notion of justification isn't essentially tied to actual world or home world reliability. And these points apply *mutatis mutandis* to other externalist friendly approaches to epistemic support such as reliability, proper functioning, safety, sensitivity, etc.

#### 6.4 Epistemic Support: A Keynesian Approach

If we can't appeal to externalist notions such as reliability, proper functioning, safety, sensitivity, etc. in an account of epistemic support how else might we understand the relation? I want to suggest that the failure of the subjectivist and externalist accounts strongly suggests that epistemic support is *sui generis* and *irreducible* relation.

The reliabilist notion of epistemic support closely resembles a frequency theory of probability. Consider the claim that the proposition *that all previously observed swans are white* makes probable *that the next observed swan will be white*. According to the frequency theory of probability this just means that this pair of propositions belongs to a set of ordered pairs such that whenever the first proposition is true *usually* the latter proposition is also true.<sup>245</sup>

We might simply define the reliabilist notion of epistemic support in terms of frequencies: F epistemically supports P for S iff (a) S has a belief forming process  $\Phi$  that has F as input and the belief that P as output and (b) when it's true *that a belief B that P is produced by  $\Phi$  it is usually also true that P*. However, there is a well-known controversy concerning the interpretation of probability that pre-dates the internalism/externalism debate. The frequency theory of probability isn't the only game in town when it comes to *objective* probability statements. Keynes (1921) offered a competing logical theory of probability.

---

<sup>245</sup> Such a view of probability faces the same generality problem that the reliabilist faces. The proposition pair *all previously observed swans are white/the next observed swan will be white* belongs to an infinite number of different sets of ordered proposition pairs. It is part of the set of ordered proposition pairs where the first is a claim about swans and the second is a claim about swans. It's a part of the set of ordered proposition pairs where the first attributes whiteness to an object and the latter does so as well. It is part of the set of ordered proposition pairs where the first is about the past and the latter is about the future. Each of these sets would give us a different result pertaining to how probable the past observation of white swans makes the claim that the next observed swan will be white. Presumably we want the relevant set of order proposition pairs to be something akin to the set where the former attributes a property  $\Phi$  to all past observations of something with property  $\Psi$  and the latter attributes the  $\Phi$  to the  $\Psi$  we will observe next. But by what principle are we to type the proposition pairs in a non-arbitrary way that gives the correct intuitive verdict?

My goal here isn't to argue that one of these interpretations is to be identified as *the interpretation* of probability. I'm a pluralist about probability statements. I take it that the subjectivist, frequentist, propensity, and logical theories of probability all capture one or another notion of probability that is often expressed in ordinary discourse. Each theory is useful in different contexts. Nevertheless, I want to suggest appealing to the Keynesian theory of probability for understanding the epistemic support relation.

#### 6.4.1 Keynesian Probability

Keynes (1921) suggests understanding probability relations between propositions on the model of entailment. In fact, we can think of entailment as the upper limit of the making probable relation. When a proposition  $P$  makes probable  $Q$  without entailing  $Q$  we can think of this probability relation as a kind of *partial* entailment. And there are four key features to keep in mind concerning entailment: (i) it's an objective and mind-independent relation; (ii) facts about entailment cannot be reduced to *any* facts about relative frequencies or nomological relations; (iii) entailment is an internal relation; and (iv) it's knowable a priori. Let's consider each of these in turn.

##### 6.4.1.1 Objectivity

Entailment relations between propositions are objective and mind-independent.  $P$  entails  $Q$  means that it's absolutely impossible for  $P$  to be true while  $Q$  is false.  $P$ 's truth guarantees the truth of  $Q$ . Entailment is a relation connecting the *truth-values* of propositions in the sense of being necessarily truth-preserving. However, the existence of such a truth connection doesn't depend on any beliefs or attitudes that you, me, or any other subject takes towards  $P$  and  $Q$ .<sup>246</sup>

---

<sup>246</sup> It's important to note the kind of mind-independence at issue regarding the relation  $R$  between  $P$  and  $Q$ . It's possible that the relational fact  $P R Q$  is mind dependent simply because propositions are mind-dependent. For instance, if one equates propositions with thoughts then surely the existence of the propositions  $P$  and  $Q$  will be mind-dependent and, in turn, so will the existence

If entailment is construed as the upper-limit of the making probable relation then we should attribute this mind-independence to Keynesian probability as well.  $P$ 's standing in a making probable relation to  $Q$  means that  $Q$ 's truth is probable relative to  $P$ 's being true. The existence of this probabilistic connection between the truth of the two propositions is independent of anyone attitudes or beliefs about the relation between  $P$  and  $Q$ .

#### 6.4.1.2 Irreducibility

Facts about entailment are also irreducible. The fact that  $P$  entails  $Q$  isn't reducible to some fact about the frequency with which the truth of  $P$  is correlated with the truth of  $Q$ . It might happen to be the case that whenever  $P$  is true  $Q$  is also true but it still might be the case that  $P$  fails to entail  $Q$ . Moreover, a lawful connection between the kind of fact that makes  $P$  true and the kind of fact that makes  $Q$  true might exist while, nevertheless,  $P$  does not entail  $Q$ .

The best one could do in an attempt to reduce facts about entailment is to try to reduce them to facts about relative frequencies *across all possible worlds*. But even this proposed reduction will fail. The fact that  $P$  entails  $Q$  doesn't obtain *because* of the fact that across all possible worlds  $Q$  is invariably true whenever  $P$  is true. We must remember that talk about possible worlds is metaphorical; the actual world is the one and only world and we must find truth-makers for claims about possible worlds in this world.<sup>247</sup> So what fact about the actual world is the truth-maker for the claim that across all possible worlds  $Q$  is invariably true whenever  $P$  is true? The best candidate is the fact that  $P$  entails  $Q$ . Thus, if anything, the relevant fact about relative frequencies across possible worlds holds *because* of certain facts

---

of the relation fact  $pR_Q$ . What is important is that the existence of this relational fact doesn't depend on anyone *representing* such a relation to exist.

<sup>247</sup> I'm being somewhat cavalier in my dismissal of modal realism here. See Lewis (1986) for a defense of modal realism. I'll simply note that I hold the controversial view that modal realism is quite implausible. Discussing the issue, however, is tangential to my goals in this dissertation.

about entailment relations. If anything, the relevant facts about relative frequencies are reduced to facts about entailment not the other way around!

This point that the fact the  $P$  entails  $Q$  grounds the fact that for every possible world where  $P$  is true and  $Q$  is true there is no possible world where  $P$  is true and  $Q$  is false is crucial and underappreciated. If Keynesian probability relations are to be modeled on entailment then we should say something similar about all probability relations. Facts about a Keynesian probability relation between  $P$  and  $Q$  should also be expected to ground certain facts about the ratio of possible worlds where  $P$  is true and  $Q$  true to possible worlds where  $P$  is true and  $Q$  is false. *And this will be absolutely critical to solving the problem of easy justification in chapter eight.* At this point, however, we should just note the general feature of the irreducibility of Keynesian making probable relations.

The frequency theorist and reliabilist want to ground facts about epistemic support relations in contingent facts about certain frequencies across possible worlds. What is it for a belief forming process  $\Phi$  to be reliable? You might think this is reduced to the fact that  $\Phi$  produces more true beliefs than false beliefs. This, however, is inadequate. We can imagine that there is a belief-forming process that only ever produces one belief which happens to be true but is such that if it were to have been used more regularly would have produced false beliefs more often than not. As such, a frequency theory of probability and a reliabilist theory of justification will quickly appeal to the frequency with which the process produces true beliefs in nearby possible worlds. But remember that possible world talk is metaphorical. What is the truth-maker for the claim about relative frequencies in nearby possible worlds? The reliabilist is likely to appeal to certain nomological facts about causal connections in the actual world. In fact, it seems that this is the unifying feature of externalist accounts of epistemic support. What are the truth-makers for facts about safety, sensitivity, etc? Presumably they are nomological facts that obtain in the actual world. So the externalist reduces facts about epistemic support to facts about relative frequencies in nearby

possible worlds which are in turn reduced to nomological facts. And this is where a Keynesian inspired theory of epistemic support sees things differently.

The Keynesian making probable relation is modeled on entailment relations. As such, the obtaining of Keynesian probability relations, like entailment relations, would hold even if the relative frequencies in nearby possible worlds and the nomological relations that act as truth-makers for these claims failed to obtain. When a Keynesian probability relation obtains it will be true that across *all (not just nearby) possible worlds*  $Q$  will usually be true when  $P$  is true.<sup>248</sup> The truth-maker for this, however, just is the existence of the making probable relation that exists in the actual world. Facts about Keynesian probability relations are not being reduced to facts about relative frequencies across possible worlds. If there is a reduction it is going the other way. Keynesian probability relations are as irreducible and fundamental features of the world as the entailment relations that constitute the upper limit of the relation.<sup>249</sup>

#### 6.4.1.3 Internal vs. External Relations

A related point is this: entailment relations are internal relations. Internal, as opposed to external, relations are those for which the intrinsic nature and existence of the rela

---

<sup>248</sup> Henderson and Horgan (2007) and (2011) defend a view that they call transglobal reliabilism. According to this view a belief that  $P$  is justified iff it is produced by a process that is reliable in *all possible worlds*. As Fumerton (2011) points out, the difficulty for this view is that there are an infinite number of worlds where the process is reliable and an infinite number where the process is unreliable. So how can we determine whether a process is reliable *across all possible worlds*? I think we can actually make sense of the notion that Henderson and Horgan are using but only by locating a truth-maker in the actual world. What could the truth-maker be? Here I've been suggesting that just as entailment relations that obtain in the actual world make true claims about frequencies across *all possible worlds* the probability relations modeled on entailment will also make true certain claims about frequencies across all possible worlds. As such, it seems that the intelligibility of the crucial transworld reliability notion to which Henderson and Horan appeal ultimately depends on the intelligibility of the more internalist friendly construal of epistemic support I've suggested.

<sup>249</sup> I come back to these ideas in chapter eight's discussion of the problem of easy justification.

necessitates the existence of the relational fact.<sup>250</sup> Using “aRb” to signify a relational fact where two particulars *a* and *b* are related by relation R, we can make the internal/external relation distinction more precise:<sup>251</sup>

**INT** A relation R is an internal relation with regards to its relata *a* and *b* iff the natures and existence of *a* and *b* necessitate the existence of the relational fact aRb.

**EXT** A relation R is an external relation with regards to its relata *a* and *b* iff the natures and existence of *a* and *b* do not necessitate the existence of the relational fact aRb.

The relation BEING FIVE FEET APART is an external relation. Two books might be five feet apart from one another but the nature and existence of these books doesn't necessitate this fact; the two books could have existed and been ten feet apart instead. Alternatively, blackness stands in the DARKER THAN relation to yellowness simply in virtue of the natures and existence of these two properties.

Entailment is also an internal relation. If *P* entails *Q* then these two propositions have an intrinsic nature such that the mere existence of these propositions necessitates the fact that *P* entails *Q*. And if this is just the upper limit of the making probable relation then we should also construe the making probable relation as an internal relation.

However, there is an important difference between entailment and *mere* making probable relations that must be noted at this point. If *P* entails *Q* then any conjunction of propositions that includes *P* will also entail *Q*. Such is not the case with probability relations. *P* might make probable *Q* even though *P* & *R* makes probable  $\neg Q$ . But this shouldn't be taken to threaten the claim that probability relations are necessary and internal relations. We

---

<sup>250</sup> This is a claim about metaphysical not logical necessity. Giving a definition of metaphysical necessity is next to impossible but I take it that there are a host of examples. I think the best examples concern relations between determinates and determinables. I'm especially fond of color examples. Necessarily, all red things are colored. But this doesn't seem to be a *logical* necessity. We can't explain the necessity simply in terms of some sort of syntactic structure. Similarly, a certain shade of blue might necessarily be darker than a certain shade of yellow. I'm highly skeptical that this could ever be shown to be a *logical* necessity.

<sup>251</sup> The distinction between internal and external relations is *not* to be understood in terms of these relations either being internal or external to the subject's perspective.



can accept that  $P \not\leftrightarrow R$  fails to stand in an internal relation  $R$  with  $Q$  despite the fact that  $P$  stands in this relation to  $Q$ . What explains the difference is just that  $P$  has a different nature than  $P \leftrightarrow R$ . Thus, the nature and existence of the former might necessitate the existence of a relation between it and  $Q$  even though the *different* nature of the latter does not. This fact will be important for deflecting an objection to a Keynesian inspired account of epistemic support.

#### 6.4.1.4 *A Priori Probability*

Finally, that entailment is an internal relation that holds between propositions accounts for its being knowable a priori. I haven't said much about how I understand the distinction between a priori and a posteriori knowledge/justification. I ultimately think that the distinction between a priori and a posteriori justification is to be given in terms of the kind of supporting evidence of which a subject is aware. A priori justification for  $P$  is justification that is due to awareness of the facts concerning conceptual relations that support  $P$ 's truth and an appreciation of this evidential connection. Why is my justification for believing that all bachelors are unmarried a priori? Because the supporting evidence for this proposition of which I'm aware is the fact that my concept UNMARRIED is a proper part of my concept BACHELOR.

Similarly,  $P$ 's entailing  $Q$  will be due to various internal relations such as the structural relations between their constituent concepts that are essential to being the propositions that they are. I can therefore have a priori justification for believing that  $P$  entails  $Q$  in virtue of my awareness of these relational facts and my appreciation of the evidential connection between these facts and the claim that  $P$  entails  $Q$ . Notice that I might also have a posteriori justification for this proposition since I might also be aware of supporting evidence other than these conceptual relations such as a subject's testimony that  $P$  entails  $Q$ . In fact, one's justification for a proposition  $P$  could be simultaneously partially a

priori and partially a posteriori if one is both aware of conceptual facts that support  $P$  and non-conceptual facts that support  $P$ .

#### 6.4.2 A Keynesian Inspired Theory of Epistemic Support

Modeling a theory of epistemic support on Keynes' logical theory of probability is incredibly attractive. Such relations are internal *objective* truth-indicating relations that are not reducible to facts about frequencies or nomological relations. And these features are precisely what seem to be suggested by the new evil demon thought experiment. Before we get to this, however, there is a small glitch for appropriating Keynesian probability relations as an account of epistemic support. Keynesian probability relations are logical relations defined over propositions. In chapter four I argued that evidential relations hold between fact-proposition pairs. I want to suggest that the Keynesian theory can easily be modified into a theory of epistemic support between such pairs. In fact, I've already shown how this can be done (see chapter 4 section 4.2.2.3).

To explain how this notion can be used to illustrate the existence of an analogous internal relation between a fact-proposition pair I'll presuppose a correspondence theory of truth. The correspondence theory is controversial but defending it would take us too far afield—thankfully, it doesn't seem that it's likely to lose its status as a main contender for a theory of truth any time soon.  $P$  is true iff  $P$  stands in a relation of correspondence to a fact  $F$ . Correspondence between a proposition and a fact is also an internal relation; the nature and existence of the fact <all swans are white> and the proposition *that all swans are white* necessitates the relational fact that the latter corresponds to the former.

Assume that there exists a fact < $P$ > that makes true  $P$  (I'm readopting my convention of using an italicized capital letter to indicate a proposition and brackets around the same letter to pick out the fact that makes this proposition true). If  $P$  entails  $Q$  then it's absolutely impossible for  $P$  to be true while  $Q$  is false. This is an internal relation between  $P$  and  $Q$ . But  $P$  is true just when there is some fact or facts that serve as a truth-maker for  $P$ .

Thus, it's also absolutely impossible for  $\langle P \rangle$  to exist while  $P$  is false. Therefore, it's absolutely impossible for  $\langle P \rangle$  to exist while  $Q$  is false. We can refer to this relation as an entailment relation between  $\langle P \rangle$  and  $Q$ . Moreover, if both  $P$ 's corresponding to  $\langle P \rangle$  and  $P$ 's entailing  $Q$  are internal relations, then it seems obvious that the relation I've picked out between  $\langle P \rangle$  and  $Q$  is also an internal relation. We now have a candidate for the *conclusive* epistemic reason relation that identifies this as an internal relation analogous to the entailment relation between propositions.

Now we need only extend this reasoning to non-conclusive epistemic support. If  $P$  makes (Keynesian) probable  $Q$  this means that it's probable that  $Q$  is true relative to  $P$ 's truth. But what would make  $P$  true?  $P$  is true iff it stand in the internal relation of correspondence to a fact  $\langle P \rangle$ . Therefore,  $Q$ 's truth is probable relative to  $\langle P \rangle$ .

Picking out the epistemic support relation in this way we can see that it will share the features of being an objective truth-conducive internal relation that cannot be reduced to facts about frequencies or nomological relations (which are in turn truth-makers for claims about frequencies in nearby possible worlds). However, we must modify the claim that epistemic support relations are knowable a priori.

Epistemic support is a real relation that holds between a fact and a proposition. This relation only exists when its relata exist. Therefore, the existence of an epistemic support relation between a set of facts  $F$  and a proposition  $P$  depends on the existence of the fact. In order to know or have justification for believing that some epistemic support relation exist we need to have knowledge that or at least justification for believing that some fact exists. Certain facts, however, can only be known (or justifiably believed) to exist a posteriori. When this is the case, one's knowledge that an epistemic support relation exists will inherit this a posteriori nature. For instance, my putative knowledge that the fact that all previously observed swans have been white epistemically supports the claim that the next observed swan will be white depends on my knowing that such a fact obtains. Knowledge of this fact, however, can only be a posteriori.

Nevertheless, we can make a claim about a kind of a priori status that epistemic support relations have that is related to its Keynesian nature. Consider the following claim: the black ball in the bin is darker than the yellow ball in the bin. Can this be known a priori? No. In order to know this claim you need to have justification for believing that the black ball and the yellow ball exist, and this can only be had a posteriori. Insofar as this claim has existential import, it cannot be known a priori—the same can be said about the claim that all bachelors are unmarried if you think that this has existential import. However, one can still know a priori (i.e. in virtue of one's awareness of certain conceptual facts) that *if* there is both a black and a yellow ball in the bin *then* the black ball would be darker than the yellow ball. A similar conditional pertaining to epistemic support relations can be known a priori. We can know a priori that *if* a fact *F* exists *then* *F* would epistemically support *P*. I'll refer to this idea by saying that epistemic support relations are *minimally* a priori. What prevents our knowledge of epistemic support relations being *fully* a priori is simply that we cannot know a priori that the antecedent of the conditional is true; claims about epistemic support relations holding between *F* and *P* have existential import. Consider an example. Right now there is no fact that corresponds to the proposition that I'm in pain (though when discussing this passage during my dissertation defense I might be in emotional pain). Therefore, I cannot know (let alone know a priori) that the fact that I'm in pain epistemically supports the claim that someone is in pain. However, in virtue of my awareness of the Keynesian relation that holds between the proposition that I'm in pain and the proposition that someone is in pain, I can know a priori that if the fact that I'm currently in pain were to exist then this fact would epistemically support the claim that someone is currently in pain. In order to know the existence of the epistemic support relation I need only learn that the relevant fact does exist. So there is still an important kind of a priori element to knowing that certain epistemic support relation holds.

Finally, the last point to make about epistemic support is that this needs to be modeled on the notion of *relevant* entailment rather than the classical notion of entailment.

Every proposition classically entails any necessary truth. That *my name is Sam* entails that  $9 \times 9 = 81$ . That *there is a beer in the fridge* entails that *either an even number of dinosaurs existed or it's not the case that an even number of dinosaurs existed*. It's absolutely impossible for the first proposition to be true while the latter proposition is false but only because the latter proposition is necessarily true. Our notion of classical entailment doesn't capture any sort of connection between the *contents* of these propositions. And, similarly, a notion of entailment holding between fact-proposition pairs modeled on classical entailment wouldn't capture any sort of significant connection between the *factual content* and the *propositional content*. It's merely a syntactic relation.

We don't want our theory of epistemic support to have the result that the fact that I ate scrambled eggs for breakfast to epistemically support Goldbach's conjecture (assuming this is true). We also want to leave open the possibility of the existence of certain facts offering less than conclusive support for necessary truths. The fact that an expert mathematician has provided testimony that Goldbach's conjecture is true epistemically supports Goldbach's conjecture but not conclusively. Moreover, we want to make sense of the idea that different facts could offer different degrees of evidential support for necessary truths. The expert mathematician's testimony offers more epistemic support for believing Goldbach's conjecture than the novice's testimony (and the proof of Goldbach's conjecture offers even better support). In order to accommodate these ideas it seems that we must understand our notion of epistemic support between fact-proposition pairs as a kind of relevant entailment and construe *relevant* entailment as the upper limit of this epistemic support relation. The fact that I ate scrambled eggs for breakfast doesn't *relevantly* entail Goldbach's conjecture. Various facts cited in the proof of the conjecture do *relevantly* entail Goldbach's conjecture. And facts about expert testimony don't *relevantly* entail the conjecture's truth but such testimony is relevant and improves the epistemic credibility of such a proposition. Facts about novice testimony may also improve the epistemic credibility of such a proposition but they are as relevant as the expert testimony. Epistemic support

relations between fact-proposition pairs should therefore be understood as mirroring probability relations between propositions with relevant entailment as the upper limit. Modeling relevant entailment between propositions is incredibly difficult and I won't discuss the formalisms here.<sup>252</sup> In the end I think we have a grasp on the notion but such a relation as I've said is a *sui generis* and *unanalyzable* internal relation between some fact and a proposition and our concept of it is primitive. We can point to such a relation in the same way we can point to the darker than relation but, unfortunately, in the case of primitives someone awaiting an analysis in more familiar terms is going to be waiting a long time.

#### 6.4.3 What the New Evil Demon Really Shows

Let's reconsider the new evil demon thought experiment. I want to suggest that what this thought experiment actually does is support the non-reductive theory of epistemic support I've suggested.

Consider again the new evil demon thought objection to externalism. The way the objection works is as follows. We consider some fact-proposition pair such that we're inclined to describe the fact as offering epistemic support for believing the proposition, i.e. that the fact lends epistemic credibility to the proposition. When I think about the experiential state I'm in at this precise moment, it strikes me that the fact that I have an experiential state of this sort (including my entire state of visual, tactile, auditory, proprioceptive, memorial, etc. experiences) and my awareness of these facts provide me with justification for believing that there is a computer in front of me right now. Now I think to myself, "what if down the road I were to become convinced that unbeknownst to me I am victim to some skeptical scenario (such as Descartes' evil genius) that robs this *taking experience at face value* of its reliability, proper function, safety, sensitivity, etc." Reflecting on

---

<sup>252</sup> See Anderson and Belnap (1975), (1992); Mares (2004); and Priest (2008) for work on relevance logics.

this it's intuitive that even if this were to happen one should not feel threatened to retract the idea that my awareness of these facts about my experiential life provided me with justification for my belief that there was a computer in front of me. I could no longer trust my experiences but I was perfectly justified in trusting them previously.

What this shows is that our concept of epistemic support is such that these experiential facts epistemically support the claim that there is a computer in front of me independently of the obtaining of various external factors. By focusing on something other than the nature of the relevant facts and the proposition that there is a computer in front of me one isn't focusing on the kind of thing that determines epistemic support. We understand epistemic support as a necessary relation, and in order to get us to retract a claim about epistemic support one must challenge our understanding of the intrinsic nature of the experiential facts and the proposition that we suppose to stand in the relation. How much epistemic support for a believing a proposition is provided is something we can determine simply by examining the nature of the relata. In other words, our consideration of the new evil demon objection just is direct support for the idea that our concept of epistemic support is a concept of an *internal* relation that can be known in the minimally a priori sense I've indicated.

This, however, falls short of showing that subjects aware of subjectively indistinguishable experiences thereby have the same degree of justification for their beliefs. For one thing, two subjects S and S\* might be aware of subjectively indistinguishable experiences while S but not S\* is aware of the epistemic support provided by those experiences for believing *P*. We saw in the last chapter that this appreciation of the support relation makes a difference to justification.<sup>253</sup>

---

<sup>253</sup> Unlike many internalists I do *not* think a consideration of new evil demon cases suffices for showing that subjects in ordinary and demon worlds with subjectively indistinguishable experiences thereby have the same degree of justification for their perceptual beliefs. One reason has been given in the main text. Another reason is that I think such a claim is true but it requires independent arguments against metaphysical direct realism and disjunctivism about perceptual

### 6.5 Objections to the Keynesian Approach

Up to this point I have developed and argued in favor of a non-reductive theory of the epistemic support relation that models epistemic support on Keynesian probability relations. Defending such a view mainly comes via negative support by pointing to the difficulties with other views. After providing intuitive examples of fact-proposition pairs that either stand or fail to stand in epistemic support relations I turned to the difficulties with subjectivist and externalist accounts of this relation. Subjectivist views led to a kind of epistemological anarchy that is inconsistent with how we think about epistemic support. Externalist views do better in this regard since they make the appropriateness of our epistemic practices an objective and mind-independent matter. The difficulty for such views is that they cannot accommodate the idea that certain facts that intuitively offer epistemic support for various claims still do so even if certain skeptical hypotheses were to obtain.

The non-reductive theory of epistemic support avoids both of these worries. By modeling epistemic support on a kind of Keynesian probability with entailment as the upper limit we avoid the epistemological anarchy of subjectivism. That  $P$  entails  $Q$  is an objective and mind-independent matter. A fact about  $P$ 's entailing  $Q$  holds independently of anyone's entertaining a representation of such a relation. Moreover, such a view avoids the new evil

---

experience. The new evil demon objection presumes that subjects aware of subjectively indistinguishable experiential states are aware of the same experiential facts. Presuming that this holds we realize that these facts would offer just as much epistemic support for various claims whether or not various skeptical hypotheses held. We therefore conclude that epistemic support consists of a necessary internal relation between a set of facts and propositions. In order to show that subjects who are aware of subjectively indistinguishable experiences have the same degree of justification we would have to vindicate the claim that such subjects are aware of the same experiential facts. This assumption has been challenged by metaphysical direct realists and disjunctivists about perception. According to these views one can be in an experiential state that includes a physical object as a constituent in successful cases of perception and hallucinatory experiences share no common element with such perceptual experiences. However, I ultimately do reject metaphysical direct realism on the grounds that I think the best explanation of the subjective indistinguishability of perception and hallucination is that they share a common core of which we're aware. But this goes beyond what the new evil demon objection shows. My idiosyncratic views on this, however, aren't very important and would take us too far afield.



demon objection since these relations are irreducible to contingent facts about relative frequencies or nomological relations that obtain in this or nearby possible worlds. All that remains for defending the non-reductive view is to parry various objections. However, my discussion here will be admittedly brief and much more could be said.

OBJECTION 1: The non-reductive view of epistemic support modeled on Keynesian probability makes epistemic support relations it a necessary and essential property of some fact that it is evidence for *P*. However, certain evidential relations are species relative and therefore contingent. Consider, for instance, Michael Bergmann's billiard-ball creatures. Bergmann (2006) imagines creatures such that were purposefully designed in a way such that when they are holding a billiard ball in their hands this *reliably* causes in them the same kind of olfactory experience humans have when they are smelling a meadow full of flowers. Moreover, his imagined creatures have an *unlearned* seeming that such an olfactory experience reliably indicates that they have a smallish round hard object in their hand (and of course this seeming is true in the environment for which they were designed).<sup>254</sup>

Bergmann suggests that for these kinds of creatures, forming the belief that there is a smallish round hard object in their hand would be the epistemically appropriate response to the olfactory experiences humans usually have while smelling a meadow full of flowers. Clearly, however, such a response to these olfactory experiences would be epistemically inappropriate for humans. Bergmann doesn't put it this way but the claim seems to be that the olfactory experience constitutes evidence that epistemically supports the claim that there is a smallish round hard object in one's hand *for this imagined species* but doesn't support this claim *for humans*. In which case epistemic support relations are not necessary internal relations holding between a fact-proposition pair as required by the non-reductive view.

---

<sup>254</sup> See Bergmann (2006) chapter 5.

Epistemic support relations depend on contingent facts about the species of the subject who is aware of the evidence.

REPLY: My reply to this “species-relative support” objection is two-fold. First, I want to mention that I just don’t have the intuition that Bergmann has. I don’t have the intuition that forming the belief that there is a smallish round hard object in one’s hand is the *epistemically* appropriate response to the olfactory experience even for the cognizers he imagines. I think that there is an intimate and essential connection between the content of the experiences normal humans have and the properties attributed to the billiard ball in the thought that there is a smallish round hard object in one’s hand such that the former provides epistemic support for the latter. When I have the experience as of pressure when squeezing the billiard ball and the proprioceptive experience as of my hand making a spherical shape the properties of my experience are intimately tied to the properties attributed in my thought. This same connection between the properties of the olfactory experience and the properties attributed in the same thought doesn’t intuitively hold. And so I have a clash of intuitions with Bergmann, even for such cognizers it strikes me that the olfactory experiences wouldn’t epistemically support the claim that there is a round hard object in one’s hand. Olfactory experiences of this sort could only constitute evidence for claims about various objects instantiating primary properties when taken in conjunction with inductive evidence correlating such experiences with the small round hard objects being in one’s hand.

