

THE IMPORTANCE OF BEING RATIONAL

ERROL D. LORD

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
PHILOSOPHY
ADVISER: THOMAS KELLY AND MICHAEL SMITH

SEPTEMBER 2013

© Copyright by Errol D. Lord, 2013.
All rights reserved.

Abstract

This work systematically defends the view that what it is to be rational is to correctly respond to the reasons that one possesses. It is broken into two parts. In Part I—entitled Coherence, Possession, and Correctly Responding—I explicate the details of my positive account. I begin my defense in chapter 2 by contrasting my view with a coherentist account of rationality popular in the literature on practical reason. I show that my view can explain the data that motivates the coherentist view without incurring its implausible implications. I then argue in chapter 3 that what it is to possess a reason r to ϕ is to be in a position to know r and be disposed to treat r as a reason to ϕ . Chapter 4, the final chapter in Part I, defends a view of what it is to correctly respond to the normative reasons you possess. According to this view, you correctly respond to reasons to ϕ when you ϕ for those reasons. I defend a novel causal account of what it is to ϕ for normative reasons. I also extend the core idea of the account to an account of the broader notion of ϕ -ing for reasons.

Part II—entitled Foundationalism, Deception, and the Importance of Being Rational—confronts traditional problems for my type of view. Chapter 5 defends the claim that foundationalists about epistemic rationality should think that foundational beliefs are based on sufficiently strong possessed reasons. In Chapter 6 I argue that my view can explain why the systematically deceived can have rational beliefs and perform rational actions despite the fact that they don't possess the same reasons as those who are not systematically deceived. The final chapter deals with recent skepticism about the claim that rationality is deontically significant. I first make clear just what it would take for a view to meet the skeptical challenge and argue that the standard views in metaethics can't do what it takes. I then argue that my view can meet the challenge. In fact, I argue that what we ought to do just is what we are rationally required to do. Thus, being rational turns out to be very important indeed.

Acknowledgements

They say it takes a village. It certainly has in my case of becoming something resembling a philosopher. I have many people to thank. I'm sure I'll forget some people. I apologize to those people.

I first started thinking about many of the issues in this dissertation during my last semester as an undergraduate at Arizona State. That semester I took two seminars that greatly shaped my philosophical trajectory. The first was a seminar taught by Doug Portmore. It was for this class that I wrote my first paper on the connection between rationality and reasons. In fact, I still remember the exact moment I began thinking about an inchoate version of the main idea of this dissertation—it was on a late winter's night as I walked my dog around Chaparral lake. Doug gave me a lot of excellent feedback on that paper. He also played a pivotal role in my early life as an aspiring philosopher. I thank him for being both encouraging and realistic. He was a paradigm advisor for an ambitious but largely ignorant undergraduate like me.

The second seminar was taught by Stew Cohen. It was on Williamson's *Knowledge and its Limits*. If you take anything away from this dissertation, it will be that *Knowledge and its Limits* has played a monumental role in my thinking. And Stew's class played a monumental role in the way I think about Williamsonian epistemology. It was in this class that I first appreciated just how consistently great a seminar can be. I thank Stew for making epistemology seem as exciting as it in fact is. Without that I might not have ever become an epistemologist.

After I graduated from ASU, I went to study at the University of Nebraska-Lincoln. My plan was to work on normative ethics. This changed about a month into my first semester. I gave my first talk as a graduate student at the in-house talk series for graduate students. UN-L has a great tradition of inviting all first year faculty to these talks. I was lucky enough to have Dave Sobel and Jan Dowell there since they were first year faculty. I gave the paper I had written for Doug the previous spring. I had actually wanted to talk about contractualism, but didn't have anything I was happy enough with. I'm glad I didn't, because the discussion of my paper that day was exhilarating enough to cause a serious obsession with the connection between reasons and rationality. I have to thank Dave especially for pushing me just hard enough to very badly want to solve the obvious problems. I was so worked up that I was madly writing notes on the back of scrap paper on my bus ride home. Dave was quite surprised when I pulled out these same pieces of paper that night at a department party. I continued to work on these issues for much of my time at UN-L. I was helped by many at UN-L. I owe a very large debt to Mark van Roojen, who was my main advisor. Without Mark, this dissertation would not exist. I would not have ended up at Princeton without Mark's tireless support. I also thank Al Casullo, Jan Dowell, Cullen Gatten, John Gibbons, Reina Hayaki, Cliff Hill, Tim Loughlin, Jennifer McKittrick, Joe Mendola, Dave Sobel, Steve Swartzter, and Adam Thompson.

I then moved to Princeton. Princeton has been an absolutely fabulous place for me to work. The program gave me the freedom to work on the projects I found interesting and to move at my own pace. More importantly, the intellectual climate of the department has been fantastic. After some hemming and hawing over dissertation topics, I decided to pursue the topic that's captivated me from the beginning. I ultimately decided to write this particular dissertation because I thought I had some argumentative debts to pay. I had been relying on many views that I hadn't properly worked through and defended. It was time to do this.

I have many people to thank for playing various helpful roles. I must start by thanking my official advisors, Tom Kelly and Michael Smith, and my unofficial advisor, Gideon Rosen. All three

have been truly fantastic to work with. They all have provided invaluable feedback on countless occasions, both about my work and about life. They have been great colleagues and great friends. I'm honored to know them.

Thanks next to John Broome, Mark Schroeder and Timothy Williamson. Their work serves as a touchstone for the kind of work I hope to do and for the best versions of all sorts of views about all sorts of things. I end up disagreeing with all of them, but the dissertation is undoubtedly better because these three are the three that are most often in the back of my mind. I thank Mark in particular for also being a mentor and friend even though he had no obligation to do so.

Many, many people have helped make the content of the dissertation better. I've presented material that ended up in the dissertation at the 2008 UT Austin Graduate Conference, the 12th Annual Oxford Graduate Conference, the Cologne Summer School in Philosophy, the CRNAP Networkshop at the L'Institut Jean Nicod, the Northern Institute of Philosophy at the University of Aberdeen, Franklin & Marshall College, Vassar College, and many times at both the University of Nebraska and Princeton. I thank members of the audiences at all of these places. I especially thank Neil Sinhababu for formally commenting in Austin, John Broome for formally commenting in Oxford, Jens Ziska for formally commenting in Paris, and Andrew Huddleston for formally commenting in Princeton.

In addition, I thank Ashley Atkins, Sam Baker, Nick Beckstead, Bob Beddor, John Brunero, Tim Campbell, Dave Chalmers, Richard Yetter Chappell, Ryan Cook, Aaron Cotnoir, Anthony Cross, Simon Cullen, Jordan Delange, Billy Dunaway, Steve Finlay, Daniel Fogal, Nate Gadd, Josh Gillon, Dan Greco, Mark Harris, Jonathan Jenkins Ichikawa, Frank Jackson, Karen Jones, Rishi Joshi, Boris Kment, Raffi Krut-Landau, Barry Lam, Alida Liberman, Kathryn Lindeman, Eden Lin, Susanne Mantel, Corey Maley, Jimmy Martin, Brennan McDavid, Aidan McGlynn, Tristram McPherson, Angela Mendelovici, Carla Merino-Rajme, Lisa Miracchi, Shyam Nair, Cory Nichols, David Nowakowski, Rachel Parsons, Kristin Primus, Joe Rachiele, Jake Ross, Vanessa Schouten, Andrew Sepielli, Nate Sharadin, Derek Shiller, Sam Shpall, Holly Smith, Justin Snedegar, Ernie Sosa, Jack Spencer, Noel Swanson, Gerard Vong, Ralph Wedgwood, Daniel Wodak, Jack Woods, and Helen Yetter Chappell.

My life in graduate school would have been much worse if it were not for my friends. Five deserve special thanks for their vigilant support and feedback, both about philosophy and about the world beyond it. These are Barry Maguire, Whitney Schwab, Kurt Sylvan, Jonathan Way and, most especially, Andrew Huddleston. I hope to know all of you for the rest of my life.

Last but not least, I thank my wonderful family. I thank my mom and dad and my mother-in-law and father-in-law. Thanks especially to my mother-in-law, Julia Thomsen, who provided much needed family support in the final weeks of work on the dissertation. I thank my kids, Thomsen, Lavinia, Copley, and Hammond, for being constant sources of wonder and fun. And I thank my wonderful wife Anne, from whom all blessings flow. It is with all my heart that I dedicate this work to you.

EL
Princeton
April 2013

To Anne

I never thought this life was possible. You're the yellow bird that I've been waiting for.

Into the caverns of tomorrow, with just our flashlights and our love, let us plunge, let us plunge, let us plunge. And then we'll get down there, way down to the very bottom of everything. And then we'll see it, we'll see it, we'll see it.

Contents

Abstract	iii
Acknowledgements	iv
1 Introduction	1
I Coherence, Possession, and Correctly Responding	6
2 The Coherent and the Rational	7
2.1 Introduction	7
2.2 The Debate as it Currently Stands	8
2.3 A Better Way	10
2.4 Broome's Challenge	13
2.5 What about Coherence?	18
2.5.1 The Myth of the Coherent	23
2.6 Back to the Beginning	24
3 What it is to Possess a Reason	26
3.1 Introduction	26
3.2 A Working Hypothesis and Two Desiderata	27
3.3 The Epistemic Condition	28
3.3.1 A Taxonomy	28
3.3.2 Against Holding Views	30
3.3.3 Against Low Bar Views	31
3.3.4 Against Non-Factive Views	36
3.3.5 Against $P \rightarrow TPJ$	44
3.4 The Treating Condition	48
3.5 The View	54
4 Deviance in Action and Belief	55
4.1 Introduction	55
4.1.1 Where We've Been and Where We're Going	55
4.1.2 The Lay of the Land	55
4.2 Acting for Reasons, Believing for Reasons, and Being Deviant	57
4.3 Acting and Believing for Normative Reasons	59
4.4 What it is to Act and Believe for Normative Reasons	63

4.4.1	Why This is, Alas, Not Enough to Get All that We Want	67
4.5	What it is to Act and Believe for Reasons	69
4.5.1	Again, This is Not Enough	73
4.6	Acting for Reasons and Acting for Normative Reasons	74
4.6.1	The Transparency of Acting for Reasons	74
4.6.2	The Transparency of Acting for Normative Reasons	75
4.7	Further Upshots	76
4.7.1	The Relationship Between <i>Ex Post</i> and <i>Ex Ante</i> Justification	76
4.7.2	Speckled Hens and the Epistemology of Perception	77
4.7.3	The Causal Efficacy of the Normative	78
4.8	Conclusion	80
 II Foundationalism, Deception, and the Importance of Being Rational		81
 5 Foundations for the Foundational		82
5.1	Introduction	82
5.2	The Regress(es), with Possible Replies	83
5.3	Foundationalism, Reasons, and Basing	86
5.4	The Argument from the Nature of Defeat	87
5.5	Arguments in Favor of Answering No to the Propositional Question	88
5.5.1	From Mentalism about Reasons to a No Answer: Lyons's Argument	88
5.5.2	Dispatches from the Counterexample Wars	93
5.6	From Propositional Foundationalism to Doxastic Foundationalism	96
5.7	The Duplicate Reasons Problem	97
5.8	Summary and Conclusion	99
 6 Defeating the Externalist's Demons		103
6.1	Introduction	103
6.2	Preliminaries and a Statement of the Problem	104
6.3	A Solution to the New Evil Demon Problem	107
6.4	But Wait, There's More: The New New Evil Demon Problem	114
6.5	Solving The New New Evil Demon Problem	117
6.5.1	The Inferential Case	119
6.5.2	The Non-Inferential Case	120
6.6	Conclusion (With A Reason to Reject $P \rightarrow K$)	121
 7 What You're Rationally Required to Do and What You Ought to Do		123
7.1	Introduction	123
7.2	Why Be Rational?	125
7.2.1	Preliminaries	125
7.2.2	Coherence, Rationality, Deontic Significance	126
7.3	A Better Way	130
7.4	Ignorance and Obligation	132
7.4.1	Ignorance and Envelopes	133

7.4.2	Ignorance and Acting for the Right Reasons	138
7.5	Why You Ought to Be Rational	145

Chapter 1

Introduction

WE devote a good portion of our lives striving to figure out what it makes sense to do, believe, intend, desire and hope for given our often limited information and abilities. In other words, we devote a good portion of our lives striving to be rational. This dissertation seeks to understand what exactly we are striving for. In it I explicate and defend a general and unified account of rationality.

The thesis of the dissertation is that what it is to be rational is to correctly respond to the objective normative reasons you possess. In this introduction I briefly explicate what I take this claim to be and summarize how I argue for it in the pages to come.

Objective Normative Reasons, What

Let's start with *objective normative reasons*. Objective normative reasons are facts that count in favor of various attitudes and actions. Usually they are facts about the world (as opposed to facts about one's mental states). Some examples. The fact that the dissertation needs to be defended by June 30 is a reason for me write this introduction as soon as possible. The fact that it is going to rain in the city I'm traveling to next week is a reason for me to intend to bring an umbrella. The fact that you yawned is a reason to believe you didn't get enough sleep last night. The fact that you would be better off if you got a job is a reason for you to desire to get a job.

One project I do *not* want to take up here is the project of figuring out *what it takes* for some fact to be an objective normative reason. I am taking it for granted that there are objective normative reasons. That said, I do think there are some earmarks of normative reasons that we can use as heuristics for determining whether some fact is a normative reason.

One earmark is that objective normative reasons are the types of things that determine whether an attitude or action has strict normative statuses like justification, rightness, or correctness. If some fact f can explain why some action ϕ is justified, right, or correct, then (it is very likely that) f is a normative reason to ϕ .

Another earmark is being able to justifiably ϕ for f . When one ϕ s for f , f is the consideration that moves one to ϕ . If one can justifiably ϕ with f as one's reason for ϕ -ing, then it is very likely that f is a normative reason to ϕ .

A third earmark is the possibility that the relevance of f is mitigated by some other facts. Reasons are *defeatable* by other considerations. The reason to write this introduction isn't as weighty when I have multiple chapters that also need to be written than when it is the last thing I have to

write. The fact that I have multiple chapters to write weakens the reason to write this introduction. Thus, if it seems like the relevance of f to ϕ -ing can be mitigated (or enhanced) by adding or subtracting facts, it is very likely that f is a normative reason.

At various points in the dissertation I will use these earmarks to help determine whether some fact is a normative reason. In some cases, I will use the earmarks to argue that facts that many take to not be objective normative reasons really are objective normative reasons. I thus think there are more normative reasons than most moral philosophers. It will be important for the plausibility of my view about rationality that there are more reasons than is traditionally thought. So the views in the dissertation do depend on taking some stands on issues having to do with what it takes to be a normative reason. All of my arguments for my extensional claims will rely on the earmarks. Since the earmarks seem to me to be absolutely central to the notion of a normative reason, I am moved by my arguments. It's important for me to confess at the outset though that I will not be directly engaging with systematic theories of what it takes for some fact to be an objective normative reason. There are only so many battles I can fight.

It is important for me to say something at this point about the methodology of the dissertation. I am interested in making the best case for analyzing rationality in terms of objective normative reasons. A methodological reason I am interested in doing this is that it is fairly uncontroversial that there are objective normative reasons. Indeed, it is popular to think that normative reasons are the basic units of normativity—i.e., it's popular to think that all normative properties can be analyzed in terms of normative reasons.¹ Since most think objective normative reasons exist and have some theoretical importance, it would be nice if we could understand rationality in terms of normative reasons. It would allow us to understand one part of the normative landscape in terms of something we are already comfortable positing.

Possessed Normative Reasons, What

There are a lot of normative reasons out there. Many of them I don't know anything about. The ones I don't know anything about are reasons for me to act, believe, intend, and desire, but they aren't reasons to act, believe, intend, and desire that I *possess*. In order to possess a reason, I need to stand in some special epistemic relation to the fact that constitutes that reason. For example, I possess the reason to intend to take my umbrella next week because the weather report informed me that it is going to rain in the city I am traveling to. Similarly, I possess the reason to believe you didn't get enough sleep because I saw you yawn.

I think that the only reasons relevant to rationality are the possessed reasons. This is because it is implausible that reasons completely outside of your ken can affect what it is rational for you to do. If I have no clue that you just yawned, the fact that you yawned doesn't have any effect on which stance it is rational for me to take on the question of how much sleep you got last night. On the other hand, if the fact that you yawned is within my ken, it does seem to bear on which stance I should take on that question.

¹See, to name a few examples, Scanlon (1998), Parfit (2011), Schroeder (2007), Skorupski (2011), Dancy (2009).

Correctly Responding to Possessed Normative Reasons, What

Suppose the fact that you yawned is a good enough reason to believe that you didn't get enough sleep last night—i.e., it is a strong enough reason to make it rational to believe you didn't get enough sleep last night. And suppose further that I both possess that reason and believe that you didn't get enough sleep last night. Is this sufficient for my belief to be rational? In a word: No.

The reason why this is not sufficient for my belief to be rational is that it is possible for all of that to be true even though I don't *correctly respond* to the reason provided by the fact that you yawned. In the simplest kind of case, this is because there is no connection at all between my belief and that reason. Suppose I believe that you didn't get enough sleep because it's Tuesday and I believe (without good reason) that you never get enough sleep on Monday nights.

Because of this, I think that in order for token attitudes or actions to be rational, they have to be held or performed for the reasons that make them rational. This is what I think it is to correctly respond to the reasons you possess.

Thus, on my view, a token action or attitude is rational when (i) one possesses normative reasons to perform that action or hold that attitude that are sufficiently weighty, and (ii) one performs those actions or holds those attitudes for those reasons.

Let's turn to a summary of how I go about defending the view.

Part I: Coherence, Possession, and Correctly Responding

The dissertation is split up into two parts. Each part contains three chapters. In Part I I develop my positive account of rationality.

Chapter 2 deals with the relationship between rationality and coherence. Many metaethicists think that rationality is primarily about avoiding certain types of incoherence. According to a popular view—a view popular mostly due to the excellent work of John Broome—rationality merely requires coherence. A major cost of this view is that it commits one to thinking that rationality never requires one to perform particular actions or hold particular attitudes. Despite this, this view seems best given the current understanding of what the options are. I argue that my account offers a plausible alternative to this coherentist account while avoiding the main problems of competing accounts. Moreover, I argue, my account can also explain why it is that you are always irrational when you are incoherent—it is because you are always failing to correctly respond to some possessed reasons when you are incoherent. Along the way we learn some lessons about the structure of any plausible account that holds that rationality is a matter of correctly responding to normative reasons.

Chapter 3 defends a view about what it is to possess a reason. I argue that not only does one need to have epistemic access to the fact that constitutes the reason (a point everyone agrees on), one also needs to see, at least in some broad sense, the normative connection between the fact that constitutes the reason and the attitude or action it counts in favor of (a point that no one has made). Thus, I argue, any adequate theory has two tasks: It needs to tell us what the epistemic condition is and it needs to tell us what the so-called treating condition is. I argue for a view about both. I think the epistemic condition is being in a position to know. I argue for this view via process of elimination. I then argue that the so-called treating condition is being disposed to treat the fact as a reason (in the way it is a reason). I explicate a view of what it is to treat something as a reason. The central idea behind this account is that when you treat r as a reason to ϕ , you possess certain

dispositions to use r in deliberation in various ways. I argue that what it is to possess a reason is to meet the epistemic and treating conditions.

In Chapter 4 I turn to what it is to correctly respond to reasons. As I said above, I think this is a matter of performing actions and holding attitudes for the normative reasons that rationalize the acts and attitudes. I defend a causal account of what it is to ϕ for normative reasons. According to this account, ϕ -ing for normative reasons involves having dispositions that are sensitive to the normative facts as such. When you hold some attitude for a normative reason r , you are disposed to drop that attitude if r is defeated. When you act for a reason r , your action is the product of a disposition to act in that way when r is a normative reason. I argue that this account solves famous deviancy problems for causal accounts of ϕ -ing for normative reasons. I also generalize the basic idea to an account of the broader notion of ϕ -ing for reasons and argue that this account solves various deviancy problems for causal accounts of that phenomenon.

Part II: Foundationalism, Deception, and The Importance of Being Rational

Sometimes people are initially puzzled about why it would take a dissertation to defend my view. It seems so obviously correct when stated at a particular level of abstraction. Who could deny that what it's rational for you to do is what it makes sense for you to do given the information available?! Sometimes this puzzlement is meant to be something of an objection to the project; it puts the burden on me to explain why my view needs defending. For all of the people who give the puzzlement *qua* objection, I give you Part II.

Part II is dedicated to solving some menacing problems for my view. Chapters 5 and 6 focus on problems traditionally discussed by epistemologists, while Chapter 7 focuses on a problem discussed by metaethicists.

Chapter 5 is about foundationalism. Foundationalism is the view that the rational status of some rational beliefs doesn't depend on one having any other rational beliefs. Most epistemologists are foundationalists. An interesting question for the foundationalist is whether all foundational beliefs need to be held for sufficient normative reasons. Many epistemologists argue that they don't need to be or even that they can't possibly be. In Chapter 5 I argue that some of the more plausible arguments for these conclusions are misguided. They are misguided either because they are confused about normative reasons—e.g., they think normative reasons are constituted by mental states—or because they are confused about what it is to hold a belief for a reason—e.g., they think this has to be inferential. I argue that my views about these matters allow me to escape the force of these challenges. Moreover, my positive view ends up making it very plausible that foundational beliefs are held for sufficient normative reasons.

Chapter 6 is about the New Evil Demon problem for externalist views of rationality. Externalist views hold that the rational status of our attitudes and actions doesn't supervene on our non-factive internal states—i.e., on the states that we can be in even if the content is false. Externalist theories are plagued by the New Evil Demon problem. In order to see the problem, reflect on the following two characters. Sam inhabits the actual world. She thus sees lots of sights, hears lots of sounds, and learns lots of things via inference. Now consider Pam. She is Sam's non-factive internal state duplicate. That is, Sam and Pam have all the same non-factive internal states. The catch is that Pam is radically deceived by an evil demon. None of her beliefs about the external world are true. Many

have the strong intuition that Sam is rational just in case Pam is. It is hard to see how the externalist can capture this thought since the externalist thinks that Sam and Pam have much different rationalizers (since they are in much different positions with respect to the external factors relevant to rationality).

I am an externalist. This is because I think that the reasons you possess are facts that you are in a position to know. Thus, on my view, Sam has many more reasons than Pam since she's in a position to know many more facts about the external world. So I have the New Evil Demon problem. I argue that my view can solve the problem. That is, I argue that on my view Sam is rational just in case Pam is. This is so despite the fact that they don't possess the same reasons.

I also argue that there is another problem, the New New Evil Demon problem, that is even worse than the New Evil Demon problem. The New New Evil Demon problem is anchored in the thought that even if Pam always possesses sufficiently strong reasons, it is implausible that she always correctly responds to those reasons. Thus, it is implausible that her token attitudes and actions are rational. I argue that my view about ϕ -ing for a normative reason gives us the resources to solve this problem. Thus, I think that my view can solve both Evil Demon problems.

The final chapter is about the deontic importance of rationality. Despite a long history of thinking that rationality has a tight connection with what we ought to do, it has recently become popular to think that it is not the case that we ought to be rational. The final chapter takes up this challenge. I argue that my view can vindicate the claim that we ought to be rational. Indeed, I argue that what we ought to do *just is* what we are rationally required to do. Thus, rationality has ultimate deontic significance.

I argue for this by arguing that what we ought to do is a function of the normative reasons we possess. The anchor of the argument is that in order for some reason to obligate, it has to be potentially action guiding in a certain sense—it has to potentially be the reason for which one acts. The rub is that a reason can be potentially action guiding in this sense only if we possess the reason. At the end of the day, I argue, both our full stop obligations and the requirements of rationality are a function of the reasons we possess. That is, our full stop obligations just are what rationality requires of us.

I thus end up thinking that rationality is centrally important to the task of figuring out how to live. The importance of being rational is that by being rational we are living up to the standards we are obliged to live up to.

I should note before we begin that each chapter was primarily written as a stand alone essay. I have gone through with an eye of unifying the chapters in a way more fitting of a book. Despite this, there is some repetition, mostly in the form of reintroducing key notions multiple times without any hint that those notions were introduced in earlier chapters. This is perhaps annoying to readers who are reading the work from cover to cover (I suspect the number of readers doing this can be counted on one hand). But it will be a good thing for those who are selectively dipping into the work (perhaps I'll make it to a second hand counting these folks).

Part I

Coherence, Possession, and Correctly Responding

Chapter 2

The Coherent and the Rational

Facts about human life are the ones we should call on to develop an account of practical rationality: that is an account of goodness in respect of the recognizing and following of reasons for action as it is for creatures situated as human beings are. Moral virtues are needed for this practical rationality, as is a modicum of self-love, and the rightful pursuit of more desire-dependent ends. Someone acts well in all these respects only if he recognizes what he has reason to do in each area, and allows this to influence his conduct. Wherever exactly we want to use the word "irrationality," we must admit that someone is *deficient in practical rationality* if he fails to recognize moral reasons as reasons for acting; as also if he thinks he has no reason to look after himself, or to take trouble to attain his more particular ends. (Foot, 1997, pp. 322)

2.1 Introduction

SOMETIMES you are rational only when you have particular beliefs, intentions, desires, or perform particular actions. To give some examples, given that the evidence overwhelmingly supports believing that the earth is older than 300 years (and I'm considering whether it is), it would be irrational for me to not believe that the earth is older than 300 years. Similarly, given my (rational) intention to provide for my children's basic needs and my (required) belief that in order for me to provide for them I must make money, it seems like I would not be rational if I didn't intend to make money. Finally, given that I (rationally) desire to eat my supper by 6 and I know that in order to do that I have to go home in 7 minutes, it seems like I would not be rational if I didn't go home in 7 minutes. It's clear: Sometimes rationality requires me to have particular attitudes and perform particular actions.

Or maybe not. As it happens, the most popular view of rationality in the literature on practical reason denies that rationality ever requires one to have particular attitudes or perform particular actions.¹ This, I think, is extremely *prima facie* unintuitive. Despite this, there are few explicit defenses of the view that rationality sometimes requires particular attitudes and the performance of

¹Indeed, on the view of the most famous defender of these views, John Broome, rationality *never* requires one to perform any actions at all. See, e.g., (Broome, MS, ch. 1).

particular actions.² This chapter will offer a defense of such a view in the context of the literature on practical reason.

I will start by explicating some of the main contours of the debate over rationality in the literature on practical reason. We'll begin this explication by getting clear on just what people are trying to explain when they're offering theories of rationality.³ I will then explain the view that I wish to defend here. The bulk of the chapter will be concerned with rebutting a challenge to the type of view I defend. This challenge is, I conjecture, the main reason why the view receives little support in the literature on practical reason. I think this is unwarranted given the relative ease of meeting the challenge. Finally, at the end, I will consider again the initial motivations of the debate and explain how my view captures the motivations.

2.2 The Debate as it Currently Stands

The starting point for the current debate in the literature on practical reason is the fact that you are irrational if you are incoherent in certain ways. The paradigm is means-end incoherence. You are means-end incoherent when you intend to ϕ , believe that in order to ϕ you must ψ , but fail to intend to ψ . Other examples include akrasia—believing that you ought to ϕ but failing to intend to ϕ —and having inconsistent intentions. For the most part, the whole debate about rationality is about what best explains why you are irrational when you are incoherent in these ways.

Given this motivation, the debate has focused on *conditional requirements*. This is because the incoherence data is about which *combinations of attitudes* are irrational. For example, not having the intention to ψ when one intends to ϕ and believes that in order to ψ one must ϕ seems to be a paradigmatically irrational combination. Conditional requirements are tailor made to explain why it's bad for certain combinations of facts to obtain. This is because they require things only given some other facts. Given these thoughts, we might be tempted to endorse something like Means-End N:

Means-End N: If you intend to ϕ and believe that in order to ϕ you must ψ , then you are rationally required to intend to ψ .

Means-End N explains why you are irrational when you are means-end incoherent. For every time you are means-end incoherent, you will lack an intention that rationality requires you to have—viz., the intention to ψ . Notice also that, according to those who endorse Means-End N, rationality sometimes requires you to have particular intentions. So far, so good.

Despite its ability to account for the incoherence data, most ethicists think that Means-End N is obviously false. This is because, they have thought, it gives rise to objectionable bootstrapping. The basic thought behind the worry is that it's implausible to think that you are required to intend to ψ , for any ψ , simply in virtue of having some other intentions and beliefs. For example, it's implausible that you are required to intend to become the King of France simply in virtue of the fact that you

²This is true for different reasons in the metaethics literature than it is in the epistemology literature. It's true in the metaethics literature because most think that the view is false. It's true in the epistemology literature because most people think it's obviously true—at least since the death of a particularly naïve form of coherentism.

³As I say above, the context will be the literature on practical reason. That said, a good bit of the chapter is about epistemology. This is in line with the literature on practical reason since the default view is that the practical and epistemic will be treated in similar ways (see, e.g., Broome (MS), Kolodny (2007b), Schroeder (2008)).

intend to be a European monarch and you believe that in order to become a European monarch you have to be the King of France. After all, it seems like it would be more rational for you to resolve the incoherence by giving up your intention to become a European monarch than to form the intention to become the King of France. Giving up your end certainly seems *permitted* by rationality. Means-End N seems to rule this out.

This problem is even worse if you think rationality is deontically significant in some sense—i.e., bears on what you ought to do.⁴ There is a spectrum of ways that rationality could be deontically significant. On one extreme, necessarily you ought to do what rationality requires of you. Means-End N seems hopeless if you think rationality is deontically significant in this sense. It is uncontroversial that it is not the case that you ought to intend to be the King of France simply because you intend to be a European monarch and believe that you must be the King of France in order to be a European monarch. On the other end of the spectrum of deontic significance, it's just that there is always a reason for you to do what rationality requires. It doesn't follow from this view that you ought to always do what rationality requires. Often times the reason you have to be rational will be outweighed by stronger reasons. This still seems implausible to most.⁵ That is, most ethicists think it is implausible that you can bootstrap into existence a reason to intend to be King of France simply by intending to be a European monarch and believing that you must be the King of France in order to be a European monarch.

These objections have led nearly everyone to reject narrow-scope requirements like Means-End N. This has not led them to reject the incoherence data we started out with, though. For the most part, the literature hasn't doubted that the primary concern for a theory of rationality is to account for the incoherence data. Since the (purported) problem with Means-End N is that it forbids too much, philosophers have thought that the fix is to propose requirements that *merely* ban incoherence. They have thought, that is, that instead of 'rationality requires' taking narrow-scope over the consequent of the conditionals, it rather takes wide-scope over the whole conditionals. So, the thought goes, we should replace Means-End N with Means-End W:

Means-End W: You are rationally required to [intend to ψ if you intend to ϕ and believe that in order to ϕ you must ψ].

Means-End W doesn't give rise to any objectionable bootstrapping. This is because you comply with it when you drop your intention to be a European monarch. Nevertheless, Means-End W provides an explanation for why you are always irrational when you are means-end incoherent. You are irrational when you are means-end incoherent because you violate Means-End W when you are means-end incoherent—i.e., a conditional you are rationally required to make true is false.

A striking feature of Means-End W is that it doesn't ever require you to have any particular attitudes. You comply with it as long as you are coherent. This is the feature that allows the view to avoid bootstrapping. You comply with Means-End W, and are thus instrumentally rational, when you drop your intention to be a European monarch and/or drop your belief that in order to be a European monarch you must be the King of France.

This is, at least in a rough way, the direction the debate about practical rationality has taken in the last decade. The focus has been on how to explain the irrationality of the incoherent states,

⁴In §7.2 of chapter 7 I will go through the arguments briefly explicated here in more detail.

⁵Schroeder (2004, 2005b) argued that this isn't in fact implausible. He seems to have later given that up (see especially Schroeder (2009b)).

and people have been driven to the wide-scope view by these types of considerations. Moreover, if you focus solely on the irrationality of the incoherent states, it seems like the same type of arguments in favor of wide-scoping will crop up in each case. This leads to the acceptance of whole-hog wide-scoping, or something near enough. This has come to fruition especially in the work of John Broome. Broome is, as far as I can tell, a whole-hog wide-scoper (although see note 7 below).

However, avoiding the bootstrapping objections in this way comes with distinctive costs. Most importantly for my purposes, the coherentism of Means-End *W* makes bad predictions about other cases. To take the case from the opening lines of the chapter, it doesn't seem right to say dropping my intention to provide for my children's basic needs is on a par when it comes to instrumental rationality with my forming the intention to make money.⁶ Coherence doesn't always seem to be *sufficient* for rationality.⁷ Even if it is a fact (and it is) that I'm irrational when I'm means-end incoherent, it's not right to say that I'm instrumentally rational so long as I'm coherent. And this point seems to apply *mutatis mutandis* for all of rationality.

This seems like a classic case of philosophical impasse. Both views capture some of the data while making some bad predictions. On the one hand, the narrow-scoper captures the idea that coherence isn't always sufficient for rationality, but (purportedly) predicts that it's always irrational to drop ends once we have them. On the other hand, the wide-scoper captures the idea that it's (sometimes) rational to drop ends, but predicts that it's always rational to drop ends.

Perhaps the debate between the wide- and narrow-scoper is at an impasse. I don't think that this entails, however, that there is an impasse over the debate over *rationality*. I think there is a third, better, kind of view about rationality. It is this kind of view that I will defend here. Our next step is to see what this kind of view is.

2.3 A Better Way

Talk of reasons has come to dominate all corners of normative philosophy. This, I think, has not been a mistake. Reasons are, at the very least, some of the most fundamental normative entities. Indeed, I'm very tempted by the thought that reasons are the *most fundamental* normative entities there are. All of normativity reduces to reasons. Or so a common thought goes.⁸

⁶Wide-scopers might try to temper this bad result a bit by arguing that, at least sometimes, I won't be able to give up some of the antecedent attitudes. They could then follow Greenspan (1975) and hold that if you ought (ϕ or ψ) and you can't ϕ , then you ought ψ (for further discussion of this principle, see Setiya (2007a), Schroeder (2009b)). The pure coherence result would thus not follow in cases where you can't give up one of the antecedent attitudes. Indeed, the wide-scope requirements would collapse into the narrow-scope ones if it was never possible to give up the antecedent attitudes. Because of this, nearly everyone since Greenspan has thought that this is a *objection* to wide-scoping. But in our context it helps with the pure coherentism charge (of course, it only helps insofar as you think the narrow-scope requirements are plausible). Thanks to Michael Smith for raising this possibility.

⁷As it happens, even Broome agrees that coherence isn't always sufficient for a particular token attitude to be rational. This is because he thinks there are some negative 'basing' requirements. These requirements forbid having or dropping attitudes for particular reasons. The paradigm example is dropping your belief that you ought to ϕ *because* you realize you aren't going to ϕ , anyway. This seems irrational. Broome explains this by appealing to the basing principles. The basing principles never require you to have particular attitudes, though. They only forbid you for having or dropping particular attitudes for particular reasons. Thus, on Broome's view, it's always permissible to drop your belief that you ought to ϕ , just so long as you don't do it for particular reasons (cf. Way (2011) for a similar view).

⁸See Scanlon (1998, MS), Parfit (2011), Schroeder (2007), Dancy (2000, 2009), Skorupski (2011).

Suppose you start with this assumption—viz., that all of normativity reduces to reasons. And then add the extremely plausible claim that rationality is normative. You would then be committed to thinking that rationality reduces to reasons somehow. Since the two starting points seem so plausible, it seems like this hypothesis is something we should take seriously. This is only reinforced by the seeming impasse between the wide- and narrow-scopers.⁹

'Ok,' you might think 'let's start thinking about whether rationality reduces to reasons *somehow*. But how?' I have a particular type of reduction in mind, and it will be fruitful to make it clear how the views I'm interested in are different than others that have been discussed in the literature. The reductive views I'm interested in hold that *what it is* to be rational is to *correctly respond to reasons*. These views hold that what is rational is *determined by* reasons. I'm interested, then, in defending a view committed to Determination:

Determination: What it is to be rational is to correctly respond to reasons.

Those who accept Determination will automatically think that there are reasons to be rational. After all, what's rational is determined by what reasons there are.

A word is in order about which reasons I'm talking about. The term 'reason' is used in a lot of different ways in both philosophy and on the street. The type of view I'm interested in holds that what it is to be rational is to correctly respond to *normative* reasons. Normative reasons are the type of reasons that justify actions and attitudes. To give some examples, the fact that the road is icy is a reason not to drive; the fact that it is cold is a reason to wear a coat; the fact that my wife is loyal is a reason to love her.

For the most part in the literature on normative reasons, normative reasons are assumed to be *objective*. Objective normative reasons have two important features for our purposes. First, they are constituted by *facts*. Second, whether a fact is an objective normative reason is, at least in paradigm cases, independent from our beliefs. In the end I'm interested in defending the thought that what it is to be rational is to correctly respond to objective normative reasons.¹⁰ However, here I want to be neutral between objective reason views and certain subjective reason views. Subjective normative reasons are determined by our beliefs. For example, according to Schroeder (2009b), *r* is a subjective reason for *A* to ϕ just in case (i) *A* believes *r* and (ii) *r* is an objective reason to ϕ if true. I think that all of my arguments here could be adopted by someone who thinks that what it is to be rational is to correctly respond to subjective normative reasons.¹¹ I'm only interested in views that hold that what it is to be rational is to correctly respond to normative reasons, either objective or subjective.¹²

⁹I think I can motivate my view using a weaker and even more plausible assumption. Namely, that the *deontic* reduces to reasons. It's plausible—recent skepticism notwithstanding—that rationality is deontically significant (that rationality affects what you ought to do). It's very plausible that the deontic reduces to reasons because deontic facts seem to be grounded in contributory facts. The most natural normative contributory facts are reasons.

¹⁰I argue for this in chapters 3 and 7.

¹¹Although I caution the reader that I will sometimes speak as if these views are false.

¹²Note that by being officially neutral between whether the normative reasons that are relevant are objective or subjective I avoid problems having to do with the rational significance of false beliefs. On views that hold that only objective reasons matter, the contents of false beliefs cannot be the type of thing that affects the rational status of ϕ -ing (since objective normative reasons are constituted by facts). These views thus face a problem with cases where it seems as if the contents of false beliefs are playing a role in determining the rational status of ϕ -ing. Views that hold that subjective reasons matter don't face these problems since subjective reasons can be false. At the end of the day

It's important to contrast Determination with other views that hold that there are reasons to be rational. According to a type of view that has received an enormous amount of attention from John Broome and Niko Kolodny, there are reasons to be rational.¹³ However, on this view, what is rational is not determined by reasons. Instead, rationality is determined by facts about coherence.

According to the type of view discussed by Kolodny and Broome, we can discover what the requirements of rationality are independently from whether there are reasons to be rational. Once we do this, we can then ask if there are any reasons to comply with the coherence requirements. Broome focuses most of his energy thinking about whether the fact that rationality requires ϕ -ing is itself a reason to ϕ . Kolodny (and Broome to some extent) explores different possibilities. However, both discussions assume that the requirements of rationality are a function of something other than reasons.

I am not interested in defending the type of view that Broome and Kolodny discuss.¹⁴ I am not, that is, interested in whether we have reasons to comply with coherence requirements like Means-End W.¹⁵ I am instead interested in whether the requirements of rationality are themselves determined by reasons.

It's perhaps helpful to say a few words about what I take the *requirements* of rationality to look like on a view that accepts Determination. I think there will be requirements like Action Schema, Belief Schema, and Intention Schema:

Action Schema: If there is a set of decisive reasons S to ϕ , then you are rationally required to ϕ .

Belief Schema: If there is a set of decisive reasons S to believe p , then you are rationally required to believe p .

Intention Schema: If there is a set of decisive reasons S to intend to ϕ , then you are rationally required to intend to ϕ .

Of course, one could accept instantiations of these schemata and not hold Determination. Those who accept Determination accept instantiations of the schemata *and* hold that when, e.g., there is a set of decisive reasons to ϕ , you are rationally required to ϕ because of the reasons in the set.

To anticipate, I will be arguing that the set of reasons that determines what is rational is the set of reasons that you *have or possess*. In order to possess a reason, you have to stand in some type of epistemic relation to the reason. Moreover, as I'll argue, you also have to treat the proposition that is the reason as a reason. I will argue that this type of view can avoid the main objections raised in

I do have this problem since I ultimately hold that only objective reasons matter. But the views defended in this chapter do not have this problem. For my solution to the problem, see chapters 5 and 6 (chapter 6 is devoted to this problem).

¹³This view is discussed most prominently in Kolodny (2005, 2007a,b) and Broome (2005a,b, 2007c, 2008a).

¹⁴More carefully, I'm not interested in the type of view Broome discusses in the papers cited in the last note. I am interested in a type of view Broome discusses in Broome (2007a).

¹⁵It is important to note that this is not equivalent to the claim that we are always irrational when we are incoherent. Below I will argue that whenever we are incoherent, we are doing something we possess reasons not to do. But this isn't to immediately say that we have reasons to comply with Means-End W. This is just to say that we have reasons to do something that guarantees that we comply with Means-End W. In order to have reasons to comply with Means-End W, we need to have reasons to be coherent that are over and above our reasons to have particular attitudes. This is spelled out more fully in chapter 7, §7.2.2.

the literature on practical reason to views that accept Determination. I will also argue that my view can account for the initial motivation—to explain why one is irrational when incoherent—while avoiding the major problems for both the narrow-scope view and the wide-scope view.

It turns out that there hasn't been much discussion of views that embrace Determination, at least in the literature on practical reason. This is reflected by the fact that most of the debate has been about the wide- and narrow-scope requirements. There has been *some* critical discussion of Determination, most notably Broome (2007a, MS). Broome argues that what's rational is not determined by reasons. My defense of the idea, then, will start by defending the view against Broome's criticisms. By the end, we'll have a more fleshed out version of Determination.

2.4 Broome's Challenge

In Broome's terminology, the views under consideration hold that rationality consists in correctly responding to reasons. He wants to know whether it does. He argues that it doesn't via counterexample.

Broome is interested in showing that Equivalence is false:

Equivalence: Necessarily, you are rational if and only if you correctly respond to reasons.

Equivalence is weaker than Determination. It merely says that correctly responding to reasons is a necessary and sufficient condition for being rational. Determination entails this, but it also makes a claim about what determines what's rational. Since I hold Determination, I also hold Equivalence.

Broome argues against Equivalence by arguing against the necessity claim—i.e., the claim that if you are rational, then you correctly respond to reasons. Here's his case from Broome (2007a).¹⁶ Suppose Jane goes to her local seafood restaurant. She orders her favorite dish. Unbeknownst to her, the fish contains salmonella. The fact that the fish contains salmonella is, let's stipulate, a decisive reason not to eat the fish. However, it seems like Jane is rational to eat the fish given her ignorance of the salmonella. Indeed, Broome adds, sometimes Jane can even know about the salmonella and still be rational in eating. These are cases where even though Jane knows the descriptive fact that the fish contains salmonella, she doesn't know that that fact is a strong reason not to eat the fish. Jane's case(s), then, seem like cases where someone is rational even though they don't do the thing that the reasons decisively support.

It's at this point that we need to start getting into the details of particular ways to flesh out requirement schemata like Action Schema. Jane's cases are counterexamples to certain instantiations, but not others. The crucial difference between these views is *which reasons* go into set S . According to one view, all the reasons go into S . This gets us Action All:

Action All: If the set S of all the reasons decisively supports ϕ -ing, then rationality requires you to ϕ .

Jane's cases are certainly counterexamples to Action All. It's very implausible—and Jane's cases bring this out nicely—that you're always irrational when you fail to do things that all the reasons

¹⁶In Broome (MS) he doesn't give a specific case, but merely the schema for the type of case he provides in the earlier paper.

support.¹⁷ This is because it's obvious that (at least sometimes) when you are ignorant of some reason r to ϕ , r does not affect the rational status of ϕ -ing.

Jane's cases, however, are simply not counterexamples to more plausible instantiations of Action Schema. The type of view I wish to defend holds that rationality is a function of the reasons you *possess*. In order to possess a reason, you have to stand in some type of epistemic relation to that reason. For example, it's plausible that you meet the epistemic condition when you know the reason.¹⁸ That is to say, knowing some reason r seems sufficient for meeting the epistemic condition for possessing r .¹⁹ The important feature of all of the views that appeal to possessed reasons is that they all agree that when you are ignorant of some reasons, those reasons don't affect what it is rational to do.²⁰ So the first version of Jane's case is not a counterexample to Action Possess:

Action Possess: If the set S of the reasons you possess decisively supports ϕ -ing, then rationality requires you to ϕ .

The epistemic relation part of the possession views easily accounts for the first version of Jane's case—the version where she is ignorant of a descriptive fact. Given her ignorance, she doesn't possess the reason provided by the fact that there is salmonella in the fish. Since she doesn't possess that reason, it doesn't bear on whether she's rational. Moreover, given the other things she knows, it seems rational to eat.

It doesn't immediately solve the problem posed by the second version of the case, which is the version where she knows the descriptive fact but doesn't know that the descriptive fact is a reason. Notice that, in that version of the case, Jane doesn't *treat* the fact that there is salmonella in the fish as a reason not to eat even though she knows there is salmonella in the fish. This seems relevant to whether she possesses that fact *as a reason to not eat the fish*. This, in turn, suggests that standing in the epistemic relation is only a necessary condition for possessing the reason. You also need to meet some *treating* condition. So in order to know if Action Possess explains the second case, we need to know what the *treating* condition comes to.

¹⁷This, like all claims of interest to philosophers, is controversial. Some, most notably Gibbons (2010), hold that only facts within your perspective can be normative reasons. His argument for this, roughly, turns on the thought that in order to be a normative reason ϕ , a fact must make ϕ -ing reasonable. But, he claims, only facts within one's perspective can make ϕ -ing reasonable. I think this is to throw the baby out with the bath water. Facts outside of our perspective can have the hallmark features of normative reasons. We can explain their added powers once they are in our perspective by appealing to the distinction between normative reasons there are and the normative reasons one possesses.

¹⁸As we'll see, it's not enough to know the descriptive fact that is the reason. You have to know the descriptive fact and treat it as a reason.

¹⁹It's very controversial whether this is also necessary. Chapter 3 will be devoted to seeing what the best of view is of what it takes to possess reasons. All views discussed there, though, hold that knowing is sufficient for possession.

²⁰You might worry that this is fine as far as it goes when it comes to non-culpable ignorance, but won't adequately account for cases of culpable ignorance. It's plausible that sometimes when you are culpably ignorant of some reason r to ϕ , r does affect whether it's rational to ϕ . I agree with this, and that's why I argue in chapter 3 that the epistemic condition on having reasons is a *non-holding* condition—i.e., you can meet it without holding beliefs. It's worth pointing out that some cases of culpable ignorance are cases where you are culpable because you do something that you were required not to do that prevents you from meeting the epistemic condition (even if it is a non-holding condition). In those cases I think you don't possess the reason you are culpably ignorant of, but you've obviously still done something you were required not to—viz., the thing that prevented you from meeting the epistemic condition.

At first pass, you might think that you meet the treating condition just in case you believe p is a reason. This is the view that Broome seems most sympathetic to in Broome (2007a).²¹ It also seems to be the view that T.M. Scanlon endorses in Scanlon (1998, 2007). I don't think believing p is a reason is a necessary condition for treating something as a reason, however. I think this would underdetermine the extension. We'll stick to the easy cases. It seems like creatures that lack enough of the concept of a reason to have beliefs about reasons can still treat certain things as reasons. My 4 year old son, for example, certainly treats the fact that he won't get his dessert unless he eats his dinner as a reason to eat his dinner. But I don't think he believes that that is a reason to eat his dinner. Similarly, my dog treats the fact that the food is stored under the sink as a reason to pace incessantly around the sink until she's fed, but she doesn't believe that that fact is a reason to pace.

So what are the less conceptually sophisticated members of my family doing in these cases when they treat those facts as reasons? They are, I think, manifesting certain dispositions—some of which are the same kinds of dispositions one has when one believes that p is a reason. The most salient disposition in these cases is the disposition to use p as a premise in reasoning. My dog is disposed to move from the judgement that the food is kept under the sink to the act (or intention to act) of incessantly pacing by the sink. Similarly, my son is disposed to use the claim that in order to get dessert he must eat his dinner as a premise in a piece of reasoning that culminates in him eating his dinner (or intending to eat his dinner).

Perhaps an example involving theoretical reasoning will crystallize the point. Imagine I believe it's 2 pm on Tuesday and I believe that if it's 2 pm on Tuesday, then I should be writing. If I'm rational, I will treat the contents of those beliefs as reasons for thinking I should be writing. Most importantly, this means that I am disposed to form the belief that I should be writing when I believe the other propositions. It seems like an overintellectualization to say that I don't treat those premise beliefs as reasons unless I believe they are reasons. It seems like I am doing all that needs to be done to treat them as reasons by being disposed to use them in reasoning in certain ways.