However, even if one shares Bergmann’s intuitions about this odd species, one can hold on to the claim that epistemic support relations are necessary internal relations between fact-proposition pairs. We need to distinguish between a notion of ‘species-relative evidential *support*’ and ‘species relative evidence’. Even if evidential support isn’t species relative surely the kind of evidence of which different species are aware is different. For instance, I see no reason to deny that creatures with different sensory modalities might be aware of certain kinds of facts that humans aren’t but that still *necessarily* make probable various claims. Bats

sonar for instance might give rise to certain kinds of experiential facts that epistemically support various claims about the objects around them even though humans will never be aware of this evidence. This doesn't show that evidential support is species relative, it just shows that what kind of evidence a subject is aware of might be species relative. Similarly, if you share Bergmann's intuitions about the odd species he imagines, one could admit that such a species has evidence of which we're unaware. The catch is that even for the odd species he considers, the fact that they have the olfactory experience doesn't epistemically support the claim that there is a smallish hard round object in one's hand. Adding the unlearned seeming that such experiences are reliable indicators doesn't thereby turn the experience into evidence for such a claim. Adding the experience adds new evidence. My view of epistemic support can claim that it's an essential feature of the conjunction of the fact that one has the olfactory experience with the fact that one has the relevant unlearned seeming epistemically supports the claim that one has a smallish hard round object in one's hand. In this regard it's important to remember that on the Keynesian inspired view a fact  $F$  might fail to epistemically support  $Q$  even though the set of facts  $\{F \text{ and } F^*\}$  does support  $Q$ . Adding the fact  $F^*$  doesn't thereby make  $F$  into evidence for  $Q$ ;  $F^*$  is one of the relata of the evidential support relation. Similarly adding the unlearned seeming doesn't turn the olfactory experiences into evidence for the corresponding claim about a billiard ball. Rather, the unlearned seeming is itself part of the relata of an evidential support relation.

OBJECTION 2: Consider the case of litmus paper. Surely the fact that the litmus paper turned pink when submerged in the solution is evidence for the claim that the solution is acidic. But this evidential relation cannot be known a priori even in my minimal sense. We can only discover such a relation a posteriori. Moreover, there are plenty of these kinds of evidential relations that we discover through science all the time.

REPLY: First, I've admitted that there are different categories of evidence that might be distinguished by the kind of truth-indicating relation they stand in to various claims. On a natural understanding of the concept of scientific evidence it consists of facts that are a

reliably correlated with the truth of some proposition. As a matter of fact there is a reliable correlation between facts like that a solution turned litmus paper turned pink and the truth of the proposition that the solution is acidic. Thus, I can admit that this is scientific evidence for the relevant claim. And talking about this kind of evidence is surely useful in scientific contexts and in various other contexts.

Nevertheless, what we've been interested in is the nature of the truth-indicating relation constitutive of justificatory evidence. We been investigating the kind of relation a fact must stand in in order to provide epistemic reason for believing a proposition. But doesn't the fact that the litmus paper turned pink constitute a reason to believe that the solution is acidic and therefore justificatory evidence for this claim? No. Various types of reasoning are enthymematic. The fact about the litmus paper is a salient part of the facts that constitute evidence that a solution is pink. Its salience is due to the fact that we all take for granted that most people have evidence that there is a reliable correlation between litmus paper turning pink and its being acidic. We assume that most people are either aware of other facts such as testimony from a 5<sup>th</sup> grade teacher, a science textbook, performing previous experiments, etc. that support this claim about the reliable correlation. What we don't assume is that someone has in fact submerged a piece of litmus paper in *this* solution. So when we ask for a person's evidence that *this* solution is acidic they cite what is a very salient piece of a larger body of total evidence. Nevertheless, none of us think that the fact that litmus paper turned pink is an epistemic reason to believe that the solution is acidic *all by itself*. It's only in conjunction with a whole host of other facts that this constitutes a reason for belief.

OBJECTION 3: Presumably, much of the value of justification derives from its connection to truth. We desire justified beliefs because forming justified beliefs is intimately tied to getting at the truth. Epistemic justification is a function of forming beliefs on the basis of evidence that epistemically supports those beliefs. You, however, get out of the new evil demon problem by stripping epistemic support relations of any empirical content (other

than the existence of the relata that is). This seems to be a double edged sword. Doesn't this make the truth-connection as mysterious on your view as on the subjectivist view? On your view, a subject could possess a high degree of epistemic support for all her beliefs about contingent truths even though *all* of these beliefs are false. So why should we even care about forming epistemically justified beliefs? It seems we would do better to stick with the externalist notion of evidential support (i.e. RST) and thereby gain a better truth-connection.

REPLY: My reply to this objection is going to end up anticipating some of the claims I'll make in my resolution of the problem of easy justification. Despite the fact that there is no *necessary* connection between forming justified beliefs and forming mostly true beliefs on my theory of epistemic support it is still true that forming justified beliefs is likely to be truth-conducive.

But what does this mean? The objector might complain that the claim that justified beliefs are likely to be truth-conducive appears to be code for the claim that these are epistemically likely to be true. And this is just code for the claim that they're justified. But then this truth-connection appears utterly trivial and no better than the kind of truth-connection that subjectivist theories of epistemic support can get.

I, however, think I can get something more substantial. Remember our discussion that the fact that *P* entails *Q* itself entails that *P*'s truth is perfectly reliably correlated with *Q*'s truth across all possible worlds. The fact about entailment itself entails a reliability claim. In chapter 8 I'll argue that the solution to the problem of easy justification is that something similar holds in the case of Keynesian probability relations. The fact that *P* makes probable *Q* itself makes probable (but does not entail) that *P* is reliably correlated with *Q*. The existence of a Keynesian probability relation makes Keynesian probable the existence of a statistical probability relation. In this way we can garner a more substantial truth-connection. Justified beliefs are *epistemically likely* to be statistically likely to be true relative to one's evidence.

Of course it's still *possible* that all of a subject's beliefs are entirely justified beliefs and is a demon victim. It's still *possible* that justified beliefs are statistically likely to be false. If you

want a closer truth-connection I can't give it to you. At this point, however, I would make a kind of partners in guilt response. It isn't clear that externalist theories can get a much more intimate truth-connection either. Consider reliabilism again. Surely we don't want to characterize reliability in terms of *actual* frequencies. Reliability will be characterized in terms of *counterfactual* frequencies so as to deal with the possibilities of a type of process that is only used once, delivers a true beliefs, but would have produced mostly false beliefs had it been used more often. However, as soon as one goes counterfactual one must admit *the possibility* that a subject form all justified beliefs that are, nevertheless, all false (or at least all the beliefs concerning contingent truths).

## CHAPTER SEVEN

### AN ACQUAINTANCE THEORY

Evidence for (or against) *P* consists of facts. Justificatory evidence for (or against) *P* consists of facts that stand in a *sui generis* and *unanalyzable* probability relation to *P*. Finally, a set of facts *F* can only play a justificatory role for a subject *S* relative to *P* if *S* has connected possession of evidence *E* for *P*—i.e. *S* is aware of *F* and aware of the connection between *F* and *P*. Now we need to bring together these insights into a full theory of justification.

I begin by presenting and motivating a contemporary theory of justification that respects these requirements concerning the relations between justification, awareness, and the relevant evidential concepts: Richard Fumerton's acquaintance theory of foundational justification and his corresponding version of inferential internalism. At this point in our investigation all the pieces have been put in place so as to apply this theory to a resolution of the puzzle of justification presented in chapter 2. However, I'll argue that Fumerton's accounts of foundational and inferential justification face a variety of objections and so applying the theory to our puzzle would be in vain if these objections cannot be overcome. I therefore show how we can modify Fumerton's accounts of foundational and inferential justification in order to arrive at a more plausible epistemology without losing its initial motivations. In the last two sections of the chapter I respond to objections to my modified version of the acquaintance theory. This will leave us with the most defensible version of the acquaintance theory and thereby place us in the best position for utilizing the theory to solve our original puzzle in chapter 8.

#### 7.1 Fumerton's Acquaintance Theory

Richard Fumerton has recently defended an acquaintance theory of foundational justification. According to Fumerton, a subject *S* has foundational justification for believing a proposition *P* iff (i) *S* is directly acquainted with the fact that *P*, (ii) *S* is directly acquainted

with the thought that *P*, and (iii) *S* is directly acquainted with the correspondence between the thought and the fact.

Such an account satisfies the desiderata set down by our discussions of evidence, evidence possession, evidential support, awareness and their relation to justification. In order to understand how Fumerton's account satisfies the constraints set on a theory of justification produced by our discussions in chapters 4-6 we need to investigate the crucial concepts of acquaintance and correspondence.

### 7.1.1 Acquaintance

Fumerton offers a nice starting point for discussing acquaintance:

Acquaintance is *not* another intentional state ... Acquaintance is a *sui generis relation* that holds between a self and a thing, property, or fact. To be acquainted with a fact is not *by itself* to have any kind of propositional knowledge or justified belief... One can be acquainted with a property or fact without even possessing the conceptual resources to *represent* that fact in thought, and certainly without possessing the ability to linguistically express that fact.<sup>255</sup>

This passage identifies four important features of acquaintance: (i) acquaintance is not an intentional state; (ii) acquaintance is not a representational state; (iii) acquaintance is not an epistemic state; and (iv) acquaintance is *sui generis*.

The claim that acquaintance is non-intentional is not meant to entail that acquaintance isn't directed at an object. Acquaintance with the fact that I'm in pain is a kind of awareness I have that is directed at this fact. Fumerton means to contrast acquaintance with states like fear, belief, and desire. We would naturally describe these states as being *directed* at an object. My fear of spiders is directed at spiders. But what we might identify as the object of these intentional states needn't actually exist. I can fear ghosts, believe that unicorns live in Scotland, or desire drinking from the fountain of youth. We want to admit the existence of these states without being committed to the view that ghosts, unicorns, or

---

<sup>255</sup> Fumerton (1995), p. 74.



the fountain of youth exist. Acquaintance with a state, property, or fact, however, isn't like these other states. Acquaintance is a relation that requires the existence of its relata. Acquaintance is a *factive* awareness; it cannot be directed at an entity that fails to exist.

Related is the point that acquaintance isn't a representational state. Acquaintance with a property, state, or fact isn't mediated by one's entertaining a proposition or representation of the relevant property, state, or fact. Acquaintance is the relation we stand in to things when they're directly *given* to consciousness. It's likely that young children and even unsophisticated animals could be acquainted with a red patch in their visual fields or with the fact that they're experience extreme pain despite their lacking the conceptual resources for representing these in thought.

This leads to the third feature. If acquaintance doesn't involve the subject's entertaining any kind of representation that can be assessed as true/false or accurate/inaccurate then it doesn't appear that it's any kind of epistemic relation. Being acquainted with a fact isn't to know a truth nor is it to have a justified belief about that fact.

Our characterization of acquaintance has only told us what acquaintance isn't. You might hope for a more positive characterization, but Fumerton insists that acquaintance is *sui generis*. Acquaintance is primitive and can't be given an analysis in any more familiar terms. Some philosophers might find the notion utterly mysterious.

Nevertheless, I think acquaintance is something we're all incredibly familiar with, and we can help someone understand the notion of acquaintance by attempting to "point" to various acts of acquaintance. I can't provide an analysis of phenomenal pain but if someone is wondering what I mean by "pain" I might get them to understand by kicking them really hard in the shin, asking them if they noticed any difference in their experience after I kicked them, and explaining that this change is what I mean by pain. Similarly, an advocate of an acquaintance theory might ask the same person to focus their attention on the very intimate relation they stand in to their own pain, visual experiences, beliefs, etc. And in this way we can help someone gain a concept of acquaintance despite its being *sui generis*.

I prefer the view that acquaintance is a kind of non-conceptual awareness that grounds our ability to selectively attend to or “mentally point” to various states, experiences, properties, features, facts, etc. Consider two ways we can become doxastically aware of a feature in the periphery of our visual fields. You might come to a justified belief about a feature’s presence on the basis of a neuroscientist’s testimony that some configuration of neurons is firing which has been correlated with such a feature being present in the periphery of people’s visual field. Needless to say, *this* way (and more realistic but analogous ways) of becoming doxastically aware of an experiential feature is rare. More often than not we become doxastically aware of the feature by focusing our mental attention on or mentally pointing to the feature (a kind of *non-conceptual* noticing) in our periphery and thereby form the belief that the feature is present. The ability to attend to a state, experience, property, or fact in this non-conceptual sense is how I think we ought to understand acquaintance.

This explains the idea that normal humans can be acquainted with the fact that they have an experience as of a 3-speckled hen, and perhaps infallible foundational justification for the corresponding belief, but not with the fact that they have an experience as of a 47-speckled hen.<sup>256</sup> Even assuming both of these features *in some sense* show up in consciousness, normal humans have the ability to attend to the feature of 3-speckledness but not to the feature of 47-speckledness. Consider figure 4:<sup>257</sup>

---

<sup>256</sup> See Chisholm (1942), Feldman (2004); Fumerton (2005); Markie (2009); Poston (2007); and Sosa (2003) for discussions of the problem of the speckled hen. It’s interesting to note that empirical evidence, see Pylyshyn and Strom (1988), suggest that subjects can simultaneously track up to five dots that are moving in random paths. This appears to track our intuitive judgments pertaining to which beliefs can be infallibility foundationally justified. It’s easy to identify 3, 4, or 5 speckles/dots, and a degree of confidence approaching certainty seems appropriate in these cases. Around 7, 8, or 9 speckles or dots it gets more difficult to identify the feature (without counting). In so far as one is convinced that one is acquainted with the determinate fact in the former cases but not in the latter I think this offers good reason to identify acquaintance with a kind of attentive awareness.

<sup>257</sup> This figure is inspired by and loosely modeled on a figure of lines from Tye (2010).




Figure 4: Attending to Dots

While focusing on the plus sign of figure 4 you're able to mentally ostend each of the three dots in your experience to the left. In fact, it strikes me that I can simultaneously attend to each of the dots individually. The latter fact seems to be necessary for attending to my experience being as of three dots to the left of the plus sign. Thus, if we identify acquaintance with the ability to mentally ostend then this would explain why I can be acquainted with the fact that my experience is as of three dots to the left of the plus sign. One of the crucial conditions for foundational justification is in place. However, you likely cannot mentally ostend to your experience as of a dot in the figure to the right of the plus sign that is 4 columns to the right and 5 rows down. Even if you change your focus and find the dot so you know which dot you're trying to ostend, when you go back to the plus sign you likely cannot mentally attend to this particular dot. In this case it seems you can only attend to the group as a whole. You can attend to your experience being as of many dots but you cannot simultaneously attend to each dot. You cannot attend to your experience being as of 36 dots to the right of the plus sign. Thus, one of the critical conditions for foundational justification for believing you're having an experience as of 36 dots is absent.

What is important here is that, when the reader focuses on the plus sign in figure 4 and tries to mentally point to various features of this experience, this activity can serve to help understand the critical concept of acquaintance. One is aware of each individual dot to the left of the plus sign in a way that one isn't aware of each individual dot to the right of the

plus sign. Reflecting on the difference in your experience can draw your thought to the critical relation of acquaintance. You're acquainted with each of the dots to the left of the plus sign in your experience but not so in the case of your experience of the dots to the right. This is meant to illustrate a phenomenon that is familiar to all of us and thereby make the notion of acquaintance less mysterious.<sup>258</sup> Nevertheless, I can't give a definition of acquaintance in the same way that I can't give a definition of phenomenal pain. This, however, doesn't mean that the notion is mysterious. I take it that we're all incredibly familiar with this notion of attentive awareness in a way that makes acquaintance understandable. The previous discussion is meant to make this familiarity explicit and give phenomenological motivation for positing the existence of the acquaintance relation.<sup>259</sup>

---

<sup>258</sup> It's important to be clear here what I'm actually trying to do. The considerations above about the direct givenness of searing pain and the exercise involved in considering one's experience when looking at figure 4 are attempts to put the reader in a state of being acquainted with acquaintance. It's this state that will provide reader with an understanding of the critical notion. Moreover, if I'm asked why I believe that there is such a relation as acquaintance I'll give the unhelpful answer that I'm acquainted with it. Such an answer is obvious question begging but, as Fumerton 1995 himself points out, given the theory you wouldn't expect an acquaintance theorist to say anything else. Remember, however, our lesson from Peter Markie. Begging the question is a dialectical failure and this doesn't necessarily imply any epistemological failure. If the acquaintance theory is the correct theory of justification then acquaintance with acquaintance would provide the best justification you could get for believing that such a relation exists (provided that you meet the other conditions on justification as well).

<sup>259</sup> We might also give theoretical motivation for positing the existence of the acquaintance relation. First, I've already mentioned that it seems that unsophisticated people and even animals could have a kind of awareness of a red patch in their phenomenal field or of the fact that they're in pain even if they lack the conceptual resources to represent these in thought. Second, various philosophers have been attracted to the idea that we can think about phenomenal features using what Chalmers (2003) calls a direct phenomenal concept. It seems I can have a concept of phenomenal redness where the content of this concept is *exhausted* by the phenomenal character of redness. I can think about using a concept that contains no information beyond the phenomenal character itself. I needn't think about phenomenal redness via some sort of indirect description such as the kind of experience that people or even I usually experience when looking at Washington apples. How do I acquire this *phenomenal* concept of phenomenal redness? Surely it isn't enough that I just have an experience as of phenomenal redness. A person could have the experience without forming the concept of that property. It seems that acquiring the phenomenal concept will require that the subject have a phenomenally red experience and then somehow picks this feature out. Clearly actual pointing can't do the work here but a kind of mental attention could play this theoretical role, giving further theoretical motivation for positing an acquaintance relation. A final theoretical motivation for positing acquaintance is simply the role it can play in preventing the epistemological regress considered below.

### 7.1.2 Correspondence

Fumerton's acquaintance theory is also developed within a controversial correspondence theory of truth. I've already indicated my acceptance and conception of the correspondence theory and so my comments here will be brief. A proposition is a psychological entity that is to be identified with a *type* (not a token) of thought. Thoughts are non-relational properties of the mind that represent the world as being some way or another. Thoughts are the primary bearers of truth-value. A thought is true when it corresponds to or "pictures" a fact and is false otherwise.

Correspondence, like acquaintance, is a *sui generis* relation. Thus, there isn't much more we can say about the nature of correspondence. We can, however, try to "point" to the relation by considering examples. Consider my true thought that there is a computer two feet in front of me. Call this my computer thought. Now consider my false thought that there is a trampoline in this room. Call this my trampoline thought. My trampoline thought fails to stand in the kind of relation to the world that my computer thought does. My computer thought stands in a special relation to a fact consisting of various worldly entities (i.e. the computer and myself) that exemplify a certain special relation (i.e. being two feet apart) and is thereby true. This special relation is the correspondence relation.<sup>260</sup>

### 7.1.3 The Epistemic Regress Argument for Foundationalism

Before considering the reasons for adopting Fumerton's theory of foundational justification we must first consider the motivation for accepting foundationalism.

---

<sup>260</sup> While this correspondence theory has been challenged in various ways it still remains a dominant (if not *the* dominant) theory of the nature of truth amongst philosophers I accept such a theory as correct. Nevertheless, it's interesting to note that my modifications to the acquaintance theory will only rely on the metaphysics of fact-proposition evidential relations outlined in chapter 4 and 6. As such, my version of the acquaintance theory will be at least consistent with rejecting a correspondence theory of truth.

It's natural to think that many of our supposedly justified beliefs are based on and justified by our further beliefs. Part of what provides me with justification for believing my fiancé will pick me up from work later is my belief that, when she dropped me off at work this morning, she told me that she would pick me up. My justification for the former belief is partly constituted by my having the latter belief. Part of my justification for believing that it snowed last night is my belief that the roads are covered in fresh snow. My justification for the former belief is partly the latter belief. Each of these cases illustrates the notion of *inferential* justification where my having justification for a certain belief is partly constituted by my having some further belief.

However, we don't think that I have inferential justification for believing that it snowed last night *merely* in virtue of my having the further belief that the roads are covered in fresh snow. This latter belief might be completely irrational and unjustified. I might believe that the roads are covered in fresh snow because my Ouija board said as much. If so, then this belief doesn't provide me with any justification for believing that it snowed last night.

Our discussion of the subject's perspective objection (SPO) in chapter 5 provides a deeper explanation for why one must have justification for believing that *P* in order to have inferential justification for believing the further proposition *Q* on the basis of *P*. We discussed the idea that a subject *S*'s having justification for believing a proposition requires that the truth of that proposition not be an accident from *S*'s own perspective. But "accidentality" can be inherited. That a proposition *P* makes probable *Q* means that *Q*'s truth is probable relative to *the truth of P*. It is *P*'s truth that makes *Q*'s truth non-accidental. If *P*'s truth is nothing more than a cosmic accident from my perspective then, even if I properly infer *Q* from *P*, *Q*'s truth is also nothing more than an accident from my perspective. I therefore lack justification for believing *Q* (at least on the basis of my belief that *P*).

All of this motivates the following requirement on inferential justification (IJ1):

**IJ1** S has inferential justification for believing that  $Q$  on the basis of S's further belief that  $P$  only if S has justification for believing  $P$ .

IJ1 is a rather uncontroversial principle accepted by epistemologists with wildly different theories. Process reliabilists, proper functionalists, evidentialists, mentalists, access internalists, coherentists, infinitists, advocates of each of these views will likely accept IJ1.

IJ1 is the cornerstone of an argument for foundationalism. If all justification were inferential then having justification for a belief B1 requires that I have a further belief B2. According to IJ1, B2 can only provide justification for B1 if I have justification for B2. Having justification for B2 in turn requires that I have some further belief B3. Once again, IJ1 entails that I must have justification for B3. This continues *ad infinitum*. Unfortunately people can't complete an infinitely long chain of reasoning. So, if all justification were inferential, no one would have any justification for believing anything at all. But such a radical form of skepticism is absurd. Therefore, there must be some beliefs whose justification is not (even partly) constituted by the subject's having further beliefs. In other words, there must be foundational justification.

This epistemic regress argument for foundationalism is far from uncontroversial. A skeptic might actually use the potential regress of justification as an argument for skepticism rather than foundationalism. Such a skeptic will claim that the so-called "absurdity" of radical skepticism isn't a good reason to reject this regress argument for skepticism.

One element of "absurdity" that philosophers often point to is that the regress would lead to a radical skepticism such that, if true, no one has justification for accepting the conclusion (nor the premises of the argument leading to the conclusion). The belief in radical skepticism is necessarily unjustified!

In chapter 2 I argued that these aren't adequate grounds for dismissing a skeptical argument but I also showed sympathy for the Moorean strategy that claims that the proposition *I'm justified in believing I have hands* or that *I'm justified in believing I'm in pain* will always be more plausible than the skeptic's abstract epistemological principles. Such a

strategy, however, only claims to give us reason to think that the argument must go wrong somewhere. A philosophically satisfying response requires that we identify where the argument goes wrong, why it goes wrong there, and that we have good reasons for making these claims. Only this kind of diagnosis of the argument will lead us toward the kind of understanding we seek in philosophy.

Foundationalism takes it that the regress argument for skepticism goes wrong in its initial assumption that all justification is inferential. Such a move seems plausible when we consider some examples of what are often considered paradigmatic foundationally justified beliefs. Consider your belief that  $2+2=4$ . Presumably you're justified in this belief but it doesn't seem as if your justification requires your *inferring* it from some further belief. Rather, your justification for believing that  $2+2=4$  seems to arise merely from your understanding the concepts involved. Even more convincing are various introspective beliefs. Consider a case where you believe that you're in pain while undergoing an experience of searing pain. Such a belief appears to possess the best justification possible but you needn't infer this from some further belief. These kinds of examples give prima facie reason to accept the foundationalist's identification of where the regress argument for radical skepticism goes wrong. It's incumbent on the foundationalist, however, to explain how this foundational justification is possible. What distinguishes your belief that  $2+2=4$  and your belief that you're in pain (while experiencing searing pain) from merely arbitrary beliefs for which you lack any support? I'll address this question in 7.1.5. First, let's briefly consider some anti-foundationalist responses to the epistemic regress.

Implicit in the regress argument is a commitment to the idea that justification is *linear*, i.e. in order for my belief B1 to be justified I must have prior justification for believing B2. Coherentists reject this assumption and adopt the metaphor of a web of belief for describing the structure of justification.<sup>261</sup> Justification is due to a belief's being a member

---

<sup>261</sup> See Bonjour (1985) for a paradigmatic statement of the coherence theory.



in a set of beliefs that cohere with one another. Justification for any belief depends on a subject's having justification for further beliefs but this doesn't require that the subject complete an infinitely long chain of reasoning. It isn't as if there are beliefs that are *first* justified and this justification then *transmits* to the further beliefs. Rather, justification for an entire set of beliefs arises simultaneously in virtue of their coherence with one another.

Alternatively, Peter Klein has attempted to defend the idea that the regress initiated by IJ1 and the claim that all justification is inferential isn't vicious. Klein develops a version of infinitism that attempts to fend off the skeptical results by requiring that the subject S has the *capacity* to defend her belief via an infinite chain of further beliefs rather than that the subject *actually* complete that infinite chain.<sup>262</sup>

Objections to coherentism and infinitism are well known. Because I don't have anything particularly new to add to this discussion I'll only briefly outline what I take to be the most pressing objections and assume that these objections are cogent when developed properly. And so foundationalism will be the only option left for parrying the skeptical result threatened by the epistemic regress.

Let's consider some well-known difficulties with coherentism. First, there are difficulties in even giving an account of this critical concept of coherence. It seems that coherence would minimally require logical consistency. However, there seems to be clear examples of justified but inconsistent sets of beliefs. I can be justified in believing of each individual ticket in the lottery that it will lose and justified in believing that not all of the tickets will lose even though these beliefs are inconsistent.<sup>263</sup>

Second, coherentism faces what has been dubbed the "isolation" objection. It doesn't appear that mere coherence is enough to make our beliefs likely to be true. Rather,

---

<sup>262</sup> See Klein (1999), (2000), (2005), and (2007).

<sup>263</sup> See Foley (1979) for this kind of point.

there needs to be some sort of input from the world that makes it likely that our coherent set of beliefs is more than mere confabulation and represents the world as it is.<sup>264</sup> Bonjour 1985 famously responded to this objection by explaining that the world constrains our beliefs by acting as a causal input to our belief system. The problem with such a response is that the isolation objection isn't that a coherent system of belief might be *causally* isolated from the world but *rationally* isolated. Causal laws could be set up such that my experience of searing pain causes a remarkably coherent set of beliefs that includes the belief that I'm not in pain but euphoric pleasure. Here a coherent system of beliefs would be causally responsive to the world but what we need in order to ensure that a coherent system is likely to represent the world as it is some sort of *rational* constraint from the world.

Finally, there is the problem that coherence can be gained too cheaply. This was briefly discussed in Chapter 3 in connection with Sosa's particular notion of cross-level coherence. For instance, imagine I start with an unjustified belief  $P$ . Remembering some elementary truths about logic I can easily arrive at a set of beliefs such that every belief is actually *entailed* by the remaining members of the set. Imagine that I'm caused to form the unrelated beliefs that  $P$  and that  $Q$ . According to coherence theories,  $P$  and  $Q$  are unjustified for me since they're unrelated and haven't been incorporated into a coherent set of beliefs. But as Richard Fumerton (2006b) points out, by simply using the rules of elementary logic I can easily construct set of beliefs that mutually entail one another by using the following procedure: form the beliefs that  $(P \text{ or } \neg Q)$  and that  $(Q \text{ or } \neg P)$ .  $P$  entails the first disjunction.  $Q$  entails the second disjunction.  $Q$  and the first disjunction jointly entail  $P$ . Finally,  $P$  and the second disjunction jointly entail  $Q$ . Thus, I arrive at a minimally coherent set of beliefs and therefore gain a miniscule amount of justification. But this gain seems illicit. Coherence gained *in this way* seems to be epistemically worthless.

---

<sup>264</sup> The cogency of the isolation objection depends on rejecting a coherence theory of truth in favor of the correspondence theory.

The biggest concern with infinitism is that the capacity to appeal to an infinite chain of propositions in defense of  $Q$  won't be sufficient for conferring justification on a belief that  $Q$ . Consider a subject  $S$  who believes that there are three invisible gnomes dancing around outside the window. Now imagine that I ask  $S$  why she has this belief and she tells me she defends her belief by appealing to her further belief that there are four invisible gnomes dancing around outside the window. I then ask  $S$  what reasons she has for this belief. She defends this further belief by appealing to yet a further belief that there are five invisible gnomes dancing around outside. Imagine that  $S$  has the disposition and ability to continue this way forever. Surely we don't want to say that this is sufficient for  $S$ 's having justification for believing that there are three invisible gnomes dancing around outside. It's only if the premise beliefs have some positive epistemic status that the infinite chain would justify  $S$ 's belief. An infinitist needs some way of distinguishing the infinite chains that do and those that do not confer justification, and this turns out to be a daunting task.<sup>265</sup>

Such difficulties with both coherentism and infinitism suggest that foundationalism is our best bet for fending off the skeptical results of the epistemic regress. However, while I think presenting these more formal objections to these alternative ways of construing the structure of justification are *dialectically* important, I think the main reason for rejecting them in favor of foundationalism is simply the intuitive plausibility of foundationalism. Consider your belief that you're in pain formed while you're in searing pain. It seems obvious that your justification derives somehow *directly* from the match between the content of your experiential state and your belief. Your belief that you're in pain needn't be incorporated into a wider set of coherent beliefs in order to be justified. Nor do you need the ability to continue giving an infinite chain of reasons supporting your belief that you're in pain.<sup>266</sup>

---

<sup>265</sup> See Moser (1991) for a development of an objection to infinitism along these lines.

<sup>266</sup> Similar remarks apply to one's beliefs about what one believes. A case of perceptual belief might appear more controversial. In the case of a perceptual belief it might strike you that the convergence (a kind of quasi-coherence) among your different perceptual modalities is important for

#### 7.1.4 The Conceptual Regress Argument for Foundationalism

The epistemic regress has traditionally been the most discussed argument for foundationalism. Richard Fumerton, however, has suggested that there is another type of regress that can be used to motivate foundationalism. It's tempting to take IJ1 to be analytically true. Our concept of inferential justification just is the concept of a belief whose justification is partly constituted by the subject's having *justification* for believing some intermediary proposition. But this suggests that we couldn't even have the concept of inferential justification without also having a concept of foundational justification.

The foundationalist has an incredibly plausible way of preventing the conceptual regress threatened by IJ1. We can model our analysis of justification on the familiar notion of a recursive definition. An account of foundational justification will serve as a base clause that doesn't utilize epistemic concepts and then IJ1 will constitute at least part of a recursive clause in our analysis of justification.

#### 7.1.5 Acquaintance, Evidence, and Foundational Justification

Given the admission that there is foundational justification, how is such justification possible? Why aren't these "foundations" merely arbitrary beliefs? Why makes the truth of these beliefs anything more than just a cosmic accident?

One answer is that foundational beliefs are justified but not in virtue of the subject's being aware of any supporting evidence. Foundational beliefs are justified simply because they were formed in a way that makes them objectively probable and thereby distinguishes them from merely arbitrary beliefs. This is how the radical externalist views discussed in chapter 5 stop the epistemic regress. According to these theories there are beliefs that are

---

your justification. However, it seems that it isn't the coherence that constitutes the epistemic support relation itself. Rather, the fact that the different perceptual modalities provide coherent information about a chair being in front of you is part of your evidence that makes probable that there is a chair in front of you.

foundationally justified in virtue of the belief-forming processes possessing certain properties (e.g. being reliable, functioning properly, being intellectually virtuous, etc.) of which the subject needn't be in any way aware.

Appealing to the reliability, proper functioning, etc. of a belief-forming process makes it look as if foundational beliefs are non-arbitrary. This, however, is an artifact of our moving to a 3<sup>rd</sup>-person perspective on these beliefs. Such beliefs might still be arbitrary from the 1<sup>st</sup>-person perspective. As I argued in chapter 5, if a subject isn't aware of what the belief has going for it then the truth of the belief would be an accident from her own perspective. Getting a true belief in this manner is like grasping gold in the dark. Even if one succeeds you won't have any kind of assurance that you did. And so our discussion in chapter 5 gives us reason to reject this account of foundational justification. The lessons of the SPO are completely general and illustrate that foundational justification (like all justification) must include an awareness of evidence that supports the truth of the relevant belief and an appreciation of that support. Supporting evidence for *P* consists of facts (a la chapter 4) that stand in a *sui generis* and *irreducible* epistemic support relations to *P* (a la chapter 6). In the case of foundational justification, a subject *S*'s awareness of such a fact must be non-doxastic, i.e. it must not consist of *S*'s belief that such a fact obtains. So what kind of awareness of evidence can generate *foundational* justification?

For largely Sellarsian reasons, the awareness of supporting evidence can only provide foundational justification if it's a non-judgmental act.<sup>267</sup> Awareness of a fact is judgmental if

---

<sup>267</sup> My discussion of the Sellarsian worry is heavily inspired by Hasan's (2013) discussion. The Sellarsian dilemma is often developed where the two horns involve the awareness being propositional or non-propositional. I adopt Hasan's development in terms of one's awareness being judgmental vs. non-judgmental since simply being propositional wouldn't make something capable of conferring justification. For instance, my belief that *P* and my desire that *P* both have the propositional content *that P*, but only the former could ever provide justification for believing a proposition *Q* that is made probable by *Q*. The reason is that a proposition *P*'s making probable *Q* just means that *Q* is probable relative to *P*'s truth (*Q* isn't probable relative to *Q*'s truth). As such, it seems that a propositional attitude directed at *P* could only provide reason to believe *Q* if it represents *P* as true. A belief is a kind of assertive attitude since it represents its content as being actualized. A desire, however, isn't assertive and doesn't represent its content as being actualized.

it involves “the assertion or acceptance of a proposition or thought, or at least the categorization of some sensory item or the application of some concept to experience.”<sup>268</sup> Awareness of a fact is non-judgmental if it involves neither the acceptance of a proposition nor the application of concepts. Notice, however, that even if the awareness itself is a non-judgmental act, the *object* of this awareness might be judgmental. For instance, one might hold that one can have a non-judgmental awareness of one’s belief that *P*. The belief that *P* is judgmental since it asserts the truth of *P* but this doesn’t entail that one’s awareness of the belief is also judgmental. I might be aware of this belief without asserting the *higher-order* proposition *that I believe that P* in which case I would have a non-judgmental form of awareness of a judgmental item. Keeping clear the distinction between features of the *objects* of awareness and the *vehicle* of awareness is crucial.

With these distinctions in place we can present the Sellarsian dilemma. A judgmental awareness could clearly provide justification for believing a proposition but a judgmental awareness would appear to stand in need of justification itself. On the other hand, if the awareness is non-judgmental then it isn’t clear how it could provide a reason for/generate justification for believing a proposition.<sup>269</sup>

Assume that my awareness of supporting evidence in the case of supposed foundational justification is judgmental. My awareness of supporting evidence *E* for *P* would involve my asserting or assenting to the proposition *that E* (perhaps my awareness of *E* is just a non-doxastic seeming that *E*). This proposition is a representation and, like all representations, is characterized by its ability to represent the world accurately or inaccurately.<sup>270</sup> But remember that a proposition *P* makes probable *Q* just means that *Q* is

---

<sup>268</sup> Hasan (2013).

<sup>269</sup> For nice statements of the Sellarsian dilemma see Bonjour (1985), (2003), (2010) and Hasan (2013). For the original source of the dilemma see Sellars (1956).

<sup>270</sup> This is a bit tricky in the case of necessary truths but I think the same general thought that I’m attempting to get at here holds it’s just a lot more difficult to state.

probable relative to  $P$ 's truth (i.e.  $Q$  isn't probable relative to  $P$ 's falsity). When I move from a proposition (whether the content of a doxastic or non-doxastic state!) *that E* to a belief that  $P$  I thereby commit myself to the truth of the former. It's only *the truth* of the proposition *that E* that what would make  $P$ 's truth non-accidental. The mere truth of this proposition, however, is insufficient to provide justification for believing that  $P$ . If *that E*'s truth is itself an accident from your own perspective then  $P$ 's truth will also be an accident from your perspective. In order for  $P$ 's truth to be non-accidental from your own perspective you must have some reason to believe the proposition *that E* is true. However, this entails that the justification for believing  $P$  wasn't foundational after all!

Foundational justification (if we're to respect the considerations of the previous chapters) therefore requires a non-judgmental form of awareness. Acquaintance seems the perfect candidate. Acquaintance is a form of awareness that isn't mediated by a subject  $S$ 's asserting or assenting to some representation that may or may not represent the world accurately. If you're acquainted with some evidence, this evidence is directly present to your mind. You can't be acquainted with some evidence if that evidence doesn't exist. If acquaintance doesn't involve any kind of representation then it isn't the kind of thing that can be justified or unjustified. As such, acquaintance appears to be the perfect candidate for a kind of awareness that can stop both the epistemic and conceptual regress.

In making one's awareness of supporting evidence a non-propositional and non-conceptual acquaintance we encounter the other horn of the Sellarsian dilemma. However, why should we think that one's awareness being non-propositional or non-conceptual necessarily prevents it from providing one with a reason or justification for believing a certain proposition? Traditionally the worry has just been that such non-propositional and non-conceptual items couldn't stand in the relevant evidential relations. But notice that acquaintance is the *vehicle* rather than the *object* of awareness. Acquaintance is an awareness that *provides* a subject with epistemic reasons. The reasons delivered to the subject consist of

the object or content of this act of acquaintance. As such, it's the content of this state that must stand in the evidential relations and not the awareness itself.

Doesn't the Sellarsian dilemma now apply to how we conceive of the object of acquaintance? If the object of one's acquaintance is propositional or conceptual then it can evidentially support belief but one would need reason to believe that the relevant proposition is true or that the concepts applied are accurate. Alternatively, if the object of one's acquaintance is non-propositional then it doesn't seem as if it can stand in the relevant evidential relations. A solution falls out of Fumerton's appeal to acquaintance with facts. Facts can stand in a relation of correspondence or *making true* to various propositions and if the content (in this case the content is factual content, i.e. the exemplification of various properties and relations) of one's acquaintance is the truth-maker for *P* then this content constitutes conclusive reason for believing said propositions.<sup>271</sup>

While this deals with the Sellarsian worry we must also remember the lessons from our consideration of the closely related SPO from chapter 5. Merely being aware of some fact *F* that *is* an evidential reason for believing *P* is true isn't sufficient for making *P*'s truth non-accidental from a subject's own perspective. If one is completely oblivious to the evidential connection (even assuming the relation is one of correspondence!) then it's hard to see how acquaintance with *F* would make any difference to my perspective on *P*'s truth. And this motivates the idea that foundational justification requires acquaintance with a fact *F* that corresponds to *P* and an appreciation of this correspondence (i.e. a kind of matching between the factual and propositional contents). Sellarsian worries, yet again, motivate construing this as a kind of non-conceptual awareness of or acquaintance with correspondence rather than any sort of judgmental awareness.<sup>272</sup> If justification for

---

<sup>271</sup> I'm unsure whether Fumerton would accept this characterization of the fact with which one is acquainted as acting as a reason or evidence for one's belief.

<sup>272</sup> Bonjour (2000) and (2003) appears to construe the awareness of correspondence as a conceptual awareness.



believing *P* requires a judgmental act *that* there is a correspondence between the factual content with which one is acquainted and the content of one's thought that *P*, such awareness would be in need of justification. In which case the belief that *P* isn't foundationally justified after all. Moreover, this would initiate an infinite regress of the sort discussed in chapter 1 in connection with access internalism where justification not only requires further justification for an infinite number of beliefs but an infinite number of beliefs of infinitely increasing complexity (see chapter 1).

All of this leads us to Fumerton's view that a subject *S* has foundational justification for believing *P* iff (i) *S* is acquainted with the fact that *P*, (ii) *S* is acquainted with the thought that *P*, and (iii) *S* is acquainted with the correspondence between the fact and the thought. The Sellarsian worry about non-judgmental awareness is circumvented by the idea that the object of one's acquaintance (a non-judgmental form of awareness) can stand in relations of correspondence/truth-making. Moreover, the acquaintance with the correspondence itself makes the truth of *P* non-accidental from one's perspective. "[W]hen everything constitutive of a thought's being true is immediately before consciousness, there is nothing more that one could want or need to justify a belief."<sup>273</sup> In the case of a hunch or guess one is merely aware of the thought that *P*. When a subject meets Fumerton's conditions on foundational justification, however, one directly confronts the proposition *P* *along with its property of being true!* What more could someone want from an account of foundational justification?

#### 7.1.6 The SPO, Acquaintance, and Bergmann's Challenge

Fumerton's theory of foundational justification respects our lessons from the previous chapters. The evidence for foundational beliefs consists of facts (a la chapter 4). The relation between this evidence and the propositional content of foundationally justified beliefs is *sui generis* and *unanalyzable* (a la chapter 6). Finally, foundational justification requires

---

<sup>273</sup> Fumerton (1995), p. 74-75.

that the subject S have *actually connected* awareness of this evidence, i.e. must be aware of both the evidence and appreciate its connection to the relevant proposition ( a la chapter 5).

Michael Bergmann, however, has recently argued that any theory where the required awareness of an evidential connection is achieved via a non-judgmental form of awareness such as acquaintance inevitably falls prey to the SPO. For the sake of argument Bergmann grants the possibility that a subject be acquainted with the fact that *P*, the thought that *P*, and the correspondence between the two.<sup>274</sup> He then offers the following case concerning a subject Jack's belief that he is appeared to redly, call this belief B2, as a supposed illustration that B2's truth will still be an accident from Jack's own perspective:

Because the direct acquaintance is nonconceptual, Jack can be directly acquainted with the fact that he is being appeared to redly [and]... directly acquainted with the relation of correspondence holding between his thought that he is being appeared to redly and the fact that he is being appeared to redly *even if he has no idea* that the relation of correspondence holds between these two items (again, this is because nonconceptual awareness is the sort of thing that can occur without the application of *any* concepts). Thus, Jack's belief B2 can satisfy Fumerton's requirements even if he conceives of his being appeared to redly as no more relevant to B2 than is the mild pain in his left knee. It is, therefore, exceedingly difficult to see how these direct acquaintances improve things *from Jack's subjective perspective*.<sup>275</sup>

#### 7.1.6.1 "He Has No Idea"

What exactly does Bergmann have in mind when he claims that a subject could be acquainted with the correspondence between her experience and her thought and still "have no idea" that this relation of correspondence holds? He can't be suggesting that a subject might be acquainted with the correspondence yet fail to be *aware* of the instantiation of this relation. Acquaintance with the correspondence relation between the experience and the thought *just is* a form of such awareness. A subject can't be aware of the correspondence

---

<sup>274</sup> Though Bergmann does seem to implicitly reveal skepticism about whether he really understands what it would be to be acquainted with something like the correspondence relation.

<sup>275</sup> Bergmann (2006), p. 30.

relation and be *completely oblivious* that such a relation exists. Bergmann could only be suggesting that one might be acquainted with the correspondence yet fail to have an awareness of the correspondence that involves the subject's asserting or assenting to the existence of such a relation.

Bergmann, however, cannot appropriately infer from this that the truth of the belief is therefore an accident from the subject's perspective. Such a move would obviously beg the question against the acquaintance theorist who suggests that this acquaintance does bring the truth into the subject's perspective. Moreover, the acquaintance theorist has a plausible explanation of why this would be sufficient. As Fumerton is adamant to point out, when these three acts of acquaintance are in place "everything constitutive of a thought's being true is immediately before consciousness, there is nothing more that one could want or need to justify a belief."<sup>276</sup>

At this point we might sympathetically reconstruct Bergmann's argument as a challenge that there is indeed something more one could want or need to justify a belief: the actual judgment that one's appearance is relevant to the truth of one's belief before one's mind! Here, however, I think the correct response on behalf of the acquaintance theory (to *some* extent) anticipates our solution to the problem of easy justification. We shouldn't impose the requirement that one actually judge that such a correspondence exists since this, as Bergmann is well aware, initiates a vicious regress. More importantly, however, we needn't impose such a requirement because adding the active conceptualization or assertion that the correspondence exists doesn't add anything to one's perspective on *P*'s truth that wasn't already present prior to such a conceptualization or judgment. When a subject is acquainted with the correspondence relation between a fact <P> and the thought that *P*, adding an

---

<sup>276</sup> Fumerton (1995), p. 74-75.

application of the concept RELEVANT TO THE TRUTH OF P or the judgment that <P> corresponds to *P* is redundant.

Presumably any classical foundationalist view such as Fumerton's is going to tell a story such that certain primitive and sui generis concepts obtain their content via mental ostension (think back to our discussions of the notion of acquaintance and correspondence). Perhaps the story will start with something like Tim McGrew's "I'm experiencing *this*" where the demonstrative concept inherits the content of the sensory perception that is picked out by the ostensive act, i.e. the content of the state picked out is constitutive of the concept.<sup>277</sup> The story of how we move from demonstrative concepts to concepts capable of reapplication is sure to be very complicated but, how it goes, it appears that the possible content of these primitive and sui generis concepts will be parasitic upon the content of items with which we're directly acquainted. Now consider our concept TRUTH. How would we gain a concept of this primitive but ever important relation? Presumably this concept would obtain its content via our ostending in thought to instances of correspondence with which we're acquainted. Thus, our concept TRUTH is parasitic upon the factual content of which we're aware when acquainted with a fact corresponding to a proposition. Our application of the concept "*corresponds with P*" to the fact <P> is redundant when we're already acquainted with the correspondence relation. It's as if all you're adding to your perspective when you apply such a concept is the proposition that *the truth of P is relevant to the truth of P*. Bringing this proposition before one's mind doesn't bring anything into our perspective on *P*'s truth that wasn't *already* present before its application! The reader may be wondering what my point is. The point is that Bergmann seems to admit that if we were to apply such a concept then the object-level proposition's truth would be non-accidental from our perspective (though it would lead to a vicious regress). If, however, I'm right and the

---

<sup>277</sup> See McGrew (1995) and Chalmers (2003).

application of this concept doesn't add any *new* content that would be relevant to our perspective on the proposition's truth then we shouldn't lose anything by taking it away either. Therefore, if we take away the assertion of the proposition that *this (i.e. the fact that I'm appeared to redly) correspondence to this (i.e. the thought that I'm appeared to redly)* yet acquaintance with the correspondence relation remains intact then the proposition's truth should still be non-accidental from our perspective. All of the materials relevant to the truth of one's beliefs are still within one's awareness.<sup>278</sup>

#### 7.1.6.2 Changing the Scope

So far I've explained why the acquaintance theory's three acts of acquaintance can remove the accidentality of a proposition's truth from a subject's perspective absent any judgment that the correspondence holds. However, in the last sentence of his objection to the acquaintance theory Bergmann moves from a discussion of "not conceiving" to "conceiving not." "Not conceiving of X as Y" means that one does not apply the concept Y to X but "conceiving of X as not Y" means that one does apply the concept NOT A Y to X.

Such a move creates a new challenge to the idea that meeting Fumerton's three conditions is *sufficient* for avoiding the SPO that one might find a bit more disturbing. Bergmann is pointing out that, since acquaintance is a non-conceptual form of awareness, one could meet Fumerton's three conditions for foundational justification while actively conceiving of the fact that she is appeared to redly as being "no more relevant to my thought's (that I'm appeared to redly) being true than is the pain in my left knee." Initially, this does strike me as quite disturbing.

---

<sup>278</sup> I need to be clear here that I don't mean to suggest that *nothing* changes when one moves up a level. In fact, moving up a level may constitute an incredible epistemic achievement. Moving up a level requires conceptual sophistication that might contribute to a kind of *understanding* that involves seeing how various notions are related. My point here is that nothing about one's perspective on *P* changes when one moves up a level and conceptualizes the correspondence relation with which one was already acquainted.

I can imagine someone trying to deny the possibility that Bergmann is suggesting. One might just dig one's feet in and say that one couldn't be acquainted with the fact that one is appeared to redly, the thought that one is appeared to redly, and the correspondence holding between them while *conceiving* of this fact as no more relevant to the truth of the thought that one is appeared to redly than is the pain in one's left knee. However, insofar as the acquaintance with correspondence and the *conceiving* are distinct states it seems to me that one cannot deny that Bergmann has identified a logical possibility.

For the most part, I think the correct response to Bergmann is to simply dig in one's heels and suggest that even when the subject *conceives* of her experience as irrelevant to the truth of her belief she still has foundational justification. Her belief's truth is nevertheless non-accidental from her own perspective. In order to see this we need only imagine how such a case might arise. Consider Paul who has become utterly convinced that eliminative materialism is true and accepts that there are no beliefs, no desires, no phenomenal pains, etc. Now imagine that Paul is directly acquainted with the fact that he believes that there is a table in front of him, his thought that a table is in front of him, and the correspondence between the two. However, Paul now remember the arguments for eliminative materialism and starts to think to himself that there are no beliefs and therefore *this* (the fact that he has the relevant belief) is irrelevant to the truth of his thought that he has the belief that there is a table in front of him. Or consider a case where Paul is acquainted with the fact that he is in *searing* pain, the thought that he's in (phenomenal) pain, and the correspondence. However, remembering the arguments of eliminative materialism he thinks to himself that there are no (phenomenal) pains and so *this* (the fact that he is in *searing* pain) is irrelevant to his thought that he is in pain.

When I consider these kind of case I'm convinced that this false philosophical theory about the mind does *nothing* to remove Paul's justification for believing that he believes there is a table in front of him or that he is in pain. His acceptance of a false philosophical theory about the mind doesn't thereby make the truth of these beliefs an accident from his

perspective. Paul is acquainted with the thoughts *and their property of being true!* What is irrational is his acceptance of the higher-order belief that these facts are irrelevant to the truth of his beliefs. And why should we think an irrational belief affects the rationality of his belief that he is in pain? I must admit one caveat, however. It's possible that moving up a level in this way might act as a defeater for Paul's *doxastic* justification since it might prevent Paul from *basing* his beliefs on his acquaintance with the fact and with the correspondence.

### 7.2 Moving to Fallible Foundations

Fumerton's acquaintance theory is closely aligned with traditional versions of classical foundationalism that sought to ground justification in *infallible* foundations. On Fumerton's account it isn't infallible beliefs that are doing the foundational work but rather beliefs for which one has infallible justification. A subject S couldn't have the justification she has for *P* while *P* is false.

Fumerton's version of the acquaintance theory respects our findings from the previous chapters; meeting his three criteria is *sufficient* for foundational justification that includes a kind of 1<sup>st</sup>-person assurance for the truth of one's belief. However, due to his infallibilism, I don't think Fumerton has identified *necessary* conditions for justification.<sup>279</sup>

A theory of justification grounded in *infallible* foundations is overly demanding. We need to loosen Fumerton's requirements so as to allow foundational justification when a subject is acquainted with a set of facts *F* and acquainted with *F*'s *making probable* the thought that *P*. There are three reasons for making this modification. First, widening the class of foundationally justified beliefs in this way leads to a more phenomenologically plausible view. Second, this explains cases of *false* beliefs that are intuitively foundationally justified. Lastly, this accommodates the intuition that there can be foundationally justified true beliefs

---

<sup>279</sup> Fumerton himself seems to waver at times on whether these conditions are necessary and sufficient or just sufficient. He often suggests that acquaintance with a fact similar to the truth-maker for one's belief might provide justification. I consider and reject this idea below.

that are nonetheless epistemically *uncertain*. The key to vindicating this idea, however, is showing that all of this can be done while holding onto the initial motivations for the acquaintance theory. I need to show that this can be done without ceding the claim that justification requires that a subject be aware of how a belief's status is different from that of a mere guess or arbitrary conviction.

### 7.2.1 Phenomenological Considerations

Consider your perceptual beliefs about the various objects around you: perhaps you have the belief that there is a table in front of you, that the table is red, that there is a loud crying child in the background, etc. And consider some of your memorial beliefs: perhaps you have the belief that you ate scrambled eggs for breakfast, that you grew up in a suburb of Minneapolis, that you sent in your taxes for the year, etc. There is some sense in which these beliefs are spontaneous. It isn't as if you go through some kind of conscious process of inferring these beliefs from beliefs about your perceptual and memorial experiences. You don't first form beliefs about the character of your experience, beliefs about how these experiences have been correlated with physical objects having certain properties or certain kinds of events having occurred, beliefs about how these experiences cohere with one another, and then subsequently *infer* the perceptual and memorial beliefs. You seem to form many of your most basic perceptual and memorial beliefs spontaneously and without the mediation of any kind of inference. More importantly though, when you reflect on the beliefs you had a moment ago, it seems as if they were formed immediately *and appropriately*. It seems as if, even prior to your reflection on the character of your experiences, the probable truth of these perceptual and memorial beliefs was directly revealed to you.

Any view that requires foundationally justified beliefs to be infallibly justified is not capable of accommodating this phenomenological point.<sup>280</sup> It doesn't take very much

---

<sup>280</sup> An infallibilism that accepted a form of metaphysical direct realism and disjunctivism about perceptual experience such as that of Brewer (2002) and (2013) could accommodate this



reflection to see that our perceptual and memorial beliefs are fallible. We need only consider the possibility of dreams, illusion, hallucination, and massive deception (e.g. the Matrix, Descartes' evil demon, Russell's hypothesis that a God-like being created the earth 5 minutes ago, etc.) to realize that we could be directly acquainted with the same facts in situations where our perceptual and memorial beliefs are false. Our evidence doesn't guarantee the truth of these beliefs. If a theory only allows infallible foundations then our perceptual and memorial beliefs could only be justified inferentially. In fact, since doxastic justification would require that we *actually* infer these beliefs from beliefs about our experiences, it would turn out that the majority of our beliefs are doxastically unjustified!

An advocate of infallible foundations might attempt to parry this worry by admitting a kind of degenerate justification for these beliefs. Our perceptual and memorial beliefs usually lack *ideal* justification since they aren't consciously inferred from infallibly justified beliefs about appearances. Nonetheless, these beliefs enjoy a kind of degenerate justification since they are still caused by facts about our experience (facts with which we're acquainted) that are the truth-makers for propositions from which we *could* perform the relevant inferences.

Moving to a characterization of certain beliefs as having a kind of degenerate as opposed to ideal justification is legitimate and even called for at times in response to the skeptical concerns that plague internalists. But such a move is inappropriate in this particular instance. Such a response fails to appreciate the nature of the phenomenological objection. When I reflect on the perceptual and memorial beliefs that I often form it isn't merely that it appears as if I failed to rely on any conscious inference. It also appears as if these beliefs were perfectly *appropriate* despite the fact that they were not inferred from infallible beliefs about features of my experience. Forming the relevant appearance beliefs and performing

---

phenomenological point. However, I find metaphysical direct realism and disjunctivism about perceptual experience incredibly implausible.

the relevant inference doesn't seem to actually improve my perspective on the truth of these beliefs. Phenomenologically it seems as if the probable truth of *some* perceptual and memorial beliefs is directly given to me without any *need* of inference.

By modifying the acquaintance theory to allow for probabilistic foundations we can vindicate this intuition. With the metaphysics of fact-fact, fact-proposition, and proposition-proposition pairs developed in chapter 4 we're in an ideal position to do this without losing the initial motivations for the acquaintance theory. Foundational justification might sometimes hold in virtue of a subject's acquaintance with a set of facts *F* and her acquaintance with this set of facts' making probable *P* (though we'll see below that moving to these probabilistic relations will move us to introduce a kind of no defeater clause in order to formulate sufficient conditions for foundational justification). Such a subject would be directly confronted with a proposition's property of being *probably true* and so would still "see" what the belief has going for it. As such, one's justification would provide the subject with a *degree* of assurance that her belief is true in proportion to the *degree* of probability with which they are acquainted. What the belief has going for it in such a case would be less than in cases where Fumerton's conditions are met but this just reflects the idea that assurance itself comes in degrees. Meeting Fumerton's conditions is therefore best construed as just a special case of foundational belief where such justification is particularly strong.<sup>281</sup>

### 7.2.2 Foundational Justification, Fallibility, and Uncertainty

Fumerton's infallibilist version of the acquaintance theory also suffers from two difficulties concerning the related possibilities of foundationally justified true but *epistemically uncertain* beliefs and of foundationally justified *false* beliefs.

---

<sup>281</sup> In chapter 5 I suggested that we might understand someone who has potentially connected supporting evidence as having a kind of degenerate as opposed to ideal justification. I'm not making that distinction here. I take it that one can have *ideal* foundational justification for fallible beliefs and that meeting Fumerton's infallibilist conditions is just a case where a subject has a particularly *high* degree of ideal justification.

Consider first the possibility of foundationally justified true beliefs that are nevertheless *epistemically uncertain*. There seem to be contrasting pairs of cases where acquaintance provides foundational justification for a true belief but where one ought to have certainty in one of the cases but one ought to have something less than certainty in the other case. It's not merely that there are some cases of foundational justification where one is *psychologically* uncertain that a proposition is true but also where psychological certainty is *epistemically* inappropriate. The infallibilist acquaintance theory doesn't have the resources for explaining this *epistemic* difference. The infallibilist theory holds that all cases of foundational justification involve acquaintance with a truth-maker and acquaintance with the correspondence relation between the truth-maker and the thought. If I'm acquainted with everything constitutive of the thoughts being true then it seems I have the best possible justification I could ever have. The truth of the belief is directly before my mind. So wouldn't epistemic certainty be the epistemically appropriate attitude in all cases of foundational justification?

The threat of foundationally justified but *false* beliefs would pose for the infallibilist theory is much more obvious. Clearly, one can't have infallible justification for a false belief. One's justification cannot guarantee the truth of a belief that is in fact false. In order for an acquaintance theory to allow foundational justification for false beliefs one must appeal to acquaintance with a fact other than the truth-maker for the thought and (in order to satisfy the SPO) acquaintance with a relation other than correspondence. As such, the infallibilist acquaintance theory cannot accommodate the possibility of foundationally justified false beliefs.

But why should we think that these are genuine possibilities? Each possibility can be motivated by appealing to either: (i) continua cases such as those discussed by Williamson in

his famous anti-luminosity argument<sup>282</sup>; or (ii) cases like those used to motivate the so-called “problem of the speckled hen.”

Consider the continue type cases. Imagine that you begin in a state of extreme pain that slowly begins to decrease in intensity until finally the pain disappears and, perhaps, slowly turns into an itchy sensation. Or consider a case where you’re experiencing a paradigmatically red patch in your visual field that slowly changes into a paradigmatically orange patch. Williamson contends that there are miniscule changes in the character of your experience that occur despite the fact that you’re incapable of noticing these changes, and such a contention is strikingly plausible.

Now imagine that you’re experiencing a pain or a patch of red that is on the continuum “right next to” an indistinguishable (to you) itch or orange patch. It’s intuitive that you would have *some* justification for believing that you’re in pain or that you’re appeared to redly. Clearly your awareness of the character of your experience would give you more justification for believing that you’re in pain than that you’re having a euphoric sensation (and you would have more justification for believing that you’re appeared to redly than that you’re appeared to greenly). In these cases, the experience of which you’re aware corresponds to the relevant thought and your belief is true. But it seems obvious that psychological *certainty* that you’re in pain or that you’re appeared to redly in these cases would be epistemically inappropriate; you clearly have *less* justification for these beliefs than you do in the cases where you’re acquainted with severe pain or a paradigmatically red patch in your visual field.