It's worth noting that we usually, perhaps always, treat p as a reason when we believe that p is a reason. This is true because of very general features of what it is to have a belief. The most important feature of belief is that we are often (perhaps necessarily) disposed to use the propositions we believe in reasoning in certain ways. When it comes to our beliefs about reasons, the things we are disposed to do in virtue of believing that those things are reasons are often the same (or at least very similar) to the things we are disposed to do when we treat propositions as reasons. So the normative-belief view has something going for it. It just puts the cart before the horse.²²

Now back to Jane. An important factor for which things you treat as reasons is the attitudes you already have. You are unlikely to treat the fact that smoking is pleasurable as a reason to smoke if you think that pleasure is worthless. And Jane is unlikely to treat the fact that the fish contains

²¹There he says that the reason why the fact that there is salmonella doesn't affect the rationality of eating is because Jane doesn't believe that the salmonella constitutes a reason not to eat. In a different context (namely, Broome (FC)) he argues that believing p is a reason to ϕ is not necessary for treating p as a reason to ϕ , for reasons similar to the ones articulated below.

²²For reasons that I make clear in chapter 3 I don't think actually treating p as a reason to ϕ is what you need to do in order to meet the treating condition. Instead, I think you meet the treating condition just in case you are *disposed* to treat p as a reason to ϕ . This doesn't matter much here, for my explanation of what's going on in Jane's case doesn't turn on this difference. I'll assume for simplicity that you meet the treating condition just in case you treat p as a reason to ϕ .

salmonella as a reason to not eat the fish if she doesn't believe that salmonella is harmful, or that getting very sick is undesirable.

In order to fully evaluate the second version of Jane's case, we need more details. For it's clear that sometimes we irrationally treat things as reasons, or we irrationally fail to treat things as reasons. It's not enough, that is, to find out that Jane does not believe that the fact that there is salmonella in the fish is a reason not to eat. For this to be a compelling case, we need to know that she is non-culpably ignorant of that fact. Moreover, as I hope I've motivated, it's not enough for her not to believe that the fact that there is salmonella in the fish is a reason not to eat the fish. It also has to be the case that she doesn't treat that fact as a reason.

In order to have a clear view of the terrain, let's flesh out two versions of the case:

Culpable Ignorance

Jane has a basic knowledge of food science. She knows that salmonella causes serious illness. She doesn't believe that being sick on Tuesdays is bad for her, though. This is despite the fact that she has normal evidence about the the relation between illness and Tuesdays. That is, she has good reason to think that the days of the week are arbitrary when it comes to the pain of being ill. Moreover, she has excellent reason to believe that the pain of illness will be just as bad on Tuesday as on any other day. Since today is Tuesday, she doesn't care about becoming sick today. So she doesn't treat the fact that there is salmonella in the fish as a reason to refrain from eating the fish.²³

Non-Culpable Ignorance

Jane knows that salmonella is a bacteria that is often found in food. However, she also believes that salmonella is one of the many forms of bacteria found in food that is harmless to humans. Moreover, this belief is rational. A renowned food scientist told her this bacteria is harmless. Thus, she doesn't treat the fact that the fish contains salmonella as a reason to refrain from eating the fish.

If you are anything like me, you have different reactions to the two versions of the case. In Culpable Ignorance, Jane seems clearly irrational. She is irrational for not treating the fact that there is salmonella in the fish as a reason to refrain from eating. This is because she has an irrational belief about whether being ill on Tuesday is bad for her. Of course, since she has that irrational belief, she doesn't actually possess the reason to refrain from eating the salmonella. So *eating* might be rational. But she is still irrational for not treating the fact that there is salmonella in the fish as a reason not to eat. Moreover, if she were to rectify this mistake—if she started to treat the fact that there is salmonella as a reason not to eat—then she would possess that reason and would be rationally required not to eat.

This is not what's going on in Non-Culpable Ignorance, though. In that case, Jane does seem rational. And this is so even though she does know a fact that constitutes a strong reason not to eat. The reason why she is rational, I claim, is that she doesn't possess that fact as a reason. She doesn't possess that reason because she doesn't treat that fact as a reason to refrain from eating. Moreover, in Non-Culpable Ignorance, unlike in Culpable Ignorance, Jane is rational in not treating the salmonella as a reason to refrain. So, for all we know, she is rational *tout court*.

²³This case is inspired by Parfit (1984)'s famous case of future Tuesday indifference.

I take it the best type of case for Broome's purposes is something like Non-Culpable Ignorance. But that case isn't a counterexample to a suitably refined version of Action Schema. According to that version, you are irrational when you ϕ just in case you *possess* decisive reasons not to ϕ . You possess some reason r to ϕ just in case you stand in some epistemic relation R to r and you treat r as a reason to ϕ . This handles both of Broome's purported counterexamples. It handles the first because in those cases you fail to possess the relevant reason by not standing in the epistemic relation to that reason. It handles the second because in that case Jane fails to possess the relevant reason because she fails to treat it as a reason.

One might object at this point that making this move threatens my attempt to avoid the problematic features of the narrow-scope view.²⁴ Recall that narrow-scope views are objectionable because they predict illicit bootstrapping. One can bootstrap into a rational requirement to do heinous things just by having certain antecedent combinations of attitudes. This is implausible on its own, but is even worse if you think that rationality has some connection with what we ought to do, full stop.

This last point is the important point for this objection to my view. For consider Jane again. On my view, Jane is rationally permitted to eat the contaminated fish, *even though* the balance of all the reasons decisively supports not eating the fish. Thus, my view doesn't guarantee that rationality co-travels with what all of the reasons support. One might think of this as a form of bootstrapping akin to the bootstrapping that plagues the narrow-scope view.

Of course it is true that my view doesn't guarantee that rationality co-travels with what *all* the reasons support. This is so by design—otherwise Jane's case is a counterexample to Determination. Moreover, it is true that if you think that our obligations are determined by all the reasons, my view of rationality will come apart from obligation. With these concessions in hand, I have two replies.

First, these concessions aren't particularly damaging in this context. This is because most theorists working on these issues now think that it is an open question whether rationality affects obligation.²⁵ This is an assumption made by wide- and narrow-scopers alike. Nevertheless, those working on this topic are still confident that rationality exists and it requires things of us (even if the things it requires aren't always the things we ought to do).

Even given this background contextual information, there is a bootstrapping problem for the narrow-scope view. That is, even if you give up on the idea that rationality co-travels with obligation, the narrow-scope view predicts objectionable bootstrapping.²⁶ We can see this by thinking of cases where one's antecedent attitudes are themselves irrational.²⁷ Take our example from §2.2 of the poor fool who intends to be the King of France. Given plausible background assumptions, it is irrational for him to intend to be a European monarch and irrational for him to believe that in order to be a European monarch, he needs to be the King of France. Given that those attitudes are irrational, it is implausible that they *make it the case that* he is rationally required to intend to be the King of France.

This is a bootstrapping problem for narrow-scope views even if we give up that rationality always affects obligation. Moreover, it is clearly a problem that my view doesn't have. My view doesn't predict that our poor fool is rationally required to intend to be the King of France. In fact,

²⁴I thank an anonymous referee for raising this objection.

²⁵A classic early statement of this view is Parfit (1997). See also Brunero (2010), Reisner (2009), Broome (MS).

²⁶I'm not the only one who thinks this. Most of the theorists cited in the last note also think this (save Parfit).

²⁷These are the cases that many wide-scopers lean on in recent literature. See especially Brunero (2010, 2012), Way (2011).

my view can give a nice explanation of why it is true that, given normal background conditions, those antecedent attitudes are irrational. They're irrational because he doesn't possess sufficient reason to hold them. Further, my view can nicely explain why it is that forming the intention is rationally subpar—it's because he is required by the reasons he possesses to give up the antecedent attitudes. Thus, my view does avoid the objectionable bootstrapping that recent theorists have been concerned with.

As we have seen, there is another type of bootstrapping objection in the neighborhood. According to this objection, a view of rationality predicts bootstrapping if it predicts that what we are rationally required to do comes apart from what the set of all the reasons supports. Given the context just explicated, it is tendentious whether there is any *objectionable* bootstrapping here.

This, however, is not my reply to this second bootstrapping objection. Instead, I think something more ambitious. As it happens, it is very controversial whether *all* the reasons determine what we are obligated to do. Call the view that all the reasons determine what we ought to do objectivism about obligation. Many prominent theorists deny objectivism.²⁸ Most who deny objectivism hold that only facts within one's *perspective* can determine what one ought to do—call this view perspectivalism. If objectivism is false and perspectivalism is true, then it is much more plausible that my view *does* co-travel with obligation. After all, my view of rationality is perspectival. The reasons you possess are all within your perspective. Moreover, it seems as if the reasons you possess are poised to explain all the data perspectivalists want to explain.²⁹

Thus, I think that my view stands a good chance of avoiding this second bootstrapping worry, as well. Even though it is clear that on my view rationality and the verdicts of all the reasons come apart, it is not so clear that on my view rationality and *obligation* come apart. This is what a fully ambitious view of rationality would want to secure.

This feels like progress. While Broome's cases aren't counterexamples to all instantiations of Action Schema, reflection upon them has helped us flesh out some details that any plausible instantiation must have. With this view in hand, let's return to the initial motivations of the debate between the wide- and narrow-scooper.

2.5 What about Coherence?

In §2.2 I agreed that it is a datum that you are irrational when you are incoherent in certain ways. I have yet to say how the view I'm advocating accounts for this. I have said, of course, that I reject certain views that do capture this. In particular, I reject views that hold that being coherent is both necessary *and sufficient* for being rational. I do think being coherent is necessary, and I'll now explain why.

I'll offer the explanation via examples. In the second example, I'll give a first pass explanation that is, as we'll see, not strictly speaking adequate. There is a hitch which will be dealt with in due course.

²⁸See, among many others, Scanlon (2008), Jackson (1991), Ross (2012), Zimmerman (2008), Kiesewetter (2011).

²⁹I will defend this claim in detail in chapter 7. See also Ross (2012), Gibbons (2006, 2010).

Suppose that you believe p , believe if p then q , but fail to believe q .³⁰ You are closure incoherent, and, I think, irrational. Why are you irrational? Recall that on my view, Belief Possess is a (schema for a) requirement of rationality:

Belief Possess: If the set of reasons you possess decisively supports p , then you are rationally required to believe that p .³¹

If you combine this with the very plausible Closure Transmission

Closure Transmission: If the reasons you possess decisively support p and the reasons you possess decisively support if p then q , then the reasons you possess decisively support q .

then we can explain why you are always irrational when you are closure incoherent.

I think you are irrational because Reasons-Coherence, Closure is true:

Reasons-Coherence, Closure: If you are closure incoherent, then you are not correctly responding to all of the reasons you possess.

Reasons-Coherence, Closure is true because it will either be true that you possess decisive reason to believe p , decisive reason to believe if p then q and thus by Closure Transmission possess decisive reason to believe q , or you won't. If you do possess decisive reason to believe p , decisive reason to believe if p , then q , then you'll possess decisive reason to believe q , and thus you are irrational when closure incoherent because you don't believe q . If you lack decisive reason to believe p or lack decisive reason to believe if p , then q , then you will be irrational when closure incoherent because you hold one of those beliefs and you are rationally required not to. Either way, you will be irrational when closure incoherent.

We can say similar things about means-end incoherence. Suppose you intend to ϕ , believe that in order to ϕ you must ψ , but fail to intend to ψ .³² Recall that, on my view, Intention Possess is a (schema for a) requirement of rationality:

Intention Possess: If the reasons you possess decisively support intending to ϕ , then you are rationally required to intend to ϕ .

If you combine this with the very plausible Means-End Transmission

Means-End Transmission: If the reasons you possess decisively support intending to ϕ and the reasons you possess decisively support believing that in order to ϕ you must ψ , then the reasons you possess decisively support intending to ψ .

³⁰And suppose that you are considering q , q is a matter of importance for you et cetera.

³¹This is not the most general requirement applying to doxastic stances towards propositions. It will get us all we want for belief regarding p . This is because it will explain why you are required to believe p when you possess decisive reason for p and why you are required to believe $\neg p$ when you possess decisive reasons for $\neg p$. However, this leaves out withholding. A more general principle says something like: For any stance S towards p , if you possess decisive reasons to take S towards p , then you are rationally required to take S towards p .

³²And your plans to ϕ are relevant to current planning, and you have access to your intention to ϕ and your means end-belief, et cetera.

then we can explain why you are always irrational when you are means-end incoherent.

I think you are irrational because Reasons-Coherence, Means-End is true:

Reasons-Coherence, Means-End: If you are means-end incoherent, then you are not correctly responding to all the reasons you possess.

To see that Reasons-Coherence, Means-End is true, consider the fact that the reasons you possess will either decisively support intending to ϕ and decisively support believing that in order to ϕ you must ψ , or they won't. If they do, then it follows from Means-End Transmission that the reasons you possess decisively support intending to ψ , and you are irrational when means-end incoherent because you lack an intention—the intention to ψ —that you are rationally required to have. If they don't, then when you are means-end incoherent you are irrational because you either have an intention—the intention to ϕ —that you are rationally required not to have, or you have a belief—the belief that in order to ϕ you must ψ —that you are rationally required to not have, or both. In any case, you are irrational when you are means-end incoherent.

This explanation of Reasons-Coherence, Means-End is too quick. This is because oftentimes there are multiple options that are *merely permissible*. That is, there are often multiple possible actions one can perform that are all permissible even though none of them are required. Right now it's permissible for me to work on this chapter, but it's also permissible for me to do some reading, or send some emails. Moreover, none of those things are required. Imagine I decide to write and I know that in order to write I have to turn on my computer, but I don't intend to turn on my computer. Intuitively, I'm irrational. But why? It doesn't look like I can use either of the above explanations. It's not true that I possess decisive reason not to have the end or not to have the means-end belief. So those aren't the attitudes that are irrational. But it also seems like I don't possess decisive reasons to have the end, and hence it doesn't seem like Means-End Transmission can help explain why I'm required to take the means. Puzzle.³³

The exact same puzzle arises for views about what all the reasons require (as opposed to what the reasons one possesses require). The puzzle of mere permissibility has been discussed at great length in that context.³⁴ The puzzle of mere permissibility has proven hard to solve. There are several different candidate solutions. Getting into the details of these views would take us too far afield. However, I will briefly sketch and motivate my preferred solution to the puzzle.³⁵

It is helpful to contrast my view with another view, which we can call Intentions Provide Reasons.³⁶ According to Intentions Provide Reasons, when you intend to ϕ , the fact that you intend

³³Some think that a similar puzzle arises in the case of belief. It should be noted that although it's uncontroversial that there are cases of mere permissibility when it comes to practical matters, it's very controversial whether there are cases of mere permissibility in the epistemic case. Many explicitly reject that such cases are possible. See, e.g., Kolodny (2007a), White (2005), Feldman (2003, 2007). Some defenders of the claim that there are cases of epistemic mere permissibility include Harman (1986), Kelly (2010), Rosen (2001), Ballantyne & Coffman (2011), Schroeder (FC). Although I think it's plausible that there are cases of mere permissibility in the epistemic case, the details would take us too far afield here. I don't think it makes a difference to my explanation of Reasons-Coherence, Closure.

³⁴See, e.g., Broome (2001), Brunero (2007), Schroeder (2009b), Scanlon (2004), Kolodny (2007a, MS), Way (2010a, 2011).

³⁵This is fleshed out more fully in Lord (MSd).

³⁶Schroeder (2009b) tentatively defends a solution to the mere permissibility problem by appealing to Intentions Provide Reasons. I think that Way (FC) raises additional worries to the ones I'll focus on here, some of which I think are damning.

to ϕ is a reason to ϕ . It's important to note that the reason provided by one's intention need not be a very weighty reason. Thus, it won't be possible to bootstrap into a *weighty* reason to ϕ just by intending to ϕ .

However, Intentions Provide Reasons can still help with mere permissibility as long as the reasons provided by intentions have *some* weight. This is because, it seems, in cases of mere permissibility, we just need a little more weight in favor of the end one actually intends to pursue to make the reasons in favor of that end decisive. The thought is that the reason provided by one's intention can provide this extra bit of weight. And once one has decisive reason to pursue the end, we can use a principle like Means-End Transmission to explain why you are required to take the means.

There are several problems with this explanation. I will just mention the one I think most damning (this is developed nicely in Way (FC)). It is important that the reasons provided by one's intentions are usually quite light weight. This is needed to avoid the bootstrapping worry. However, if they are light weight, then it is implausible that they will be able to explain all cases of mere permissibility, for it's not true that all cases are cases where the reasons there are for the various permissible ends are *roughly equal* in weight.

In some cases the merely permissible options are incommensurable. I might have the choice to take a job as a philosopher or to take a job as a lawyer. It might be that both are permissible, but not because the reasons there are for each option are commensurable and of roughly equal weight. The reasons there are for each option might be sufficiently weighty to make it permissible to pursue that job even though we can't directly compare the weights of the reasons. Moreover, in cases in incommensurability it is well known that adding a light weight reason in favor of one side doesn't necessarily give one decisive reasons to act in that way.³⁷ To see this, imagine that the law firm, in an effort to recruit me, offers me a \$500 signing bonus. This gives me an extra reason to become a lawyer, but it doesn't seem to make the reasons in favor of becoming a lawyer decisive.³⁸ Thus, establishing that there is a light weight reason provided by one's intention to be a lawyer isn't sufficient for establishing that one has decisive reasons to become a lawyer. Without this, the Intentions Provide Reasons theorist has nothing.

Despite the fact that Intentions Provide Reasons is implausible, it is very plausible that choosing one of the ends makes a difference to which means you should pursue. Before you make a decision there are no particular means you ought to pursue. After you make a decision there are. One might think we need to abandon this idea if we abandon Intentions Provide Reasons.

Fortunately, there are at least two ways for some fact f to affect the weight of reasons in favor of ϕ -ing. One way is for f to provide a reason to ϕ or provide a reason against ϕ -ing. Thus, providing additional reasons for the end is one way in which an intention to ϕ might increase the weight of one's reasons to pursue the necessary means (via some transmission of weight principle). However, it is not the only way. For another way f can affect the weight of reasons in favor of ϕ -ing is by *intensifying* or *attenuating* the *other* reasons there are for ϕ -ing. For example, the fact that you just took a pill that makes some red walls appear blue isn't itself a reason to disbelieve that the wall that seems to you blue is blue. Instead, it just weakens or attenuates the reason provided by the appearance. Similarly, the fact that today is the last day to see a special exhibition at the local art gallery doesn't provide a further reason to go to the gallery. It *intensifies* the weight of the other reasons I have to go see the special exhibition.

³⁷Cf. Chang (1997).

³⁸See Hare (2009, 2010) for arguments that these reasons are decisive

Even though it's implausible that intentions affect the weight of our reasons to take necessary means in the first way, it is quite plausible that they affect the weight of our reasons to take the means in the second way. That is, it is quite plausible that our intentions *intensify* the weight of the reasons we already have to take the means. Moreover, I claim, when ψ -ing is necessary for ϕ -ing and one has sufficient reason to ϕ , the fact that one intends to ϕ intensifies one's reasons to ψ so that those reasons are decisive. I'll briefly argue for each claim in turn.

The argument I will provide here for the claim that intentions intensify—let's call this Intentions Intensify—will be anchored in the fact that Intentions Intensify avoids the main problems with Intentions Provide Reasons. Given the antecedent plausibility of the claim that our intentions make some difference to what we should do (at least in these cases), this will be a serious virtue of the view.

First things first, there is no threat of bootstrapping. In a case where there are no reasons to pursue the end, there won't be any reasons to pursue the means (because of the end). Thus, intending the end won't intensify the reasons to take the means because there won't be any reasons to take the means.³⁹ One might worry that one can bootstrap into sufficient or decisive reasons to take the means in cases where one lacks sufficient reasons to pursue the end. That is, it might be that the fact that you intend intensifies your reasons to take the means to a point where they are sufficient or decisive even though you lack sufficient reasons to pursue the end.

I don't think this is possible. If your reasons for pursuing the end are insufficient, then they will be defeated by some reasons. Those reasons will also defeat your reasons to take the means, even if the fact that you intend to pursue the end intensifies your reasons to take the means. Since there are defeaters of your reasons to pursue the end, the intensification will not be enough to make your reasons to take the means sufficient. Why is this? The intuitive idea is that the defeaters of the end trump the intensification. If the end is impermissible, the fact that the means are necessary can't intensify the reasons to take the means enough to make the reasons sufficient.

Of course, you might have other reasons (or there might be other intensifiers) that make your reasons to take the means sufficient or decisive. But as long as your reasons to pursue the end are insufficient, the fact that you intend the end won't intensify your reasons to take the means enough.

It avoids bootstrapping but also avoids the problem with incommensurability. The problem there, recall, was with the claim that the reason provided by the intention made one's reasons to pursue the *end* decisive. In cases of incommensurability, adding a light weight reason to one side doesn't always make the reasons in favor of that side decisive. According to Intentions Intensify, one's intentions don't affect the reasons in favor of pursuing the end. They affect the reasons to take the means. Moreover, they do this not because they add a light weight reason, but by intensifying.

There is one other virtue of the account worth mentioning here. It avoids commitment to the claim that once you intend to pursue the end, you ought to pursue the end. The explanation given by Intentions Provide Reasons is committed to this. This is because that explanation turns on one gaining decisive reasons to pursue the end once one intends to pursue the end. This is implausible because in many cases it seems permissible to drop the end after one forms an intention to pursue it.⁴⁰ Intentions Intensify avoids this problem because it holds that one's intention affects the weight of one's reasons to take the means. One's intention does nothing to the weight of the reasons to

³⁹It's also plausible that intending the end won't intensify other reasons to take the means. The fact that ψ -ing is necessary for ϕ -ing doesn't seem to be a good candidate for explaining why reasons to ψ that are completely independent of ϕ -ing are intensified.

⁴⁰This objection is pushed in Broome (2001) and Brunero (2007).

pursue the end. Thus, it might be permitted to drop one's intention after one forms it, although it is not permitted to keep one's intention but fail to take the means.

While this isn't a complete defense of Intentions Intensify, it is a good start. Showing Intentions Intensify isn't enough to fully solve the problem of mere permissibility. In order to fully solve the problem, I have to explain why this intensification always makes the reasons to ψ decisive. The short answer, which I don't have the space to fully defend, is that there being sufficient reason to ψ before you intend to ϕ guarantees that there won't be defeaters that prevent the intensification from making your reasons to ψ decisive. In order to see this, it is important to ask *why* one's reasons are not decisive before one chooses one of the ends. It is very plausible that it is because one hasn't made up one's mind.⁴¹ This is the only feature that is present across all cases of mere permissibility. The rub, of course, is that that feature won't be present after you make up your mind. Thus, the intensification will make one's reasons decisive.⁴²

I think that I can adopt an analogous view for the weights of the reasons-possessed. This will hold that the fact that you intend to ϕ intensifies the reasons you possess to ψ such that the reasons you possess to ψ are decisive. This fills the hole in my explanation above, and it vindicates and Reasons-Coherence, Means-End.

I think that the basic structure of these two explanations will hold across the board of intuitively irrational incoherences. With the structure of my explanation in hand, we should go back to reevaluate the initial motivations for the debate between the wide- and narrow-scopers. First, though, we should pause to compare the view advocated here with so-called myth views advocated by Joseph Raz and Niko Kolodny.

2.5.1 The Myth of the Coherent

The explanation offered above of why one is always irrational when incoherent (and why one is always coherent when rational), is very similar to things that have been said by Joseph Raz and Niko Kolodny (see, e.g., Raz (2005b), Kolodny (2005, 2008a,b)). Both Raz and Kolodny are skeptical about coherence requirements of rationality—requirements that directly require coherence. Despite the similarities between my view and theirs, I think there is a very important dissimilarity that should be explicitly pointed out. I'll explain this after briefly spelling out Raz and Kolodny's views.

Both Raz and Kolodny are skeptics about coherence requirements of rationality. However, they see the appeal of the coherence requirements, and thus a big part of their debunking strategy is an explanation of why we are tempted to think that there are coherence requirements. Their explanation holds that we are tempted by coherence requirements because it is true that, whenever we comply with all the requirements of reason, we are coherent, and that whenever we are incoherent, we don't comply with the requirements of reason. Thus, the coherence facts fall out of more

⁴¹Of course, this is only given the fact that there are other acts that you also have sufficient reason to perform. That is, if there weren't any other acts you had sufficient reason to perform, the fact that you haven't made up your mind wouldn't make the reasons to take the means indecisive. However, it's important not to infer from this that what is happening in all cases of mere permissibility is that the reasons there are for the other permitted acts *defeat* the decisiveness of the reasons to perform a particular permitted act. This won't be generally true since some cases will be cases of incommensurability. Thus, I think the only thing that will be generally true is that you haven't made up your mind yet.