Fumerton (2002) actually defends the idea that correspondence and (thereby) truth can come in degrees. I find such a view of truth to be implausible but this may give Fumerton wiggle room for allowing his version of the acquaintance theory to deal with the

---

<sup>282</sup> See Williamson (2000), chapter 4.

epistemic uncertainty of foundationally justified beliefs in these continua cases. The idea would be that the very concepts PAIN, ITCH, REDNESS, and ORANGENESS are vague concepts that fail to specify determinate truth-conditions for borderline cases. Thus, when one is acquainted with an experience that is a borderline case the thought that one is in pain or that one is appeared to redly is “kind of true” in the sense of having some degree of truth or correspondence with the fact. Moreover, one might be acquainted with this *degree of correspondence* and thereby have justification for taking the thought to be “kind of true.” In the case of paradigmatic pain or redness, however, one might be acquainted with a full degree of correspondence and thereby have justification for taking the thought to be true simpliciter. Perhaps this makes sense of the contrast between foundational justification that licenses epistemic certainty and those that don’t *in these cases*. Notice, however, that this response commits one to a substantive account of vagueness that locates this in our concepts. If one adopts an epistemicist approach to vagueness where there is a sharp line between pains and itches, redness and orangeness, and bald and not-bald but we just cannot *know* where this line is, then this response is out of the picture. It would be dialectically helpful if we could find a solution that didn’t commit us either way in a substantive debate about vagueness.

More importantly, however, this solution to foundational justification with epistemic uncertainty is extremely implausible in speckled hen cases. The traditional problem of the speckled hen is as follows. Imagine that you have a conscious experience as of a hen with exactly 47-speckles. Even if you were acquainted with such an experience and happened to form the belief that you’re having an experience as of a 47-speckled hen, this belief would not be justified. The acquaintance theory outlined above actually has an obvious response to this problem: “Yes, one might be acquainted with an experience as of a 47-speckled hen but lack foundational justification for the corresponding thought since one could be acquainted with the experience while lacking acquaintance with the correspondence relation itself.”<sup>283</sup>

---

<sup>283</sup> See Poston (2007).

This isn't my preferred response to the traditional problem of the speckled hen (I indicated earlier that I think my construal of acquaintance as a kind of selective attention is sufficient to avoid this problem) but here I'm concerned with a further problem for the acquaintance theory raised by the speckled hen case. Clearly, I have *more* justification for believing that I'm having an experience as of a 47-speckled hen than I have for believing that I'm having an experience as of a 13-speckled hen or of a 272-speckled hen. But this implies that I have *some* justification for believing that it's an experience as of a 47-speckled hen. Clearly one lacks epistemic certainty in this case but one doesn't lack justification all together. And it isn't clear how the infallibilist version of the acquaintance theory can accommodate this. We can't appeal to vagueness or degrees of truth. It's plausible that the concepts PAIN, ITCH, REDNESS, and ORANGENESS lack precise boundaries. But there doesn't seem to be vagueness in the concept of "47 speckles in one's phenomenal field." When I'm acquainted with an experience with the relevant character and I form the thought that I'm having an experience as of a 47-speckled hen my belief is true simpliciter. My thought isn't just "kind of true." Moreover, in order to explain the existence of this justification (in line with the considerations of the SPO in the previous chapter) we must appeal to the subject's awareness of the connection between the fact and the thought. If one construes this connection as correspondence then one cannot explain the epistemic appropriateness of uncertainty in this case. If S isn't aware of the connection then she has no justification for her belief. If S is aware of the correspondence relation then she's aware of a relation of full correspondence thereby making certainty the epistemically appropriate attitude. Fumerton's version of the acquaintance theory cannot accommodate the middle ground required in this case.

These cases also help to illustrate the further worry concerning the possibility of foundational justification for *false* beliefs. In the continua cases I might have an experience as of an itch that is right next to a pain on the continuum or as of an orange patch that is right

next to redness on the continuum. Surely I might have some justification for my belief that I'm in pain or that I'm appeared to redly even though these beliefs are false.

By hypothesis, that state [i.e. the itch] is right next to a state that is pain, and I can't tell the difference between the two... I might surely have *some* level of justification for believing that I am in pain even when I'm not, but where I am instead in a state that is only very similar to pain.<sup>284</sup>

This becomes even more compelling when you consider the fact that the character of the experience surely provides the subject with *more* justification for the false belief that they are in pain than for the false belief that they are having a sensation of euphoria.

The appeal to continua cases in defending the possibility of foundationally justified false beliefs rests on the assumption that there is a precise point where the state is no longer a pain. A rejection of this assumption and an appeal to degrees of correspondence, degrees of truth, and vagueness in our concepts provides wiggle room for defending the original version of the acquaintance theory in these continua cases.

However, yet again, this strategy won't deal with cases of foundationally justified false beliefs in the case of the speckled hen scenarios. When I (a normal human, not a person with Rain Man like abilities) have an experience as of a 47-speckled hen I have *more* justification for believing that I'm having a 49-speckled hen experience than for believing that I'm having an 11-speckled hen experience. If I have *more* justification for this belief, I have *some* justification. Foundational justification for this false belief cannot be accounted for on the traditional acquaintance theory since, by hypothesis, I'm not acquainted with a truth-maker for my belief. Nor am I acquainted with a fact that corresponds to some degree and makes my thought *kind of true* (as might be plausible in the continua cases).

These continua and speckled hen cases illustrate the possibilities of both: (i) foundationally justified true beliefs that are nonetheless epistemically uncertain, and (ii)

---

<sup>284</sup> Fumerton (2010), p. 380.

foundationally justified false beliefs. An infallibilist version of the acquaintance theory cannot accommodate these possibilities. We must, therefore, modify the theory so as to allow for fallible foundations. Moreover, if we're going to allow for fallible foundational beliefs about experiential states I see no reason not to allow for fallible foundational perceptual and memorial beliefs so as to accommodate the phenomenological point mentioned in 7.2.1.

### 7.2.3 Moving from Correspondence to Probabilistic Relations

The key to allowing for epistemic uncertainty and fallibility in foundational justification is a modified version of the acquaintance theory that holds that the relevant evidential relation is a probability relation even in cases of foundational justification. And the metaphysics of fact-proposition probability relations developed in chapter 4 along with my discussion in chapter 6 allows just this. On this modified theory a subject *S* has (defeasible) justification if and only if *S* is acquainted with a set of facts *F* that make *P* probable and acquainted with the probability relation between *F* and *P*.<sup>285</sup>

It is straightforward how this modification allows the acquaintance theory to accommodate the phenomenological point of 7.2.1. My acquaintance with a set of facts about my perceptual experiences, apparent memories, etc. can provide foundational justification for perceptual or memorial beliefs in virtue of making these thoughts probable and my being acquainted with this probability relation.

How this modification allows for epistemic uncertainty and fallibility in the case of introspective foundational beliefs, however, is a bit more complicated. There are actually three possibilities we must consider. Epistemic uncertainty and fallibility might be due to the factual content with which one is acquainted, the clarity of one's acquaintance with the evidential support relation, or even a combination of both.

---

<sup>285</sup> The defeasibility modifier is crucial, see 7.4.4.



Consider the continua cases where a normal human S forms a *true* foundationally justified belief about a property of S's experience that is right at the edge of the continuum. Consider, specifically, the red-orange continuum case. In this situation S has a red experience such that the redness is right at the edge of the continuum and a change to an orange experience right next to the red experience would go unnoticed. S has *some* foundational justification for believing that S is appeared to redly but this belief is still epistemically uncertain. What might explain the uncertainty on my modified acquaintance theory?

The first possibility is that S isn't acquainted with the correspondence between the fact that she is appeared to redly and the thought that she is appeared to redly due to the character of her experience being right at the edge of the red-orange continuum. And so S lacks the epistemic certainty provided by meeting the requirements of the traditional acquaintance theory. Nevertheless, S might also be acquainted with the fact that her experience has some more determinable color property, something like redish-orangishness, which makes it at least somewhat probable that her experience has the more determinate color property of redness. And S might be acquainted with this probability relation and thereby have *some* degree of justification for believing that she is appeared to redly. Minimally, S would have *more* justification for believing that she is appeared to redly than she does for believing that she is appeared to greenly.

The second possibility that would account for a subject S's having foundational justification that falls short of epistemic certainty pertains to the nature of S's awareness of the connection between the truth-maker and her thought. We're all familiar that our perceptions and memory can be "fuzzy." Some perceptions and memories are clearer or more vivid than others. And these differences in the vivacity of our experiences can make epistemic difference. In an extremely valuable discussion of probability, Russell (1948) explores the idea that we can similarly have a "dim awareness of logical connection." He explains that:

[O]ur perceptions of the logical connections between propositions, like our sense perceptions and our memories, can be ordered by their degrees of credibility: in some, we see the logical connection so clearly that we cannot be made to doubt it, while in others our perception of the connection is so faint that we are not sure whether we see it or not.<sup>286</sup>

Russell is noting an incredibly important phenomenon. Sometimes we might be aware of some evidence  $E$ , the thought that  $P$ , but only have a kind of “fuzzy” awareness of the evidential connection between  $E$  and  $P$ . In such a scenario we might only be acquainted with the fact that the connection has certain determinable features without being acquainted with any determinate degree of evidential support for  $P$  offered by  $E$ . We might, for instance, be acquainted with the fact that  $E$  makes  $P$  more probable than not, that  $E$  makes  $P$  probable to a degree that is somewhere between .6 and .8, or that the extent that  $E$  makes  $P$  probable is at least 90%.<sup>287</sup> Nevertheless, there are situations like these, where our acquaintance with facts about determinable properties of the evidential connection between  $E$  and  $P$  would only make probable to some degree claims about a more determinate evidential connection.

A person can be acquainted with an epistemic support relation in virtue of her being acquainted with the relation’s instantiating a certain property (i.e. with a certain fact about that relation). Consider again a distinction between a certain kind primary and secondary seeing. There are many cases where I see your body even though the bottom half of your body is blocked by a fence. In such a case I see your body in virtue of seeing a part of your body (e.g. I see your whole body in virtue of seeing your torso, arms, and head). I see your torso, arms, and head in the primary sense and see your whole body in a secondary sense. Similar remarks can be made about acquaintance. I might be acquainted with an experience that *has* 47 dots in virtue of being acquainted with that experience’s instantiating the determinable property of being many speckled. In such a case I’m acquainted with the

---

<sup>286</sup> Russell (1948), p. 395

<sup>287</sup> It makes matter considerably more complicated but there is also the possibility that the edges of the probability interval with which we’re acquainted might themselves be fuzzy.

experience as of 47 dots but not with the *fact* that my experience has 47 dots. I'm acquainted with the *fact* that my experience is as of many speckles. I'm acquainted with the experience in virtue of my acquaintance with the instantiation of a property of that experience. Finally, this same observation applies to the evidential support relation. I can be acquainted with the evidential connection between  $E$  and  $P$  (which is a relation of a determinate degree of support) but in virtue of my acquaintance with the *fact* that the relation is one of making  $P$  more probable than not (i.e. acquaintance with the instantiation of a determinable relation). In such a situation, one shouldn't adopt a *precise* attitude regarding  $P$ . One's credence should be distributed along the interval of determinate probabilities that have the relevant determinable feature. In other words, a subject in this kind of situation should adopt what some have called a "mushy" credence.<sup>288</sup> Notice that this discussion suggests an important difference between acquaintance with an experience, state, or property and acquaintance with a fact. It's only acquaintance with a fact that is epistemically relevant (this ties in with the fact ontology of evidence discussed in chapter 4).

How does all of this tie into epistemically uncertain but true foundationally justified beliefs? Imagine that a subject is acquainted with the fact  $F$  that they're appeared to redly. However,  $S$  is only acquainted with the entailment relation between  $F$  and the thought  $P$  that she is appeared to redly in the *secondary sense*.  $S$  is acquainted with this entailment relation but not with the fact that the entailment relation holds.  $S$  is acquainted with the entailment *in virtue of* being acquainted with the fact that the relation between  $F$  and  $P$  is a relation of making probable that falls somewhere on the interval between .5 and 1.  $S$  would only be justified in holding an imprecise credence in the proposition that she is appeared to redly

---

<sup>288</sup> See Joyce (2010).

which is spread out along the probability interval that includes the determinate probability relations that fall under the relevant determinable.<sup>289</sup>

A similar strategy can be used in speckled hen type cases. First, the epistemic uncertainty might be due to the fact that your only source of foundational justification is your acquaintance with the fact that your experience has certain determinable features (e.g. that it is many speckled, the speckles take up roughly such-and-such phenomenal space, the speckles collectively take up roughly such-and-such portion of your visual field, the speckles are distributed in such-and-such fashion, etc.) and your acquaintance with the fact that these features make it probable that your experience is as of 47 speckles. Second, the epistemic uncertainty might be due to the fact that your source of foundational justification is the result of your acquaintance with the fact F that your experience is as of 47-speckles but where you're only acquainted with the fact that there is some determinable evidential relation between F and the thought that you're having an experience as of 47-speckles.

The possibility of a subject S being acquainted with the fact that S's experience has certain determinable properties and acquainted with the fact that this makes probable that S's experience has more determinate properties also explains the possibility of foundational justification for false beliefs. Consider the continua case where I'm acquainted with an orange experience right next to a red experience on the continuum. It's intuitively possible that this would provide me with some justification for my belief that I'm appeared to redly. My modified version of the acquaintance theory can accommodate this by pointing to the possibility that I'm acquainted with the fact that my experience has the determinable feature

---

<sup>289</sup> Moreover, if the subject were to try to move up a level, while remaining acquainted only with this determinable fact about the support relation, she wouldn't be able to tell whether she is acquainted with the more determinate property of her experience (i.e. redness) that would be a truth-maker for her belief or whether she was acquainted with a less determinate property (reddish-orangeness) that merely makes it probable that her experience instantiates the more determinate property. Thus, there are ways that epistemic uncertainty at the object-level will produce epistemic uncertainty at the meta-level as well.

of a kind of redish-orangeness and acquainted with the fact that this makes probable that I'm having a red experience. At the very least I might be acquainted a stronger degree of evidential support between the fact about this determinable feature and the claim that I'm having a red experience than between this fact and the claim that I'm having a green experience. Similar remarks apply *mutatis mutandis* to speckled hen cases where I form a false belief about the number of speckles that is near the correct number.

### 7.3 The Inferential Internalist Insight

I've argued for a modified version of the acquaintance theory of foundational justification. Now we need to consider how we can move beyond our foundations. We need an account of inferential justification. Yet again, Fumerton's inferential internalism provides a nice starting point for motivating my preferred view.

Earlier I introduced the relatively uncontroversial principle of inferential justification IJ1. Fumerton appeals to the way we might challenge a subject's claims to inferential justification in order to argue for a more demanding restriction on inferential justification.

We often challenge a subject S's justification for believing  $Q$  by challenging her justification for believing the inferential basis of this belief. Imagine that I come to you and claim that you're going to experience excruciating pain tomorrow. Having thought that you were going to be spending tomorrow at a relaxing spa, you ask me for my reasons. I offer as my *putative* justification the claim that you're going to be kidnapped and tortured by government spies who think you have some sensitive information about a terrorist plot. We both realize that if this proposition is true then there would be good reason for believing that you'll experience excruciating pain tomorrow. But imagine you ask why I believe this latter proposition and you discover that I formed the belief on a whim and have absolutely no justification for believing the proposition offered as my putative evidence. You would rightly conclude that this belief fails to provide me with inferential justification for believing that you'll experience excruciating pain tomorrow. All of this supports IJ1.

However, Fumerton points to another way to challenge the justification of a belief.

The astrologer makes all sorts of predictions about your life based on your birthday and the positions of celestial bodies. Almost all of us think that the astrologer's predications are comically irrational. Why? It is not that we doubt his knowledge of the stars. Rather, we doubt that he has any reason to believe that the positions of celestial bodies [have] anything to do with the affairs of human beings.<sup>290</sup>

We don't challenge the justification of the astrologer's beliefs in the way that we challenged the justification of my belief that you'll experience excruciating pain tomorrow. In most cases we grant that the Astrologer has justification for her beliefs about your birthday and the positions of celestial bodies. Nevertheless, we don't think the astrologer has inferential justification for her belief because we think the astrologer lacks justification for believing these facts make probable her predications. All of this leads Fumerton to accept IJ2:

**IJ2** A subject S has inferential justification for believing *Q* on the basis of her belief *P* only if: (i) S has justification for believing *P* and (ii) S has justification for believing that *P* makes probable *Q*.

Michael Huemer argues against clause (ii) of IJ2 by arguing that IJ1 alone can explain the astrologer's lack of justification solely since the astrologer's reasoning must be construed as *enthymematic*. If an astrologer were to form the belief (2) that there will be prosperity next year on the basis of her belief (1) that Jupiter will align with Mars, surely the astrologer's inference would utilize various background beliefs.

The astrologer does not merely believe (1); she has a host of background beliefs about the characteristics of the planets, how they influence human affairs, and the like. As a simple example, the astrologer might believe

(3) The alignment of Jupiter with mars causes people to be more productive.

and then it would be the *conjunction* of (1) with (3) that the astrologer would base her belief in (2) on. She would not simply infer (2) from (1).<sup>291</sup>

---

<sup>290</sup> Fumerton (2006), p. 102.

<sup>291</sup> Huemer 2000.

Huemer's move here is to claim that astrologer's belief that (2) is at least partially based on her belief that (3). IJ1 thereby explains why, intuitively, the astrologer lacks inferential justification for believing that there will be prosperity next year since we think that she lacks justification for believing (3).

This objection to IJ2 fails. Huemer is right that we need to be careful to recognize reasoning that is and reasoning that is *not* enthymematic. But once all of the implicit premises are in place a subject would still need to "see" the connection between the premises and the conclusion. Otherwise, as we've seen in chapter 5, the belief's truth would be an accident from the subject's perspective. Fumerton's motivation for IJ2 is just another version of the SPO. Imagine we have justification for a belief that  $P$  which provides the best possible support for  $Q$ : entailment. If I have justification for believing  $P$  I might still lack justification for believing  $Q$  if the entailment relation is too complicated for me to "see" and appreciate. If I fail to see the relevant connection between  $P$  and  $Q$  then my justification for believing  $P$  does nothing to distinguish, from my own perspective,  $Q$ 's truth from a random guess or arbitrary conviction. I therefore lack justification for believing  $Q$ .

However, I do think we ought to reject IJ2 for two reasons. First, IJ2 is too strong and falls prey to Lewis Carroll's (1895) infamous regress. IJ2 therefore leads to skepticism about inferential justification. Second, IJ2 is designed to avoid a kind of SPO objection to inferential justification but it does so in a way that makes an unprincipled distinction between what is required to avoid such an objection in the cases of foundational and inferential justification.

Consider the regress first. It's important to note that the problematic regress initiated by the second clause of IJ2 isn't the epistemic regress. Fumerton solves this regress by allowing that beliefs about probability relations can be foundationally justified. Our difficulty involves the nature of inference. Let's stick with the most straightforward case of entailment. Assume that  $P$  entails  $Q$ . According to IJ2,  $S$ 's won't have inferential justification for believing  $Q$  on the basis of  $P$  merely in virtue of  $S$ 's justification for believing  $P$ . Also

relevant is S's justification for believing that  $P$  entails  $Q$ . Now it seems that S's inferential justification for believing  $Q$  depends on two basis beliefs: the belief that  $P$  and the belief that  $P$  entails  $Q$ . Applying IJ2, however, has the result that S won't have justification for believing  $Q$  on the basis of these two beliefs unless he also has justification for believing that they jointly entail  $Q$ ; S must have justification for believing that  $P$  and  $P$  entails  $Q$  jointly entail  $Q$ . And now it seems that S's inferential justification for believing  $Q$  depends on three basis beliefs. Such a process continues *ad infinitum*. IJ2 appears to have the result that inferential justification requires a subject to have justification for an infinite number of beliefs about entailment relations of *infinitely increasing complexity*. People, however, cannot even entertain the beliefs past the even the 5<sup>th</sup> or 6<sup>th</sup> iteration. And it's hard to convince yourself that you have justification for believing a proposition that you cannot even entertain in thought.

Fumerton suggests avoiding the Lewis Carroll regress by admitting that one must have justification for believing the evidential connection between  $P$  and  $Q$  holds but that this proposition doesn't constitute a premise from which you must infer  $Q$ ; in which case the proposition about the connection doesn't constitute a *basis* of one's inferential justification that  $Q$  and so IJ2 doesn't apply. I, however, am at a loss for what the rationale for such a claim might be (other than the hope of avoiding the regress but then the move just looks *ad hoc*). This judgmental awareness of the connection between  $P$  and  $Q$  is playing an essential role in one's justification for believing  $Q$  according to Fumerton. S's judgmental awareness of the connection is part of what provides S with assurance for  $Q$  and thereby makes  $Q$ 's truth non-accidental from S's perspective. If a judgmental awareness is playing this role then it strikes me that the proposition asserted by the judgment must be playing the role of a premise. And this is relates to the next worry.

When we compare Fumerton's account of foundational justification with his inferential internalism it appears that there is an illegitimate bifurcation in how he avoids SPO like objections. Consider again the case of entailment motivating Fumerton's acceptance of IJ2. Contrast two subjects S and S\* who both know that  $P$ .  $P$  entails  $Q$ . This



entailment is so complicated that  $S$  is completely oblivious to the relation. However, given her superior cognitive abilities,  $S^*$  actually sees the connection. In this kind of case it seems that  $Q$ 's truth is an accident from  $S$ 's but not  $S^*$ 's perspective. As such, it's only when one can see the connection that one's knowledge that  $P$  provides a subject with inferential justification.

A nice example of this phenomenon is the contrast between the mathematical dilettante and the expert. Both may know that certain axioms are true but the expert will likely see various entailments that the mathematical dilettante cannot, and so only the expert will have justification (at least on the basis of her knowledge of the axioms). It's only when the subject can see the connection that a justified belief that  $P$  could make the truth of another proposition  $Q$  non-accidental from a subject's perspective. The problem is the move from this observation to inferential internalism. Our considerations show that inferential justification isn't present when the subject is completely oblivious to the inferential connection. But one needn't adopt as demanding a condition on inferential justification as clause (b) of IJ2. We might instead require only that a subject be acquainted with the relevant connection. This is a middle ground between a theory of inferential justification that merely requires the existence of a connection (and falls prey to the SPO style objections) and inferential internalism (which appears to lead to the Lewis Carroll regress).

Fumerton allows that acquaintance with correspondence is sufficient for removing accidentality from a subject's perspective in the case of foundational justification. One needn't have justification for believing that the correspondence obtains. When everything constitutive of  $P$ 's truth is directly before the mind what else could one want for justification? Adding the proposition about correspondence doesn't appear to improve one's perspective on  $P$ . Similarly, when one is justified in believing  $P$  (and, therefore,  $P$ 's truth is non-accidental from one's perspective) and one is acquainted with  $P$ 's making probable  $Q$  one has directly before one's mind everything constitutive of  $P$ 's probable truth. What more could one want from inferential justification? Yet again, it doesn't seem that moving up a

level to a belief that the probability relation obtains makes a difference to one's perspective on  $Q$ 's truth.

The basic objection here is that Fumerton illegitimately treats foundational and inferential justification differently. We ought to understand the nature of the required awareness of the connection the same in both the foundational and inferential case. If acquaintance with the relevant connection is sufficient to make a proposition's truth non-accidental at the foundation then it should be sufficient at the inferential level. And we've already seen that the antecedent is true. I, therefore, suggest weakening IJ2 in this way.

Notice that this nicely unifies how we think about both foundational and inferential justification. In each case a subject must be aware of how her ultimate evidence (i.e. a set of facts with which she is directly acquainted) makes her belief probable in a way that is ultimately grounded in acts of direct acquaintance with probability relations (for both foundational and inferential justification). In this way a subject must always "see" what her belief has going for it. However, sometimes a subject cannot directly see the evidential connection between a set of facts and a proposition. And the direct acquaintances with probability relations holding between propositions acts as a way to extend our awareness of what certain facts make probable.

We can get a better grasp of this by considering a Cartesian theory of inference. Descartes explains an important distinction between intuition and deduction:

Very many facts which are not self-evident are known with certainty, provided they are inferred from true and known principles through a continuous and uninterrupted movement of thought in which each individual proposition is clearly intuited. This is similar to the way in which we know that the last link in a long chain is connected to the first... [W]e can have knowledge of the connection provided we survey the links one after the other... [W]e are distinguishing mental intuition from certain deduction on the grounds that we are aware of a movement of a sort of sequence in the latter but not in the former.<sup>292</sup>

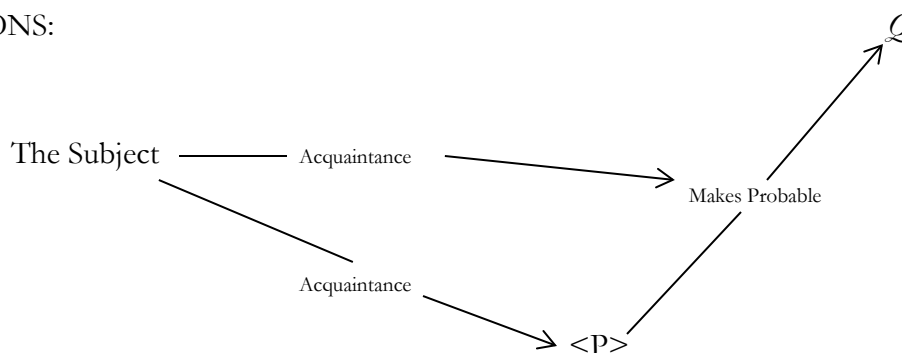
---

<sup>292</sup> Descartes ([1628] 2007), p. 15

Descartes was concerned with certainty and so he required that the foundations be certain and that each entailment relation be known with certainty. Nevertheless we can appropriate this idea to my version of the acquaintance theory that allows fallible foundations and fallible inferences.

In many cases we can see immediately and directly see (i.e. be acquainted with) the evidential connection between a set of facts and a proposition. Direct awareness of the evidential connection as depicted in figure 5 gives us *prima facie foundational* justification.

PROPOSITIONS:



FACTS:

Figure 5: Acquaintance and Foundational Justification

In this case we are immediately aware of what a belief has going for it. Justification that holds in virtue of the subject's being directly aware of evidential relations in this way is *not* constituted by conditions that provide justification for believing any other proposition.

Alternatively, it might not be immediately obvious to a subject *S* that an evidential relation holds between a fact or set of facts *F* and a proposition *Q*. A subject might be acquainted with a set of facts but fail to be acquainted with the relation obtaining between *F* and *Q*. Nonetheless, a subject might be aware of the evidential connection via a chain of inferential steps. *S* might be acquainted with *F*'s being supporting evidence for *P* and acquainted with *P*'s making probable *Q*. In this way a subject can *indirectly* grasp the connection between *F* and *Q*. As Descartes explained, "we know that the last link in the chain is connected to the first... provided we survey the links one after another." This is

easiest to understand in the case where  $F$  is the truth-maker for the proposition  $P$  that makes probable  $Q$  since this relation is necessarily isomorphic to a probability relation holding between  $F$  and  $Q$  (a la chapter 4). A subject will have a kind of indirect awareness of a probability relation holding between a fact and a proposition  $Q$  mediated by her awareness of  $P$ 's truth. Such a case is illustrated in figure 6.

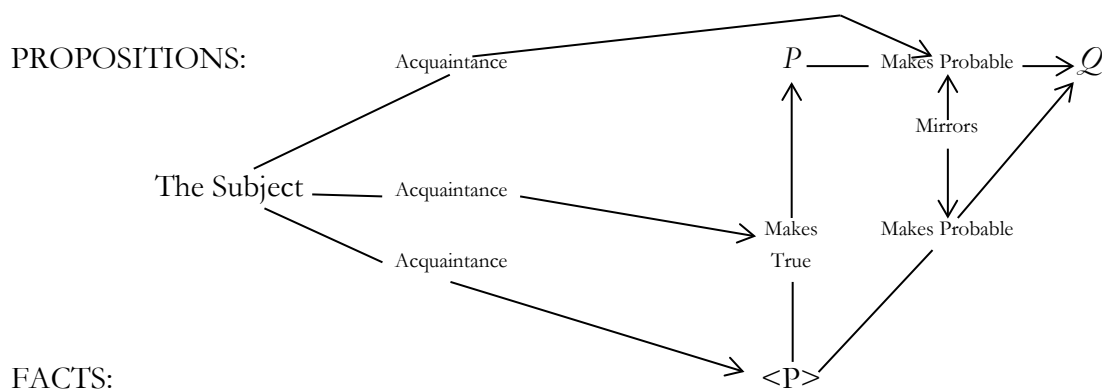


Figure 6: Acquaintance and Inferential Justification

In this kind of case where one is indirectly aware of an evidential connection between  $F$  and  $Q$  one's justification essentially depends on  $S$ 's meeting conditions that provide justification for believing the intermediary propositions (e.g.  $P$ ). As such,  $S$ 's justification for believing  $P$  is inferential.

#### 7.4 Having Evidence and Justification: Stating the View

At this point I've presented and argued for each aspect of my version of the acquaintance theory of justification (though there is one addition that I make with my fourth condition and explain in 7.4.4). Now we need only put the pieces together into a full theory of justification. In stating the overall view that we've arrived at I'll adopt a few notational conventions. A capital "S" will refer to the subject. A capital "F" will be a variable that only ranges over facts or sets of facts. Different facts and sets of facts will be distinguished by attaching a "\*" to the capital "F." And, finally, any italicized capital letter, e.g. " $P$ ", will be a

variable that only ranges over propositions. Having adopted these conventions, my view can be presented with the following principle of justification (PJ) which is meant to capture both the notion of foundational and the notion of inferential justification:

- PJ** S has justification for believing that  $P$  iff:
- (i) S is acquainted with a set of facts  $F$
  - (ii)  $F$  makes probable that  $P$
  - (iii) S is aware of  $F$ 's making probable  $P$  in one of the two following ways:
    - a. *Directly* aware of the fact that  $F$  makes probable  $P$  in virtue of being acquainted with this relation
    - b. *Indirectly* aware of the fact that  $F$  makes probable  $P$  in virtue of either: (i) being acquainted with the fact that  $F$  makes probable  $Q$  and acquainted with the fact that  $Q$  makes probable  $P$ ; or (ii) being acquainted with the fact that  $F$  makes probable  $Q$ , acquainted with the fact that  $Q$  makes probable  $R$ , and acquainted with the fact that  $R$  makes probable  $P$ ; or (iii) ... continued *ad infinitum*
  - (iv) There is no set of facts  $F^*$  such that:
    - a.  $F \subseteq F^*$
    - b. S is acquainted with  $F^*$
    - c.  $F^*$  does *not* make probable  $P$  (i.e.  $F^*$  is either neutral between  $P$  and  $\neg P$  or  $F^*$  makes probable that  $\neg P$ )
    - d. S is either directly or indirectly aware of the fact that  $F^*$  does *not* make probable  $P$ .

I'll explain PJ by considering each condition individually and explaining how it relates to the evidential concepts: evidence, evidence possession, and evidential support.

#### 7.4.1 Facts and Evidence: Conditions (i) and (ii)

Chapter 4 developed and defended the view that evidence consists of facts. These are non-linguistic complexes composed of an entity or entities instantiating various properties and relations. Facts are neither true nor false but they are truth-makers and probability makers. Chapter 6 argued that a fact, or set of facts, is justificatory evidence for or against a proposition  $P$  in virtue of standing in a *sui generis* and *unanalyzable* probability relation to  $P$ . The property of being an evidential reason for believing  $P$  is a fundamental and irreducible feature of reality. Our main motivation for such a view came from the new evil demon thought experiment.

#### 7.4.2 Evidence Possession: Conditions (i) and (iii)

Chapter 5 argued that possession of evidence was to be understood in terms of awareness and that this awareness of supporting evidence is necessary for justification. Our main motivation for this was the cases of the brain lesion patient, Mr. Truetemp, Norman, and the related SPO. As we've seen in this chapter, however, for largely Sellarsian reasons this awareness must ultimately be a non-conceptual and non-judgmental awareness known as acquaintance. Facts of which we are acquainted are directly present in consciousness without mediation from the application of concepts or the entertaining a representation of the relevant fact. A subject's justificatory evidence is exhausted by facts with which they're acquainted. This is the idea behind condition (i) of PJ.

If a subject *S* is aware of a fact *F* and *F* is evidence for *p* (i.e. *S* meets the first two conditions of PJ) then there is clearly *a sense* in which *S* *has* supporting evidence for *P*. Our discussion, however, has shown that the SPO (see 5.3.2) motivates the idea that a subject must possess supporting evidence *E* for *P* in a more demanding sense in order for *E* to play its justificatory role. Justification requires that a subject *see* what a proposition has going for it (or against it) and *see* how its status differs from that of an arbitrary guess or hunch. Evidence *E* can only play a role in *S*'s justification for believing *P* when *S* is aware of *E* while also aware of the evidential connection between *E* and *P*. However, we've seen that this awareness of the evidential connection must also be non-conceptual (refer back to 7.1.5 and 7.1.6). And condition (iii) of PJ captures this notion. Clauses A and B of (iii) then explain a direct and indirect way in which a subject might grasp an evidential relation between a set of facts and a proposition non-conceptually which corresponds to the distinction between foundational and inferential justification.

#### 7.4.3 The No Defeaters Clause: Condition (iv)

PJ allows that acquaintance with a fact or set of facts that *merely* make probable *P*, along with direct awareness of the probability relation, provides a subject with foundational

justification. Requiring that these facts do anything less than entail or make true a subject's belief, however, creates complications for formulating *sufficient* conditions for justification.

If a subset of the evidence I possess makes true or entails  $P$  then my entire evidence set makes true and entails  $P$ . This, however, doesn't hold true in the case of probabilistic relations. A part of my evidence might make  $P$  probable despite the fact that my total evidence makes probable  $\neg P$ . The fact that I'm having a perceptual experience as of a green lamp makes it likely that there is a green lamp in front of me.<sup>293</sup> Nevertheless, the larger set of facts {that I'm having a perceptual experience as of a green lamp, that I'm having an apparent memory as of taking a hallucinogenic drug, that I'm having a perceptual experience as of my friend testifying that there isn't a green lamp, that I'm having an apparent memory as of my friend refusing to take the hallucinogenic drug but agreeing to watch over me} makes it likely that there isn't a green lamp in front of me. Because of this, the move to probabilistic relations (rather than truth-making relations) forces us to include a kind of defeater clause. This is the purpose of condition (iv) in PJ.

I have foundational *prima facie* justification for believing  $P$  when I'm acquainted with a set of facts  $F$ , the thought that  $P$ , and  $F$ 's making probable  $P$ . I have defeating evidence for this justification if I'm acquainted with a larger set of facts  $F^*$  that includes  $F$  as a subset and such that I'm either directly or indirectly aware of the fact that  $F^*$  fails to make probable  $P$ . If there fails to be any such large set of facts  $F^*$  that meets these conditions, my *prima facie* justification is undefeated and I have *ultima facie* justification for believing  $P$ . This completes

---

<sup>293</sup> Strictly speaking, this is probably false. I'm very tempted toward the view that a perceptual experience as of a green lamp in front of me doesn't *by itself* make probable that there is a green lamp in front of me but, rather, it is only in conjunction with a coherent set of experience (including memorial experiences) that the perceptual experience would make this probable. This isn't important for the point I'm making in this paragraph and so I relegate this point to a footnote for the sake of easy presentation. However, this point will be incredibly important to my discussion of my solution to the problem of easy justification that arises in chapter 8.

my theory of justification and its relation to the notions of evidence, evidence possession, and evidential support.

### 7.5 Objections to the Acquaintance Theory

There are a number of objections philosophers might pose to the theory developed here. I want to very briefly consider what I take to be two of the most prominent objections: an over-intellectualization objection and etiological objections. My comments, unfortunately, must be brief.

#### 7.5.1 The Over-Intellectualization Objection

Even though I've defended the idea of fallible foundations where these may even include perceptual and memorial beliefs, the view I've defended will be closely associated with various classical foundationalist views. I still require that all justification has its foundation in a kind of factive awareness of, i.e. direct acquaintance with, various facts. Even in the cases of perceptual and memorial beliefs it's our acquaintance with fact about our non-doxastic perceptual and memorial experiences that make probable the relevant propositions that accounts for our justification. Moreover, I've argued that a subject will only have justification when they are aware of the evidential support relation itself. Such a view does strike one as quite demanding and many philosophers might accuse me of over-intellectualizing justification (despite my rejection of the need for infallible foundations).

My first form of response to this worry is to stress again the nature of acquaintance as a non-conceptual form of awareness. A subject S can be acquainted with a fact F even if S lacks the conceptual resources to represent that fact in S's thought. Consider unsophisticated children for instance. It's plausible that some younger children might be acquainted with the fact that they have certain visual, tactile, auditory, memorial, and a whole host of other experiences even if they cannot conceptualize in thought such facts. Children might be acquainted with these facts even if they haven't conceptualized the distinction between appearance and reality. Of course, lacking such a distinction it's unlikely that they have a



concept of a physical object and therefore couldn't form beliefs let alone justified beliefs about such physical objects. Nevertheless, the point illustrates that forming a belief on the basis of one's acquaintance with certain facts doesn't require that the subject actually form a representation of those facts.

In this same vein it's important to remember that even acquaintance with facts about evidential support relations can take place independently of the conceptual resources required to entertain a representation of that fact in thought. Now consider an unsophisticated child who has acquired concepts of shapes, colors, chairs, tables, etc. but who has not developed various epistemic concepts of justification, probability, epistemic support, evidence, etc. Such a child might still be acquainted with facts about her experiential life and acquainted with the fact that these make probable that there is a red table in front of her and thereby have justification for believing that there is a red table in front of her on my theory. Notice that the child need only be acquainted with the particular *instance* of the epistemic support relation holding between facts about her experience and the claim that there is a table in front of her. She needn't even possess the ability to conceptualize the facts about experiences as making her belief probable. Moreover, she needn't even be acquainted with anything like a *general* epistemic support relation that holds between experiences of this *kind* and propositions of this *kind*. In the same way various person might be acquainted with an entailment relation that holds in a particular case of modus ponens without being acquainted with the *general* validity of modus ponens. In fact, I imagine that acquaintance with particular instances of epistemic support is much more common than acquaintance with anything like the validity of general epistemic principles. This is why it's so helpful to beginning logic students to "see" the validity of an inference by inserting particular propositions for the "P"s and "Q"s.

So I think my theory does avoid certain forms of the over-intellectualization challenge. First, my theory allows for the possibility of a subject's having justification for beliefs about the external world and about the past that doesn't essentially depend on

entertaining representations of one's experiential states from which we must infer these beliefs. Second, the required acquaintance with evidential support need only be acquaintance with the support present in a *particular* instance. A subject needn't be acquainted with anything like the legitimacy of a general rule. And, finally, the acts acquaintances involved in my account of justification are all acts which I think unsophisticated subjects, including younger children, could and often do have and so even these subjects might have justification for many ordinary beliefs on my account.

### 7.5.2 Etiological Objections

A number of etiological objections have also been leveled at views similar to my own. On the developed theory, justification is determined solely by the facts with which the subject is *currently* acquainted. Various philosophers have suggested that such a view fails to appreciate how certain etiological facts can be positively and negatively relevant to a belief. I take this to be one of the most serious objections to my view but, unfortunately, my discussion must be brief. Here I'll simply indicate the kind of reply I'm inclined to give.

Suppose that a subject S is acquainted with various facts that provide her with justification for believing that humans haven't contributed to global climate change. Now imagine three possibilities: (i) S is only acquainted with this evidence because she was too lazy to do the research we think is required of people; (ii) S only has this evidence because (given her desires) she only read the studies that she already *knew* would provide her with supporting evidence for the belief she wanted to form; or (iii) S originally forms this belief and the basis of bad evidence, subsequently forgets the origin of her belief, and so now has *current* memorial evidence supporting her belief. Critics maintain that these possibilities show that a person's justification can be negatively affected by etiological facts even when her belief is supported by the facts with which she is *currently* acquainted.<sup>294</sup>

---

<sup>294</sup> See Baehr (2011); DeRose (2000), (2011); and Goldman (1999) for development of this kind of objection. Baehr and DeRose focus on the kind of irresponsible inquiry mentioned in the

The first thing to do in response to these kinds of worries is to note that we must be supposing that the person has no evidence to suggest that their belief has a bad etiology. In the first kind of situation it we must suppose that S either has no idea that she was lazy or no idea that this might affect the likelihood of her belief. In the second case we must suppose that the person is *unaware* of the fact that there are prominent scientists arguing the other side of the debate such that she would be justified in believing that reading their work would give her evidence that her belief is false. Finally, in the third case, we must suppose that the subject has no awareness that the belief was originally formed on a bad basis. If we don't assume these then responding is easy since the subject would have higher-order evidence that would defeat any justification she has.

Having made this caveat, I now want to suggest that the intuition that such a subject would lack justification arise out of inappropriately taking a 3<sup>rd</sup>-person as opposed to a 1<sup>st</sup>-person perspective on the beliefs. Consider what the beliefs would be like from the 1<sup>st</sup>-person perspective. The subject is directly acquainted with facts supporting her belief and directly acquainted with these making probable her belief. From her perspective the belief is incredibly likely (by hypothesis). What attitude other than belief ought she take towards such a proposition. Should she believe the opposite? Surely the answer is no. Should she withhold assent? I see no reason to think that this would be the right response given her perspective on the proposition since she has every reason to think the belief is likely to be true!

It's of course true that from *our* perspective the beliefs are not likely to be true. *We* are aware that the beliefs are formed via a procedure extended over time that is unlikely to get at the truth. But this means that we aren't justified in accepting these beliefs. Or, at the very least, we cannot treat this other person's belief like a thermometer in this case. *We* can't take S's believing *P* as evidence *for us* to believe *P* in these kinds of situations. Similarly, if S

---

first two possibilities. Goldman focuses on the kind of forgotten bad basis mentioned in the third possibility.

was aware of these facts about the etiology of her belief she would have a defeater but this would be determined by facts with which she is currently acquainted.

Finally, a related point is that the objections' 3<sup>rd</sup>-person evaluation of S's beliefs presupposes the 1<sup>st</sup>-person notion of justification I've been developing. Much of the time when we're considering other person's beliefs we're attempting to determine the likelihood of the beliefs so as to treat them as a kind of thermometer. If I have reason to think your beliefs will reliably get at the truth then I can use your beliefs as evidence that various claims are true. However, why would one be interested in this 3<sup>rd</sup>-person evaluation of another subject's beliefs? It seems that this would be in the service of *my* gaining 1<sup>st</sup>-person guidance or assurance for the beliefs that *I* form. My evaluation of your beliefs as likely to be true in some 3<sup>rd</sup>-person objective sense requires that I *currently have good evidence for this evaluation*. I must rely on facts *with which I'm currently acquainted* so as to arrive at any justification for believing that beliefs formed via certain kinds of inquiry aren't likely to be true. Far from challenging my theory of justification I think making sense of the purpose of these kinds of 3<sup>rd</sup>-person evaluations presupposes the kind of theory of justification I'm presenting here. These replies to the etiological objections are rather underdeveloped but expanding on these points and considering the relation between inquiry and 1<sup>st</sup>-person assurance could take up an entire dissertation. Here I've only indicated the general direction I would take in response to these (and similar) etiological objections to a theory that focuses on the subject's current states.

## CHAPTER EIGHT

### EASY JUSTIFICATION AND SKEPTICISM REVISITED

We've reached the culmination of this dissertation. In chapter 2 I presented a puzzle that any adequate theory of justification must resolve. Chapter 3 illustrated just how difficult our puzzle is by considering and rejecting various proposed solutions. Chapters 4-7 left our puzzle by the wayside so as to independently develop and defend an evidentialist theory of justification that ultimately culminated in my version of the acquaintance theory presented in the previous chapter. Various parts of these chapters have anticipated and hinted at the kind of solution that I'll offer to our puzzle. Now, however, the time has come to put these pieces together so as to illustrate the unique position that this acquaintance theory is in for resolving the issues raised by our initial puzzle.

There are two key pieces involved in my attempt to apply my version of the acquaintance theory in attempt to resolve our puzzle. First, evidential support is a *sui generis* and *irreducible* relation such that we can know certain truths about evidential support a priori (see chapter 6). Second, justification must always involve an awareness of the evidential support relation itself (see chapter 5) but it's important that this awareness be construed as something weaker than justified belief (see chapter 7). After briefly reviewing our puzzle I'll illustrate how these features help us navigate our way through our puzzle while avoiding impalement by either horn of the supposed dilemma.

#### 8.1 Reviewing Our Puzzle

##### 8.1.1 Easy Justification

Consider a belief source such as perception. How is it that a perceptual experience as of *P* can thereby provide us with justification for believing *P*? Must we have justification for believing that perception is reliable *prior* to its having the ability to confer justification on our

beliefs? Such a question can be asked about any belief source and so we must consider the status of the following general principle:

**JR** Necessarily, a belief source  $\Phi$  provides S with justification for believing  $P$  only if S already has justification for believing that  $\Phi$  is reliable.<sup>295</sup>

As explained in chapter 2, most theories of justification—internalist and externalist alike—reject JR. I refer to any belief source that constitutes a counterexample to JR as a “basic justification source.” Any theory that allows basic justification sources (i.e. denies JR) appears committed to allowing illegitimate gains in justification.

Assume that perception is a basic justification source. A perceptual experience as of a red table can thereby provide me with justification for believing that there is a red table in front of me without my possessing any prior justification for believing that such a perception would reliably be accurate. But now it seems that I need only introspect on my perceptual experience and I gain justification for believing that my perception was accurate in the current conditions. I need only reason as follows:

PERCEPTUAL CLOSURE REASONING

1. There is a red table in front of me (Justified by my perception as of a red table).
2. I have a perception as of a red table (Known introspectively).

---

3. Therefore, my perception is accurate on this occasion.

The evidence that provides me with justification for both premises is just my perceptual experience. These two beliefs then combine so as to provide justification for believing that this very perception is accurate. However, since the only evidence that provides justification for these premises is just the perceptual experience itself, it appears as if the perception is vouching for its own accuracy. It strikes us as incredible that a subject’s perceptual experience can be a source of information that that very perception is *contingently* accurate! It’s analogous to claiming that we might have no reason to believe that a person’s testimony

---

<sup>295</sup> “Already” indicates epistemic rather than temporal priority.

would reliably be accurate but that if she testifies that she is a reliable testifier we magically gain justification for believing her testimony is reliable.

Problems become more severe when we iterate this type of reasoning. It appears that, by iterating the reasoning, a set of perceptual experiences can justify a modal conclusion about the general reliability of perception. The structure of the supposed justification would be as follows:

#### EASY BOOTSTRAPPING ON PERCEPTION

1.  $P_1$  (justified by a perception as of  $P_1$ )
2. I have a perception as of  $P_1$  (known by introspection)
3. My perception as of  $P_1$  is accurate (from 1 & 2)
4.  $P_2$  (justified by a perception as of  $P_2$ )
5. I have a perception as of  $P_2$  (known by introspection)
6. My perception as of  $P_2$  is accurate (from 4 & 5)

[Repeat for  $P_3$  to  $P_N$ ]

---

7. Therefore, my apparent perception is reliable

Gaining justification for the general reliability of one's perception in this way seems all too easy and illegitimate.

Both the closure and bootstrapping style reasoning can be generalized to any supposed basic justification source: memory, testimony, etc. Criticizing the resulting structures of justification as "too easy", however, is too vague to be of much use. Such a generic statement fails to get at the heart of what is objectionable about this procedure and fails to pinpoint precisely what we want a theory of justification to avoid. As an example, we've already seen that the problem cannot be understood merely in terms of circularity since there are circular arguments that are (at the very least) *less* problematic (refer back to chapter 2 section 2.2.3.1).

When we consider the ways in which a subject is reasoning to the conclusion that a certain belief source is accurate or reliable in these easy justification reasoning scenarios, we realize that one could know in advance that this method would *inevitably* lead to a favorable conclusion. And this leads us to the conclusion that either (i) when the premises are justified in this way they don't provide the subject with justification for the conclusion or (ii) such reasoning is only available when the subject is aware of evidence for the conclusion that didn't essentially depend on the reasoning (otherwise it seems as if the person is "pulling a rabbit from the hat").

The problem with the first option is that the lack of justification for the conclusion when a subject goes through this kind of reasoning cannot be contained; if the subject goes through this reasoning and lacks justification for the conclusion then she lacks justification for the premises as well (refer back to my discussion of Vogel and Weisberg in chapter 3 section 3.5). The problem with the second option is that it appears that, where a basic justification source is concerned, what is required of the subject's perspective in order to gain this inductive evidence for the source's reliability doesn't in any way discriminate between a source being reliable or unreliable. So what evidence could the subject possibly have for the conclusion when such reasoning is available to her?<sup>296</sup> Moreover, it initially appears that this is to simply reject the possibility of basic justification sources—we'll see later that this impression is mistaken.

---

<sup>296</sup> We saw that the admission that the perceptual experience doesn't make the conclusion probable all by itself is Weisberg's way of defending the former option that a subject can go through the reasoning, have justification for the premises, yet fail to have justification for the conclusion—refer back to chapter 3 section 3.5.2 for a review of Weisberg's proposal and the problems with it.



### 8.1.2 The Problem of the Criterion

One response to the problem of easy justification is to simply accept JR. Such a move, however, appears to lead to global skepticism when conjoined with another plausible epistemic principle:

**JT** Necessarily, a subject S has justification for believing a source  $\Phi$  is reliable only if S already has justification for believing that particular deliverances of  $\Phi$  are true.

Whether a source  $\Phi$  is reliable is a contingent matter.<sup>297</sup> It seems straightforward that the only way we could gain justification for believing that  $\Phi$  is reliable is by inferring this from justified beliefs correlating  $\Phi$ 's delivering certain verdicts and those verdicts being true. How else could you gain justification for believing a source is *contingently* reliable?

How do JR and JT lead to skepticism? Consider a particular case of visual experience. Can a visual experience provide me with justification for believing  $P$  if JR and JT are both true? If we accept JR then vision only provides justification provided that I have epistemically prior justification for believing my vision is reliable. According to JT I can only gain justification for believing vision is reliable if I already have justification for believing that particular deliverances of vision are true. Clearly one must appeal to some further belief source if one is to acquire this justification *prior* to vision's having any justification-conferring ability. A natural suggestion would be to appeal to auditory experiences. But by JR, my auditory experience can only provide justification for believing that the deliverances of vision are true if I have *prior* justification for believing that particular deliverances of hearing are reliable. Since we've accepted JT, I'll need some third source  $\Phi$  to provide me with justification for believing that particular deliverances of my hearing are true. But the same reasoning will apply regarding  $\Phi$ 's justification-conferring ability. Clearly this reasoning will proceed *ad infinitum*. Even if I had an infinite number of belief sources, such a regress would

---

<sup>297</sup> This is complicated by the case of a priori justification but the complications don't affect my main points in this chapter.

still be vicious. If JR and JT are both true then there could never be an initial source that could, so to speak, “get the justificatory juices flowing.” JR and JT cannot both be satisfied. As such, accepting both principles leads to a global skepticism regarding justification.

### 8.2 Getting Re-Acquainted with the Acquaintance Theory

At this point I want to briefly recap the theory of justification presented in chapter 7. A subject S will have foundational *prima facie* justification for believing *P* iff S is acquainted with supporting evidence E for *P* and acquainted with the evidential connection between E and *P*. A subject S will have inferential *prima facie* justification for believing *Q* iff (i) S is acquainted with supporting evidence for *Q* and (ii) acquainted with a chain of probabilistic relations connecting F and *Q*, a chain which includes various intermediary propositions (refer back to 7.2 through 7.4).<sup>298</sup>

S will have either foundational or inferential *ultima facie* justification for believing *P* provided she has *prima facie* justification for believing *P* which is undefeated. My “no defeaters” clause isn’t essential to my resolution of our puzzle and so I ignore it in the remainder of my discussion (see 7.4.3 for a statement of the no defeaters clause).

### 8.3 Basic Inferences, Deduction, and Easy Justification

A solution to our puzzle becomes apparent when we consider how the apparent problem arises concerning the ability of deductive inference to provide justification. Having explained why the easy justification reasoning doesn’t lead to *problematic* justification for the conclusion in cases of deductive inference we will then be able to generalize the solution to foundational sources and to defeasible evidence.

---

<sup>298</sup> The idea here is that this justification is inferential because S’s awareness of the connection between F and *P* is an awareness that is essentially constituted by acts of acquaintance that provide *prima facie* justification for the intermediary propositions.

### 8.3.1 Basic Deductive Inferences

A basic justification source is any source that constitutes a counterexample to JR. This leaves open the possibility that a source might be both *inferential* and basic in the relevant sense.

Consider an application of enumerative induction from premises  $P_1$ - $P_N$  to the conclusion  $C$ . This will only confer justification on  $C$  if one has justification for believing  $P_1$ - $P_N$ ; in other words, applications of enumerative induction are a source of *inferential* justification. Nevertheless, this is consistent with holding that enumerative induction is a basic source of justification in the sense defined earlier. You can insist that you need justification for believing the premises of an inductive argument are true while also admitting that you needn't have justification for believing that an application of enumerative induction to true premises is reliable. On such a view, enumerative induction would be an inferential but basic justification source. We can refer to an inference pattern that can confer justification absent prior justification for believing that its application to true premises is reliable (i.e. absent justification for believing this source to be conditionally reliable) as a basic inference.

Our basic worry concerning easy justification arises just as easily if we allow for the possibility of basic inferences. I illustrated how the worry arises in the case of induction in chapter 2 and explained why this worry arises for any basic inference that is non-deductive (see 2.2.2 for a more detailed discussion). However, if there are no non-basic non-deductive inferences then it seems that a far reaching (though not global) skepticism looms. It seems that any justification for the conditional reliability of non-deductive inference must rely on non-deductive reasoning itself. And if there are no *basic* non-deductive inferences then we couldn't ever get the *prior* justification for the conditional reliability of a non-deductive inference required for such inferences to confer justification.<sup>299</sup> Our basic puzzle can be

---

<sup>299</sup> Here we are running up against Hume's ([1748] 1975) infamous problem of induction.

applied equally well to both foundational and inferential sources (though the key notion in the latter case is conditional reliability rather than reliability simpliciter).

It may be startling but the same apparent problem applies equally well to the case of deductive inferences! Navigating our way out of the puzzle in the case of deductive inferences illuminates a path to a more general solution.

### 8.3.2 Modus Ponens and Easy Justification

Consider one of the most plausible candidates for a deductive inferential justification source: an application of Modus Ponens (MP). Can an application to justified premises provide me with justification for believing  $Q$  absent prior justification for believing that MP is conditionally reliable? If yes, then MP is a basic inference.

MP's being a basic inference is extremely plausible. Even children might reason in accordance with MP and thereby extend their justified beliefs despite lacking the conceptual resources to even entertain the belief that MP is conditionally reliable. Let's grant that my nephew Jackson—who is five years old—knows that his Mom is working tonight and that he knows that if his mom is working then his Dad will stay home with him and his brother Max. On the basis of his knowledge, Jackson can come to know that his Dad will stay home with him and Max. Nevertheless, it would be a gross over-intellectualization to suggest that Jackson can even contemplate the structure of this inference let alone form a belief about its conditional reliability.

I take this to provide prima facie reason to think MP is a basic inference. Now imagine that I form a belief  $Q$  on the basis of an application of MP to foundationally justified beliefs. MP can provide justification for  $Q$  even if I lack justification for believing that MP is conditionally reliable. But now all I need to do is find out that my belief that  $Q$  is the result of an application of MP to true premises and this will provide me with me justification for believing that such a procedure didn't undetectably lead me astray on this

occasion. At which point I can iterate this reasoning and gain inductive evidence for the general reliability of applying MP to true premises:

MP EASY BOOTSTRAPPING REASONING

1.  $Q_1$  (Justified by an application of MP to justified premises)
2. I formed my belief that  $Q_1$  on the basis of an application of MP to true premises<sup>300</sup>
3. Forming my belief that  $Q_1$  belief via an application of MP to true premises didn't undetectably lead me astray on this occasion (From 1 & 2).
4.  $Q_2$  (Justified by an application of MP to justified premises)
5. I formed my belief that  $Q_2$  on the basis of an application of MP to true premises.
6. Forming my belief that  $Q_2$  belief via an application of MP to true premises didn't undetectably lead me astray on this occasion (From 1 & 2).

[Repeat for  $Q_3$ - $Q_N$ ]

---

7. Therefore, forming a belief via applications of MP to true premises is generally reliable.

A person who reasoned in this way would surely strike us as odd since, for most of us, the conclusion is knowable a priori. A person's using an empirical track-record argument is silly given that we have such a better means of justifying the conclusion. It would be like someone who gains justification for believing that all bachelors are unmarried by doing a survey of bachelors and asking how many of them are unmarried. Even provided that the inductive argument supports and can provide justification for the conclusion, if you have the ability to even entertain the conclusion in thought then there is no need for such an

---

<sup>300</sup> Notice that if you have justification for believing the premises of an application of MP then you only need to gain justification for believing that a belief  $Q$  was formed by an inference that is an instance of MP in order to gain justification for believing that  $Q$  was formed on the basis of an application of MP to true premises. Moreover, MP is an inferential source and, therefore, you'll always have justification for believing the premises of an application of MP when that application provides you with justification for believing the conclusion.

inductive argument; this inductive reasoning could never provide you with any more justification for the conclusion than was available to you a priori.

Nevertheless, however odd such reasoning appears, *if* MP is a basic inference (and it appears we should accept this for the reasons considered above) and induction can confer justification *then* we're committed to the idea that the MP EASY BOOTSTRAPPING REASONING generates justification for its conclusion. And at this point I want to make three points that I think any discussion of bootstrapping on MP ought to respect.

First, as I've mentioned, we should not deny the justificatory force of inductive reasoning. Something makes us uneasy with the idea that someone could gain justification for believing MP is conditionally reliable via this bootstrapping reasoning but we shouldn't attribute this uneasiness to an uneasiness concerning induction in general. Surely if we had independent confirmation for the deliverances of MP (e.g. imagine a case where we had perceptual evidence providing justification for  $Q_I$ - $Q_N$ ) we would want to admit that the inductive argument could generate justification for the conclusion (provided we had justification for believing the premises).

Second, as with all examples of easy justification, our uneasiness derives from the feeling that the subject is somehow spinning her wheels. It isn't just that the justification is circular. What makes us uneasy with the idea that such reasoning can provide justification that wasn't present prior to the reasoning is that an attempt to gather inductive evidence regarding the reliability of MP *in this way* is guaranteed to produce premises that support the favorable conclusion. Therefore, provided that the application of MP really does provide justification for believing  $Q$ , there must have been *some sense* in which any justification produced by this reasoning was already available to the subject when she formed her initial belief that  $Q$  via MP.

Finally, we don't want to commit ourselves to the idea that a subject must have a prior justified belief that MP is conditionally reliable in order for an application of MP to be capable of providing justification. One reason was already mentioned in connection with

children being able to use MP to extend their knowledge; such a move appears to over intellectualize justification. More importantly, however, is that such a move invites the Lewis Carroll style regress concerning inferential justification. Our discussion of inferential internalism in 7.3 has already suggested that this but it will be helpful to remind ourselves how the regress arises.