⁴²Of course, other defeaters might come about after one forms the intention. But those, I claim, will also defeat the sufficiency of one's reasons to pursue the end.

basic requirements of reason. Since all we wanted to account for were the coherence facts, there is no reason to posit requirements that directly require coherence.

This is obviously very similar to my view. However, it must be stressed that for Raz and Kolodny, the requirements of reason are completely independent of *rationality*. The requirements of reason, on their view, are determined by *all of the reasons*. They assume that these requirements are the one's that determine what we ought to do and believe. They think that rationality is about coherence. They thus become skeptics about rationality by being skeptics about coherence.

A skeptic about rationality I am not. Indeed, the main aim of this chapter is to show that Determination can plausibly account for the data important to recent discussions of rationality without incurring the most damning problems for rival accounts. A big part of my approach is a rejection of the type of coherentism advocated by wide-scopers. In this way I am like Raz and Kolodny. Moreover, like Raz and Kolodny, I think that there are more basic requirements that explain the coherence facts. Finally, like Raz and Kolodny, I think that these more basic requirements are a function of normative reasons. However, and this is the rub, unlike Raz and Kolodny, I think that these more basic requirements *just are* the rational requirements. Raz and Kolodny's mistake is that they have too narrow a conception of what the options are. We can embrace their skepticism about coherence requirements without embracing skepticism about rationality.

Finally, it is worth pointing out that the explanation given in the last section *doesn't* immediately entail that the wide-scope requirements are false. In fact, if we hold a certain kind of deontic logic (standard deontic logic happens to be a member of this kind), then it turns out that requirements like Action Possess will *entail* the wide-scope requirements. Let me quickly explain.

The relevant principle holds that if you are required to ϕ and ϕ entails ψ , then you are required to ψ —call this Inheritance. If Inheritance is true, then the requirements appealed to in the previous section will entail the wide-scope requirements.⁴³ To see this, just note that in all the worlds where one complies with all the requirements appealed to above, the conditionals that the wide-scope requirements scope over are true. This fact, together with Inheritance, entails the wide-scope requirements—i.e., entails we are required to make those conditionals true. This shows two things. First, Raz and Kolodny must reject Inheritance in order for their hostility to the truth of wide-scope requirements to make any sense at all.⁴⁴ Second, and more importantly, my explanation of the fact that if you are rational, then you are coherent isn't hostile to the truth of the wide-scope requirements. Indeed, it might be positively friendly to the truth of the wide-scope requirements.

If I were to embrace Inheritance and thus embrace the thought that requirements like Intention Possess entail the wide-scope requirements, then I can answer a popular objection to my views. This objection starts by asking us to imagine someone that is incoherent—means-end incoherent, say. Suppose it is our monarch loving fool. The thought behind the objection is that he is making not one but two rational mistakes. He is irrational for having that end and having that means-end belief, but he is *also* irrational for being incoherent. If he were to form the intention to become the King of France, he would be rectifying one of his rational mistakes. He'd be more rational than he was before. This doesn't seem compatible with my views because they seem to predict that his only rational mistakes are having the end and the means-end belief.

⁴³Just to be clear, I'm not endorsing Inheritance. It leads to lots of problems. I mention it here because I think it is important to notice that the so-called myth view requires supplementation if it is going to help explain why the wide-scope requirements are false.

⁴⁴I draw out this point more in §7.2 in chapter 7.

However, as we've seen, myth views aren't incompatible with the truth of wide-scope requirements. In fact, the requirements that myth theorists hold dear to their hearts might actually entail the wide-scope requirements. If they do, then the intuition that drives the objection can be captured. It's true that our poor fool makes three mistakes and it's true that he is more rational after he forms the intention. However, the myth theorists adds, he is still irrational for having the end and the means-end belief. So even if he can comply with *a* requirement of rationality by forming the intention, he isn't complying with all of the requirements.

2.6 Back to the Beginning

Recall what the debate between the wide- and narrow-scooper is about. In the first place, it is about what the best explanation is of why you are irrational when you are incoherent. That is to say, it is about what the best explanation is of the fact that it is necessary to be coherent in order to be rational.

Narrow-scoopers think that being coherent is necessary, but not sufficient. However, as the wide-scoopers are keen to point out, the narrow-scope view seems to deny sufficiency in an implausible way. It doesn't seem like you are always rationally required to intend to ψ when you intend to ϕ and believe that in order to ϕ you must ψ . Sometimes you are rationally required to give up one of the antecedent attitudes instead.

Wide-scope requirements are compatible with saying that sometimes you should give up the antecedent attitudes. The wide-scope view—the view that holds that all the requirements are wide-scope—holds that the narrow-scope view is implausible because, in fact, being coherent is both necessary and *sufficient* for being rational. The failures of the narrow-scope view directly motivate the claim that being coherent is both necessary and sufficient for being rational.

This itself is implausible for a variety of reasons. Most fundamentally, it's implausible because it seems obvious that sometimes rationality requires us to hold particular attitudes or perform particular actions. Without some alternative account of why being coherent is insufficient, perhaps we should throw up our hands, be wide-scoopers, and call it a day.

This is no longer our epistemic position. For I've offered another account of why being coherent is insufficient. It's because what's rational is determined by the reasons you possess, and the reasons you possess sometimes require you to hold particular attitudes or perform particular actions. Moreover, my view lacks the problems that the traditional narrow-scope view has. My view not only is compatible with the claim that sometimes you are rationally required to give up the antecedent attitudes, it can give a principled explanation of what needs to obtain for this to be the case.

My view can also explain the initial data, which is that being coherent is necessary for being rational. This is because it's plausible that when you are incoherent you are not correctly responding to some of the reasons you have. All in all, I think the considerations adduced here provide strong support for my third way in the debate about the nature of rationality.

Chapter 3

What it is to Possess a Reason

3.1 Introduction

THERE are reasons for believing things, intending things, and doing things. For example, the fact that the baby is crying is a reason to believe she is hungry, the fact that the 2014 World Cup is in Brazil is a reason to intend to be in Brazil in 2014, and the fact that the light is green is a reason to drive through the light. There are also reasons for believing things, intending things, and doing things that *I have or possess*. For example, I possess a reason for believing the baby is hungry because I *hear* her crying, I possess a reason for intending to be in Brazil in 2014 because ESPN *informed me* that the 2014 World Cup is in Brazil, and I possess a reason to go through the light because I *see* that the light is green.¹ I must have some type of access to the facts that are the reasons that I possess, even though I don't need to have access to facts that are merely reasons for me to believe, intend, and do things.

It's very plausible that only the reasons you possess bear on the rationality, justification, or reasonableness of your actions and attitudes. If r is a reason to ϕ that you don't possess, r can't contribute to the rationality of ϕ -ing. If I have no access to the fact that the baby is crying but I believe she is hungry, my belief is not made rational, if it's rational at all, by the fact that the baby is crying. Conversely, it's plausible that if I possess some reason, r , to ϕ , then r affects the rationality of ϕ -ing. The reasons you possess seem to be *difference makers* when it comes to what it's rational for you to believe, intend, and do.

Even if we take it for granted that only the reasons we possess contribute to the determination of which attitudes and actions are rational, it is still an open question *what it is to possess a reason*. I will offer an answer to this question in this chapter. In the end, I will argue that what it is to possess a reason r to ϕ is to be in a position to know r and to be disposed to treat r as a reason to ϕ .

The plan is this. In §3.2 I will introduce some desiderata a successful view must meet. In §3.3 I will introduce a taxonomy for the various views about which epistemic condition must be met in order to possess a reason. There will be three different divisions between the views. I will argue against my opponents' views by arguing for one side of each division. In the end, only one view will be left—viz., that you possess r only if you are in a position to know r . In §3.4 I will then argue for

¹The fact that all of the epistemic relations I cite here entail that I know the facts in question is no accident. As we'll see, everyone agrees that I meet the epistemic condition on possessing reasons when I know the reason. It's very controversial whether knowing the reason is also necessary.

a view about what it is to be disposed to treat something as a reason. When you combine my view about the epistemic relation with my view about what it is to be disposed to treat something as a reason, you will get an analysis of what it is to possess a reason.

3.2 A Working Hypothesis and Two Desiderata

In order to focus discussion, it's important to be clear about what theories of possession have to account for. This is the only way to assess the virtues and vices of the various theories. As we saw in the introduction, it's plausible that there are at least two tight connections between the reasons you possess and the attitudes and actions that are rational. On the one hand, it's plausible that if some reason r to ϕ contributes to the rationality of ϕ -ing, then you possess r . If I'm irrational for not taking into account some reason to ϕ , then I better have some access to that reason. And, on the other hand, it's plausible that if I possess r to ϕ , then r makes a difference to the rational merits of ϕ -ing. Of course, this doesn't mean that possessing r always *changes* the rational merits of ϕ -ing. It's just that it has an effect on the rational merits—i.e., the merits of ϕ -ing from the point of view of rationality—of ϕ -ing.² I propose we accept this as a working hypothesis.

Working Hypothesis: You possess reason r to ϕ iff r affects the rational merits of ϕ -ing.

Working Hypothesis provides two ways to test the plausibility of views of what it is to possess a reason. First, it's a major strike against a theory if it's possible to satisfy the conditions that the theory proposes for possessing r to ϕ even though, intuitively, r doesn't affect the rational merits of ϕ -ing. Second, it's a major strike against a theory if it's possible for r to intuitively affect the rational merits of ϕ -ing even if you don't satisfy the conditions the theory proposes for possessing a reason. These will be the main tests we will use to evaluate theories.

We can refine our focus even more by cataloging two desiderata that a successful theory must account for. First, it's clear that in order to possess a reason, you have to meet some epistemic condition. If r is completely outside your epistemic ken, then you don't possess r . Which relation is involved is the primary topic of debate in the literature on possession. There is a wide-range of options, and views diverge along several dimensions. Nevertheless, everyone agrees that in order to possess a reason, you have to stand in some epistemic relation to the reason.

Almost all focus has been placed on the right epistemic condition. This is a shame, because it seems equally clear that meeting the epistemic condition, whatever it is, is not sufficient for possessing r .³ This is because it's common for one to stand in the correct epistemic relation to some proposition, p , that is in fact a reason to ϕ even though, intuitively, p does not affect the rational merits of ϕ -ing. Here's one case that shows this:⁴

Lois' Fish

Lois has ordered fish for supper. She finds out that the fish contains salmonella. However, she reasonably believes that salmonella is one of the many types of bacteria found in food that is harmless to eat. She thus intends to eat the fish.

²It does this, I think, by contributing to the weight of the set of possessed reasons to ϕ .

³I myself make this mistake in my Lord (2010). I was convinced I was wrong to do this by John Broome and Kurt Sylvan.

⁴This case comes from Broome (2007a). It was discussed extensively in chapter 2.

The fact that there is salmonella in Lois' fish is an excellent reason for her not to intend to eat the fish. Moreover, Lois knows the proposition that constitutes this reason. Thus, by anyone's lights she meets the epistemic condition. However, given Lois' false but reasonable belief that salmonella isn't harmful, she doesn't *treat* that fact as a reason to refrain from intending to eat the fish. It seems like this makes an important difference as to whether Lois possesses that reason. I think that since she doesn't treat that fact as a reason, it doesn't affect the rational merits of intending to eat the fish. And thus, I think she doesn't possess that reason for not intending to eat the fish even though she meets the epistemic condition.

Cases like this point to a second desideratum. The correct theory will also have to some type of *treating* condition. In addition to meeting the epistemic condition, it has to be that you meet some type of treating condition.

The plan now is this. First, we'll tackle the epistemic condition. I will argue by elimination for the view that you meet the epistemic condition just in case you are in a position to know the proposition that constitutes the reason. I will then turn to the treating condition. I will argue you meet the treating condition just in case you are disposed to treat the reason to ϕ as a reason to ϕ .

3.3 The Epistemic Condition

3.3.1 A Taxonomy

Recently there has been a vigorous debate on what the epistemic condition needs to be in order to possess a reason or piece of evidence. Many views have been defended. I will present the views by explicating three key divisions that separate the various views. Each view takes a stand on each division.

High Bar v. Low Bar

The first division is between High Bar and Low Bar views of the epistemic condition.⁵ High Bar views hold that in order to possess a reason, you have to stand in some positive epistemic relation with the reason. Two popular High Bar views hold, respectively, that you possess r only if you know r and you possess r only if you justifiably believe r . The High Bars are knowledge and justification. To make reference easier, I will call the first view $P \rightarrow K$ and the second view $P \rightarrow JB$.⁶ Other possible high bar views will be introduced below. For now it's only important to know what it is to be a high bar view.

Low Bar views, on the other hand, place a low bar on the epistemic condition. They don't hold that you have to necessarily stand in some positive epistemic relation to r in order to meet the epistemic condition for possessing r . The most obvious low bar view holds that you possess r only if you believe r . For reasons that Schroeder (2008, 2011) makes vivid, the best low bar view seems to be that you possess r only if you have a presentational state with r as its content. Presentational states are states that represent their contents as being true. Schroeder holds that not

⁵This terminology comes from Schroeder (2011), but the debate goes back to at least Feldman (1988).

⁶In general, I'll use ' $P \rightarrow x$ ' to refer to views that hold that you possess r only if you stand in the x relation to r . Those in the literature tend to use '=' instead of ' \rightarrow '. This is because they have apparently thought that standing in the epistemic relation is both necessary and sufficient for possessing a reason. As we've seen, there are compelling reasons to think this is false (in the last section I used Lois' Fish to illustrate this). Hence my use of ' \rightarrow '.

only are beliefs presentational states, but perceptual states are as well. We'll call these low bar views $P \rightarrow B$ and $P \rightarrow \text{Pres}$, respectively.

Holding v. Non-Holding

A second division has to do with whether the epistemic relation essentially involves having some mental state with the reason as the content. Theorists who are committed to Holding views think that the epistemic relation essentially involves having some mental state with the reason as the content. All of the views we've looked at so far, both High Bar and Low Bar, are holding views.

There are possible non-holding views, though. Consider the familiar distinction between doxastic and propositional justification. You are propositionally justified to believe that p just in case p is supported by sufficient justifiers. However, this isn't enough to have a doxastically justified belief that p . In order to have a doxastically justified belief, there must be an appropriate connection between your belief and the justifiers of the belief. The important point here is that p can be propositionally justified for you even if you don't believe it.⁷ So one possible (high bar) non-holding view is that you possess r only if r is propositionally justified for you— $P \rightarrow PJ$.

There is an analogous distinction to be made with knowledge. Just as you can have propositional justification for p without doxastic justification for p , you can be in a position to know p without knowing p .⁸ This opens up the door to another non-holding view. Namely the view that holds that you possess r only if you are in a position to know r — $P \rightarrow PTK$.

Factive v. Non-Factive

The final division in our taxonomy is between factive views and non-factive views. Factive views hold that the epistemic relation you must bear to a reason is factive. An epistemic relation is factive just in case you can only stand in that relation to p when p is true. Non-factive views hold that you can stand in the epistemic relation necessary for possession to propositions that are false.

So far, only $P \rightarrow K$ and $P \rightarrow PTK$ are factive views (given widely held assumptions about justification). $P \rightarrow B$, $P \rightarrow \text{Pres}$, $P \rightarrow JB$, and $P \rightarrow PJ$ are all non-factive views. However, you can still think justification is as high as the bar goes even though one can only possess truths as reasons. For you could hold that the epistemic relation is justified true belief, or propositional justification in a truth— $P \rightarrow TJB$ and $P \rightarrow TPJ$.⁹

The following table represents our taxonomy.

⁷Consider this trivial example. I have propositional justification that I'm typing right now. *A fortiori*, I have propositional justification to believe that I'm typing right now or I'm Elvis. At least before I thought of the example, I didn't believe the disjunction even though I was propositionally justified in believing it.

⁸The example in n. 7 can equally well illustrate this distinction.

⁹Another route, one that avoids some of the objections I'll press below, is to hold that justification itself is factive. This is argued for by Unger (1979), Littlejohn (2009, 2012), Sutton (2007).

Table 3.1: Epistemic Condition Taxonomy

View	High Bar v. Low Bar	Holding v. Non-Holding	Factive v. Non-Factive
$P \rightarrow B$	Low Bar	Holding	Non-Factive
$P \rightarrow \text{Pres}$	Low Bar	Holding	Non-Factive
$P \rightarrow JB$	High Bar	Holding	Non-Factive
$P \rightarrow PJ$	High Bar	Non-Holding	Non-Factive
$P \rightarrow TJB$	High Bar	Holding	Factive
$P \rightarrow TPJ$	High Bar	Non-Holding	Factive
$P \rightarrow K$	High Bar	Holding	Factive
$P \rightarrow PTK$	High Bar	Non-Holding	Factive

I will argue that the correct view has to be a High Bar, Non-Holding, Factive account by arguing first against Holding views, then against Low Bar views, and then against Non-Factive views. Since there are two High Bar, Non-Holding, Factive views— $P \rightarrow TPJ$ and $P \rightarrow PTK$ —I will end by arguing for $P \rightarrow PTK$ over $P \rightarrow TPJ$.

3.3.2 Against Holding Views

Holding views of the epistemic condition all share a commitment to the view that the epistemic relation being in a mental state that has the reason as its content. I think that Holding views are too restrictive. Consider the following case.

Out of the Ordinary

Each morning I casually peruse a magazine while I eat my breakfast. I am doing this on Monday morning. As I'm doing this, my wife tells me that she has an unusual schedule that day. Given her unusual schedule, I need to pick up my son from school (this is something she almost always does because we have agreed that that is one of her daily tasks). She speaks clearly when she tells me these facts, and she is a mere 10 feet from me. However, I am engrossed just enough in the Newsweek I'm reading to not process the information. Given the fact that picking up my son is not one of my usual tasks, I believe that I don't need to pick him up that day. My son doesn't get picked up on time.

Focus on the rationality of my belief that I don't need to pick up my son that day immediately after my wife tells me this information. Intuitively, this belief is not rational. After all, my wife just told me some very pertinent information that decisively supports thinking I need to pick him up. However, I don't believe that my wife can't pick him up. Thus, according to Holding views, I don't possess that reason. By Working Hypothesis, then, that reason shouldn't affect the rational merits of my belief that I don't need to pick him up. Moreover, it seems that all the other reasons I possess sufficiently support the proposition that I don't need to pick him up. Holding views seem to entail, then, that my belief that my schedule is not unusual is rational. My wife would disagree.

Non-holding views can easily account for why I'm not rational. It's because I possess the reasons provided by my wife's testimony. $P \rightarrow PJ$ (and $P \rightarrow TPJ$) holds that I possess those reasons because those propositions are propositionally justified for me. $P \rightarrow PTK$ holds that I possess those

reasons because I'm in a position to know them. On non-holding views, inattentiveness does not get one off the hook for taking into account reasons that one has what we'll call *broad access* to. For now I won't analyze what broad access comes to. We've seen two proposals—viz., being propositionally justified and being in a position to know. What's important right now is that it's intuitive that reasons you merely have broad access to affect the rational status of your beliefs (and other attitudes and actions). Non-holding views can account for this, holding views cannot.

3.3.3 Against Low Bar Views

I have not given you enough information yet to evaluate the plausibility of low bar views. For I have not given you a low bar analysis of what the necessary and sufficient conditions are for something to be a reason you possess. It's certainly false that all the contents of all your beliefs or presentational states are reasons you possess for everything whatsoever. That is to say, I've told you the epistemic condition that must be met, but I haven't told you what condition must be met to make the content a possessed reason.

While there hasn't been a tremendous amount of discussion on this score, it's clear what the intuitive idea is that low bar theorists are attracted to.¹⁰ The basic idea is that the content p of one your presentational states is a reason to ϕ only if p is an objective reason to ϕ if p is true.¹¹ Given Working Hypothesis, low bar views thus hold that the contents of all of your presentational attitudes that are objective reasons if true affect the rationality of your attitudes and actions.

This prediction seems obviously false. Let's just think about the epistemic case. Suppose r is a reason to believe p . Plausibly, what it means for r to affect the rationality of believing p , in this case, is for r to put you in a better evidential position with respect to p when you possess r . But it's incredibly plausible—indeed, some call it a truism¹²—that you aren't put in a better evidential position with respect to p when you *unjustifiably* believe r . So if Working Hypothesis is true, then low bar accounts are in trouble. Since Working Hypothesis is intuitively plausible, low bar accounts are in trouble from the beginning. Moreover, I think that Working Hypothesis can be bolstered by thinking about the relationship between the objective weight of reasons—i.e., the weight reasons have when weighed against all of the reasons—and the subjective weight of reasons—i.e., the weight reasons have when weighed against all of the reasons that are had.

The intuitive idea driving the argument is something like this: Reasons for belief are the types of things that rationalize beliefs. The point of coming to possess reasons is to come to be in a better position with respect to what you're rational in believing. Minimally, gaining p is to come to be in a

¹⁰See Schroeder (2007, 2009b), Way (2009, 2010a, 2012), Vogelstein (FC), Parfit (2011).

¹¹There is some disagreement about whether the analysis should be in terms of counterfactuals or some other type of conditional. All extant views, with the exception of Vogelstein's, at least sometimes cash things out in terms of counterfactuals. Vogelstein explicitly rejects this, but for dubious reasons (his strategy is to first provide an analysis of what needs to obtain in order for some subjective reason to exist, and then provides an analysis of what it is to be a subjective reason. His objection against counterfactual views comes at the first stage—he provides an objection to of a counterfactual analysis of what needs to obtain in order for some subjective reason to exist. But this objection doesn't hold any water against the best counterfactual analysis of what it is to be a reason and the existential analysis that results from this analysis). Nonetheless, I strongly suspect that counterfactual analyses are false for general reasons. In personal communication Schroeder has expressed serious doubts about whether a counterfactual analysis can work.

¹²For example, Feldman (1988), Schroeder (2011).

better position with respect to p itself. This is a fundamental part of the role that possessed reasons play.

Here's the argument:

- (1) If p is an objectively sufficient reason to believe p and you possess p , then you possess a subjectively sufficient reason to believe p .
 - (2) If p , then p is an objectively sufficient reason to believe p .
 - (3) If p , then if you possess p , then you possess a subjectively sufficient reason to believe p .
 - (4) If you possess a subjectively sufficient reason to believe p , then you are propositionally justified in believing p .
- (C) If p , then if you possess p , then you are propositionally justified in believing p .

If this argument is sound, then low bar views are false. For suppose p . Suppose that you believe p , but p is not propositionally justified. The low bar accounts hold that you possess p . But it follows from the contraposition of the conclusion of the above argument that if you are not justified in believing p when p is the case, then you don't possess p .

Recall the basic idea behind the argument is that a fundamental role that possessed reasons play is putting one in a better position with respect to what it's rational to believe. The argument extends and precisifies this basic idea. It does this by making claims about the connection between the subjective weight of certain reasons that are possessed and the objective weight of those reasons. The objective weight of some reason r is just the weight it has when weighed against all the reasons. Few reasons always have sufficient objective weight. But a special class of reasons for belief do. Namely, all of the truths. For any true proposition p , p is an objectively sufficient reason to believe p . This is not to say that there aren't objective reasons to not believe truths. It is to say, though, that the objective reasons to believe falsehoods are always objectively insufficient. This seems correct. If you were omniscient, you would only possess sufficient reason to believe the truths. You would never possess sufficient reason to believe something false, even though you would possess lots of reasons to believe false things (and to withhold belief in true things).

So when p is true, p will be an objectively sufficient reason to believe p . If this is right, then it seems very plausible that if you were to come to possess p , you would gain a sufficient reason to believe p . It's hard to see what good comes from gaining p as a reason if you can't even be justified in believing p !

You might not be convinced by this because you might be suspicious of the claim that p is a sufficient reason to believe p —i.e., you might doubt (2). After all, you can never become justified in believing p by *inferring* p from p . But why not if p is a sufficient reason to believe p ? The answer is that in order for the output of a pattern of inference to be justified via the inference, one must be justified in believing the inputs. But if you are already justified in believing p , then certainly no new justification will be *transmitted* via the inference. And if you're not already justified in believing p , then justification won't be transmitted via inference because there is no justification to transmit in the first place. But this very plausible story is compatible with p being sufficient reason to believe p . It just shows that you can't get a justified belief that p via an inference from p .

You still might not be satisfied with this. For suppose p , suppose you possess p and suppose the argument above is sound. Thus, it follows that p is propositionally rational for you. What are

your sufficient reasons for believing p ? Suppose the answer to this is simply p . If this is right, then it's very plausible that you can come to believe p because of p . This follows from the plausible principle that if the members of some set of reasons you possess S sufficiently support p , then you can come to believe p because of the members of S . If this is right, then in this type of case it must be that you can believe p because of p . Some find this to be implausible.

This is an interesting challenge that I can't fully deal with here without further resources that would take us too far afield. I talk about these issues extensively in chapters 5 and 6 using resources defended in chapter 4. I'll say a few words here about how I think we should meet this challenge and how this relates to the current dialectic. I think the answer is that you can base your belief that p on p , even though you can't do this via inference. Such basing is always uninferential.¹³ In the paradigm case, you come to believe p because of p when you see p . Roughly, you base your belief on p not because you infer p from p , but because you come into perceptual contact with the fact that p , which stands in the right causal relationship with your belief to become the basis of the belief. If this is right, then you can base your belief that p on p , even though this is never inferential.¹⁴

So, roughly, this is how I see the terrain. p is a sufficient reason to believe p . You can never be justified in believing p because you infer p from p . This isn't particularly problematic because we have a principled explanation for why not. If you are going to be justified in believing p because you infer it from something, you'll have to be justified in believing the things you infer it from.¹⁵ But you can still form a rational belief that p because of p . This happens in cases of uninferential belief forming processes like perception. As I see it, the main burden of the current challenge is to motivate (i) that p is a sufficient objective reason to believe p (when p) and (ii) that when you possess p you can sometimes base a belief that p on p . I believe I have met this burden.

The final comment I'll make about this challenge is about our dialectical situation. I think that low bar theorists only really gain an advantage over high bar views if they agree with the essentials of the view about basing I just briefly sketched. The best low bar view, $P \rightarrow \text{Pres}$, gains an advantage over high bar views precisely because it holds that you possess p whenever you perceive p . And this allows a straightforward explanation of why it is that you possess the same reasons for belief in both the case of veridical perception and the case of non-veridical perception.¹⁶ But this works

¹³Uninferential here just means not-inferential. You might not like this terminology (I don't). I use this terminology in order to avoid confusions that inevitably arise when you use the more common non-inferential. I get into this particular issue in chapter 5, section 5.2.

¹⁴I think there is uninferential basing even in cases where you don't come into the kind of contact with the fact that p that you do when you see p . To give the most extreme case, I think you can believe p on the basis of p even when your belief was formed inferentially. This is because I think that what it is to base a belief that p on some reason r is for r to sustain the belief. What it is for r to sustain A 's belief that p is for A to be disposed to drop the belief if A loses r . I think you can be disposed to drop your belief that p if you lose the reason constituted by p . And when you are, you base your belief that p on p (at least partially). I think this is actually a virtue of the view, for it allows one to explain certain otherwise troubling cases of memorial knowledge. I discuss this virtue in chapter 5. This view of basing is defended in chapter 4, and further discussion of inferential v. uninferential belief forming processes can be found throughout chapters 5 and 6.

¹⁵As I say in note 13, it might be that *once* you justifiably infer p from some other propositions you're justified in believing, you are then in a position to uninferentially base that belief on p .

¹⁶Of course, it's not clear at all if this is that much of a virtue, or even if it's a virtue at all. (Williamson, 2000, ch. 8) argues that it's a bug that leads to skepticism. I agree even though I also think it's a virtue in some ways. But certainly it's been popular to think it's a virtue since at least 1983 when Cohen & Lehrer (1983) appeared. This is the main topic of chapter 6.

only if you think that you can come to have a justified belief that p by basing a belief that p on p in some of those cases. Moreover, it's very implausible that you can gain justification by basing a belief that p on p by inferring p from p . So my opponent will have to follow the same line I just tried to follow.¹⁷

(3)—if p , then if you possess p , then you possess subjectively sufficient reason to believe p —follows from (1) and (2). (4)—the claim that if you possess sufficient reason to believe p , then you are propositionally justified in believing p —is very plausible.¹⁸ Since I think (1) and (2) are quite plausible, I think the conclusion is quite plausible as well.

As it happens, the main defender of the low bar views, Mark Schroeder, explicitly recognizes that he has to deny Working Hypothesis (see Schroeder (2011)). So my defense is incomplete without considering what he has to say about it. I don't think anything he says should give us much pause.

Schroeder is interested in *making room* for low bar accounts by casting doubt on what he sees to be the best two routes to high bar views. His case is broken up into two parts. In the first part he considers an indirect argument for a high bar account. The argument is indirect because it runs through an auxiliary premise—viz., that positing a high bar is the best way to explain why one is not put in a better evidential position with respect to the consequences of p when one irrationally believes p . He resists this claim by arguing for an explanation of the fact that you aren't in a better evidential position with respect to p 's consequences when you irrationally believe p that is compatible with low bar views. The second part of his case involves a direct argument for a high bar view. It's direct because it appeals directly to intuitions about whether certain characters possess reasons. He argues that such intuitions are unreliable. With this précis in mind, let's turn to more of the details of each part.

The Indirect Argument

Schroeder, following Feldman (1988), thinks that it's a truism that when you're propositionally irrational in believing p , you aren't in a better evidential position with respect to the consequences of p in virtue of believing p . Feldman at least flirts with the thought that you can get from this and something like Working Hypothesis to some high bar view. Schroeder's target isn't exactly the inference to a High Bar view from something like Working Hypothesis and our truism. Rather, Schroeder is concerned with showing that positing a high bar isn't needed to *explain* the truism. Schroeder's strategy is to explain the truism in a way that's compatible with low bar views. By doing this, he shows that low bar views are compatible with the truism, and thus we shouldn't think accounting for the truism weeds us to high bar views.

Schroeder's explanation is that whenever you irrationally believe p , p is guaranteed to be *completely* and *generally* defeated. The defeat will be general because it will be irrational to treat p as a reason for anything (even though it will still be a reason you possess for lots of things). And the defeat will be complete because p will have no weight whatsoever (even though it will still be a reason). If this is right, then our truism is explained—when you believe p irrationally, you won't be in a better evidential position with respect to p 's consequences. This explanation is compatible with low bar views because it's compatible with you possessing p when you irrationally believe p . It's just that despite the fact that you possess p , it will be irrational to give p any weight whatsoever

¹⁷This is exactly how it goes in Schroeder (2011).

¹⁸I defend this view directly in chapter 4.

when it comes to anything, including the question of whether p . Thus, Schroeder denies that if you possess a reason r to ϕ , r affects the rational merits of ϕ -ing, which is the left-to-right direction of Working Hypothesis.

I'm not that interested in explaining the truism. In fact, I suspect Schroeder's explanation is correct. But it also doesn't speak against Working Hypothesis at all. This is because, in giving his explanation, Schroeder doesn't have to make any claims about the relationship between the reasons you possess and your evidential position. So while his explanation is compatible with low bar views, it is also compatible with high bar views.

While it's obviously necessary for low bar theorists to offer some explanation of the truism that's compatible with their view, it is very, very far from being a sufficient defense of their view. Working Hypothesis still seems like an incredibly plausible hypothesis about the connection between the reasons you possess and rationality. It's hard to see at this point why we should sever the connection between possessing p to believe q and p affecting the rationality of believing q .

Think of it this way: The claim that whenever you irrationally believe p , p will be completely and generally defeated doesn't speak to the relationship between irrational beliefs and possession. It is compatible with this explanation of the truism that you do in fact possess p when you believe it irrationally. But mere compatibility is not good reason to sever possession from a high bar. At best Schroeder's argument shows that a reason to wed possession with a high bar does not obtain.

The Direct Argument

Schroeder anticipates this, and thus tries to argue that we shouldn't trust our intuitions about the relationship between possession and a high bar. In particular, he thinks we shouldn't trust our intuitions about negative existential reason claims. He thinks we shouldn't trust our intuitions about these claims because they are unreliable.

His favorite illustrative example of this is Lehrer & Paxson (1969)'s Tom Grabit case. In that case, you see someone who looks just like Tom Grabit steal a library book. The fact that you saw this is a reason to believe Tom stole the book. However, Tom happens to have a twin brother, Tim, who is a kleptomaniac. This fact is, intuitively, a defeater of the first reason. The orthodox opinion about this case is that the fact that Tom has a kleptomaniac for a brother makes the fact that you saw someone that looks just like Tom steal a book no reason at all to believe Tom stole the book.¹⁹ Schroeder rightfully points out this is a mistake. For imagine that Tom and Tim have another identical sibling, Tam, who is also a kleptomaniac. In that case, it seems you have even less reason to think Tom stole the book. So you must have some reason in the case where it's just Tom and Tim.²⁰

Once again, I agree with Schroeder. I think we should be wary of our negative existential intuitions about reasons in particular cases. We are quite likely to get it wrong for roughly the reasons Schroeder elucidates. However, I think that (at least in the larger dialectic) Schroeder is misapplying these lessons here.

¹⁹This opinion is seen as orthodox, I think, because of the tremendous influence of Pollock & Cruz (1999)'s discussion of *undercutting defeat*. They define undercutting defeaters as defeaters that make it the case that the original reason is not a reason at all once defeated. The Tom Grabit example is supposed to be a paradigm of this kind of defeat.

²⁰Schroeder also offers an debunking explanation of why we are unreliable. I don't go into that here because I agree with it and it's irrelevant to the point I want to make.

The rub is this: The current debate doesn't turn on the intuitions Schroeder shows are problematic. Our task is *not* to evaluate whether there is a reason in some particular case. We *already know* that Schroeder's low bar view predicts there will be some cases where you will possess p as a reason to believe q even though p has *no weight whatsoever*. This itself is implausible.²¹ We needn't worry about our unreliability at picking which cases are cases of this type. This is what our unreliable negative existential intuitions are trying to track. They're trying to track when it is that something has no weight whatsoever. We aren't very good at this, as the Grabit examples show. Let Schroeder pick the cases, then. We can then ask whether, in those cases, you possess p to believe q even though p has no weight whatsoever. It doesn't seem like you do. It seems like, if we have any grip on possessing reasons at all, you can possess p only if it affects the rationality of believing p 's consequences.²²

For these reasons I am not moved to abandon a high bar view because of what Schroeder says. He doesn't attack Working Hypothesis. He merely argues for an explanation of the truism that is compatible with denying Working Hypothesis and argues that our intuitions about when there aren't reasons are unreliable when applied to particular cases. I agree with the main claims of both parts of Schroeder's argument. The catch is that they're neither here nor there.²³

3.3.4 Against Non-Factive Views

Non-factive views hold that one can stand in the epistemic relation with propositions that are false. Non-factive views fail because they cannot account for the fact that if you possess some reason r to ϕ , then you can ϕ because of r .

Let me be clear that we are interested in a particular kind of *because* explanation. As we'll see, there are all kinds of *because* explanations. The type we are interested in here is what I call a *normative because* explanation. The best way to glob onto this type is by example. I believe the heat is on in my office. It's also warm in my office and cold outside. There is a tight relationship between my belief that the heat is on and these other facts. I believe the heat is on in my office *because* it's warm in my office and cold outside. The *because* claim is true partly because I correctly inferred that the heat

²¹It is also worth noting that it seems like this commits Schroeder to denying standard conditional probabilistic accounts of evidence. This is because, for Schroeder, reasons for belief are always evidence. But if we can possess a reason r to believe p that has no weight whatsoever, then it's implausible that the conditional probability of p given r will be greater than the probability of p on its own. After all, if it were greater, then it would be plausible that the reason would have some weight. Thanks to Tom Kelly for pushing me to mention this.

²²Here's another way to put the point for those who are familiar with the details of Schroeder's paper. He tries to show that if you don't argue for a high bar view with an indirect argument through the truism (which he shows won't work because his explanation of the truism is compatible with low bar accounts), then the only way to argue for it is through appeal to negative reason existential intuitions. This conditional is false. The Working Hypothesis gives us an independent way to test the theories. Moreover, its plausibility doesn't rest on negative reason existential intuitions. Negative reason existential intuitions are thus neither here nor there.

²³It must be noted that Schroeder makes explicit in both Schroeder (2008, 2011) that his main motivation for going low bar is to account for some vexed data in the epistemology of perception. Specifically, he holds that low bar views can give the most elegant explanation of why both the deceived and the non-deceived have justified perceptual beliefs. I agree that his low bar view provides a nice explanation. Importantly for our purposes here, Schroeder at least implicitly recognizes the intuitive power of Working Hypothesis because he sees it as a major obstacle for providing his explanation. Despite the elegance of his view about the epistemology of perception, I am unmoved by what he has actually says with respect to Working Hypothesis. I won't go into the issues in epistemology of perception because it would take us too far afield. Much of chapter 6 is about this issue.

is on in my office from the claims that it's warm in my office and cold outside. Similarly, I intend to exercise 3 times a week. It's also true that exercising 3 days a week is very good for my health. There is a tight relationship between my intending and the fact about my health. Indeed, I intend to exercise 3 times a week *because* it is very good for my health. Although it's hard to say precisely when a consideration is an explanans in a normative because explanation, the phenomenon is clear enough.²⁴ Here is an argument for why non-factive views cannot account for the tight connection between possessing a reason r to ϕ and being able to ϕ because of r :

The Argument from Acting for Reasons

- (1) If you possess reason r to ϕ constituted by some proposition p , then you can ϕ because of p .
 - (2) If you can ϕ because of p , then p .
- C. Thus, if you possess reason r to ϕ constituted by some proposition p , then p .

(1) is intuitively plausible, and for the moment I'll take it for granted. This is because we are assuming that the reasons you possess are the reasons that determine which attitudes and actions are rational. If we assume this and decide that in some particular case that ϕ -ing is rational in virtue of some reason you possess r , then r can rationalize you in ϕ -ing. But if r rationalizes your ϕ -ing, then it seems that you can base your ϕ -ing on r . It's plausible that when you base your action on r when you ϕ , you ϕ (at least partially) because of r .²⁵ I'll break the following discussion up into two subsections. I'll first defend (2), given (1). I'll then come to defend (1).

For (2)

I'll take it for now, then, that if you possess r to ϕ , then you can ϕ because of r , where this is a normative because explanation. Given this, (2) is where the action is at. (2) entails that 'x because y' statements are factive. That is to say, 'x because y' can be true only if y (and x for that matter). This seems quite plausible. In general, *because* explanations are factive. Moreover, at first blush it seems like normative because explanations are also factive. Consider some cases:

Normal Colloquium

Today there is a department colloquium. It is at the usual time for such events, 4:00 pm. I notice that it is 3:55, and is thus time to walk over to the building where the colloquium is taking place. On my way I see my friend Dan. He asks me why I'm going to Robertson Hall. I tell him it's because colloquium starts at 4:00.

Unusual Colloquium

²⁴Hard, but I don't think impossible. I provide a view about in chapter 4 (see also Lord & Sylvan (MS)).

²⁵Proponents of non-factive views might deny this. They might hold, that is, that there isn't any tight connection between the fact that you acted on r when you ϕ -ed/ r rationalized your ϕ -ing and the claim that you ϕ because of r . For various reasons I think this might be the best route for the non-factive theorist to take. Actual non-factive theorists, though, do not opt for this view. We'll return to this thought below.

Today there is a department colloquium. This talk, however, is not at the usual time. It is at 4:30. I am unaware of this fact (despite receiving numerous emails that said so). I believe it is at the usual time. So at 3:55 I start walking towards Robertson Hall. I see Dan, who asks me why I'm going to Robertson. I tell him it's because colloquium is at 4:00.

My explanation of why I'm going to Robertson Hall makes sense in Normal Colloquium, but is puzzling in Unusual Colloquium. Suppose Dan reads his emails more carefully than I do, and he calls me out in Unusual Colloquium. He says, 'No. Colloquium is not at 4:00 today. It's at 4:30.' It would be odd, I think, for me not to retract my claim after Dan says this. It would be odd, that is, for me to reply like this: 'Really. Well I guess I can go back to my office. Still, I was going to Robertson because colloquium is at 4:00.' Instead, it would be much more natural for me to retract my earlier claim and say 'Well, I thought colloquium was at 4:00. Thanks for setting me straight.'

As it happens, many disagree with my diagnosis of what's going on in Unusual Colloquium.²⁶ These people think that it would be very odd for me to decide that the reason why I walked to Robertson isn't that the colloquium is at 4:00 even after I find out that the colloquium isn't at 4:00.

The issue is vexed. Those who deny (2)—let's call them non-factualists—do not deny that I would retract once Dan corrected me in Unusual Colloquium. But they think there is an independent explanation of this compatible with non-factualism. The explanation, they say, is that once I learn that the colloquium is at 4:30, I don't want to implicate that I still believe it's at 4:00. Thus, I'll weaken my report by saying things like 'I thought colloquium was at 4:00.' Why I would do this, claim the non-factualists, can be explained on basic Gricean grounds. I don't want to implicate that I think p when I think $\neg p$. Thus, I use the weaker 'I thought...' construction. This explanation of the retraction data is compatible with non-factualism.²⁷

At first blush, this doesn't seem to be enough. This is because these implications should be cancelable. Of course I do implicate that I believe that colloquium starts at 4:00 when I say 'I'm going to Robertson because colloquium starts at 4:00.' And I don't when I say 'I'm going to Robertson because I thought colloquium starts at 4:00.' But if it's really the case that I'm going to Robertson because colloquium starts at 4:00, then I should be able to say as much *while also* canceling the implication. I should be able to felicitously say, for example, 'I no longer believe that colloquium starts at 4:00, but I was going to Robertson because colloquium starts at 4:00.' To my ear this still sounds awful.

Suppose I'm right and these things are infelicitous. This is some reason to think that non-factualism is false—perhaps even a very strong reason. But, the non-factualist might contend, this isn't the end of the story. For we know that there are some true sentences that are never felicitous. Perhaps the most famous examples are Moorean absurdities. For example, it's never felicitous for me to say (nor is it ever rational for me to believe) 'I believe I'm writing, but I'm not writing.' This

²⁶There has been extensive discussion of these issues. Much of it is inspired by Dancy (2000). See, for example, Dancy (2004b), Lord (2008), Hornsby (2007), Mele (2007), Miller (2008), Millar (2004), Littlejohn (2012). The best discussion I know of takes place in the comment thread of a Pea Soup post written by Clayton Littlejohn. See <http://peasoup.typepad.com/peasoup/2010/02/thoughts-as-motivating-reasons>.

²⁷Here's Mark Schroeder in the aforementioned Pea Soup post: 'The fact that when I don't believe that p , I can't aptly describe Sue's reason for ϕ -ing by saying that her reason for ϕ -ing was that p , doesn't show that her having the motivating reason for ϕ -ing that she did doesn't consist in her bearing the had-as-her-motivating-reason relationship to the proposition that p . It just shows something about the further commitments of saying 'her motivating reason was that p .'

very well might be true, though. So the non-factualists might take recourse in the fact that Moorean absurdities are infelicitous even though they can still be true. Perhaps our normative because explanations are like that, at least in the bad cases.

Perhaps. But as Littlejohn (2012) has pointed out, this is too quick. For there is at least one test for truth that Moorean absurdities pass that our normative because explanations don't seem to pass. Namely, when it comes to Moorean absurdities, there is nothing wrong with third-personal reconstructions. There is nothing wrong with saying 'EL believes he is writing, but he's not.' However, it still seems infelicitous to say 'EL was going to Robertson because colloquium starts at 4:00, but colloquium starts at 4:30.' This, I think, makes the prospects of saving the non-factualists by drawing some analogy with Moorean absurdities much dimmer.

To bolster my factualist view, let me consider some of the standard objections. I think the core intuition driving non-factualism is a sort of internalism about the reasons we base our actions and beliefs on. The thought is that the reasons for which we act and believe supervene on our non-factive internal states. To make the pull of internalism vivid, consider my counterpart in the close possible world where today's colloquium is actually at 4:00. The internalist idea that drives non-factualism is that the mere fact that colloquium starts at 4:00 in my counterpart's world can't change whether we act on the same reason when we start walking to Robertson at 3:55. Nor does it matter that my counterpart has a true belief and I have a false belief, and likewise for the fact that my counterpart knows that it starts at 4:00 and I don't. None of these facts make a difference to which reasons I act on when I start walking towards Robertson.

One way to motivate the internalist view is to think about things first-personally. It *seems* like I have unproblematic access to the reasons for which I act and believe. When I start walking towards Robertson, it seems bizarre to think that I could be easily mistaken about why I was leaving my office, walking down the stairs, etc.²⁸ But, the thought goes, this is only true if the internalist view is true. For it will seem like I know what my reasons are whether my beliefs are true or not. So in order to hold onto the thought that we have unproblematic access to the reasons for which we act and believe, we better be internalists.²⁹

This general line of argument is often bolstered by the thought that to deny the internalist view would be embrace the thought that one doesn't act/believe for any reason in the bad cases.³⁰ And this suggests that when one is in the bad case one's actions and beliefs are simply inexplicable.³¹ This thought, I think, is what Dancy is getting at when he says that non-internalist views are 'too harsh'.³²

I think both of these results would be bad for factualists. But I don't think either result is forced upon the factualist. Let's start with the inexplicability point first. As David Lewis might say, there are explanations and there are explanations. In a particular context, my action in Unusual Colloquium is perfectly explicable by the claim that I was going to Robertson because I thought collo-

²⁸Of course, the linguistic data—especially the retraction data—seems to cut against this. I'm probably blinded by theory, but I have a hard time feeling the pull of these thoughts. In general, I think we are apt to be mistaken about all sorts of things internalists think we have unproblematic access to.

²⁹Williams (1981) famously says that the truth-value of the beliefs we reason from don't affect the structure of the explanation given. This idea runs rampant in the philosophy of action. As we'll see, I think it's true for a particular kind of action explanation, but not the kind theorists of rationality should be primarily interested in.

³⁰Some non-internalists embrace this conclusion (see especially Littlejohn (2012)).

³¹Although Littlejohn thinks you don't act for a reason in the bad case, he doesn't think your actions are completely inexplicable.

³²This was in conversation. For more, see Littlejohn (2012).

quium started at 4:00. This is a perfectly good explanation for many purposes. If the question is 'Why was EL going to Robertson?', this would often be a perfectly good answer. It is, after all, the most natural answer I would give upon being informed that colloquium starts at 4:30. This is, as far as it goes, compatible with factualism, which is a view about what it takes for something to be a reason that you potentially base your action or belief upon.

Moreover, I think there are other true, non-factive, explanations for why I'm going to Robertson. For example, it seems like it's true that my rationale for going to Robertson was that colloquium starts at 4:00. *Rationale* explanations seem non-factive. It seems felicitous to say 'EL's rationale for going to Robertson was that colloquium starts at 4:00, but he was mistaken about that.' Why is this? It is because rationale explanations are in the business of making attitudes and actions *intelligible*. And intelligibility comes cheap. Some consideration can make an attitude or action intelligible even if it is false.

The rub, though, is that rationale explanations *don't* entail normative because explanations. It might be that my rationale for walking to Robertson is that colloquium starts at 4:00; it doesn't follow that I walked to Robertson because colloquium starts at 4:00.³³ This is because intelligibility comes cheap. The acts and attitudes of even the most deluded amongst us can be made intelligible. That is, their acts and attitudes have rationales. But it is implausible that they act and hold attitudes because of certain facts. At the very least, this is a substantive issue, and thus non-factualists need to argue that the rationale explanations entail the normative because explanations.³⁴

Back to (1)

At this point it's natural for my opponent to become skeptical about (1)—i.e., the claim that if you possess r to ϕ , then you can ϕ because of r . Everyone agrees that there is some Acting On condition when it comes to reasons that are possessed. That is, everyone agrees that there is some important connection between the reasons we possess and rationally salient explanations of our actions. To put it in neutral terms, everyone agrees that there is a tight connection between the reasons we possess and the reasons *we can act on*. My view is that this amounts to a tight connection between the reasons we possess and the reasons that can provide normative because explanations. However, given that even I think there are some normative looking explanations that are non-factive, why think the connection is between the reasons we possess and normative because explanations? Why not think that having p merely entails being able to ϕ with p as your rationale? I think there are two reasons for preferring the stronger view—i.e., the view that having p entails being able to ϕ because of p .

The first reason is that it's plausible that there is also a Reverse Acting On condition. That is, it's plausible that if you meet the Acting On condition for some proposition p —on my view this is being able to ϕ because of p —then you possess p . The Reverse Acting On condition is plausible if the Acting On condition is being able to ϕ because of p , but is obviously false if the Acting On condition is having p as your rationale.

³³It should be noted that Dancy (2004a) makes the exact opposite inference. He infers that normative because explanations aren't factive because something like rationale explanations aren't factive and they entail the normative because explanations. He doesn't argue for this beyond reporting intuitions I (and others) simply don't have.

³⁴For those keeping track of the connections between this chapter and the next, I think that the rationale explanations go together with merely ϕ -ing for reasons and the normative because explanations go together with ϕ -ing for normative reasons.