How does such a move lead to the Lewis Carroll regress? The move, remember, is that the justified beliefs that  $P$  and *if  $P$  then  $Q$*  cannot *by themselves* provide a subject with inferential justification for believing  $Q$ . A subject must also have a justified belief that forming the belief that  $Q$  on the basis of  $P$  and *if  $P$  then  $Q$*  is conditionally reliable.<sup>301</sup> But now it seems that the subject's inferential justification for  $Q$  actually depends on inferring this from (a)  $P$ , (b) *if  $P$  then  $Q$* , and (c) *MP is conditionally reliable*. Label this new inference pattern MP+. Even if we admit that (c) can be known a priori we run into the same problem concerning the justificatory power of MP+. If MP+ is a basic inference then the easy justification reasoning can be used to produce justification that MP+ is conditionally reliable. And there seems no principled reason to respond differently to the challenge of easy justification in the cases of MP and MP+. Therefore, the same move would be required in this latter case and so we should also claim that MP+ only provides justification for  $Q$  if the subject has justification for believing it's conditionally reliable. And so our reasoning repeats *ad infinitum*. If we hope to avoid the Lewis Carroll regress we must hold out hope for a vindication of the possibility of basic inferences.

---

<sup>301</sup> Our framing of the Lewis Carroll regress in the context of this discussion is slightly different than the one considered in connection with Fumerton's inferential internalism. The reason for this is that here we are challenging the idea that a justified higher-order belief about the reliability of an inference is necessary for inferential justification whereas in connection with Fumerton's theory we were concerned with the idea that a justified higher-order belief about the inferential basis making probable (where this relation is *stripped* of all empirical content) the inferred proposition is necessary. These differences, however, are not important for the main thrust of the argument.

Getting out of our puzzle requires showing how to make these second and third points consistent. What our investigation of the puzzle of easy justification illustrates is that we need a theory of justification that simultaneously vindicates each of the following claims:

- (A) An application of MP can provide S with justification for believing  $Q$  absent prior justification for believing MP is conditionally reliable.
- (B) Necessarily, if an application of MP does in fact provide S with justification for  $Q$  then justification for the reliability of MP is *in some sense* already available to S.

The key to our solution requires providing a satisfactory interpretation of the sense in which the justification is “already available” to the subject in a way that falls short of denying (A).<sup>302</sup>

### 8.3.3 Getting Acquainted with the Way Out

In explaining my way of resolving our puzzle in the case of MP I want to begin by revisiting this idea that the reasoning strikes us as odd since the subject clearly could have had better *a priori* justification for the conditional reliability of MP than any justification she might gain via this empirical track-record argument for the same conclusion.

As I said earlier, the need for an empirical track record argument is silly given that we have a better means of justifying the conclusion. In fact, the *degree* of support provided by this a priori means is the *limit* of the degree of support that would be garnered as the number of instances in the inductive argument approaches infinity! It would be like someone who gains justification for believing that all bachelors are unmarried by doing a survey of bachelors and asking how many of them are unmarried. Surely the inductive argument supports and can provide justification for the conclusion. However, if you can form a belief about bachelors then you have the concept of a bachelor, and if you possess the concept of a

---

<sup>302</sup> Wright’s notion of entitlement or non-evidential justification solution and Sosa’s attempt to claim that there was *implicit* animal but not reflective justification for the reliability of a source attempt to do just this. Refer back to 3.3 and 3.4 respectively. However, we’ve already seen that the particular ways in which they utilize these notions in an attempt to vindicate the consistency of (a) and (b) encounter a host of problems.



bachelor then there is no need to use such empirical inductive reasoning since it cannot provide any justification that you can't get for the conclusion a priori by simply examining your concept of a bachelor. If you can entertain the justified belief in one of the track-record argument's premises that so-and-so is a bachelor then you're already aware of a priori evidence that supports the conclusion.

Something similar happens in the case of the easy justification reasoning on MP. A subject S's justification for the empirical premises involved in easy justification reasoning produces justification for the conclusion *but only to the extent that S's justification for the deliverances of MP already includes S's awareness of a priori evidence for this conclusion*. As such, the same justification could have been provided without combining one's belief that  $Q$  with the further belief that one used an application of MP in forming this belief. Combining one's beliefs that  $Q_1$ - $Q_N$  with the justified beliefs about what applications of MP testify to (i.e. the 2<sup>nd</sup>, 5<sup>th</sup>, etc. premises of the bootstrapping reasoning) doesn't produce any new support for the reliability of MP; it only illustrates the support provided by a priori evidence for this conclusion of which the subject was already aware *when forming the beliefs that  $Q_1$ - $Q_N$* .<sup>303</sup>

Notice that most theories which deny JR cannot vindicate this idea. A common view of how an application of MP would confer inferential justification is that simply in virtue of having justification for believing  $P$  and *if  $P$  then  $Q$*  and the mere existence of the entailment relation provides S with justification for  $Q$ . One reason you might accept such a view is that requiring a further justified belief about the legitimacy of this inference leads to the Lewis Carroll regress and seems like a gross over-intellectualization of justification. This denial of JR, however, cannot accommodate the idea that forming a justified belief that  $Q$  *necessarily*

---

<sup>303</sup> In a case where one has *independent* justification for  $Q_1$ - $Q_N$  the combination can generate new justification for the conclusion that wasn't *in some sense* already available to the subject in forming the justified beliefs  $Q_1$ - $Q_N$ . Since, by hypothesis, the justification for  $Q_1$ - $Q_N$  didn't result from an application of MP the subject wasn't required to be aware of the entailment relation. The ties into the point of the next paragraph that views that don't require awareness of the entailment relation cannot accept such a solution.

includes an awareness of evidence for MP's conditional reliability. The only thing you need be aware of on the views currently under consideration is  $P$  and *if  $P$  then  $Q$* . But surely these propositions aren't to be construed as evidence for MP's conditional reliability!

The choice between the denial of JR just sketched and an acceptance of JR, however, overlooks a middle road that my version of the acquaintance theory makes use of in its account of inferential justification. Requiring a justified belief about the legitimacy of an inference is too demanding for a theory of inferential justification and will ultimately lead to vicious regress. But one shouldn't rashly jump to the conclusion that *no* awareness of the legitimacy of this inference is required. We can adopt the weaker requirement that the justified beliefs that  $P$  and *if  $P$  then  $Q$*  only provide inferential justification for  $Q$  if you're acquainted with the fact that the first two propositions jointly entail the latter.<sup>304</sup> Acquaintance isn't a representational state. Nor is it an epistemic relation. Acquaintance is a real relation of direct awareness that holds between a subject and a fact, state, property, or event (refer back to chapter 7 section 7.1.1). It's possible that even children stand in this relation to entailment relations. And it's quite plausible that while they might not be acquainted with *general* facts about the legitimacy of certain *general* forms of inferences, children are often aware of *particular* instances of entailment relations holding between *particular* propositions. It's only with such a theory that the response I've briefly sketched (and develop further below) to the problem of easy justification becomes available.

Notice that the fact that the propositions  $P$  and *if  $P$  then  $Q$*  together entail  $Q$  is itself conclusive evidence that MP is conditionally reliable. Therefore, on my version of the acquaintance theory applied to inferential justification, *necessarily* an application of MP provides a subject S with justification for believing  $Q$  only if S is aware of evidence for the conditional reliability of MP. And so this seems a clear sense in which principle (B)

---

<sup>304</sup> Surely a justified belief would be sufficient. The key here is that making acquaintance sufficient prevents the Lewis Carroll regress.

presented at the end of 8.3.2 is true on my acquaintance theory. The justification for MP's conditional reliability is available in the sense that the subject is already aware of the a priori evidence that supports such a claim. However, it's important to note that this is consistent with the idea that MP is a basic inference. Awareness of evidence for *P* is insufficient for justifying *P*. Justification only arises when one is also aware of the evidential connection. And this is the case even when the proposition under consideration is the reliability of some evidence. Thus, the fact that MP only confers justification on *Q* if a subject is *aware of evidence* for MP's conditional reliability is consistent with the idea that a subject needn't have *justification* for believing that MP is conditionally reliable.

Extending this solution to the case of foundational sources of justification where the foundational evidence *entails* the truth of the relevant belief is straightforward. Consider a case where my belief that I'm having a color experience is based on my awareness of the *fact* that I'm having a perceptual experience as of a red ball. Acquaintance with the fact F <that I have a perceptual experience as of a red ball> and my acquaintance with the fact F\* <that F entails that I'm having a color experience> provides me with *foundational* justification for my belief about the more determinable property.<sup>305</sup> These two acts of acquaintance are sufficient for having justification for my belief. I needn't, in addition, have a prior justified belief that the fact that one has a perceptual experience as of a red ball is a reliable indicator of the truth of the claim that one is having a color experience.

At this point, however, I might engage in the easy justification reasoning to gain justification for the conclusion that forming the belief that I'm having a color experience on the basis of the fact that I'm having a perceptual experience as of a red ball is reliable. But

---

<sup>305</sup> Moving from F to the belief that I'm having a color experience does amount to (*in some sense*) taking F as evidence for my belief and making a kind of "quasi-inference". The resulting justification, however, isn't inferential since it doesn't involve inference from a truth-bearer but rather a truth-maker. Since the evidence is a fact there is no need for the basis of my belief to be justified. The fact evidentially justifies my belief that I'm having a color experience even though the fact itself is beyond justification.

given the conditions for foundational justification proposed by my version of the acquaintance theory this shouldn't strike us as problematic. My foundational justification for believing that I'm having a color experience already includes *as a constituent* my acquaintance with a fact F\* that evidentially supports this conclusion. Combining my belief that I'm having a color experience with a justified belief that this belief is formed on the basis of the fact F can only provide me with a degree of justification equal to or less than the extent to which facts (i.e. F\*) already within my ken support such a conclusion. The easy justification reasoning merely illustrates that I was aware of supporting evidence for such a conclusion when forming my original belief about my having a color experience.

It will help to relate this back to my distinction from chapter 5 between two ways a subject can possess evidence for *P* (see 5.3.1). If a subject is aware of evidence E for *P* but isn't actually aware of the evidential connection then E is part of S's unconnected evidence for *P*.<sup>306</sup> If a subject is aware of E and aware of E's supporting *P* then E is part of S's actually connected evidence for *P*. My version of the acquaintance theory is a more fully developed version of the idea that S's having justification for believing *P* requires that S have actually connected supporting evidence for *P*. When this is combined with my view of epistemic support modeled on the Keynesian view of probability we avoid the problem of easy justification because S's having actually connected supporting evidence E for *P* entails that S has unconnected supporting evidence that forming a belief that *P* on the basis of E is reliable. An awareness of evidence for the conclusion about reliability is already present in S's justification for believing the object-level proposition. The easy justification doesn't involve acquiring an awareness of any new pieces of evidence but rather involves teasing out

---

<sup>306</sup> Unconnected evidence was further divided into unconnected evidence simpliciter and potentially connected evidence. This distinction, however, isn't crucial for the discussion that follows. What is important is the distinction between the more general category of unconnected evidence and actually connected evidence.

the fact that items of which we were already aware evidentially support the meta-level proposition.

In order for some evidence  $E$  to provide  $S$  with justification for believing  $P$ ,  $S$  must be aware of  $E$  and aware of  $E$ 's evidentially supporting  $P$ . That  $E$  evidentially supports  $P$  itself evidentially supports the claim that  $E$  is a reliable indicator of  $P$ . Thus, on my version of the acquaintance theory, *necessarily*, some evidence  $E$  provides  $S$  with justification for believing  $P$  *only if*  $S$  is aware of evidence that  $E$  bears an appropriate reliability connection to  $P$ 's truth. This is what we might call the Evidential Reliability (ER) principle:

**ER** Necessarily, a belief source  $\Phi$  (i.e. some evidence  $E$ ) provides  $S$  with justification for believing  $P$  only if  $S$  is aware of evidence for  $\Phi$ 's reliability.

The acquaintance theory's vindication of ER is what allows us to avoid the problem of easy justification. When you move up a level to a belief about the accuracy or reliability of your evidence via the particular kind of reasoning used in bootstrapping you gain no more assurance concerning the likelihood of your belief's truth than was already present prior to moving up a level. When I'm acquainted with a fact  $F$  and with  $F$ 's entailing  $P$  I have the very property of  $P$ 's truth being guaranteed before my mind; my perspective on  $P$  provides a kind of full assurance of  $P$ 's truth. Utilizing the bootstrapping reasoning to reach the further conclusion that forming a belief that  $P$  on the basis of  $F$  was accurate or even that this is perfectly reliable doesn't bring any information before my mind that would improve my perspective on  $P$ 's truth. In fact, any justification a subject might get for such a conclusion is parasitic upon support derived from my awareness of the entailment relation.<sup>307</sup> Just as important, however, is that none of this implies an acceptance of JR. Being aware of evidence for  $P$  (including claims about accuracy or reliability) is insufficient for having justification or believing  $P$  on the acquaintance theory I've been advocating.

---

<sup>307</sup> This isn't to say that there isn't some sort of valuable achievement in moving-up a level. Acquiring the conceptual resources to move up a level might very well be something valuable but the point is that my perspective on the truth of the first-level proposition isn't improved by moving-up a level even if other things of value are in place when I move up a level.

Our discussion shows that the illegitimate gains in justification present in cases of easy justification are not due to denying JR but rather ER. My acquaintance theory vindicates ER and thereby avoids the problem of easy justification. Nevertheless, it's important to see that most theories that deny JR do so in a way that also commits them to a denial of ER and thereby fall prey to the problem of easy justification.

Traditional reliabilist theories deny JR. On the most radical versions of reliabilism what matters is the *de facto* reliability of the belief-forming mechanism. One needn't be aware of the input to that process, aware of the reliability of that process, nor be aware of any reason to believe that such a process is reliable. Clearly such a view denies ER. And even more sophisticated versions of reliabilism such as Alston's (1988) indicator reliabilism and Comesaña's (2010) evidentialist reliabilism that require some form of awareness or access to evidence for one's belief eschew any requirement that a subject have any awareness of or evidence for the adequacy (i.e. reliability on this view) of this evidence. Such views, therefore, appear to fall victim to the problem of easy justification.

Even many theories commonly characterized as internalist deny ER along with their denial of JR and thereby fall prey to the problem of easy justification. For instance, according to phenomenal conservatism, its seeming to S that *P* provides S with defeasible justification for believing *P*. This, however, doesn't entail that the subject is aware of any evidence that seemings are reliably correlated with the truth of their contents. One could, therefore, use the closure or bootstrapping reasoning to arrive at the conclusion that a seeming is accurate or generally reliable. But since this view fails to vindicate the idea that necessarily when a source (i.e. the seeming) provides a subject with justification a subject is aware of evidence for the reliability of that source it's utterly mysterious from whence the justification for the conclusion of the easy justification reasoning derives. Such a view falls prey to the problem of easy justification.

## 8.4 Generalizing the Solution

My discussion of a way out of our puzzle thus far has only considered the justificatory power of entailing evidence and of deductive inference. It remains to be seen whether this solution can be extended to the justification conferring ability of non-deductive inference and defeasible foundational evidence. One might think that something very important changes when we begin to consider these latter kinds of cases.

### 8.4.1 Irreducible Epistemic Support's Role in Our Solution

The key to extending our solution to the justification-conferring ability of non-deductive inference and defeasible foundational evidence lies in recognizing the role of irreducible epistemic support in my proposed solution.

In chapter 6 I argued that we should model our view of epistemic support on the Keynesian interpretation of probability. The Keynesian interpretation of probability is such that entailment relations holding between propositions are merely the upper limit of the making probable relation. Similarly, an entailment relation that holds between a fact-proposition pair is the upper limit of the epistemic support relation.

Facts about entailment are most definitely not to be reduced to facts about relative frequencies, not even to facts about relative frequencies across all possible worlds. However, we've seen that facts about the existence of entailment relations do *entail* certain claims about relative frequencies. Important to our solution to the problem of easy justification as applied to Modus ponens, deductive inference, and entailing foundational evidence is the following relation:<sup>308</sup>

---

<sup>308</sup> "E" represents a fact, "P" represents a proposition, the arrows with "ENTAILS" in the middle represents the entailment relation, the arrow with a "RELIABLE" represents a reliable indicator relation, and the boxes gather the constituents together into either a fact or a proposition as indicated below the box.



Figure 7: Entailments Entailing Reliability

It's the existence of this internal relation *necessitated* by the intrinsic natures and existence of the relata that vindicates principle ER and thereby resolves the problem of easy justification in cases where a subject's justification rests on entailing evidence or a basic deductive inference such as Modus Ponens.

In order to extend our solution to cases where one's justification depends on less than conclusive evidence for one's belief we need to posit a similar relation between the existence of merely probabilistic relations and the existence of reliability relations. This, however, seems more difficult since the existence of a Keynesian probability relation *does not entail* the existence of a reliability relation. Keynesian probability relations are necessary relations that exist in every possible world where the relevant relata exist. Assume that a perceptual experience as of  $P$  makes it probable that  $P$ . This relation holds even in a world where you're the victim of a Cartesian demon and perception isn't reliable. And similar remarks can be made about supposed defeasible evidence for beliefs about the past and future.<sup>309</sup> If a fact  $F$  makes probable  $P$  but doesn't *entail*  $P$ 's truth then it's surely possible that  $F$  isn't a reliable indicator of  $P$ . The fact  $\langle E \text{ is } \textit{defeasible} \text{ evidence for } P \rangle$  doesn't entail that  $E$  bears an appropriate reliability connection to  $P$ 's truth. So how could there be a

---

<sup>309</sup> Assume that my memory experiences make probable that I previously broke my leg skiing (this is one of the most vivid memories I have and coheres greatly with my other memorial experiences). Such a relation would still hold even if Russell's 5-minute hypothesis were true. In such a world, however, my memorial evidence wouldn't be reliable. Alternatively, imagining worlds where nature fails to be uniform (though it just randomly happens to have every appearance of uniformity while I am alive) can be used to make similar points about any supposed inductive evidence  $E$  making probable  $P$ .



necessary relation connecting the fact that  $E$ 's making probable  $P$  and the truth of the claim that  $E$ 's is a reliable indicator of  $P$  that would thereby vindicate ER in the case of non-deductive inference and defeasible foundational evidence?

Remember that entailment isn't a relation distinct from the epistemic support relation. Entailment is just a determinate of this determinable: it's the upper limit of epistemic support. All epistemic support relations are like entailment in being internal relations that hold between a fact-proposition pair and that are not reducible to facts about relative frequencies (or to facts about nomological relations that make true claims about relative frequencies in nearby possible worlds). Our motivation for this claim about irreducibility comes from our discussion of the new evil demon problem discussed in chapter 6.

However, if entailment is just the upper-limit of epistemic support then, even if the existence of epistemic support doesn't necessitate the existence of an entailment relation between this fact and a claim about the existence of a reliability relation, we should construe the relation depicted in figure 7 as a particular instance of a general feature of epistemic support relations. In other words, we should construe the relation depicted in figure 7 as just a particular instance of the fact that *necessarily* the existence of a Keynesian probability relation *makes probable* (in the Keynesian sense) the existence of a reliability relation:<sup>310</sup>



Figure 8: Probability Making Probable Reliability

---

<sup>310</sup> The interpretation of figure 8 is the same as figure 7 except that the arrow with a “K-PROBABLE” subscript (a reference to its Keynesian origin) represents the more general epistemic support relation (of which entailment is the upper limit).

*Assume* that the fact that one has a perceptual experience as of  $P$  makes probable  $P$  in this Keynesian inspired sense. This itself *makes probable but does not entail* that the perceptual experience as of  $P$  is a reliable indicator of  $P$ 's truth. Like all epistemic support relations, this is an internal relation necessitated by the natures and existence of its relata (e.g. the fact that the perception makes probable  $P$  and the proposition that the perception is a reliable indicator of  $P$ ). Therefore, since my acquaintance theory holds that a necessary condition for some evidence  $E$  providing a subject  $S$  with justification with justification for believing  $P$  is that  $S$  be aware of the evidential connection between  $E$  and  $P$ , whether it be conclusive or non-conclusive, such a theory vindicates ER even when one's justification depends on defeasible foundational evidence or non-deductive inference.

In this section I have sketched how my solution in the case of Modus Ponens, deductive inference, and entailing foundational evidence depended on the idea that entailment is an irreducible form of the epistemic support as defended in chapter 6. I then illustrated how we to extend this idea to cases that would cover non-deductive inference and defeasible foundational evidence. However, all of this is a bit abstract at the moment. In the next section I illustrate how one could have a priori evidence that makes probable *contingent* reliability claims. I also illustrate and defend the defeasibility of the *a priori* evidence for contingent reliability claims. In section 8.4.3 I will then consider the particular cases of induction, perception, and memory in order better to illustrate how my proposed solution extends to these cases.

#### 8.4.2 Defending Defeasible A Priori Evidence for Reliability

In the previous section I explained that extending my solution so as to explain how to avoid the problem of easy justification in cases where one's justification depends on a basic non-deductive inference or non-entailing foundational evidence requires holding that the existence of an a priori knowable probability relation makes probable that a certain

reliability relation holds. But this would appear to give us a priori evidence supporting a contingent claim! This intuitively strikes us as incredibly counterintuitive and the reader would be right to ask for examples illustrating such a possibility.

John Hawthorne (2002) has presented a plethora of cases that he takes to establish the possibility of contingent a priori knowledge. The most interesting example he considers is a case involving inference to the best explanation. Given my focus on justification, however, I'll modify the example so as to discuss justification. Hawthorne asks us to imagine a weird subject S who, prior to any kind of sensory experience, reflects upon various possible experiential lives. He suggests that (provided S has the conceptual resources to reflect on the possible experiential lives) S could come to know a priori that a particular experiential history E would be best explained by some hypothesis H. Let's grant this for the moment. Hawthorne then contends that on the basis of this knowledge that H would be the best explanation of E, S thereby has justification for believing the *contingent* material conditional that *if E then H*. The idea here is that a necessary truth about explanatory relations can be knowable a priori and then act as evidence for a related contingent truth. Such evidence for the contingent truth, however, still derives from the subject's awareness of various conceptual necessities and is thereby a priori.

The problem with Hawthorne's example is a problem common to many attempts to justify a belief on the basis of an inference to the best explanation. There might be an infinite number of possible explanations for E. H might be more probable relative to E than any of the other individual explanations *even if H is still incredibly improbable relative to E*. Consider a situation where the competing explanations of E are  $H_1$ ,  $H_2$ ,  $H_3$ , and  $H_4$ . Now imagine that  $\Pr(h_1|e) = .4$ ,  $\Pr(h_2|e) = .2$ ,  $\Pr(h_3|E) = .2$ , and  $\Pr(h_4|e) = .2$ . Moreover, imagine that conditional probabilities are knowable a priori. A subject S might thereby know a priori that  $H_1$  is the best explanation of E (it's the most probable explanation). However, it doesn't seem that this a priori knowledge makes probable the contingent material conditional *if E*

then  $H_1$ , and so the fact that  $H_1$  is the best explanation of  $E$  cannot act as a priori evidence that can provide  $S$  with justification for believing this conditional claim.

Nevertheless, even if Hawthorne's example doesn't illustrate a case where a certain fact can provide us with a priori evidence that might justify us in believing a contingent claim I do think it illustrates the possibility of the phenomena needed to extend my solution to the problem of easy justification. Awareness of the fact that certain conceptual relations hold that make  $H_1$  the best explanation of  $E$  would still provide an awareness of a fact that constitutes a reason to place *some* confidence in the contingent conditional even if it didn't support believing the conditional outright. Hawthorne's example still shows that there could be a priori evidence that bears some epistemic relevance to the truth or falsity of various contingent claims.

Another example of the general phenomenon comes from John Turri. He offers the following example:

Sam considers whether the most unlikely possible event is not presently occurring. By 'the most unlikely possible event', Sam intends to designate whatever was, at the immediately preceding instance,  $t-1$ , the possible event most unlikely to occur at the next instance,  $t$ , which is the moment at which her deliberation occurs. Sam understands the proposition in question. Solely in virtue of that understanding...she intuits... that the proposition is true, though not necessarily so. On the basis of this intuition, she believes that the most unlikely possible event is not presently occurring.<sup>311</sup>

I take it that Turri's example illustrates the possibility of a person's having an awareness of certain evidence that could (and in this case *does*) provide her with a priori justification for a believing a contingent claim. My story, however, I imagine is very different than the one Turri would give. For one I'm highly suspicious of Sam's 'intuiting that the proposition is true.' What Sam can *intuit* is that such a proposition is probably true. What is the a priori evidence for this claim's probable truth? Here Turri identifies the key to this example: it

---

<sup>311</sup> Turri (2011), pp. 337-338.

focuses on a “proposition that is overwhelmingly likely to be true *as a matter of conceptual necessity*.”<sup>312</sup> On the acquaintance theory of justification defended it is Sam’s acquaintance with the fact that certain conceptual relations hold and her awareness of this fact evidentially supporting *but not entailing* the claim that the relevant event isn’t occurring. These conceptual relations of the proposition *necessitate* the truth of a claim about the high probability that the most unlikely possible event is not occurring. In virtue of this we can see that the conceptual relations also *make probable* this proposition minus the probability operator, which is a contingent proposition since it’s surely possible that the most unlikely event is occurring.

These examples help illustrate the possibility of a priori evidence (i.e. facts about conceptual relations) supporting contingent truths. However, we haven’t yet seen a plausible example of a contingent claim about some sort of reliable correlation. I take it, however, that it isn’t difficult to imagine extending Turri’s example so as to apply it across time. Is it not a conceptual necessity that *it’s probable that, reliably, the most unlikely event of each moment (understood as indicated in the quote from Turri) does not occur at that moment?* If this is right then don’t the conceptual relations necessitating such a truth make probable that *reliably, the most unlikely event of each moment does not occur at that moment?* This reliability claim, however, is contingent. It’s incredibly unlikely but it could turn out that the most unlikely event of each moment more often than not occurs (of course we’ll have to construe the most unlikely event in terms of counterfactual frequencies rather than actual frequencies). As such, acquaintance with the fact that this proposition stands in these conceptual relations would be an awareness of evidence supporting the contingent reliability claim.

Another example has been suggested to me in conversations with Richard Fumerton that is especially useful for connecting up the idea with the non-reductive theory of epistemic support presented at the end of chapter 6. Imagine that a God presents Siri with

---

<sup>312</sup> Turri (2011), p. 338.

an infinitely large urn. For every natural number there is one sheet of paper inside the urn with that number written on it and (being omnipotent) God makes it so that Siri can easily reach any of these numbers by reaching her arm in to the urn. God then shakes up the urn so that the natural numbers are randomly placed throughout the urn. Now imagine that Siri reaches into this infinite urn and draws a number from inside. Intuitively it seems right to say that there are more possible worlds where Siri draws a number other than 7 than there are possible worlds where Siri draw a 7. Moreover, it also seems that among the possible worlds such that Siri draws a number, places it back in the urn, and draws another number there are more where Siri reliably choose a number other than 7 than there are possible worlds where this isn't the case. It's a priori likely that these facts about the urn would be reliably associated with Siri's tending to choose a number other than 7.

Finally, consider the possibility of a priori evidence supporting the conditional claim that if a fair die is rolled 1,000 times and lands with one of the six sides face up then it will not usually land with the 6-side face up. You can know *a priori* the conditional claim that *probably* if the six sides are equally likely to land face up on a roll then if this die were to be rolled 1,000 times and land with one of the six sides face up each time it would not usually land with the 6-side face up. But if you know that this conditional is probably true can't you drop the probability operator and belief the conditional claim outright (perhaps with a credence corresponding to its probability). Intuitively, awareness of this probabilistic fact could provide me with justification for believing that if the six sides are equally likely to land face up on a roll then if this die were to be rolled 1,000 times and land with one of the six sides face up each time it would not usually land with the 6-side face up. Notice, however, that I know very well that this conditional has a non-zero probability. I know very well that it's entirely possible for a fair die to be rolled 1,000 times and usually land with the 6-side face up. In fact, it's *possible* for a fair die to be rolled 1,000 times and *always* land with the 6-side face up. Nevertheless, my awareness of the improbability of this happening provides me with justification for believing the contingent conditional. There are more possible worlds

where a fair die is rolled 1,000 times and usually lands with something other than the 6-side face up than there are possible worlds where a fair die is rolled 1,000 times and usually lands with the 6-side face up.<sup>313</sup>

Possible world talk, however, should be taken as metaphorical. We want to find a truth-maker for this claim in the actual world. What would make such a claim about possible worlds true? My suggestion from chapter 6 is that the truth-maker for such a claim are the a priori probability relation (if it exists) that holds between the propositions about God's urn and that Siri randomly drew a number from this urn make probable the proposition that she drew a number other than 7. Similar claims apply *mutatis mutandis* to the dice case. Facts about frequencies across *all possible worlds* are made true by Keynesian relations holding among propositions in the actual world. The entailment relations in the actual world (the upper-limit of this making probable relation) are the truth-makers for the claim that there is a *perfectly* reliable correlation across all possible worlds. The probability relations in the actual world are truth-makers for the claim that there is an imperfect reliable correlation across all possible worlds. The relevant probability relation that exists in the actual world is just the kind of probability relation represented in figure 8.