If you are able to ϕ because of p , then it's very plausible that you possess p . In order for p to be able to provide a normative because explanation of your ϕ -ing, it has to be that you are appropriately related to p in some way. If p is completely outside of your epistemic ken, or if you don't even implicitly see the appropriate connection between p and ϕ -ing, then it's hard to see how p could provide a normative because explanation of your ϕ -ing.

On the other hand, it's not plausible that if you are able to ϕ with p as your rationale, then you possess p as a reason to ϕ , at least not on any view about possession on the table. Take the consequent affirmer. Let's call him Jimmy. Jimmy knows if p , then q and he's just learned q . He then infers p from q and if p , then q . It's plausible that his rationale for believing p is q and if p , q . But he doesn't thereby possess q and if p , then q as reasons to believe p . On high bar views (at least as we've understood them), some reason r to ϕ is a reason you possess only if r is an objective reason to ϕ . And obviously q and if p , then q aren't objective reasons to believe p in this case. On low bar views, some reason r to ϕ is a reason you possess only if r is an objective reason to ϕ if true. But by stipulation q and if p then q are not objective reasons to believe p if true because they are true and they aren't objective reasons to believe p . So if the Acting On condition is just being a possible rationale, then the Reverse Acting On condition is false.

That's the first reason to be skeptical that the Acting On condition is the rationale condition. The second reason is that once you adopt a non-factive Acting On condition, you've got to give up on the idea that those in favorable epistemic environments are in an importantly different situation from those in the depraved epistemic environments. In particular, you're committed to thinking that the facts that, in the favorable case, you know, or truly believe, or truly justifiably believe, aren't important when it comes to the explanation for why you are rational. For you could have possessed the relevant reasons even if you didn't stand in any of those relations with the reasons that explain why you are rational. This is because the explanation of your actions and attitudes relevant to whether you're rational are just rationale explanations. But the rationale explanations are the same in the favorable conditions and the depraved conditions, even if the normative because explanations are different.

This requires a bit of unpacking, for I am making a crucial assumption. I'm assuming that the Acting On condition is important because it tells us something important about what it is to *correctly respond to reasons*. This is because it tells us the connection between possessing reasons and a certain kind of normative explanation. Plausibly, what it is to correctly respond to a reason r to ϕ when you ϕ is for r to (at least partially) explain why you ϕ -ed, where this is the kind of explanation possessed reasons are linked to via the Acting On condition. The rationale explanation will be the same whether you are in the good environment or the bad environment. And thus, knowing p , or truly believing p , or truly believing with justification doesn't matter when it comes to correctly responding to reasons.

This is objectionable because it seems like there is an important rational difference when a normative because explanation can be given rather than merely a rationale explanation. When a normative because explanation is available, I'm connected to an objective reason to ϕ in a way that I'm not in cases where merely a rationale explanation is available.³⁵ The fact that you know certainly seems relevant when you're explaining yourself to other people. At least when I know that I know

³⁵For more detailed arguments for this, see Hornsby (2007, 2008), Gibbons (2001), Hyman (2006).

p , I will often find it essential to explain myself in terms of the fact that I know p when I am really pushed to explain myself.³⁶

It's important to stress that I'm not objecting to the fact that the internalist holds that you are rational in both the good and the bad case. I agree with this.³⁷ Rather, I'm objecting to the fact that the internalist view commits us to thinking that, as far as rationality is concerned, both the deceived subject and the knowing subject are *on a par*. This is what seems objectionable to me. It seems essential to explaining the rationality of the knowing to appeal to the fact that they know, even if the deceived are also rational.

So factualism doesn't commit one to thinking that, in the bad case, one's actions and attitudes are inexplicable. But even so, it still seems plausible to think that factualism commits one to thinking that, in the bad case, one doesn't act on reasons. This is motivated by the plausible claim that if I act on any reason in the bad case, it must be the same reason I act on in the good case. After all, my *reasoning* is the same in both cases.

I think this is partially right, but I don't think it's right enough to show that factualism commits you to thinking that in the bad case one doesn't base one's actions and attitudes on reasons. One problem with the last paragraph is that it assumes that, in the good case, the only reason my action is based upon is the fact that colloquium starts at 4:00. I don't think this is true. Perhaps the easiest way to see this is to think about what my belief that colloquium starts at 4:00 is based upon. In the good case, the fact that colloquium starts at 4:00 is one of the reasons I base my actions on when I go to Robertson. But it's not the only one. I also base my action on the other reasons my belief that colloquium starts at 4:00 is based upon. One of these reasons is the fact that colloquia normally start at 4:00. This is a reason to think that today's colloquium will also start at 4:00. It's one of the reasons why I believe colloquium will start at 4:00 today. A good test for seeing this is to note that it's quite plausible that I will be disposed to not think today's colloquium starts at 4:00 if I didn't possess the reason constituted by the fact that colloquia usually start at 4:00.³⁸ If I found out that they usually don't start at 4:00, then my confidence that today's starts at 4:00 would probably go down (and rightly so).

The fact that colloquia usually start at 4:00 is also, plausibly, one of the reasons I base my action upon. Again, counterfactuals provide a nice test. Given normal background conditions, it's plausible that I'm disposed not to walk over to Robertson at 3:55 if I didn't possess the reason constituted by the fact that colloquia usually start at 4:00. For example, if I didn't know that colloquia usually start at 4:00, then I'd be less likely to go to Robertson at 3:55. This is good evidence that I do partially act on the fact that colloquia usually start at 4:00 when I go to Robertson at 3:55. If this is right, then factualism doesn't require us to say that in bad cases one doesn't act on any reasons at all.

One might think there is a tension between my view about explicability and my view about which reasons are acted upon in the bad cases. The internalist non-factualist might insist that the

³⁶Objection: Your view doesn't explain *this*. After all, you think the Acting On condition just comes to p being able to explain your action, not the fact that you know that p . Reply: This is a canard. Given plausible assumptions about assertion, when I offer up a normative because explanation of the form 'I ϕ -ed because of p ,' I am expressing my knowledge that p (or at least I think I am). Often when I am pushed harder to explain myself, I will explicitly say that I knew p . This is not to say that I wasn't expressing my knowledge all along.

³⁷Indeed, chapter 6 will be devoted to explaining why I think you are justified in the bad case.

³⁸In chapter 4 I will argue that not only is this a good test, it just is what it is to base my belief on a reason that is had. The weaker claim that it's a good test will suffice here.

fact that we can make my action intelligible by citing the fact that 'colloquium starts at 4:00' as my rationale for heading to Robertson shows that that is at least one of the reasons that I base my action upon.

Again, there are explanations and there are explanations. Here's another perfectly good explanation: I went to Robertson at 3.55 because it takes 5 minutes to walk to Robertson. In some contexts this would be a good explanation. Here's another: I went to Robertson because neurochemical event *A* caused neurochemical event *B*, which caused me to begin moving towards Robertson. If I were a neuroscientist, I might want to know why I went to Robertson based on the more fine grained brain states I was in when I left. In that context, filling in a complicated neurological story would be a good explanation of why I went to Robertson. The rub is that these explanations aren't the ones we are interested in when we are asking for the reason that I based my action upon. What we want to know is what justifies my action.

It is on this score that I think the internalist view fails—even the internalist who appeals to the rationale explanations. Take my original rendering of Unusual Colloquium. Plausibly, when Dan asks me why I'm going to Robertson, he's asking for the reasons I acted on. My original answer—I'm going to Robertson because colloquium starts at 4:00—strikes him as wrong because he knows colloquium doesn't start at 4:00. He thus corrects me. At this point it would be unreasonable for me to stick to my guns and say that I am justified in going because colloquium starts at 4:00. Instead what I'll do is retract and offer up the explanation that I thought it started at 4:00. Usually Dan won't really care why I thought that. Given that Dan thinks I'm generally right about a whole host of things, he'll probably give me the benefit of the doubt and leave it at that. However, suppose Dan is curious why I thought that and he asks me. I would not appeal to the claim that colloquium starts at 4:00. I would instead appeal to the types of reasons I mentioned before, the paradigm example being that colloquia usually start at 4:00. I would only appeal to other facts (or least things I take to be facts) in defending myself. Given normal background conditions, there will be some facts to appeal to. I think this pattern of defense is strong evidence that factualism is correct.

Here's how I see the dialectic, at this point. The initial linguistic data suggests that normative because explanations are factive, and the most obvious non-factualist stories don't seem compelling. Still, there is some intuitive force behind the type of internalism that underwrites non-factualism. This force is especially acute if the non-factualist is right to think that factualism commits one to the view that in the bad cases one's attitudes and actions are inexplicable and not based on reasons. But factualism commits you to no such thing. On the first point, there are many ways in which your action is explicable. On the second point, there are usually other reasons that you base your actions and attitudes on in the bad cases (and the good cases, for that matter). It's true that the factualist posits that you are in error about your reasons in the bad case. But this seems to be confirmed by the fact that we retreat away from the stronger claims once we find out they are false. We actually seem to think we were wrong about what reasons we based our attitudes and actions on once we find out those beliefs were false. This, it seems to me, provides a compelling case for factualism, and thus for (2).

Another way out is to deny that we need to be able to ϕ because of the reasons we possess—this is to deny (1). Perhaps it is enough that the reasons we possess can be our rationales. This is implausible, as well. This is because we want the Acting On condition to tell us something about what it is to correctly respond to reasons. Because of this, it's plausible that there is a Reverse Acting On condition. It should be that when I exercise the ability required by the Acting On condition with

respect to some consideration p , I possess p . But the Reverse Acting On condition is surely false if we only need to have the ability to act with p as a rationale. We can act with p as a rationale much more often than we possess p . This, I think, provides a compelling case for (1). It follows from (1) and (2) that non-factive views are false.

3.3.5 Against $P \rightarrow TPJ$

I've now argued against Low Bar views, Holding views, and Non-Factive views. If those arguments are sound, then the correct view about the epistemic relation must be a High Bar, Non-Holding, and Factive view. Our chart now looks like this (the views that have been eliminated via previous arguments are crossed out; the features of the views that led to their rejection are bolded).

Table 3.2: Updated Epistemic Condition Taxonomy

View	High Bar v. Low Bar	Holding v. Non-Holding	Factive v. Non-Factive
$P \rightarrow B$	Low Bar	Holding	Non-Factive
$P \rightarrow P_{res}$	Low Bar	Holding	Non-Factive
$P \rightarrow JB$	High Bar	Holding	Non-Factive
$P \rightarrow PJ$	High Bar	Non-Holding	Non-Factive
$P \rightarrow TJB$	High Bar	Holding	Factive
$P \rightarrow TPJ$	High Bar	Non-Holding	Factive
$P \rightarrow K$	High Bar	Holding	Factive
$P \rightarrow PTK$	High Bar	Non-Holding	Factive

As you can see, there are two High Bar, Non-Holding, Factive views: $P \rightarrow TPJ$ and $P \rightarrow PTK$. $P \rightarrow TPJ$ holds that you possess a reason r only if the proposition p that constitutes r is true and you have propositional justification to believe p . $P \rightarrow PTK$ holds that you possess r only if you are in a position to know p . In this final subsection I will first try to allay skepticism about the notion of being in a position to know, and then I will offer two arguments against $P \rightarrow TPJ$ in favor of $P \rightarrow PTK$.

Position to Know, What

The notion of being in a position to know has come to be a common one in contemporary epistemology. Epistemologists rarely bat an eye when it is used on the periphery. However, I've found that when you place it at the forefront of your theory, philosophers will quickly become skeptical. In this case this seems fair enough, for the notion is obviously vague. To see this, we needn't look past my paradigm case in §3.3.3.

In that case, I believe that my wife will pick up my son at the normal time, despite the fact that she just told me that she can't pick him up that day. Despite the fact that she told me this, I was engrossed in my magazine just enough not to take the information in to form the belief that she can't pick him up that day. I am, at least when she tells me this, in a position to know my wife can't pick him up that day. This is why, I claim, I'm not rational in believing that she will pick him up at the normal time. Now consider this feature of my situation: 'I ate pizza on January 3rd, 2004' is

written on a piece of paper stuck between page 192 and 193 of my copy of *American Psycho* on the bookshelf behind the kitchen table. Good question: Am I in a position to know that I ate pizza on January 3rd, 2004?

I hope not, and I think not, too. I think there is a way to precisify 'position to know' such that I am in a position to know that I ate pizza on January 3rd, 2004. But this is not the precisification used here. Roughly, in order to be in a position to know p , you have to be able to learn that p without a significant change in your epistemic situation.³⁹ What constitutes a significant change is also vague (which isn't in-itself objectionable). I'm happy to adjudicate this question in the normal way we do with vague concepts. Plus, since 'rational' itself is vague, we shouldn't be too surprised if some of the analyses of what it is to be rational are also vague.

Still, I sense hostility towards this strategy. So let me precisify some more.⁴⁰ To start, compare being in a position to know with being propositionally justified. When you have propositional justification to believe p , all of what we can call the impersonal conditions for doxastic justification are met. The personal conditions are believing p and believing p for the right reasons. The impersonal conditions are just whatever is left. There are two different kinds of cases where you are propositionally justified. The first kind of cases are *inferential*, while the second kind of cases are *uninferential*. In the inferential cases, you have propositional justification in virtue of some of your other beliefs. In the uninferential cases, you have propositional justification in virtue of some other kinds of states—e.g., perceptual states.⁴¹

Since believing p is a personal condition, it's possible to have propositional justification for p without believing p . When this happens, what is it about the facts that ground your propositional justification that puts you in a position to meet the personal conditions? Here's a stab at an answer: In the inferential cases, it's the fact that you have some beliefs with contents such that, if you could and did attend to those contents and inferentially form a belief that p in the right way, you would have a doxastically justified belief that p . In the uninferential cases, it's the fact that you have some experiences such that, if you could and did attend to certain features of those experiences and uninferentially form a belief in the right kind of way, then you would have a doxastically justified belief that p . Of course I haven't told you what the right way is, but presumably my burden isn't that heavy.

Being in a position to know works in an analogous way. You are in a position to know p when all the impersonal conditions for knowledge are met. The personal conditions are (1) believing p and (2) believing p for the right reasons. The impersonal conditions are just whatever are left over. There are also inferential cases and uninferential cases. Since believing p is a personal condition, it's possible to be in a position to know p when you don't believe p . In cases like this, what is it in virtue of which you are in a position to know p ? In the inferential cases, it's the fact that you have some beliefs with contents such that, if you could and did attend to those contents and inferentially form a belief that p in the right way, you would know that p . In the uninferential cases, it's the fact that you have some experiences such that, if you could and did attend to certain

³⁹Compare with Williamson (2000) and Gibbons (2006).

⁴⁰What follows was partially inspired by Smithies (2006). I'm confident Smithies wouldn't be all that happy with the product of my inspiration (or at least the use I put it to).

⁴¹Note that I'm using 'in virtue of' in its usual broad sense. I'm not committing myself to the claim that *the beliefs themselves* or the *experiences themselves* are the justifiers. On my view, a *background condition* on the contents of the beliefs/perceptual states being the justifiers is that you have the belief or perceptual state. I'll return to this issue in ch. 5.

features of those experiences and uninferentially form a belief that p in the right kind of way, then you would know that p . Again I haven't told you what the right way is, nor I have told you what impersonal condition or conditions must be added to go from being propositionally justified to being in a position to know. But I think that I needn't solve the problem of the basing relation or the Gettier problem in order for the average reader to glob onto what I mean by 'position to know'.⁴²

In my paradigm case, I have an auditory experience that is such that were I to attend to certain features of it—the feature of it that sounds like my wife telling me she can't pick up my son that day—and formed a belief that p in the right way, then I would know that my wife can't pick up my son that day.⁴³ It's not the case that I have any such experiences when it comes to what is written on that piece of paper inside *American Psycho*, nor do I have any beliefs that play the analogous inferential role. So I'm not in a position to know that I ate pizza on January 3rd, 2004. Voila.

Let me stress that, in some moods, I think being in a position to know is broader than this characterization. For example, it doesn't seem like having my eyes closed when I enter a room I've never been in before changes the fact that I'm in a position to know the color of the room's walls (assuming that the the room is lit in a normal way). But if this characterization is correct, then having my eyes closed would make a difference. The problem is that I don't know of any characterization specific enough to allay the doubts of a skeptic that captures these cases. This is why I'm happy enough to work with the more specific version above, at least for the time being. I think there will be plenty of progress to be made even if we fix position to know in this way.

Now let's see why we should accept $P \rightarrow \text{PTK}$ over $P \rightarrow \text{TPJ}$.

The Circularity Objection

Many think that justification can be analyzed in terms of possession. Indeed, the main goal of this dissertation is to defend the view that what it is to be rational/justified is to correctly respond to the reasons you possess. Those who think this cannot accept $P \rightarrow \text{TPJ}$. For if they did, then their analysis of justification would be viciously circular. It would say, partly, that what it is to be justified is to correctly respond to the reasons you are propositionally justified in believing that are true. Since *justified* is one of the analysans and is the supposed analysandum, it's very plausible that the analysis is viciously circular.⁴⁴

Many will think that the same objection applies to $P \rightarrow \text{PTK}$. After all, they'll say, knowledge is analyzed in terms of justification. And thus being in a position to know will be analyzed in terms of justification. So if justification is analyzed in terms of being in a position to know, *justification* will

⁴²I offer my own solution to the problem of the basing relation in ch. 4 (don't ask for a solution to the Gettier problem).

⁴³One might think this way of characterizing this case opens the door for the Holding view to resist my counterexample in §3.3.2. The thought is that one's auditory experience has the relevant reason as its content. This is not how I'm thinking of the auditory experience. As I'm imagining the case, I do not have an experience that has propositional content. It's rather that the experience puts me into a position to be in a state with propositional content if I were to attend to the right features. Much of our perceptual experiences are like this. Right now there are many things in my visual and auditory field that are such that if I attended to them, I'd have a perception with propositional content. But without attending to them I don't have such a perception. Thanks to Shyam Nair for raising this objection.

⁴⁴Beddor (MS) independently makes a similar point. He thinks the objection extends to $P \rightarrow \text{PTK}$ as well. For reasons I give in the next paragraph, I disagree.

end up being one of the analysans, albeit in a more round about way. Thus, if you accept $P \rightarrow PTK$, then you can't give a non-circular analysis of justification/rationality.

No doubt this has its *prima facie* appeal. It is an orthodox view that knowledge can be analyzed in terms of justification. Nevertheless, it is an orthodoxy that I reject. Instead, I think that it is justification that is to be analyzed by knowledge. Since I think this, it's no surprise that knowledge will one of the analysans in my analysis of justification!

Defending my knowledge-first methodology is outside the scope of this chapter.⁴⁵ I think that it's independently interesting to work out a view using this method in order to see what fruit it bears. That is, I think this dissertation has merit even if I don't defend the knowledge-first methodology. The devil is in the details, and thus progress will be made by working out the details. I'm happy to write this promissory note here. The important point for our purposes is that there's simply no way for the defender of $P \rightarrow TPJ$ to accept that justification is analyzed in terms of the reasons that are possessed without thereby accepting a circular analysis. It is possible for the defender of $P \rightarrow PTK$ to do this by denying that knowledge is analyzed in terms of justification. This gives those who think justification is analyzed in terms of the reasons that are possessed decisive reason to reject $P \rightarrow TPJ$. Given the other arguments, one who accepts that justification is analyzed in terms of the reasons that are possessed has very strong reason to accept $P \rightarrow PTK$ and the claim that justification is analyzed in terms of knowledge.

The Naturalness Objection

Although the Circularity Objection provides a decisive reason to reject $P \rightarrow TPJ$ for those who think that justification is analyzed in terms of the reasons that are possessed, its force only extends that far (or at least the *decisive* nature of its force). It would thus be nice to have something to say to those readers who aren't on board with my particular project of analyzing justification in terms of the reasons that are possessed. I should say up front to readers in that position that I'm not too concerned if you get off the boat here (of course, I will be concerned with your upstream decision to not think that justification is analyzed in terms of the reasons that are possessed!). You will agree with me on much of what is important to me in this chapter if you accept $P \rightarrow TPJ$. Moreover, if I've shown that everyone who accepts that justification is to be analyzed in terms of the reasons that are possessed must accept $P \rightarrow PTK$, then I'll consider today a pretty great day for the advancement of my view.

Those *caveats* aside, there is at least one thing that can be said in favor of $P \rightarrow PTK$ over $P \rightarrow TPJ$. Namely, the property of being in a position to know is much more natural than the property of being

⁴⁵The *locus classicus* defense of the knowledge-first approach is chapter 1 of Williamson (2000). For Williamson's remarks on how this relates to analyzing justification in terms of the evidence that is had, see §9.1 of Williamson (2000).

propositionally justified in believing a truth.⁴⁶ Think of it this way:⁴⁷ Suppose we decide that the relevant epistemic condition is factive, and then we ask what the explanation is of this fact. Both $P \rightarrow PTK$ and $P \rightarrow TPJ$ offer an explanation. The $P \rightarrow PTK$ theorist explains the factivity by appealing to the factivity of being in a position to know. The $P \rightarrow TPJ$ explains factivity by pointing out that all the truths you are propositionally justified in believing are true.

It's plausible that the former explanation appeals to more a natural or unified property than the latter. An *essential* property of the state—being in a position to know—posited by the $P \rightarrow PTK$ theorist is that you can be in that state only if the content of the state is true. It's true that you can't be propositionally justified in believing a truth without the content being true. But you can be in the same state you're in when you are propositionally justified in believing a truth even when that proposition is false. Adding the truth condition to $P \rightarrow PJ$, in other words, seems *ad hoc*. And thus the factivity of the possession relation turns out to be, in Williamson's words, 'an *ad hoc* afterthought, not an organic consequence.' But the factivity of the possession relation is an organic consequence of $P \rightarrow PTK$.⁴⁸

The fact that being in a position to know is more natural than being propositionally justified in believing a truth gives us some reason to prefer $P \rightarrow PTK$ over $P \rightarrow TPJ$. I spelled this out via illustration. It seems like the factivity of the possession relation is better explained by $P \rightarrow PTK$ than $P \rightarrow TPJ$ precisely because being in a position to know is essentially factive and being propositionally justified in believing a truth isn't. Better explanations make for a better theory.

3.4 The Treating Condition

So far I've argued that you can possess a reason only if you are in a position to know the proposition that constitutes that reason. That is to say, I've argued that being in a position to know is a necessary condition on possessing a reason. The recent literature on possession has been entirely focused on

⁴⁶Here's Williamson (2009) making the same point against a similar proposal advocated by Goldman (2009b) (Goldman's view, for what it's worth, faces the Circularity Objection): 'Such a view is a rather unnatural hybrid; the truth condition is an *ad hoc* afterthought, not an organic consequence' (311).

And here's Weatherson (MS) making a similar point: 'To put the point in Lewisian terms, it seems that knowledge is much more *natural* relation than justified true belief.' In a footnote he says 'What is important here is that on sufficient reflection, the Gettier cases show that some justified true beliefs aren't knowledge and that the cases in question also show that being a justified true belief is not a particularly natural or unified property' (7).

⁴⁷I don't think this is the only way to think of it. I give another way in note 45 below. The way of putting it in the text seems to me to be one of the more simpler ways of putting it, and that's why I use that way of putting the point.

⁴⁸I think the lesson we should take from Williamson's primeness argument (see (Williamson, 2000, ch. 3) is not that knowledge is unanalyzable (this is conclusion that many readers of Williamson draw), but that knowledge is unified in a way that composite states aren't. This doesn't show that knowledge is unanalyzable because it's possible for knowledge to be both prime and analyzable (this is forcefully shown in Brueckner (2002). I also heard the point independently made by Ernie Sosa). In general, I think we should expect the prime states to be more explanatorily basic than the composite ones. Since being in a position to know is a prime state and being propositionally justified in believing a truth is a composite one, I think we should expect the former to be more explanatorily basic than the latter. I also think this provides some independent motivation a knowledge-first methodology. Williamson does make a point close to this in (Williamson, 2000, §3.6-3.7), but he doesn't think this is the grandest conclusion one can draw from the primeness arguments.

the epistemic condition.⁴⁹ In fact, as far as I can tell, all of the participants in the recent debate think that meeting the epistemic condition is both necessary and sufficient for possessing a reason.

I think this is incorrect. Above I gave the following case as a counterexample to the claim that standing in the appropriate epistemic relation to a reason is sufficient for possessing the reason.

Lois' Fish

Lois has ordered fish for supper. She finds out that the fish contains salmonella. However, she reasonably believes that salmonella is one of the many types of bacteria found in food that is harmless to eat. She thus intends to eat the fish.

I take it that, in Lois' Fish, the fact that the fish contains salmonella is a reason to not intend to eat the fish. Moreover, Lois knows that the fish contains salmonella. Thus, she meets all of the epistemic conditions.⁵⁰ However, she intuitively doesn't possess that fact *as a reason* to not intend to eat the fish. She's not in a better position with respect to the question of whether to eat the fish now that she knows it contains salmonella. Given Working Hypothesis, she doesn't possess that fact as a reason not to intend to eat.

My hypothesis about what is missing in Lois' Fish is that Lois isn't disposed to *treat* the fact that the fish contains salmonella as a reason not to intend to eat the fish. Lois isn't disposed to think (in some broad sense) that the fact that the fish contains salmonella is a reason to not intend to eat the fish. This is why she fails to possess that reason even though she stands in the epistemic relation to the proposition that constitutes that reason. At first pass at least it seems like you can possess a reason r to ϕ only if you are disposed to treat r as a reason to ϕ .⁵¹

In order to know what it is to be disposed to treat something as a reason we need to know what is it to treat p as a reason to ϕ . Despite the prevalence of the term 'treat p as a reason,' not much philosophical attention has been directly paid to this question.⁵² So in some sense we can start fresh (or at least our start will be fresher than most).

I think that one obvious condition is that you believe the proposition that you treat as a reason. If you think $\neg p$, then you can't treat p as a reason to ϕ , although you can be disposed to treat p as a reason to ϕ .⁵³ For example, you might be disposed to treat the proposition *it is hot out today* as a reason to turn on the air conditioner when you don't believe it's hot out today, but you can't treat that proposition as a reason to turn on the air conditioner when you don't believe it's hot out today.

⁴⁹Some prominent examples include Williamson (2000), Schroeder (2008, 2011), Neta (2008). It should also be noted that all of the critical replies to Williamson's arguments for E=K that I'm aware of make this assumption, although that obviously doesn't mean all of the critics share the assumption.

⁵⁰She is in a position to know, she knows, she is propositionally justified, she is doxastically justified, she believes, and she has presentational attitude with the reason as its content. Moreover, it's vacuously true that the content is an objective reason if true because it is true and it's an objective reason.

⁵¹You might ask, 'Why isn't the condition that you have to actually treat it as a reason instead of merely being disposed to treat it as a reason?' In short, it's because in order to treat p as a reason to ϕ , you have to believe p . However, I think there are cases where you intuitively possess p even though you don't believe p . So in those cases you don't treat p as a reason, but you are disposed to. I will argue for this after I spell out what it is to treat something as a reason.

⁵²An admirable exception is Schossler (2012). Much of my discussion is heavily indebted to his paper, although I disagree on some details.

⁵³This doesn't mean you can possess p when $\neg p$ since you won't be able to meet the epistemic condition when $\neg p$.

Believing p is a necessary condition, but it's obviously not sufficient. On one view, all that needs to be added are normative beliefs. Let's call views of this type Normative Belief views. On the simplest Normative Belief view, all that is needed is that you believe that p is a reason to ϕ .⁵⁴ There are at least two different versions of this simple Normative Belief view. According to the first, you have to explicitly believe that p is a reason to ϕ in order to treat p as a reason to ϕ . This explicit belief version is very implausible. We very rarely explicitly believe that p is a reason to ϕ , yet we very often treat things as reasons. According to a second view, all that is needed is that we implicitly believe that p is a reason to ϕ . As we'll see, I think there is something to this view. However, I think even this view is both too broad and too narrow.

It's too broad because it's plausible that you can believe that p is a reason to ϕ and not treat p as a reason to ϕ . For example, it seems plausible that I could believe that the fact that Oxfam can save a child's life for $\$x$ is a reason to send $\$x$ to Oxfam *and yet* this fact could have no effect whatsoever on whether I'll send money to Oxfam. It could be that the fact that I believe that Oxfam can save a child's life for $\$x$ just doesn't make any difference—even any potential difference—to whether I'll send money to Oxfam. This is compatible, as far as I can tell, with implicitly believing that that fact is a reason to send money to Oxfam.⁵⁵ However, it doesn't seem like you really *treat* p as a reason to ϕ if you aren't any more likely to ϕ when you think p obtains.

I think this objection highlights the fact that believing that p obtains must have some type of broadly motivational upshot vis-à-vis ϕ -ing if you treat p as a reason to ϕ . Following Schossler (2012), I propose that the relevant condition is that believing p disposes you to ϕ —i.e., you treat p as a reason to ϕ only if believing p disposes you to ϕ . As we'll see, this isn't enough, either. Before we get to that, we should consider why the implicit version of the Normative Belief view is too narrow.

The implicit belief view is too narrow because some agents can treat things as reasons even if they don't have enough of the concept of a reason to have beliefs about reasons. For example, I think that my dog treats the fact that I called her as a reason to come to me. However, I don't think that my dog can even implicitly believe that the fact that I called her is a reason to come to me. I doubt she's conceptually sophisticated enough to have that belief. Similarly for my toddler daughter. I think she treats the fact that honey nut cheerios are tasty as a reason to eat her breakfast, but I'm confident that she doesn't have the conceptual capacities to have beliefs about reasons. This doesn't mean that she completely lacks the concept of a reason. I think she (and my dog) partially have the concept, just not enough to have beliefs about reasons.⁵⁶ Because I think this, I think any

⁵⁴(Scanlon, 1998, pgs. 65-6) seems to hold this view, but it's not clear to me that Scanlon is talking about the same thing I'm talking about here (Schossler (2012) does think Scanlon is talking about the same thing.)

⁵⁵Things are tricky on this point. I think Scanlon would deny this. He argues in Scanlon (1998, MS) that beliefs about reasons are actually very much like desires—in fact, he thinks we can accurately call them beliefs and call them desires. So I don't think he thinks that you can be left completely cold by a (purported) fact and still believe it's a reason to ϕ . The problem with his view is that it doesn't seem like a good account of normative *beliefs* nor does it seem like a good account of *desires*, at least in their non-stipulative senses. For convincing arguments about the desire point, see Schossler (2012).

⁵⁶I should say that I'm not strongly wedded to this point. If I thought there was compelling reason (which I don't) to think my daughter or dog aren't treating these things as reasons in a full blown sense, I wouldn't think this is a grave cost. I think the broadness point is more important. Moreover, I think it's quite plausible that those of us who have enough of the concept of a reason to have beliefs about reasons will believe p is a reason to ϕ whenever they treat p as a reason to ϕ .

view that has a normative belief component overintellectualizes the notion of treating something as a reason.

In fact, I think this point generalizes to all views that have some type of awareness condition on treating something as a reason. Schossler (2012)'s view is an example of a view that has an awareness condition that doesn't amount to a normative belief condition. On Schossler's view, in order to treat p as a reason to ϕ , I have to be aware of the fact that believing p disposes me to ϕ . It's not clear to me why he thinks this. He simply says that a condition like this is needed. It seems to me like this is just an overintellectualization.⁵⁷ Why would we need to be aware of the fact that believing p disposes us to ϕ in order to treat p as a reason to ϕ ?

So far the two conditions we've settled on for treating p as a reason to ϕ are (1) you believe p and (2) your belief that p disposes you to ϕ . Let's call the former the Belief Condition and the latter the Action Condition. Unfortunately, these two conditions are not sufficient, either. For it could be that the connection between your belief and your disposition to ϕ is the wrong kind of connection. For example, it might be that my belief that it's hot out today disposes me to turn on the air conditioner, but only because it disposes me to think I don't have any shoes on and, oddly, that disposes me to turn on the air conditioner. To give another example, suppose that a benevolent neuroscientist implants a disposition-creating chip in my brain and he makes it so that I come to be disposed to turn on the air conditioner when I believe it's hot outside. Intuitively, I might not see any justifying connection between it being hot outside and turning on the air conditioner. I might love to be hot and to save money on my electric bill. Still, the chip will make it so I'm disposed to turn on the air conditioner when I believe it's hot.

In order to solve this problem, I think we need to add something akin to the implicit normative belief. But, of course, it can't exactly be a belief with a normative content because that is an overintellectualization. I think the first part of the answer is that it needs to be that one is disposed to think, during deliberation about whether to ϕ , that p is positively relevant to whether to ϕ .⁵⁸ Let's call this the Deliberation Condition. One can satisfy the Deliberation Condition without having the conceptual resources to have beliefs about reasons. I'm quite confident that my daughter is disposed to think that the tastiness of her cheerios is positively relevant to whether eating her breakfast is the thing to do when she deliberates about whether to eat her breakfast.

This still isn't enough, though. It's true that when I meet the Normative Condition it's implausible to think that p leaves me cold vis-à-vis ϕ -ing. But this is still insufficient because it seems like these three conditions can be met, but I can still fail to treat p as a reason because there isn't the right connection between my being disposed in deliberation to think that p is positively relevant to whether to ϕ and the fact that my belief that p disposes me to ϕ . Our benevolent neuroscientist case might be, for all I've said, like this. There might not be any connection between the fact that my belief that p disposes me to ϕ and my disposition to think p is positively relevant in deliberation about whether to ϕ . In short, the Action Condition might not connect in the right way to the Deliberation Condition.

I think there is a simple solution to this problem. I think that, in order to treat p as a reason to ϕ , it has to be that your belief that p disposes you to ϕ *because* you're disposed to think p is positively

⁵⁷Moreover, it's just not at all clear what Schossler means by 'awareness.' He makes it clear that the awareness need not be occurrent, nor need it be an awareness of actual causation.

⁵⁸Note that it's quite plausible that you have this disposition whenever you believe that p is a reason to ϕ . So I don't think the normative belief accounts have a problem accounting for this feature, even though they do have a hard time accounting for the action condition.

relevant to whether ϕ is the thing to do. That is to say, the fact that the Deliberation Condition holds needs to explain why your belief that p disposes you to ϕ . When this is the case, the Deliberation Condition is connected in the right way to the Action Condition. Let's call this the Explanatory Condition.

I think that meeting the Belief Condition, Action Condition, Deliberation Condition, and Explanatory Condition is necessary and sufficient for treating p as a reason to ϕ . Thus, I think Treating is true:

Treating: A treats p as a reason to ϕ iff (1) A believes p , (2) A 's belief that p disposes A to ϕ , (3) A is disposed to think p is positively relevant to whether to ϕ in deliberation, and (4) (2) is explained by (3).

Now that we have a better idea what it is to treat p as a reason to ϕ , we are in a better position to see what it is to be disposed to treat something as a reason. A rough way to put what I think the treating condition comes to is to say that you meet the treating condition on having a reason just in case you are disposed to meet conditions (2)-(4) when you believe p .⁵⁹

A natural question to ask now is why should we think that you meet the treating condition when you are *merely disposed* to treat p as a reason to ϕ instead of when you *actually* treat p as a reason to ϕ . The answer is that we've already seen a counterexample to the claim that you have to actually treat p as a reason to ϕ . This was the following case from §3.3.4:

Out of the Ordinary

Each morning I casually peruse a magazine while I eat my breakfast. I am doing this on Monday morning. As I'm doing this, my wife tells me that she has an unusual schedule that day. Given her unusual schedule, I need to pick up my son from school (this is something she almost always does because we have agreed that that is one of her daily tasks). She speaks clearly when she tells me these facts, and she is a mere 10 feet from me. However, I am engrossed just enough in the Newsweek I'm reading to not process the information. Given the fact that picking up my son is not one of my usual tasks, I believe that I don't need to pick him up that day. My son doesn't get picked up on time.

Consider again my belief that I don't need to pick up my son. A natural explanation of this is that my wife just gave me decisive reasons to think this is false—e.g., the fact that she can't pick him up. Intuitively, it seems like I possess that reason. After all, my wife just told me that. But *ex hypothesi* I don't believe that I need to pick up my son from school. So I don't actually treat that as a reason to think my day will be unusual. I am, it's plausible to suppose, disposed to treat that as a reason to believe my day is unusual. If I were paying more attention and came to believe what my wife told me, I would treat that fact as a reason to believe that my schedule is unusual.

This isn't the case in Lois' Fish though. Lois isn't even disposed to treat the fact that there is salmonella in the fish as a reason not to intend to eat the fish. We can see this because Lois does meet condition (1) but doesn't meet the other conditions—her belief that there is salmonella in the fish doesn't dispose her to not intend to eat the fish and she's not disposed to think the fact that the

⁵⁹This is rough because it seems, strictly speaking, false. That is, I'm not sure if the disposition to treat p as a reason to ϕ comes to the claim that I put in the text. But what I put in the text is a nice way to get at the underlying idea.

fish contains salmonella is positively related whether to refrain from intend to eat the fish. These differences between me in Out of the Ordinary and Lois in Lois' Fish explain why I possess the reason in Out of the Ordinary despite the fact that I don't believe the proposition that constitutes the reason and why Lois doesn't possess the reason even though she does believe the proposition that constitutes the reason.

I must consider one objection to the way I've been thinking of the treating condition. There are cases similar to the cases that motivated our rejection of Holding views that seem to motivate the rejection of what we might call Holding views of the treating condition. Holding views of the treating condition hold that in order to meet the treating condition, you *actually* have to treat or be disposed to treat the reasons in question as reasons. Non-holding views don't think you have to actually treat or be disposed to treat things as reasons in order to meet the treating condition. You might think Holding views of the treating condition are implausible because of cases like Murderer:

Murderer

Bob is a sociopathic murderer. He knows that stabbing his next victim to death would cause her great pain. However, he does not treat this as a reason for not stabbing her to death, nor is he disposed to treat this as a reason not to stab her.

If my view of the treating condition is correct, then Bob doesn't possess the reason constituted by the fact that stabbing his next victim to death will cause her great pain. Moreover, if Bob is sufficiently depraved, then he won't possess any of the reasons not to stab his next victim to death. It thus won't be irrational for him to stab his next victim to death.

I certainly feel the pull of this objection! The first thing to say is that we need to pay careful attention to what Bob and his ilk are actually like. For the more depraved they are, the more plausible it becomes that they aren't required by rationality to refrain from doing heinous things. The reason for this is familiar. At a certain point of depravity it will be plausible that their rational capacities themselves are sufficiently compromised to open them up to the type of rational criticism that is intimately connected with rational requirements.

Of course, pointing this fact out isn't sufficient for meeting this challenge. For there will be some level of depravity that won't be sufficiently debilitating to compromise one's rational faculties. Let's suppose Bob is depraved in this way.

The most natural fix would be to move to a Non-Holding account of the treating condition—perhaps the view that you meet the treating condition for some r to ϕ just in case you are in a position to know that r is a reason to ϕ . I am attracted to this view, save one consideration. Namely, this reintroduces the type of overintellectualization I objected to above in the discussion of Normative Belief views. This is because in order to be in a position to know that r is a reason to ϕ , you must possess enough of the concept of a reason to form explicit beliefs about reasons. But it seems to me that agents that lack this conceptual sophistication can still treat things as reasons. This makes me hesitant to adopt a Non-Holding account.

Because of this, I think I should endorse another answer to the challenge. According to this view, even though it might be true that Bob lacks the reasons not to stab his victim, he does possess reasons to treat the relevant facts as reasons. Moreover, those reasons are decisive reasons to treat the relevant facts as reasons. Thus, Bob is irrational in not treating those facts as reasons. So while he might not be rationally required to treat his next victim well, he is rationally required to do something that would make it the case that he is rationally required to treat his next victim well.

I should point out that cases like this are ubiquitous. Suppose Hank is required to hold his office hours. If he held his office hours, then when Susan his student came by, he'd be required to talk to her. But he's not necessarily required to talk to her if he doesn't fulfill his requirement to hold office hours.⁶⁰ Another example: Suppose I'm required to save for my children's college educations. If I meet this requirement, then when my children go to college, I'll be required to write checks for large sums of money. However, if I don't meet this requirement, I might not be later required to write those checks. Writing those checks without the money might do much more harm than good, even to my children and to their prospects for obtaining a college education. In those cases some other course of action might be the one I'm required to pursue.

As I see it now, this is what should be said about Bob and his ilk. They might not be required by rationality to refrain from doing heinous things. But that doesn't mean they're fully rational when they do. For it's plausible that they are rationally required to treat the relevant facts as reasons. And if they did that, then they would be rationally required not to do the heinous things. Now of course it is possible for there to exist individuals so depraved that they don't even possess reasons to treat certain facts as reasons. These characters, I conjecture, are quite likely to be so depraved that their rational capacities are compromised. So it's no surprise that they would fail to possess reasons that a very diverse range of human agents possess.

3.5 The View

Now—finally—we've arrived at a view of what it is to possess a reason. There are two main conditions. The first is an epistemic condition. I argued above that you meet the epistemic condition for having some reason p to ϕ just in case you are in a position to know p . Despite the fact that the literature has spent all of its time focusing on the epistemic condition, it is not sufficient. There is also a treating condition. I argued that you meet the treating condition for having some reason p to ϕ just in case you are disposed to treat p as a reason to ϕ . Combining these, we get Possession:

Possession: What it is to possess a reason p to ϕ is to be in a position to know p and to be disposed to treat p as a reason to ϕ .

I argued in the last section that you are disposed to treat p as a reason to ϕ just in case you are disposed to meet the following three conditions when you believe that p : (1) your belief that p disposes you to ϕ , (2) you're disposed to think p is positively relevant to whether to ϕ , and (3) (2) explains (1).

I won't attempt to rehash the arguments that have led us to this point. Suffice it to say, I think that Possession is the view that faces up to scrutiny the best.

⁶⁰Suppose instead of holding office hours he goes for a walk. On his walk he comes across a pond with a drowning child in it. This happens at exactly the same time Susan stops by his office. He's not required to meet with Susan in this case, he's required to save the child, even though he would have been required to meet with Susan if he would have done what he was required to do earlier.

Chapter 4

Deviance in Action and Belief

4.1 Introduction

4.1.1 Where We've Been and Where We're Going

IN chapter 2 I initially motivated the main claim of this dissertation—viz., that what it is to be rational is to correctly respond to the normative reasons you possess. In chapter 3 I defended a view about what it is to possess a reason. In this chapter, the final chapter in Part I, I will defend a view of what it is to correctly respond to the normative reasons you possess. On my view, you correctly respond to the normative reasons you possess when you hold attitudes or perform actions *for* those reasons. Since those reasons are normative reasons, you correctly respond when you hold attitudes or perform actions for normative reasons. The primary goal of this chapter with respect to the goals of the whole dissertation is thus to defend a view of what it is to act and hold attitudes for normative reasons. Given the unorthodox view I have about the relationship between acting and holding attitudes for normative reasons and acting and holding attitudes merely for reasons, it also makes sense for me to defend an account of what it is to act or hold an attitude merely for a reason. With that, onward!

4.1.2 The Lay of the Land

Most of our actions and beliefs are performed or held for reasons. Bob cleared the snow from his driveway because he needs to drive to the office; Mary believes it is cold outside because it is snowing; and I am using examples involving snow because it is currently snowing in Princeton. Sometimes the reasons for which we act or believe are *normative reasons*. Plausibly, Bob's action is performed for a normative reason and Mary's belief is held for one—the reason for which Bob acts counts in favor of his act and the reason for which Mary believes counts in favor of (the content of) her belief.

An important project for both epistemology and metaethics is to understand what it is for token beliefs and actions to be held and done for normative reasons.¹ This is important because it is very plausible that we won't be able to fully understand what it is for token beliefs and actions to

¹The focus isn't just on beliefs and actions, of course. A common assumption—one that I will make—is that we will be able to generalize from these two paradigm cases to other attitudes (e.g., desire, intention, hoping, supposing et cetera).

be rational, justified, creditworthy, or correct without understanding what it is to believe and act for normative reasons.² This can be seen via example. Imagine that Jack has excellent reasons to believe that the temperature will be between 40 and 70 degrees tomorrow—5 reliable forecasts predict it will be between 50 and 60 degrees. Jack also believes the temperature will be between 40 and 70 degrees. However, his reasons for believing this are that tomorrow is April 12th, that the temperature of any given day will be the product of the month and the day, and that $4 \times 12 = 48$. He infers from these beliefs that the temperature will be between 40 and 70 degrees.

It is clear that Jack's token belief is not rational, justified, creditworthy etc., despite the fact that he has excellent reasons to believe what he in fact believes (I will herein focus on rationality). The lesson to be learned is that not only must one's beliefs and actions be supported by normative reasons to be rational, there must also be the right connection between one's actions and beliefs and one's normative reasons for those actions and beliefs. An extremely plausible hypothesis about what this connection amounts to is that the actions and beliefs that are rational are done and held for (at least some of) the normative reasons one has for those actions and beliefs. To figure this out is to figure out what it is to believe and act for normative reasons.

Common ground in the debate about what it takes to believe and act for normative reasons is that when one acts and believes for normative reasons, those normative reasons explain one's actions and beliefs. Ground that is far from common is whether the relevant explanatory relation is a causal relation. The central problem for causal accounts is the famous *deviancy problem*. The deviancy problem shows that not all causation by normative reasons counts as believing and acting for normative reasons. In order to solve the deviancy problem, one has to explain what 'causation in the right way' amounts to. This has proven a hard task. In this chapter I will provide an analysis of believing and acting for normative reasons that, I claim, solves the deviancy problem.

This is not all this chapter is about, though. It is also about a topic central to the philosophies of action and mind. Namely, what it is to believe and act for reasons. Most beliefs and actions are done for reasons. This, of course, is not to say that most beliefs and actions are rational. This is because our ability to believe and act for reasons far outstrips the actual normative facts. The sad fact is that we believe and act for silly reasons all the time. Nevertheless, it is an important and interesting question in the philosophies of mind and action what exactly believing and acting for reasons comes down to.³

The debate about this question is very similar to the debate about believing and acting for normative reasons. It is common ground that the reasons for which one believes and acts explain one's actions and beliefs. And it is far from common ground whether the relevant explanatory relation is a causal one. Moreover, the main problem for causal accounts is the existence of deviant causes. It turns out that many philosophers of mind and action are also after an elusive 'causation in the right way'. In this chapter I will provide an analysis of these notions that, I claim, solves the deviancy problem.

²Theorists who think the topic is important because of some connection to rationality or justification include, among many, many others Turri (2011), Schroeder (MSb), Wedgwood (2003). Theorists who think it is relevant to correctness include Schroeder (MSb), Wedgwood (2003). Theorists who think it is relevant to credit or praise/blame include Markovits (2010), Arpaly (2006).

³Indeed, according to one approach to the philosophy of action, this is the central problem. On this approach, the essential feature of actions is that they can be explained by considerations that provide some type of justification or rationale. This is often treated as an axiom of the philosophy of action, with support from the two most influential theorists in the 20th century, Anscombe and Davidson.

It is natural to be a bit puzzled at the way I have divided the terrain. This is because it is natural to think that believing and acting for normative reasons is parasitic on believing and acting for reasons. Surely what it is to act and believe for a normative reason is just a matter of acting and believing for a reason that is also normative. If this is right, then our primary task should be to figure out what it is to believe and act for reasons. We can then combine our theory of this with the correct theory of normative reasons to get the correct theory of what it is to act and believe for normative reasons.

I'll argue—show, dare I say—that this is a mistaken way to think about acting and believing for normative reasons. Any view that tries to analyze acting and believing for normative reasons in terms of acting and believing for reasons that happen to be normative reasons is going to face a new type of deviancy problem. There will be cases of acting and believing for reasons that are normative reasons that aren't cases of acting and believing for normative reasons. Thus, acting and believing for normative reasons deserves its own treatment and its own analysis. I will show how my analyses of acting and believing for normative reasons solve the deviancy problems.

The structure will be this. I will start, in §4.2, by introducing the lay of the land in the debate about what it is to act and believe for reasons. Central to this introduction will be explanatory accounts and their deviancy problem. The main task of §4.3 will be to argue that acting and believing for normative reasons is a subject in its own right. It is not parasitic on the debate about acting and believing for reasons. Thus, we will need two analyses. One of what it is to act and believe for reasons and one of what it is to act and believe for normative reasons. Both will have to solve various deviancy problems. §4.4 will be dedicated to defending an analysis of what it is to act and believe for normative reasons, while §4.5 will be dedicated to defending an account of what it is to act and believe for reasons. I will argue that these accounts solve the deviancy problems. In §4.6 I'll explain how the two accounts are related. I'll close, in §4.7, by showing how my accounts can be used to solve or dissolve standing problems in philosophy of action, practical reason, and epistemology.

4.2 Acting for Reasons, Believing for Reasons, and Being Deviant

I just bought my son an ice cream cone. Why? Because he has been well behaved today. Rose went to room 201 on May 10th, 2012. Why? Because her sociology lecture was held in that room at that time. Bart believes that Andrew Wiles proved Fermat's last theorem. Why? Because his math textbook says so. Sally always smokes cigarettes at parties. Why? Because it makes her look cool (or so she thinks).⁴ Louis believes he will be lucky this year. Why? Because his horoscope says he will be.