I want to be clear, however, that this a priori justification for the contingent reliability of evidence is defeasible. The fact <E makes K-probable P> might make K-probable (and thereby be evidence for believing) *that E is a reliable indicator of P* even though one has further evidence that E isn't a reliable indicator of P. Assume again that the fact that S has a visual experience as of P makes probable P's truth. On my view this would make probable that visual experiences as of P are reliably associated with P's truth. However, this reliability relation might not be probable relative to evidence that includes the fact that my

---

<sup>313</sup> In these last two examples I'm not committing myself to the truth of a principle of indifference. These are just nice illustrations of the fact that *if* an a priori probability relation exists *then* it will a priori make probable the existence of a reliability relation.

other sensory modalities tend to conflict with deliverances of my visual faculties. Remember, some evidence  $E_1$  can make probable  $P$  even though  $E_1 + E_2$  makes probable  $\neg P$ .

#### 8.4.3 Induction, Perception, and Easy Justification

It will help clarify how all of this works by considering particular examples of supposed basic non-deductive inference and supposed defeasible sources of foundational justification.

Consider first the possibility of basic non-deductive inferences. Enumerative induction seems like a prime candidate for such an inference. Consider an application of enumerative induction from premises  $P_1 - P_N$  concerning past observations correlating instances of  $\Phi$  and instances of  $\Psi$  to the conclusion that the next  $\Phi$  will be associated with a  $\Psi$ . This will only confer justification on  $C$  if one has justification for believing the premises. However, you can insist that you need justification for believing that the premises of an inductive argument are true while also admitting that you needn't have justification for believing that an application of enumerative induction to true premises is reliable. We should not attempt to avoid the easy justification problem by denying that enumerative induction is a basic inference in light of bootstrapping since the same problem will apply to any supposed basic non-deductive inference. And as I've stressed, if there are no non-basic non-deductive inferences then it seems that a far reaching (though not global) skepticism looms; any justification for the conditional reliability of non-deductive inference must rely on non-deductive reasoning itself. And if there are no *basic* non-deductive inferences then we couldn't ever get the *prior* justification for the conditional reliability of a non-deductive inference required for such inferences to confer justification.

Now let's consider again the problem of bootstrapping to the conclusion that induction is reliable so that we can show how the acquaintance theory avoids the problem of easy justification in the case of basic non-deductive inferences.



## EASY BOOTSTRAPPING ON INDUCTION

1.  $P_1$  (Justified by induction).
  2. I formed my belief that  $P_1$  via an application of induction to true premises.
  3. Forming my belief that  $P_1$  belief via an application of induction to true premises didn't undetectably lead me astray on this occasion (From 1 & 2).
  4.  $P_2$  (Justified by induction).
  5. I formed my belief that  $P_2$  via an application of induction to true premises.
  6. Forming my belief that  $P_2$  belief via an application of induction to true premises didn't undetectably lead me astray on this occasion (From 4 & 5).
- [Repeat for  $P_3$ - $P_N$ ]
- 

7. Therefore, applications of induction to true premises are reliable.

Now, consider what is required for an application of induction to provide me with justification for believing  $P_1$ - $P_N$ . Not only must I be justified in believing propositions about past observations I must also be acquainted with the fact F that these make probable the relevant prediction about the future. This probability relation is a priori and holds even in the irregular worlds mentioned in chapter 6. However, in the same way that an entailment relation entails perfect reliability, I've argued that F itself would be evidence for/epistemically support the claim that these the truth of these beliefs about past observations are *reliably* correlated with the truth of the predictions. This evidential support relation is again an internal and necessary relation. Therefore, in forming a justified belief via an application of induction a subject is *necessarily* acquainted with/aware of evidence for the conditional reliability of induction.

On my version of the acquaintance theory, *if* induction is a basic inferential justification source *then* the amount of justification conferred upon the conclusion of this argument is equal to or less than the degree of evidential support provided by a priori evidence of which the subject was already aware in forming the justified beliefs  $P_1$ - $P_N$ . This is

why we get the feeling that such reasoning isn't giving you anything new. In some sense, the subject was already in a position to get that justification a priori. Equally important yet again is the fact that the subject's necessarily being aware of evidence supporting the reliability of induction prior to the reasoning isn't equivalent to requiring that the subject already have *justification* for believing the conclusion. And we thereby avoid the Lewis Carroll regress.

The only difference between our solution in the case of Modus Ponens and induction is the following: in the case of MP conferring justification a subject is necessarily aware of a priori evidence that *guarantees* MP's reliability but in the case of induction conferring justification a subject is necessarily aware of a priori evidence that *makes probable* induction's reliability. One could still get further evidence that applications of induction are in fact unreliable. If you continue to utilize enumerative induction but the predictions continue to be in conflict with the deliverances of independent sources then this offers new defeating evidence for believing in the reliability of induction. Nevertheless, this doesn't threaten the claim that *necessarily* the fact <that past observations correlating  $\Phi$  and  $\Psi$  makes probable that the next  $\Phi$  is  $\Psi$ > is evidence that makes probable that past observations correlating  $\Phi$  and  $\Psi$  reliably indicate that the next  $\Phi$  is  $\Psi$ . A piece of evidence can make probable  $P$  even though one's total evidence makes probable  $\neg P$ .

Similar remarks apply to cases of supposed cases of defeasible evidence that provide *foundational* justification. Assume that a perception as of  $P$  makes probable that  $P$ . According to my acquaintance theory a perception as of  $P$  would only provide  $S$  with justification for believing  $P$  if  $S$  is acquainted with the perception and the probability relation. The existence of such a probability relation itself makes probable that perceptual evidence is reliable. Thus, if  $S$  were to go through the bootstrapping reasoning on perception outlined earlier, the reasoning could provide no more support for the conclusion than that provided by the a priori evidence of which the subject was already aware when forming beliefs via perception.

Similar remarks will apply to any supposed basic non-deductive inference (inference to the best explanation, analogical inferences, etc.) and supposed defeasible sources of

foundational justification (e.g. memory). This concludes the generalization of the solution from the case of MP so as to cover basic deductive inference, basic non-deductive inference, entailing foundational evidence, and non-entailing foundational evidence. However, I want to stress that this is a solution to a *structural* problem. I've argued that both foundational and inferential justification requires acquaintance with irreducible probabilistic relations. I've also argued that the existence of these irreducible probabilistic relations themselves make probable contingent claims about reliability or frequencies. If E makes probable *P* then this fact makes probable that E is a reliable indicator of *P*. As such, on my version of the acquaintance theory, any justification-conferring evidence E will always be accompanied by an awareness of evidence E\* that makes probable E's reliability. But such awareness falls short of justified belief and therefore falls short of accepting JR. This is what allows us to avoid easy justification without falling into skepticism. This was my main goal in this dissertation.

However, all of my cases have *assumed* that the antecedent holds in order to illustrate the solution to the structural worry. I have not argued that any of these are cases where the antecedent does *in fact* hold. And this raises closely-related issues that I want to briefly consider in the last section of this chapter.

### 8.5 Meta-Level Issues

I've stressed that my solution is meant to be a solution to the structural problem. A foundational or inferential basic justification  $\Phi$  source will only provide a subject with justification for *P* if she is aware of  $\Phi$ 's making probable *P* where this is a *sui generis* relation (modeled on entailment) that is irreducible to facts about frequencies even across possible worlds. Such a relation constitutes evidence that  $\Phi$  and *P*'s truth are reliably correlated. A difficulty arises when we consider the fact that different philosophers disagree with one another about what makes probable what. Some philosophers hold that a perceptual experience as of *P* makes probable *P*. Others hold that it's only a perceptual experience as of

$P$  taken in conjunction with a set of perceptual and memorial experiences that cohere with one another that makes probable  $P$ . Some hold that a memorial experience as of some past event by itself makes probable that the past event occurred. Others hold that it's only a memorial experience as of some past event taken in conjunction with a coherent set of memorial experiences that makes probable that the past event occurred. Some hold that the fact that  $P$  is the best explanation of some fact or event makes probable  $P$ . Others deny this.

At this point someone might worry about how to determine the probability relations with which she is actually acquainted. I think the correct answer is that precisely the same methods are legitimate at the second-level as at the first. I can't imagine how any other answer would be consistent. Now someone might worry that we'll end up begging the question at the second-level.

Imagine that I form the belief that  $P$  on the basis of evidence  $E$  and someone asks me how  $E$  justifies me in believing  $P$ . I give them my acquaintance theory and explain that if I'm acquainted with  $E$  and acquainted with the fact that  $E$  makes probable  $P$  then I'm justified in believing  $P$ . A skeptic then responds, "Okay fine, but what reason do you have for believing the antecedent holds? What reason do you have to believe that you're acquainted with  $E$ ? More importantly, what reason do you have to believe that you're acquainted  $E$ 's making probable  $P$ ?" Don't I need to have such reasons in order to truly be justified in believing  $P$ ? My comments here will be brief.

First, I think this makes a level-confusion. Having reasons for believing that I'm acquainted with  $E$  and acquainted with  $E$ 's making probable  $P$  is necessary for having justification for believing that I'm justified in believing  $P$ . Once one has offered conditions sufficient for 1<sup>st</sup>-level justification one must respond by saying that *if* those conditions are in fact met *then* I have justification for  $P$ . One needn't have justification for believing those conditions are met. Otherwise the conditions weren't genuinely sufficient for justification in the first place. All of this sounds very similar to the things externalists say but I think it's importantly different. I reject the access requirements of access internalism—though my ER

principle might remind some readers of an access requirement. However, I think that on the acquaintance theory I've defended, moving up a level and getting justification for believing that these conditions are met doesn't improve a subject's perspective on the first-level proposition's truth. When one has met these conditions one is aware of what the belief has going for it. *If* a subject is aware of what her belief has going for it (and what it has going for it is good) *then* the subject has justification for her belief and the subject's having an awareness of what the belief has going for it mentioned in the antecedent makes this very different from externalism. One doesn't gain assurance by moving up a level. My discussion of Bergmann's challenge in 7.1.6 is relevant here.

Second, even though one needn't move up a level I take it that one can use the very same methods to justify the high-level propositions. How can I get reason to believe that I'm acquainted with the fact that past observations correlating  $\Phi$  and  $\Psi$  make probable that next  $\Phi$  will be correlated with  $\Psi$ ? The same ways I can get justification for believing any other claim. I could legitimately get the higher-order justification in virtue of my acquaintance with my acquaintance with the probability relation. If I really do have these acquaintances then I really am justified in believing that I'm acquainted with the probability relation. I might even use non-deductive inference to justify my acquaintance with certain probability relations. Perhaps acquaintance with a probability relation between facts about perceptual facts and claims about the external world is the best explanation of why there is so much agreement amongst humans that perception is justification-conferring. Perhaps this is the best explanation of the persistence of Moorean and particularist responses to skepticism. If IBE were to epistemically support its conclusions at the first-level and we're acquainted with such a relation concerning 2<sup>nd</sup>-level questions then there is no reason we should deny the justificatory power in this latter case.

This of course raises the worry that we would beg the question against a skeptic, but we must remember the lesson we learned from Peter Markie in chapter 3. Begging the question is a *dialectical* failure and not an *epistemological* failure. If someone skeptical of the

existence of acquaintance asks me what reason I have for believing that it exists and I reply that I'm acquainted with it I surely begged the question. My interlocutor should not be moved by such a claim. Nevertheless, this doesn't mean that the reason I cited wasn't a good reason that actually provided me with justification. If I'm actually acquainted with my acquaintance with various facts (such as my acquaintance with pain, etc.) then this is the best reason I could have for believing such a relation exists.

Sometimes, however, we might be unsure even from the 1<sup>st</sup>-person perspective whether we're aware of certain probability relations. Here I want to suggest that the easy justification reasoning itself might be useful.

Consider again Vogel's gas gauge example. I look at what is in fact a reliable gas gauge that reads "F" at which point I directly form a belief that the gas tank is full. I then use my vision to form the belief that the tank reads "F." At which point I combine these beliefs and infer that the tank was accurate on this occasion. I repeat and eventually wind up with inductive evidence for the conclusion that the tank is reliable. Surely such an argument cannot provide me with new justification for believing that the gauge is reliable. However, assuming my solution to the problem of easy justification is correct, we should say that *if* the fact that the gauge reads "F" makes probable that the tank is full *then* if I were acquainted with this probability relation the argument would provide me with justification for the conclusion absent any confirmation of its reliability from independent sources such as a dipstick. Our uneasiness in the track-record in this case, however, isn't that it seems silly to us since one could have a priori justification for the conclusion. The gas gauge case is not analogous to a case like that of someone surveying bachelors in an attempt to get inductive reasons for believing all bachelors are unmarried. Thus the correct response to the truth of this conditional is a modus tollens: we reject the antecedent.

We aren't acquainted with any such probability relation (remember that the relevant probability relation isn't reducible to any claim about frequencies or reliability!). In fact, most of us likely think that we're acquainted with the fact that the reading on the gas gauge *doesn't*

epistemically support the claim that the tank is full. And the fact that the bootstrapping doesn't merely strike us as something silly that we could have done without the inductive argument illustrates this. The fact that the gauge reads "F" is only evidence for the claim that the tank is full when conjoined with independent facts supporting the reliability of the tank.

Something similar might be said about various epistemic support relations that various philosophers have proposed. Speaking autobiographically, when I consider bootstrapping on a single perceptual experience as of *P* I get uneasy because I don't think that the bootstrapping argument is silly in the relevant way. I take the fact that I get uncomfortable here as indicative of the fact that I'm not acquainted with a probability relation between the fact that I have a perceptual experience as of *P* and its being true that *P*.

So what evidence do I think we're aware of for our perceptual beliefs? My own view is that it consists of the convergence of different perceptual modalities and the coherence amongst perceptual experiences that all take place in the specious present. In ordinary cases, when I look at a car, I don't undergo only a single experience as of a car. Within the specious present I might begin with a visual experience as of a car, next I might have a proprioceptive experience as of tilting my head simultaneous with the visual experiential changes one would expect it to if there were an object of roughly car size and shape, next I might have proprioceptive experiences as of walking toward the car and placing my hand on it accompanied by the expected changes in my visual and tactile experiences, etc. We might add various auditory experiences to the species present as well. What I think actually stands in a probability relation to the claim that there is an object of a car size and shape in front of me isn't a single experiential state but rather experiences of different modalities within the species present converge or cohere with one another that makes probable that there is a physical object of a certain size and shape in front of me. It's these facts about my experience and their making probable my belief with which I take myself to be acquainted and thereby providing me with (foundational) justification for a belief about the existence of a physical object.

At this point, when I imagine going through the bootstrapping argument where I form the physical object belief, cite *all* of these experiential facts, infer that the set of experiences were accurate on this occasion, and repeat so as to produce an inductive argument for the reliability of converging and coherent perceptual experiences I really do feel as if this is silly in the sense that I was aware of evidence for the conclusion prior to the reasoning. It strikes me that it's a priori unlikely that such convergence amongst different modalities wouldn't reliably be correlated with the existence of a physical object. I know that it is *possible* but this doesn't threaten its being a priori unlikely.

In this way, going through the easy justification reasoning regarding *putative* evidence  $E$  for  $P$  acts as a means for testing whether we think we really are directly acquainted with  $E$ 's making probable  $P$ . If such reasoning strikes us as silly in the same way that an empirical track-record for the conclusion that all bachelors are unmarried or for the claim that Modus Ponens is reliable then this is indicative of our actually being directly acquainted with an a priori probability relation. If, however, such reasoning doesn't strike us as silly in this way but actually problematic like the gas gauge example then this is indicative that we weren't directly acquainted with an a priori probability relation between  $E$  and  $P$ . In my own case I find that when I perform this test with the dogmatist view that, by itself, a perception as of  $P$  makes probable  $P$  the reasoning doesn't strike me as merely silly. It's for this reason that I think we must find more detailed features within our fine-grained experience (as suggested above) if we want to find the features that actually make our perceptual beliefs probable. However, all of this goes beyond the solution (presented in this chapter) to the *structural* worry posed by our puzzle.



## CONCLUSION

This dissertation has focused on the nature of justification. Specifically, our discussion was centered by the nature of the relationship between the justification conferring power of a belief source and our ability to justify the reliability of that source.

We began by noting a kind of dilemma revolving around the following question: can a belief source confer justification absent prior justification for believing the source is reliable? If we answer in the affirmative we seem to allow illegitimate ways of acquiring justification for believing a source is reliable. Alternatively, if we offer a negative answer we are quickly led into a global skepticism.

After presenting the dilemma I developed and defended a modified version of the acquaintance theory of justification within an evidentialist framework. First, a subject *S* only *has* some evidence *E* for *P* when *S* is aware of that evidence *and aware of its evidential relation to P*. Second, *S*'s evidence ultimately consists of facts. Third, the evidential connection relevant to justification (i.e. the epistemic support relation) is a *sui generis* and unanalyzable probability relation. Finally, the required awareness of these facts and of the epistemic support relation must ultimately be grounded in acts of non-conceptual awareness or acquaintance. Putting these claims together via the evidentialist idea that justification is fully determined by a subject's possession of evidence and what this supports produced my theory of justification. Crudely, the thesis is that a subject *S* has justification for believing *P* iff (i) *S* is non-conceptually aware of (i.e. acquainted with) a set of facts *F*, (ii) *F* makes probable *P*, and (iii) *S* is non-conceptually aware of (i.e. acquainted with) *F*'s making probable *P*.

In the final chapter I utilized this acquaintance theory to illustrate a way out of our puzzle. The developed theory doesn't allow evidence *E* to confer justification on *P* while the subject is completely oblivious to the evidential connection between *E* and *P*. A body of evidence *E* will only confer justification for believing *P* for *S* if the subject is aware of the fact that *E* is supporting evidence for *P*. We avoid the skeptical worry, however, because the

required awareness is weaker than justified belief. Nevertheless, it may seem that the view is still susceptible to the problem of easy justification since the evidential support relation with which a subject must be acquainted isn't understood in terms of frequencies or reliability. However, I argued that the fact that E supports *P* is itself supporting evidence that E is a reliable indicator of *P* (in cases where E fails to entail *P* the existence of the relation is merely *defeasible* evidence for the reliability connection). Thus the acquaintance theory entails that evidence E only confers justification on *P* for S if S is aware of evidence for E's reliability. And so the acquaintance theory vindicates the idea that if a subject has justification for believing *P* in virtue of evidence E then prior bootstrapping reasoning a subject was already aware of a priori evidence for E's reliability. A subject cannot gain any justification for the conclusion that wasn't in some sense available without the reasoning. Therefore, we avoid the problem of easy justification.

One issue that I have not addressed, however, is whether my theory is a version of internalism or externalism. Let's consider this in relation to our discussion of various internalism/externalism controversies discussed in chapter 1. The theory I've defended is not a version of access internalism. Meeting my conditions for justification *does not* entail that a subject has justification for believing that *any* J-factors are present. It's a conceptual possibility on my view that a subject has foundational justification for a perceptual belief that *P* about the external world in virtue of her acquaintance with facts about her perceptual experiences and with these making probable *P*. This doesn't entail that she has met the conditions that would provide justification for *believing* that the experiential facts make probable *P*. In fact, this even fails to entail that the subject has justification for *believing* that the experiential facts that act as evidence for her belief obtain. Second, it isn't clear whether the theory I've developed is a version of mentalism. I've required acquaintance with a probability relation. Moreover, the acquaintance theory I've defended is at least *compatible* with the claim that we're acquainted with physical objects in the external world. I think there is good reason to reject the idea that we *are* acquainted with facts about the external world

but nothing about the theory of justification I've defended is incompatible with such a claim. Would acquaintance with probability relations and facts about the external world amount to mental states? If so, then perhaps my theory is a version of mentalism. If not, then my theory is not a version of mentalism. Finally, I've also rejected inferential internalism. One needs to be acquainted with an evidential connection between  $P$  and  $Q$  in order for one's belief that  $P$  to inferentially justify believing  $Q$  but one needn't have *justification* for believing that this evidential connection holds.

However, I do believe that it's correct to closely associate or even characterize the developed theory of justification as a version of internalism. The association is easiest to see when we consider how I've motivated my theory of justification. My motivations have stressed the 1<sup>st</sup>-person perspective throughout the dissertation. First, I used cases like Norman the Clairvoyant to motivate the requirements that a subject be aware of supporting evidence and aware of the support relation itself. Second, I used the new evil demon thought experiment to motivate the claim that epistemic support is a sui generis and internal relation holding between a set of evidence and a proposition. And I stressed that this relation cannot be understood in terms of the causal or nomological relations characteristic of the analyses of epistemic support implicit in paradigmatic externalist theories of justification. When we put these together we see that justification is fully determined by facts of which a subject is non-conceptually aware (i.e. facts with which the subject is directly acquainted). If we can understand the egocentric perspective in terms of acquaintance then it seems that the only way a fact can have an effect on a subject's justification on my theory is if it makes its way into the egocentric perspective.

I also want to briefly reconsider my theory of justification's relation to evidentialism. My theory does satisfy *a version* of evidentialist maxim I laid out in chapter 1:

**EM** The epistemic justification of any doxastic attitude holds in virtue of the subject's possession of evidence and what this evidence supports.

There is an ambiguity in EM pointed out by my distinction between unconnected and connected possession of evidence. On my theory, two subjects S and S\* might have the same evidence where this includes both unconnected and connected evidence yet differ justificationaly. EM is only true when restricted to connected evidence. There is a sense in which a subject S has evidence E for *P* merely in virtue of being aware of E but this evidence is only relevant to justification if S is also aware of E's status as evidence for *P*.

Despite satisfying EM, my theory might entail that the specific evidentialist theses put forward by Conee and Feldman (C&F) are false. First, C&F sometimes advocate a thesis which *identifies* justification with a doxastic attitude's standing in a relation of *fit* with a subject's evidence. If *fit* here is construed as fully determined by probability relations between evidence of which the subject is aware and a proposition *P*, my theory has rejected this thesis. However, if *fit* is construed as partially determined by the probability relations between evidence of which the subject is aware and a proposition *P* and partially determined by the subject's awareness of that relation then my theory vindicates this thesis.

Alternatively, C&F sometimes advocate the view that justification strongly *supervenes* on a subject's evidence. Even when construed in terms of connected evidence, this thesis is surely false on my theory of justification. In chapter 7 I explained that a subject's awareness of an evidential relation can be more and less clear, and that this can affect what attitude is justified for the subject. Consider S and S\*. Both are acquainted with E. S is acquainted with the fact that E makes *P* more probable than not. S\* is acquainted with the fact that E makes probable *P* to a degree between .8 and .9. Both S and S\* have connected possession of evidence E for *P* but different doxastic attitudes are justified for the two subjects. S\* is justified in having a more precise doxastic attitude towards *P*'s truth than S is. Justification isn't determined solely by the evidence a subject possesses but also by the nature of the subject's possession of that evidence (i.e. the clarity of her awareness of the support relation).

This concludes my discussion of my theory and its relation to internalism, externalism, and evidentialism. At this point I want to conclude with some suggestions for areas of future research. There is much more to be said about the modified acquaintance theory and its solution to the problem of easy justification. First, the metaphysics positing probability relations between facts and propositions in chapter 4 is surely controversial but the benefits illustrated in chapter 7 surely show that it is worth further investigation. Second, there is much more to be said in order to adequately defend the theory from the etiological objections mentioned at the end of chapter seven. Third, it would be interesting to continue the investigation from chapter 7 concerning the extent to which the theory I've developed can accommodate the insights of classical foundationalism without taking on many of the associated difficulties. Fourth, the solution to our problem of easy justification has suggested a connection between the existence of an a priori probability relation and the existence of a reliability relation, such a connection might be of enormous use in clarifying the nature of the truth-connection. Fifth, I think further investigation of the two distinct ways of understanding a subject's total evidence and the requirement of total evidence which I distinguished (in my discussion of Neta's objections to a fact ontology of evidence) would be incredibly fruitful. These are just the natural next steps in developing the issues involved in this dissertation that I myself find most interesting. There are many more claims put forward in this dissertation that could be further developed, defended, and applied to various epistemological issues. Unfortunately, one can only discuss a tiny portion of the issues of epistemological interest within the space of a dissertation.

## BIBLIOGRAPHY

- Alston, W. (1986). Epistemic Circularity. *Philosophy and Phenomenological Research*, 47(1), 1-30.
- Alston, W. (1988). An Internalist Externalism. *Synthese*, 74, 265-283.
- Alston, W. (1989). The Deontological Conception of Epistemic Justification. In W. Alston, *Epistemic Justification: Essays in the Theory of Knowledge* (pp. 115-152). Ithaca: Cornell University Press.
- Alston, W. (1993). *The Reliability of Sense Perception*. Ithaca: Cornell University Press.
- Alston, W. (1995). How to Think About Reliability. *Philosophical Topics*, 23(1), 1-29.
- Alston, W. (2005). *Beyond Justification: Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.
- Amico, R. (1988). Roderick Chisholm and the Problem of the Criterion. *Philosophical Papers*, 17(3), 217-229.
- Amico, R. (1995). *The Problem of the Criterion*. Lanham: Rowman and Littlefield.
- Anderson, A., & Belnap, N. (1975). *Entailment: The Logic of Relevance and Necessity* (Vol. 1). Princeton: Princeton University Press.
- Anderson, A., & Belnap, N. (1992). *Entailment: The Logic of Relevance and Necessity* (Vol. 2). Princeton: Princeton University Press.
- Armstrong, D. (1968). *A Materialist Theory of the Mind*. London: Routledge.
- Audi, R. (1988). *Belief, Justification and Knowledge*. Belmont: Wadsworth.
- Audi, R. (1991). Justification, Deductive Closure and Reasons to Believe. *Dialogue*, 30, 77-84.
- Audi, R. (2001). Doxastic Voluntarism and the Ethics of Belief. In M. Steup (Ed.), *Knowledge, Truth, and Duty* (pp. 93-111). Oxford: Oxford University Press.
- Audi, R. (2008). The Ethics of Belief: Doxastic Self-Control and Intellectual Virtue. *Synthese*, 161(3), 403-418.
- Austin, J. (1962). *Sense and Sensibilia*. (G. Warnock, Ed.) Oxford: Oxford University Press.
- Bach, K. (1985). A Rationale for Reliabilism. *The Monist*, 68(2), 246-263.
- Baehr, J. (2011). Evidentialism, Vice, and Virtue. In T. Dougherty (Ed.), *Evidentialism and its Discontents* (pp. 88-102). Oxford: Oxford University Press.
- Baldwin, T. (Ed.). (1993). *G.E. Moore: Selected Writings*. London: Routledge.
- Bayer, B. (2012). Internalism Empowered: How to Bolster a Theory of Justification with a Direct Realist Theory of Awareness. *Acta Analytica*, 27(4), 383-408.

- Becker, K., & Black, T. (Eds.). (2012). *The Sensitivity Principle in Epistemology*. Cambridge: Cambridge University Press.
- Beebe, J. (2004). The Generality Problem, Statistical Relevance and the Tri-Level Hypothesis. *Nous*, 38(1), 177-195.
- Bennett, J. (1990). Why is Belief Involuntary? *Analysis*, 50, 57-107.
- Bergmann, M. (2004). Epistemic Circularity: Malignant and Benign. *Philosophy and Phenomenological Research*, 69(3), 709-727.
- Bergmann, M. (2006). *Justification without Awareness*. Oxford: Oxford University Press.
- Bergmann, M. (2008). Reidian Externalism. In *New Waves in Epistemology*. Palgrave Macmillan.
- Bergmann, M. (Forthcoming). Externalist Justification and the Role of Seemings. *Philosophical Studies*.
- Black, M. (1958). Self-Supporting Inductive Arguments. *Journal of Philosophy*, 55(17), 718-725.
- Bonjour, L. (1980). Externalist Theories of Empirical Knowledge. *Midwest Studies in Philosophy*, 5(1), 53-73.
- Bonjour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Bonjour, L. (1999). In *Defense of Pure Reason: A Rationalist Account of A Priori Justification*. Cambridge: Cambridge University Press.
- Bonjour, L. (2000). Toward a Defense of Empirical Foundationalism. In M. DePaul (Ed.), *Resurrecting Old-Fashioned Foundationalism* (pp. 21-38). Lanham: Rowman and Littlefield.
- Bonjour, L. (2003). A Version of Internalist Foundationalism. In L. Bonjour, & E. Sosa, *Epistemic Justification: Internalism vs Externalism, Foundations vs. Virtues* (pp. 3-96). Malden, MA: Wiley-Blackwell.
- Bonjour, L. (2004). In Search of Direct Realism. *Philosophy and Phenomenological Research*, 69(2), 349-367.
- Bonjour, L. (2010a). *Epistemology: Classic Problems and Contemporary Responses* (2nd ed.). Lanham: Rowman and Littlefield.
- Bonjour, L. (2010b). The Myth of Knowledge. *Philosophical Perspectives*, 24(1), 57-83.
- Brewer, B. (2002). *Perception and Reason*. Oxford: Oxford University Press.
- Brewer, B. (2013). *Perception and its Objects*. Oxford: Oxford University Press.
- Brogaard, B. (2013a). Do We Perceive Natural Kind Properties? *Philosophical Studies*, 162(1), 35-42.
- Brogaard, B. (2013b). It's not what it seems. A Semantic Account of 'seems' and seemings. (H. Cappelen, & S. Shapiro, Eds.) *Inquiry*.

- Burge, T. (1979). Anti-Individualism and the Mental. *Midwest Studies in Philosophy*, 4(1), 73-122.
- Carroll, L. (1895). What the Tortoise Said to Achilles. *Mind*, 4(14), 278-280.
- Chalmers, D. (2003). The Content and Epistemology of Phenomenal Belief. In Q. Smith (Ed.), *Consciousness: New Philosophical Perspectives* (pp. 220-272). Oxford: Oxford University Press.
- Chisholm, R. (1942). The Problem of the Speckled Hen. *Mind*, 51, 368-373.
- Chisholm, R. (1977). The Problem of the Criterion. In R. Chisholm, *Theory of Knowledge* (2nd ed.). Eaglewood Cliffs, NJ: Prentice Hall.
- Christensen, D. (2010). Higher-Order Evidence. *Philosophy and Phenomenological Research*, 81(1), 185-215.
- Cling, A. (1994). Posing the Problem of the Criterion. *Philosophical Studies*, 75(3), 261-292.
- Cling, A. (1997). Epistemic Levels and the Problem of the Criterion. *Philosophical Studies*, 88(2), 109-140.
- Coady, C. (1992). *Testimony: A Philosophical Study*. Oxford: Clarendon Press.
- Coady, C. (1994). Testimony, Observation, and 'Autonomous Knowledge'. In B. Matilal, & A. Chakrabarti (Eds.), *Knowing from Words*. Dordrecht: Kluwer Academic Publishers.
- Cohen, S. (1984). Justification and Truth. *Philosophical Studies*, 46(3), 279-295.
- Cohen, S. (1988). How to Be a Fallibilist. *Philosophical Perspectives*, 2, 91-123.
- Cohen, S. (1998). Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery. *Australasian Journal of Philosophy*, 76(2), 289-306.
- Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives*, 13, 57-89.
- Cohen, S. (2002). Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research*, 65(2), 309-329.
- Cohen, S. (2005a). Knowledge, Speaker and Subject. *Philosophical Quarterly*, 55, 199-212.
- Cohen, S. (2005b). Why Basic Knowledge is Easy Knowledge. *Philosophy and Phenomenological Research*, 70(2), 417-430.
- Cohen, S. (2008). Ascriber Contextualism. In J. Greco, *The Oxford Handbook of Skepticism* (pp. 415-436). Oxford: Oxford University Press.
- Cohen, S. (2010). Bootstrapping, Defeasible Reasoning, and a Priori Justification. *Philosophical Perspectives*, 24(1), 141-159.
- Comesaña, J. (2002). The Diagonal and the Demon. *Philosophical Studies*, 110(3), 249-266.



- Comesaña, J. (2005). We Are (Almost) All Externalists Now. *Philosophical Perspectives*, 19(1), 59-76.
- Comesaña, J. (2010). Evidentialist Reliabilism. *Nous*, 44(4), 571-600.
- Conee, E. (1992). The Truth Connection. *Philosophy and Phenomenological Research*, 52(3), 657-669.
- Conee, E. (2010). Rational Disagreement Defended. In R. Feldman, & T. Warfield (Eds.), *Disagreement* (pp. 69-90). Oxford: Oxford University Press.
- Conee, E. (2013). The Specificity of the Generality Problem. *Philosophical Studies*, 163(3), 751-762.
- Conee, E., & Feldman, R. (1998). The Generality Problem for Reliabilism. *Philosophical Studies*, 89(1), 1-29.
- Conee, E., & Feldman, R. (2001). Internalism Defended. In H. Kornblith (Ed.), *Epistemology: Internalism and Externalism* (pp. 231-260). Oxford: Wiley-Blackwell.
- Conee, E., & Feldman, R. (2004). *Evidentialism*. Oxford: Oxford University Press.
- Conee, E., & Feldman, R. (2008). Evidence. In Q. Smith (Ed.), *Epistemology: New Essays* (pp. 83-104). Oxford: Oxford University Press.
- Coppenger, B. (2012). Achieving Epistemic Descent. The University of Iowa. Proquest, UMI Dissertations Publishing. 3526804.
- Crisp, T., Davidson, M., & Laan, D. (Eds.). (2006). *Knowledge and Reality: Essays in Honor of Alvin Plantinga*. Springer.
- Cullison, A. (2010). What Are Seemings? *Ratio*, 23(3), 260-274.
- DePaul, M. (Ed.). (2000). *Resurrecting Old-Fashioned Foundationalism*. Lanham: Rowman and Littlefield.
- DePoe, J. (2011). Defeating the Self-Defeat Argument for Phenomenal Conservatism. *Philosophical Studies*, 152(3), 347-359.
- DeRose, K. (1992). Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, 53(4), 913-929.
- DeRose, K. (1995). Solving the Skeptical Problem. *Philosophical Review*, 104(1), 1-52.
- DeRose, K. (1996). Knowledge, Assertion, and Lotteries. *Australasian Journal of Philosophy*, 74(4), 568-580.
- DeRose, K. (2000). Ought We to Follow Our Evidence? *Philosophy and Phenomenological Research*, 60(3), 697-706.
- DeRose, K. (2009). *The Case for Contextualism: Knowledge, Skepticism and Context*. Oxford: Oxford University Press.

- DeRose, K. (2011). Questioning Evidentialism. In T. Dougherty (Ed.), *Evidentialism and its Discontents* (pp. 137-146). Oxford: Oxford University Press.
- Descartes, R. ([1628] 1985). Rules for the Direction of the Mind. In *The Philosophical Writings of Descartes* (J. Cottingham, R. Stoothoff, & D. Murdoch, Trans., Vol. 1, pp. 7-78). Cambridge: Cambridge University Press.
- Descartes, R. ([1641] 1984). Meditations on First Philosophy. In *The Philosophical Writings of Descartes* (J. Cottingham, R. Stoothoff, & D. Murdoch, Trans., Vol. 2, pp. 1-62). Cambridge: Cambridge University Press.
- DiFate, V. (2007, December 30). *Evidence*. Retrieved from Internet Encyclopedia of Philosophy: <http://www.iep.utm.edu/evidence/>
- Dodd, D. (2011). Against Fallibilism. *Australasian Journal of Philosophy*, 89(4), 665-685.
- Dougherty, T. (Ed.). (2011). *Evidentialism and its Discontents*. Oxford: Oxford University Press.
- Dougherty, T. (2011). In Defense of Propositionalism about Evidence. In T. Dougherty (Ed.), *Evidentialism and its Discontents* (pp. 226-232). Oxford: Oxford University Press.
- Dretske, F. (1970). Epistemic Operators. *The Journal of Philosophy*, 67, 1007-1023.
- Dretske, F. (1971). Conclusive Reasons. *Australasian Journal of Philosophy*, 49, 1-22.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, MA: The MIT Press.
- Dretske, F. (2005). The Case against Closure. In M. Steup, & E. Sosa (Eds.), *Contemporary Debates in Epistemology* (pp. 13-26). Malden, MA: Wiley-Blackwell.
- Edward. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Oxford University Press.
- Empiricus, S. (1935). *Against the Logicians*. (R. Bury, Trans.) Cambridge, MA: Harvard University Press.
- Empiricus, S. (1996). Outlines of Pyrrhonism. In B. Mates, *The Sceptics Way: Sextus Empiricus's Outlines of Pyrrhonism* (pp. 87-218). Oxford: Oxford University Press.
- Fantl, J., & McGrath, M. (2002). Evidence, Pragmatics, and Justification. *Philosophical Review*, 111(1), 67-94.
- Fantl, J., & McGrath, M. (2007). On Pragmatic Encroachment in Epistemology. *Philosophy and Phenomenological Research*, 75(3), 558-589.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Feldman, R. (1985). Reliability and Justification. *The Monist*, 68(2), 159-174.
- Feldman, R. (1995). In Defence of Closure. *The Philosophical Quarterly*, 45, 487-494.
- Feldman, R. (2001). Voluntary Belief and Epistemic Evaluation. In M. Steup (Ed.), *Knowledge, Truth, and Dute* (pp. 77-92). Oxford: Oxford University Press.

- Feldman, R. (2004). The Justification of Introspective Beliefs. In E. Conee, & R. Feldman, *Evidentialism: Essays in Epistemology* (pp. 199-218). Oxford: Oxford University Press.
- Feldman, R. (2005). Respecting the Evidence. *Philosophical Perspectives*, 19(1), 95-119.
- Feldman, R. (2009). Evidentialism, High-Order Evidence, and Disagreement. *Episteme*, 6(3), 294-312.
- Fish, W. (2009). *Perception, Hallucination, and Illusion*. Oxford: Oxford University Press.
- Fitelson, B., & Feldman, R. (2012). Evidence of Evidence is Not (Necessarily) Evidence. *Analysis*, 72(1), 85-88.
- Foley, R. (1979). Justified Inconsistent Beliefs. *American Philosophical Quarterly*, 16(4), 247-257.
- Foley, R. (1983). Epistemic Conservatism. *Philosophical Studies*, 43(2), 165-182.
- Foley, R. (1985). What's Wrong With Reliabilism? *The Monist*, 68(2), 188-202.
- Foley, R. (1987). *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.
- Frege, G. ([1892] 1997). On Sense and Reference. In M. Beaney (Ed.), *The Frege Reader* (pp. 151-171). Malden, MA: Blackwell.
- Fumerton, R. (1995). *Metaepistemology and Skepticism*. Lanham: Rowman and Littlefield.
- Fumerton, R. (2002). *Realism and the Correspondence Theory of Truth*. Lanham: Rowman and Littlefield.
- Fumerton, R. (2004). Epistemic Probability. *Philosophical Issues*, 14(1), 149-164.
- Fumerton, R. (2005). Speckled Hens and Objects of Acquaintance. *Philosophical Perspectives*, 121-138(19), 121-138.
- Fumerton, R. (2006a). Epistemic Internalism, Philosophical Assurance, and the Skeptical Predicament. In T. Crisp, M. Davidson, & D. Laan (Eds.), *Knowledge and Reality: Essays in Honor of Alvin Plantinga* (pp. 179-192). Springer.
- Fumerton, R. (2006b). *Epistemology*. Malden, MA: Blackwell.
- Fumerton, R. (2007). Epistemic Conservatism: Theft or Honest Toil? *Oxford Studies in Epistemology*, 2, 64-87.
- Fumerton, R. (2008). The Problem of the Criterion. In J. Greco, *The Oxford Handbook of Skepticism* (pp. 34-52). Oxford: Oxford University Press.
- Fumerton, R. (2011). Evidentialism and Truth. In T. Dougherty (Ed.), *Evidentialism and its Discontents* (pp. 179-192). Oxford: Oxford University Press.
- Ginet, C. (2001). Deciding to Believe. In M. Steup (Ed.), *Knowledge, Truth, and Duty* (pp. 63-76). Oxford: Oxford University Press.
- Goldberg, S. (2006). Reductionism and the Distinctiveness of Testimonial Knowledge. In J. Lackey, & E. Sosa (Eds.), *The Epistemology of Testimony*. Oxford: Oxford University Press.

- Goldberg, S. (2008). Testimonial Knowledge in Early Childhood, Revisited. *Philosophy and Phenomenological Research*, 76, 1-36.
- Goldman, A. (1979). What is Justified Belief? In *Justification and Knowledge* (George Pappas ed., pp. 1-23). Dordrecht, Holland: D. Reidel.
- Goldman, A. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Goldman, A. (1988). Strong and Weak Justification. *Philosophical Perspectives*, 2, 51-69.
- Goldman, A. (1993). Epistemic Folways and Scientific Epistemology. In A. Goldman, *Liasons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge: Cambridge University Press.
- Goldman, A. (1999a). Internalism Exposed. *Journal of Philosophy*, 96(6), 271-293.
- Goldman, A. (1999b). *Knowledge in a Social World*. Oxford: Clarendon Press.
- Goldman, A. (2008, April 21). *Reliabilism*. (E. Zalta, Ed.) Retrieved from The Stanford Encyclopedia of Philosophy: <http://plato.stanford.edu/entries/reliabilism/>
- Goldman, A. (2009). Internalism, Externalism, and the Architecture of Justification. *Journal of Philosophy*, 106(6), 309-338.
- Goldman, A. (2011). Toward a Synthesis of Reliabilism and Evidentialism? Or: Evidentialism's Troubles, Reliabilism's Rescue Package. In T. Dougherty (Ed.), *Evidentialism and its Discontents* (pp. 254-280). Oxford: Oxford University Press.
- Greco, J. (1993). Virtues and Vices of Virtue Epistemology. *Canadian Journal of Philosophy*, 23(3), 413-432.
- Greco, J. (1999). Agent Reliabilism. *Philosophical Perspectives*, 13, 273-296.
- Greco, J. (2000). *Putting Skeptics in their Place*. Cambridge: Cambridge University Press.
- Greco, J. (2006). Virtue, Luck and the Pyrrhonian Problematic. *Philosophical Studies*, 130(1), 9-34.
- Greco, J. (Ed.). (2008). *The Oxford Handbook of Skepticism*. Oxford: Oxford University Press.
- Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
- Greco, J. (2011). Evidentialism about Knowledge. In T. Dougherty (Ed.), *Evidentialism and its Discontents* (pp. 167-178). Oxford: Oxford University Press.
- Greenough, P., & Pritchard, D. (Eds.). (2009). *Williamson on Knowledge*. Oxford: Oxford University Press.
- Gupta, A. (2008). *Empiricism and Experience*. Oxford: Oxford University Press.
- Hanna, N. (2011). Against Phenomenal Conservatism. *Acta Analytica*, 26(3), 213-221.

- Hasan, A. (2013a). Internalist Foundationalism and the Sellarsian Dilemma. *Res Philosophica*, 90(2), 171-184.
- Hasan, A. (2013b). Phenomenal Conservatism, Classical Foundationalism, and Internalist Justification. *Philosophical Studies*, 162(2), 119-141.
- Hawley, K., & Machpherson, F. (Eds.). (2011). *The Admissible Contents of Experience*. Malden, MA: Wiley-Blackwell.
- Hawthorne, J. (2002). Deeply Contingent a Priori Knowledge. *Philosophy and Phenomenological Research*, 65(2), 247-269.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, J. (2005). The Case for Closure. In M. Steup, & E. Sosa (Eds.), *Contemporary Debates in Epistemology* (pp. 26-43). Malden, MA: Wiley-Blackwell.
- Heller, M. (1995). The Simple Solution to the Problem of Generality. *Nous*, 29(4), 501-515.
- Henderson, D., & Horgan, T. (2007). Some Ins and Outs of Transglobal Reliabilism. In S. Goldberg (Ed.), *Internalism and Externalism in Semantics and Epistemology* (pp. 100-130). Oxford: Oxford University Press.
- Henderson, D., & Horgan, T. (2011). *The Epistemological Spectrum*. Oxford: Oxford University Press.
- Hobson, K. (Forthcoming). In Defense of Relational Direct Realism. *European Journal of Philosophy*. Retrieved from <http://onlinelibrary.wiley.com/doi/10.1111/j.1468-0378.2011.00462.x/abstract>
- Huemer, M. (2001). *Skepticism and the Veil of Perception*. Lanham: Rowman and Littlefield.
- Huemer, M. (2002). Fumerton's Principle of Inferential Justification. *Journal of Philosophical Research*, 27, 329-340.
- Huemer, M. (2006). Phenomenal Conservatism and the Internalist Intuition. *American Philosophical Quarterly*, 43(2), 147-158.
- Huemer, M. (2007). Compassionate Phenomenal Conservatism. *Philosophy and Phenomenological Research*, 74(1), 30-55.
- Hume, D. ([1748] 1975). Enquiry Concerning Human Understanding. In D. Hume, L. Selby-Bigge, & P. Nidditch (Eds.), *Enquiries Concerning Human Understanding and Concerning the Principles of Morals* (3rd ed.). Oxford: Clarendon Press.
- Johnston, M. (2004). The Obscure Object of Hallucination. *Philosophical Studies*, 120, 113-183.
- Joyce, J. (2010). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives*, 24, 281-323.
- Keynes, J. M. (1921). *A Treatise on Probability*. London: Macmillan.
- Klein, P. (1999). Human Knowledge and the Infinite Regress of Reasons. *Philosophical Perspectives*, 13, 297-325.

- Klein, P. (2000). Why Not Infinitism? *The Proceedings of the Twentieth World Congress of Philosophy*, 5, 199-208.
- Klein, P. (2005). Infinitism is the Solution of the Epistemic Regress Problem. In *Contemporary Debates in Epistemology* (pp. 131-139). Malden, MA: Wiley-Blackwell.
- Klein, P. (2007). Human Knowledge and the Infinite Progress of Reasoning. *Philosophical Studies*, 134(1), 1-17.
- Lehrer, K. (1974). *Knowledge*. Oxford: Clarendon Press.
- Lehrer, K. (1989). *Thomas Reid*. New York: Routledge.
- Lehrer, K. (1990). *Theory of Knowledge*. Boulder, CO: Westview.
- Leplin, J. (2007). In Defense of Reliabilism. *Philosophical Studies*, 134(1), 31-42.
- Levin, M. (1997). You Can Always Count on Reliabilism. *Philosophy and Phenomenological Research*, 57(3), 607-617.
- Lewis, D. (1970). Anselm and Actuality. *Nous*, 4(2), 175-188.
- Lewis, D. (1986). *On the Plurality of Worlds*. Malden, MA: Wiley-Blackwell.
- Lewis, D. (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, 74(4), 549-567.
- Lipton, P. (1997). The Epistemology of Testimony. *British Journal for the History and Philosophy of Science*, 29, 1-31.
- Littlejohn, C. (2009). The Externalist's Demon. *Canadian Journal of Philosophy*, 39(3), 399-434.
- Littlejohn, C. (2011). Defeating Phenomenal Conservatism. *Analytic Philosophy*, 52(1), 35-48.
- Luper, S. (Ed.). (2003). *The Sceptics: Contemporary Essays*. Burlington, VT: Ashgate.
- Lycan, W. (1987). *Consciousness*. Cambridge, MA: MIT Press.
- Lycan, W. (2001). A Simple Argument for a Higher-Order Representation Theory of Consciousness. *Analysis*, 61, 3-4.
- Lyons, J. (1997). Testimony, Induction, and Folk Psychology. *Australasian Journal of Philosophy*, 75, 163-178.
- Lyons, J. (2005). Clades, Capgras, and Perceptual Kinds. *Philosophical Topics*, 33(1), 185-206.
- Lyons, J. (2009). *Perception and Basic Beliefs*. Oxford: Oxford University Press.
- Majors, B., & Sawyer, S. (2005). The Epistemological Argument for Content Externalism. *Nous*, 39(1), 257-280.
- Mares, E. (2004). *Relevant Logic: A Philosophical Interpretation*. Cambridge: Cambridge University Press.
- Markie, P. (2005). Easy Knowledge. *Philosophy and Phenomenological Research*, 70(2), 406-416.

- Markie, P. (2009). Classical Foundationalism and Speckled Hens. *Philosophy and Phenomenological Research*, 79, 190-206.
- McCain, K. (2008). The Virtues of Epistemic Conservatism. *Synthese*, 164(2), 185-200.
- McCain, K. (2012). Against Hanna on Phenomenal Conservatism. *Acta Analytica*, 27(1), 45-54.
- McDowell, J. (1982). Criteria, Defeasibility, and Knowledge. *Proceedings of the British Academy*, 68, 455-479.
- McDowell, J. (1994). *Mind and World*. Cambridge: Harvard University Press.
- McGrew, T. (1995). *The Foundations of Knowledge*. Lanham: Rowman and Littlefield.
- McGrew, T., & McGrew, L. (2007). *Internalism and Epistemology: The Architecture of Reason*. New York: Routledge.
- Montmarquet, J. (1986). The Voluntariness of Belief. *Analysis*, 46, 49-53.
- Montmarquet, J. (2008). Virtue and Voluntarism. *Synthese*, 161(3), 393-402.
- Moore, G. (1993). Certainty. In T. Baldwin (Ed.), *G.E. Moore: Selected Writings* (pp. 171-196). London: Routledge.
- Moore, G. (1993a). Hume's Theory Examined. In T. Baldwin (Ed.), *G.E. Moore: Selected Writings* (pp. 59-78). London: Routledge.
- Moore, G. (1993b). Proof of an External World. In T. Baldwin (Ed.), *G.E. Moore: Selected Writings* (pp. 147-170). London: Routledge.
- Moore, G. (1993c). The Refutation of Idealism. In T. Baldwin (Ed.), *G.E. Moore: Selected Writings* (pp. 23-44). London: Routledge.
- Morre, G. (1993d). A Defence of Common Sense. In T. Baldwin (Ed.), *G.E. Moore: Selected Writings* (pp. 106-133). London: Routledge.
- Moser, P. (1989). *Knowledge and Evidence*. Cambridge: Cambridge University Press.
- Nagel, T. (1974). What is it Like to be a Bat? *Philosophical Review*, 10(28), 435-450.
- Neta, R. (2008). What Evidence Do You Have? *British Journal for the Philosophy of Science*, 59(1), 89-119.
- Neta, R. (2013). Easy Knowledge, Transmission Failure, and Empiricism. *Oxford Studies in Epistemology*, 4, 166-184.
- Nottelmann, N. (2006). The Analogy Argument for Doxastic Voluntarism. *Philosophical Studies*, 131, 559-582.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge: Harvard University Press.
- Papineau, D. (1992). Reliabilism, Induction and Skepticism. *Philosophical Quarterly*, 42(166), 1-20.

- Plantinga, A. (1993). *Warrant and Proper Function*. Oxford: Oxford University Press.
- Pollock, J. (1974). *Knowledge and Justification*. Princeton: Princeton University Press.
- Poston, T. (2007). Acquaintance and the Problem of the Speckled Hen. *Philosophical Studies*, 132(2), 331-346.
- Priest, G. (2008). *An Introduction to Non-Classical Logics: From If to Is*. Cambridge: Cambridge University Press.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- Pritchard, D. (2007). Anti-Luck Epistemology. *Synthese*, 158, 277-298.
- Pritchard, D. (2009). Safety-Based Epistemology: Wither Now? *Journal of Philosophical Research*, 34, 33-45.
- Pritchard, D. (2012). *Epistemological Disjunctivism*. Oxford: Oxford University Press.
- Pryor, J. (2000). The Skeptic and the Dogmatist. *Nous*, 34(4), 517-549.
- Putnam, H. (1975). The Meaning of 'Meaning'. *Minnesota Studies in the Philosophy of Science*, 7, 131-193.
- Putnam, H. (1981). *Reason, Truth, and History*. Cambridge: Cambridge University Press.
- Pylyshyn, Z., & Storm, R. (1988). Tracking Multiple Independent Targets: Evidence for a Parallel Tracking Mechanism. *Spatial Vision*, 3(3), 179-187.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Reid, T. ([1764] 1983). An Inquiry into the Human Mind: On the Principles of Common Sense. In T. Reid, R. Beanblossom, & K. Lehrer (Eds.), *Inquiry and Essays*. Indianapolis: Hackett.
- Reid, T. ([1785] 2002). *Essays on the Intellectual Powers of Man*. (D. Brookes, Ed.) Edinburgh: Edinburgh University Press.
- Riggs, W. (2002). Reliability and the Value of Knowledge. *Philosophy and Phenomenological Research*, 64(1), 79-96.
- Rizzieri, A. (2009). Evidence Does Not Equal Knowledge. *Philosophical Studies*, 153(2), 235-242.
- Rosenthal, D. (2012). Higher-Order Awareness, Misrepresentation and Function. *Phil. Trans. R. Soc. B*, 1-15.
- Roush, S. (2005). *Tracking Truth*. Oxford: Oxford University Press.
- Russell, B. ([1912] 1997). *The Problems of Philosophy*. Oxford: Oxford University Press.
- Russell, B. (1948). *Human Knowledge: Its Scope and Limits*. New York: Simon and Schuster.



- Rysiew, P. (2000). Testimony, Simulation, and the Limits of Inductivism. *Australasian Journal of Philosophy*, 78, 269-274.
- Sanford, D. (1990). The Inductive Support of Inductive Rules: Themes from Max Black. *Dialectica*, 44, 23-41.
- Schacter, D. (1999). The Seven Sins of Memory: Insights From Psychology and Cognitive Neuroscience. *American Psychologist*, 54(3), 182-203.
- Schechter, J. (2013). Rational Self-Doubt and the Failure of Closure. *Philosophical Studies*, 163(2), 428-452.
- Schmitt, F. (1992). *Knowledge and Belief*. London: Routledge.
- Schroeder, M. (2008). Having Reasons. *Philosophical Studies*, 139(1), 57-71.
- Schwitzgebel, E. (2000). How Well Do We Know Our Own Conscious Experience? The Case of Human Echolocation. *Philosophical Topics*, 28, 235-246.
- Schwitzgebel, E. (2008). The Unreliability of Naive Introspection. *Philosophical Review*, 117(2), 245-273.
- Schwitzgebel, E. (2011). *Perplexities of Consciousness*. Cambridge, MA: MIT Press.
- Sellars, W. (1956). Empiricism and the Philosophy of Mind. *Minnesota Studies in the Philosophy of Science*, 253-329.
- Siegel, S. (2011). *The Contents of Visual Experience*. Oxford: Oxford University Press.
- Skene, M. (2013). Seemings and the Possibility of Epistemic Justification. *Philosophical Studies*, 163(2), 539-559.
- Smith, A. (2002). *The Problem of Perception*. Cambridge: Harvard University Press.
- Smith, Q. (Ed.). (2008). *Epistemology: New Essays*. Oxford: Oxford University Press.
- Sosa, E. (1988). Beyond Skepticism to the Best of Our Knowledge. *Mind*, 97, 153-188.
- Sosa, E. (1991). *Knowledge in Perspective*. Cambridge: Cambridge University Press.
- Sosa, E. (1993). Proper Functionalism and Virtue Epistemology. *Nous*, 27(1), 51-65.
- Sosa, E. (1999a). How must knowledge be modally related to what is known? *Philosophical Topics*, 26, 373-384.
- Sosa, E. (1999b). How to Defeat Opposition to Moore. *Philosophical Perspectives*, 13, 141-154.
- Sosa, E. (2001). Goldman's Reliabilism and Virtue Epistemology. *Philosophical Topics*, 29, 383-400.
- Sosa, E. (2007). *A Virtue Epistemology*. Oxford: Oxford University Press.
- Sosa, E. (2009). *Reflective Knowledge*. Oxford: Oxford University Press.

- Stalnaker, R. (1978). Assertion. In P. Cole (Ed.), *Syntax and Semantics 9* (pp. 315-332). New York: New York Academic Press.
- Steup, M. (2000). Doxastic Voluntarism and Epistemic Deontology. *Acta Analytica*, 15, 25-56.
- Steup, M. (2004). Internalist Reliabilism. *Philosophical Issues*, 14(1), 403-425.
- Steup, M. (2008). Doxastic Freedom. *Synthese*, 161, 375-392.
- Steup, M. (2011). Belief, Voluntariness and Intentionality. *Dialectica*, 65(4), 537-559.
- Steup, M., & Sosa, E. (Eds.). (2005). *Contemporary Debates in Epistemology*. Malden, MA: Wiley-Blackwell.
- Swain, M. (1981). *Reasons and Knowledge*. Ithaca: Cornell University Press.
- Titlebaum, M. (2010). Tell Me You Love Me: Bootstrapping, Externalism, and No-Lose Epistemology. *Philosophical Studies*, 149(1), 119-134.
- Tolhurst, W. (1998). Seemings. *American Philosophical Quarterly*, 35(3), 293-302.
- Tucker, C. (2010). Why Open-Minded People Should Endorse Dogmatism. *Philosophical Perspectives*, 24(1), 529-545.
- Tucker, C. (2012). Movin' on Up: Higher-Level Requirements and Inferential Justification. *Philosophical Studies*, 157(3), 323-340.
- Tucker, C. (Ed.). (Forthcoming). *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. Oxford: Oxford University Press.
- Tye, M. (1998). Externalism and Memory. *Proceedings of the Aristotelian Society*, 72, 77-84.
- Tye, M. (2010). Attention, Seeing, and Change Blindness. *Philosophical Issues*, 20, 410-437.
- Unger, P. (1975a). A Defence of Skepticism. *Philosophical Review*, 80, 198-219.
- Unger, P. (1975b). *Ignorance: A Case for Scepticism*. Oxford: Clarendon Press.
- Van Cleve, J. (1984). Reliability, Justification, and the Problem of Induction. *Midwest Studies in Philosophy*, 9(1), 555-567.
- Van Cleve, J. (2006). Reid on the Credit of Human Testimony. In J. Lackey, & E. Sosa (Eds.), *The Epistemology of Testimony*. Oxford: Oxford University Press.
- Van Cleve, J. (2003). Is Knowledge Easy -- Or Impossible? Externalism as the Only Alternative to Skepticism. In S. Luper (Ed.), *The Sceptics: Contemporary Essays*. Burlington, VT: Ashgate.
- Vogel, J. (1990). Cartesian Skepticism and Inference to the Best Explanation. *Journal of Philosophy*, 87(11), 658-666.
- Vogel, J. (2000). Reliabilism Leveled. *Journal of Philosophy*, 97(11), 602-623.

- Vogel, J. (2008). Epistemic Bootstrapping. *Journal of Philosophy*, 105(9), 518-539.
- Weisberg, J. (2010). Bootstrapping in General. *Philosophy and Phenomenological Research*, 81(3), 525-548.
- Weiskrantz, L. (1990). *Blindsight: A Case Study and Its Implications*. Oxford: Oxford University Press.
- Weiskrantz, L. (1997). *Consciousness Lost and Found*. Oxford: Oxford University Press.
- White, R. (2006). Problems for Dogmatism. *Philosophical Studies*, 131(3), 525-557.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wittgenstein, L. ([1969] 1991). *On Certainty*. (G. Anscombe, & G. von Wright, Eds.) Malden, MA: Wiley-Blackwell.
- Wright, C. (1991). Scepticism and Dreaming: Imploding the Demon. *Mind*, 100, 87-116.
- Wright, C. (2004). Warrant for Nothing (and Foundations for Free)? *Aristotelian Society Supplementary Volume*, 78(1), 167-212.
- Wright, C. (2011). Frictional Coherentism? A Comment on Chapter 10 of Ernest Sosa's Reflective Knowledge. *Philosophical Studies*, 152(1), 29-41.