These are all cases of acting or believing for reasons. Interesting question: What is it in virtue of which the above *because* claims are true? What is it about me that makes it true that I bought my son an ice cream cone because he has been well behaved today? And what it is about Rose that makes it true that she went to 201 on May 10th, 2012 because her sociology lecture is held in that room on that day?

Nearly all accounts hold that in order for some consideration p to be the reason for which A ϕ s, p must explain A 's ϕ -ing. We'll call these *explanatory accounts*. A central debate in the literature is about whether the relevant kind of explanation is *causal*. According to causal accounts, the reasons for which we act and believe cause our actions and beliefs. Not just any old cause will

⁴For those keeping track of the connections between this chapter and the last, I think that Sally's rationale for smoking is that it makes her look cool.

do, though. This is because there are clearly deviant causes. Here's a version of Chisolm (1966)'s famous example:

Dead Uncle

Dan has decided to kill his uncle Fred because he is set to inherit a vast amount of money from his uncle. The thought of inheriting the money excites him so much that he loses control of his car and kills a pedestrian. The pedestrian happens to be his uncle. He inherits the fortune.

In *Dead Uncle*, some consideration—viz., that he is set to inherit a fortune from his uncle—causes Dan to do something that leads to the death of his uncle. However, it doesn't seem like the reason for which Dan kills his uncle is that he is set to inherit a fortune from him. Indeed, it doesn't seem like Dan performs any action at all. Thus, causing the outcome is not sufficient for some consideration to be the reason for which some agent acts.

Here's a version of Davidson's most famous deviancy case (originating in Davidson (1980)):

Dead Climber

Dan and Fred are climbing an arduous mountain together. Dan decides that he will only survive the climb if he rids himself of the burden Fred brings to the project. He thus believes that he must get rid of Fred in order to survive. This thought so terrifies him, though, that it causes him to loosen his grip on the rope that is supporting Fred. Fred falls to his death. Dan survives.

In *Dead Climber*, there is some consideration—that Dan will survive only if Fred dies—that causes Dan to do something that leads to Fred's death. However, it doesn't seem like the reason for which Dan loosens his grip is that he must kill Fred in order to survive. Once again, it doesn't seem like Dan performs an action at all. Thus, causing the outcome isn't enough to count as a reason for which one acts.

Cases with the same structure litter the literature on believing for reasons found in epistemology. Here's one from Pollock & Cruz (1999):

Late for Class

Joe is late for class. This causes him to quicken his pace. This causes him to slip, which causes him to fall on his back and look skyward. This causes him to believe there are birds in the tree.

In *Late for Class*, there is some consideration—that Joe is late for class—that causes him to believe that there are birds in the tree. However, it's clear that the reason for which Joe believes there are birds in the tree is not that he is late for class. So some consideration causing a belief isn't enough for that consideration to be a reason for which the belief is held.

Some have thought that deviancy problems are unsolvable and thus think we should reject causal theories of acting and believing for reasons.⁵ Others are convinced that only a causal theory

⁵See, e.g., Ginet (1990, 2008), Swain (1981).

will do and have faith that even though all proposed solutions are false, some solution must be out there.⁶ In due time I will offer analyses that I think solve these problems. First, though, I want to introduce the closely related debate about what it is to act and believe for normative reasons and argue that the notion of acting or believing for normative reasons is not parasitic on the notion of acting or believing for reasons.

4.3 Acting and Believing for Normative Reasons

Although we almost always act and believe for reasons when we act and believe, we don't always act and believe for normative reasons.⁷ We often, although not as much as one would like, act and believe for normative reasons. I might refrain from jumping off the building because it's so tall. I might believe Anne is amused because she's smiling. I might start walking at 3:55 because the talk starts at 4:00 (alas, I'm not so good at doing this).

It's important to work out what it is to act and believe for normative reasons primarily because there is a tight connection between acting and believing for normative reasons and my beliefs and acts being rational, justified, blameworthy, or creditworthy. I'm particularly interested in the notions of *ex ante* and *ex post* rationality or justification.⁸ One is *ex ante* rational for ϕ -ing when one has sufficient reason to ϕ . But this isn't enough to be *ex post* rational for ϕ -ing. In order to be *ex post* rational for ϕ -ing, one must ϕ and there must be the right connection between the reasons that *ex ante* rationalize ϕ -ing and one's ϕ -ing. One's belief that p isn't *ex post* rational if it's held because of wishful thinking, even if it is supported by the evidence. A very plausible hypothesis about what the right connection is between one's rationalizers and one's attitudes and actions is that of being held for or done for those reasons. Thus, the notion of acting or believing for normative reasons has become extremely important for the debate about *ex post* rationality and justification.

It's natural to think that when I act or believe for normative reasons, I'm not doing anything over and above acting and believing for reasons. It's just, in those cases, the reasons for which I'm acting and believing are also normative reasons. If this were right, then once we worked out both what it is to act and believe for a reason and what is for something to be a normative reason, we'd know what it is to act and believe for a normative reason. This view is nearly always assumed in the literatures on believing and acting for normative reasons. Audi (1993) eloquently gives voice

⁶See, e.g., Davidson (1980), Arpaly (2006), Pollock & Cruz (1999). One gets the sense that trying to solve the deviancy problems is a bit passé. Some evidence for this: Nearly all of the more recent work on the problem in epistemology devote large amounts of space in the introduction justifying the importance of the topic to other issues in epistemology (see, e.g., Turri (2011) and especially Evans (FC)). Apparently epistemologists need constant reminding that a lot of things hang on the right answer. Even though the problem has not been solved and most theorists are card-carrying causal theorists, solving the problem doesn't seem to be high on most people's to-do list.

⁷I will assume here that normative reasons are facts (or, if think there is a difference between true propositions and facts, true propositions). One needn't hold this view in order to accept what I say below. I think my view about acting and believing for normative reasons is compatible with any view of normative reasons. For what it is worth, I make this assumption because I think it is true, as do many others. It is the orthodox view about practical normative reasons in metaethics (see, among many others, Dancy (2000), Scanlon (1998), Parfit (2011), Schroeder (2007)). It is also held by some epistemologists, most prominently Schroeder (MSb), Williamson (2000).

⁸This more often goes under the heading of propositional and doxastic justification in the epistemology literature. The problem with this terminology, in this context, is that it can't be generalized to action or practical attitudes. This is why I, following Goldman (1979), Wedgwood (2002, FCb), use the *ex ante/ex post* terminology.

to the view. The pages preceding this quote are spent defending a view of what it is to believe for reasons.

Since believing for a good reason is believing for a reason (one that is good), the account clarifies believing for a good reason. [...] Indeed, if an indirectly (prima facie) justified belief is simply a belief held for at least one good reason, then if our conditions are supplemented with an account of what constitutes a good reason, we shall have all the materials we need to understand one of the main kinds of justified belief and, in good part, one of the main kinds of knowledge (267).

Here is Dancy (2000) explaining the connection between acting for reasons and acting for normative reasons.

It seems, then, that the explanation of action [...] can always be achieved by laying out the considerations in light of which the agent saw the action as desirable, sensible, or required. If things were as the agent supposed, there is no bar against the agent's reasons being among the reasons in favor of doing what he actually did. [...] The reasons that motivated the agent can be among or even identical with the good reasons in the case (136-37).⁹

This suggests Composite:

Composite: What it is to ϕ for some normative reason r is to ϕ for r and for r to be a normative reason to ϕ .

Despite the fact that Composite is often taken as an axiom in the debates about believing and acting for normative reasons, it is false.

The problem is that it's possible to ϕ for r and for r to be a normative reason to ϕ even though you don't intuitively ϕ for the normative reason r . Cases that show this have been around in epistemology since at least Swain (1988), even though no one has fully appreciated what they show.¹⁰ Affirming the Consequent is one such case:¹¹

⁹There are several others who explicitly endorse the thought. See, e.g., Schossler (2012), Armstrong (1973), Swain (1988, 1981). Over and above these examples, the thought often functions as a basic presupposition of the debates, as evidenced by the fact that nearly everyone thinks the key thing to explain is acting and believing for reasons even in contexts when they are searching for the connection that must hold between normative reasons and beliefs and actions in order to be rational.

¹⁰ See Turri (2010) for many cases of this kind. Millar (1991), Goldman (2012), Alston (1988) discuss the same type of cases. I will return to Turri and Millar below. See also Lord & Sylvan (MS).

¹¹An important qualification: I will assume that the only normative reasons that are relevant here are possessed normative reasons. I will be officially neutral in this chapter about what the possession relation is. I will, however, make the simplifying assumption that characters in my stories know all of the relevant facts. As we saw in the last chapter, on any plausible view about possession, knowing the fact is sufficient to meet the epistemic condition. I will say a bit more about possession in §6 below.

It is also important to note that this doesn't mean that I am assuming that Sam possesses the relevant reason in the case below. She probably doesn't. This doesn't spoil the point because she does believe for the consideration that is the normative reason. If Composite is true, then this should be sufficient for acting for a normative reason. Moreover, I think you can get cases with the same structure as Affirming the Consequent where the character possesses the reasons. These cases are not as clear as Affirming the Consequent, though. Thanks to Shyam Nair for pressing me to clarify this.

Affirming the Consequent

Sam is wondering whether Terry took the bus to work. She knows that Terry's car is in the driveway. This is, in fact, a sufficiently good reason to think that Terry took the bus. Sam also believes that if Terry took the bus, then Terry's car is in the driveway. She comes to believe that Terry took the bus by (deductively) inferring that she took the bus from her belief that Terry's car is in the driveway and her belief that if Terry took the bus, then Terry's car is in the driveway.

Affirming the Consequent is, I submit, a deviance counterexample to Composite. First things first, let's confirm that the two conditions of Composite are met. By stipulation, the fact that Terry's car is in the driveway is a sufficient reason to believe Terry took the bus. This is just to say that it is stipulated that the fact that Terry's car is in the driveway is a reason to believe Terry took the bus and that it is stipulated that there are no defeaters of this reason. Moreover, Sam does believe Terry took the bus for that reason. After all, she inferred that Terry took the bus from that proposition. Nonetheless, she intuitively doesn't believe *for a normative reason*. For one thing, she affirms the consequent! More importantly, notice that she is actually insensitive to whether the fact that Terry's car is in the driveway is a reason to believe Terry took the bus.^{12, 13}

To see this, let's add a bit to the story. Suppose that the fact that Terry's car is in the driveway isn't always a reason to believe Terry took the bus, for on Wednesdays Terry carpools. So on Wednesdays the fact that Terry's car is in the driveway is no reason to believe she took the bus.¹⁴

¹²Objection: The problem here is that Sam is also believing for another reason—viz., that if Terry took the bus, then her car is in the driveway. In order to believe for a normative reason *r*, one might think, it has to be that one doesn't also believe for a consideration that isn't a normative reason. I have two replies to this objection. First, this would underdetermine the extension of believing for normative reasons. We often believe for multiple reasons. Sometimes some of those reasons aren't normative reasons. But that doesn't automatically entail that our beliefs aren't based on any normative reasons (cf. Turri (2010)). Second, *and most importantly*, this objection can be evaded by modifying the case. Instead of having the conditional be a *premise* in Sam's inference, you can change the case so that *Terry's car is in the driveway* is the only premise and Sam follows the Affirm the Consequent inference rule. Since inference rules can't be further reasons for belief lest we posit a Carroll regress, in this modified version of the case the only thing to say is that the sole reason for which Sam believes Terry took the bus is that Terry's car is in the driveway.

¹³Objection: We can't factor out Sam's reasons like this. Instead, Sam only has one reason, which is a conjunction of all the claims that Sam's belief that Terry took the bus to work depends on (for a reply like this to a similar kind of case, see Alston (1988)). Since this conjunction isn't a normative reason to believe Terry took the bus to work, this isn't a potential counterexample to Composite. Reply: This assumes both an implausible view about normative reasons and an implausible view of believing for (normative) reasons. First, it annihilates the difference between background and foreground conditions. To take an uncontroversial example, the fact that the wall seems red is a normative reason to believe the wall is red. But this normative fact depends on other facts in certain ways. For example, the fact that it seems red is not a reason to believe it's red if you know that there is a red light shining on it. This doesn't show, though, that when it is a reason the normative reason is actually constituted by the fact that it seems red *and* the fact that it is not the case that a red light is shining on it. The latter fact is just a background condition for the former fact being a normative reason. Second, there are similar problems for the account of believing for reasons that is assumed by the objector. Sam's belief that Terry took the bus might depend on a very large amount of things. It might be that Sam will give up that belief if she comes to believe all sorts of things. It doesn't follow from this that her reason for believing that Sam took the bus is that Terry's car is in the driveway *and* she doesn't believe *p* *and* she doesn't believe *q* *and* she doesn't believe *x*...

¹⁴This is a bit controversial, but suppose it's right. One can change the details if one doesn't like these ones. We just need an undercutting defeater of the original reason.

Today is Tuesday, though, so the fact that the car is in the driveway is a reason to believe Terry took the bus. Nevertheless, Sam doesn't believe for the normative reason constituted by the fact that Terry's car is in the driveway because she would have believed as she does, and for *the very same reasons*, even if it were Wednesday and the relevant fact wasn't a reason. This is strong evidence that she is not sensitive to the relevant normative facts. This shows, I think, that despite the fact that Sam believes for a consideration that happens to be a normative reason, she doesn't believe for a normative reason.

If you don't find this immediately intuitive, fear not (if you think the fact that Sam is relying on the conditional belief is problematic, see notes 11 and 12). For I think two platitudes about ϕ -ing for normative reasons predict that Sam isn't believing for a normative reason. The first platitude is that when one ϕ -s for a normative reason r , one is sensitive to r *qua* normative reason. Call this The Sensitivity Platitude. The Sensitivity Platitude is a platitude because it is extremely plausible that attitudes and actions that are held or performed for normative reasons *track* the normative facts. Those are the features of the considerations we ϕ for that we seek to track when we ϕ for normative reasons. The rub is that Sam is *not* sensitive to the relevant fact *qua* reason. This is shown by the fact that she would continue to believe Terry took the bus, and for the same reason, even when it is carpool day. If this is right, then she doesn't believe for a normative reason.

The second platitude is that when one ϕ -s for a normative reason r , one's ϕ -ing is explained by r *qua* reason. This is plausible for similar reasons to the first platitude. When we ϕ for normative reasons, our ϕ -ing is explained by some normative facts. This, I think, follows from the fact that we are sensitive to the normative facts when we ϕ for normative reasons. Once again, though, Sam's belief is not explained by the relevant fact *qua* reason. Thus, it is not held for a normative reason.

What follows from the failure of Composite? The first thing that follows is that we need a separate analysis of what it is to act and believe for a normative reason. We can't get by just with analyses of normative reasons and acting and believing for reasons. This point has been, as far as I can see, unappreciated.

These cases not only show this, they also tell us a bit about what the right kind of view must look like.¹⁵ They show that the correct view must respect Prime:

Prime: What it is to ϕ for some normative reason r is not just a matter of ϕ -ing for r and r being a normative reason.

Prime on its own isn't very illuminating. It is, after all, just the negation of Composite. We do better by thinking of what failure the cases point to. They point to what all deviancy counterexamples point to, which is a failure of the analysis to account for some type of connection between the parts. The problem highlighted by the cases is that Composite doesn't specify what relation must obtain between the fact that A ϕ s and r being a normative reason to ϕ . Since it doesn't do this and there is more than one way one can ϕ for r , there are cases where both conditions of Composite are met that are intuitively not cases of acting or believing for normative reasons. Since it will be helpful later, let's say that an analysis is *prime* just in case the analysis holds that some connection must hold between the various parts of the analysis.

With this in hand, let's turn to the details and defense of my account of what it is to act and believe for normative reasons.

¹⁵I don't believe anyone has made this lesson explicit, but several accounts in the literature have the relevant feature (e.g., Wedgwood (2006), Arpaly (2006), Turri (2011)).

4.4 What it is to Act and Believe for Normative Reasons

The last section showed that acting and believing for normative reasons is a subject in its own right. Acting and believing for normative reasons isn't merely acting and believing for reasons plus something else—viz., the consideration for which one acts being a normative reason. In this section I will explicate and defend my preferred analyses of acting and believing for normative reasons.

The lesson we should draw from the counterexamples to Composite is that the correct analysis of acting for normative reasons holds that when A ϕ s for a normative reason r , there is some connection between the fact that A ϕ -ed and the fact that r is a normative reason for A to ϕ . I think the relevant connection is the causal *in virtue of* relation. Here are some examples of this relation. The towers collapsed in virtue of the heat (as opposed to the force of the collision). The grass grew in virtue of the rain (as opposed to the sprinklers). The computer broke in virtue of the water that was spilled on the keyboard.

I hold that Normative is the correct account of ϕ -ing for a normative reason:

Normative: What it is for A to ϕ for some normative reason r is for A to ϕ in virtue of the fact that r is a normative reason to ϕ .

The problem with Sam in Affirming the Consequent is that while she believes Terry took the bus because Terry's car is in the driveway, she doesn't believe Terry took the bus in virtue of the fact that Terry's car is in the driveway is a reason to believe Terry took the bus.¹⁶ We can see this when we notice that Sam would continue to hold the belief because Sam's car is in the driveway even if that fact were no longer a reason—e.g., when it is carpool day.

This is, I think, a very intuitive explanation of what's going wrong in cases like Affirming the Consequent. The short answer being that in these cases the subjects aren't sensitive to the fact that the fact is a reason to hold the belief or perform the action. However, this is clearly less than fully satisfying without some story of how the relevant *in virtue of* relation works. At least one theorist who defends this type of view refuses to give an account of this, claiming that it is a task for the metaphysicians (see Arpaly (2006)). This is clearly unacceptable once one realizes that the project of fleshing out what it is to act for normative reasons is *metaphysics*. We're all metaphysicians, whether we like it or not.

My view of the *in virtue of* relation is dispositional.¹⁷ Before I get to it, it's important to point out a key difference between beliefs, intentions, and desires (among other propositional attitudes) and actions. Beliefs, intentions, and desires are states, whereas actions are events. This leads to an important difference between beliefs, intentions, and desires on the one hand and actions on the other when it comes to believing (or intending or desiring) and acting for normative reasons.

¹⁶A note about the causal relation here. I am drawing on the idea that drives *aspect* theories of causation (see, e.g., Paul (2000)). Aspect theorists usually think causes are properties. While this is compatible with my claims, I don't think one needs to think this in order to hold my view. One could also think of the causes as facts (the fact that some fact is a reason to ϕ) or think in more purely dispositional terms (the disposition manifested is the disposition to...). I will use all three idioms in what follows. At the end of the day, I think that dispositions must play a role. Nevertheless, I think we can still incorporate properties or facts into the picture by having them play important roles in the manifestation conditions of the dispositions.

¹⁷There is a good reason why I don't jump directly to my dispositional account. I think that Normative would be true *even if* the following dispositional account failed. This is why I introduce Normative before explicating my dispositional theory of the *in virtue of* relation.

The difference is a difference in the type of causal relation involved. Attitudes are *sustained* by the reasons for which they are held, whereas actions are often *produced* by the reasons for which they are performed. Once my belief that p (e.g.,) is formed, any reason for which it is held will, at least partially, *sustain* it—the reason will help ensure its preservation. Actions, on the other hand, are produced (at least partially) by the reasons for which they are performed. The event that is the action is caused to come about by the reasons for which it is performed.

It's tempting to think that the way I drew this distinction is artificial, for clearly our attitudes are *sometimes* produced by the reasons for which they are held. For example, my belief that it is after 5:00 is caused by the reason provided by the fact that it's 5:45, for I inferred the belief from that reason. I obviously agree that inference has an important role to play in the correct theory of our epistemic agency. But this is because inference is a type of activity that we engage in for a reason. In this case the event of *forming the belief that it's after 5:00* is caused by the fact that the fact that it's 5:45 is a reason to believe it's after 5:00. Forming the belief is, I think, the result of a mental activity that is one engages in for a reason.¹⁸ Indeed, I think something stronger is plausible. I think it's plausible that the inferences are *mental actions*. Thus, my account of what it is to act for a reason should tell a story about it, and it does. My forming the belief that it is after 5:00 is done for the normative reason constituted by the fact that it's 5:45 just in case it is done in virtue of the fact that the fact that it's 5:45 is a reason to believe it's after 5:00. So inference is an action that is done for a reason. In typical cases, the reason for which one performs the inferential action will also be the reason for which the belief is held. But this will be true, I claim, only if that reason sustains the belief. Being the original cause of the belief via inference is not sufficient for being the reason for which the belief is held.¹⁹

Let me register recognition that it is controversial whether inference is an action or merely an activity. It is not controversial that an activity is involved. Indeed, it is central to most of the going views about inference that inference is an activity.²⁰ But it is a jump from this to the claim that inferences are actions. I will explore this in future work.²¹ In any case, I am convinced that the reasons for which we believe are sustainers and that the reasons for which we make inferences need not always be sustainers. Because of this, I think that inferences play an important but defeasible role in our believing-for-reasons.²²

Notice that I only said that our acts are *often* produced by the reasons for which they are done. This is because it is clear that our reasons for performing some temporally extended actions can

¹⁸With an emphasis here on *forming*. Beliefs that pop into one heads for wacky reasons (e.g., because one is hit in the head) are not formed, in this sense.

¹⁹For counterexamples to the claim that being the original cause of the belief via inference is sufficient for being the reason for which the belief is held, see the Lottery and Red Widgets cases in §6.5 of chapter 6.

²⁰This comes out very nicely in the excellent recent symposium on the topic in *Philosophical Studies*. See Boghossian (FC), Broome (FC), Wright (FC).

²¹Just a taste for enthusiasts. The main reason for thinking inferences aren't actions is that actions have to be done for practical reasons and one cannot perform inferences for practical reasons (this is suggested in Kelly (2002)). I am moved by this thought. However, there is something incredibly similar about the moves one makes under the veil of supposition and inference. It is hard for me to get myself to think that in one case actions are being performed and in the other actions are not being performed. At the very least, the similarities beg to be explained. I think that it is very plausible that even if we ultimately conclude that inferences aren't actions, we should think they are events done for reasons. For theorists who think inferences are actions, see Arpaly & Schroeder (2012), Hookway (1999).

²²That said, I think the reasons for which we make inferences can justify *judgments*, where judgments are a certain kind of conscious state. The relationship between judgments and beliefs is murky. This will also be a topic for future research.

change over the course of the action. For example, I might start walking home because my wife told me I needed to come home at once. This might have been the original cause of my act of walking home. But halfway through she might call back to tell me that the crisis has been averted, but that my favorite chocolate just arrived in the mail. My reason for walking home might then change. I might cease walking home because she told me to and start walking home because my favorite chocolate arrived. The reason provided by the chocolate is not the original cause of the act of walking home (or at least this is plausible). But it does become the reason for which I'm walking home.

I don't think we can have a complete understanding of these types of cases without settling other issues in the philosophy of action that I want to be neutral about here. On some views—views on which actions are individuated by the intentions they are performed for—there isn't a single act of walking home here. There are two.²³ This view thus seems compatible with holding that all the normative reasons for which we act are original causes of the actions.

I'm inclined to think that the actions that are produced by the normative reasons for which they are done are the *basic actions*—the actions that are done directly.²⁴ Non-basic actions—e.g., walking home—can be sustained by a normative reason *r* via *r* being the reason for which a series of basic actions is done. But, just to emphasize, I don't want to take a stand on these dicey issues about the ontology of actions. What is important for our purposes is that at least some actions are produced by the normative reasons for which they are done, whereas I don't think beliefs ever are. Thus, I would like to use the limited space I have here to analyze what is going on when actions are produced by the acts for which they are done.²⁵

The distinction between sustaining and production is important in this context because whether something is an attitude or an action will make a difference to what dispositions one must have in order to count as acting or believing in virtue of the fact that something is a normative reason. The dispositions involved in sustaining are different than the dispositions involved in producing. Let's stick to belief and action and take them one at a time.

Beliefs are sustained or preserved by normative reasons when they are held for normative reasons. That it is to say, what it is for my belief that *q* to be held in virtue of the fact that *p* is a normative reason to believe *q* is for it to be sustained by the fact that *p* is a normative reason to believe *q*. This just pushes the issue back onto sustaining. Luckily for me (and for you!), I have an account of how sustaining works. It is a dispositional account. To see how it works, notice that a common feature of normative reasons is that the facts that constitute them aren't always reasons in worlds in which they obtain.²⁶ This phenomenon is illustrated by Affirming the Consequent. On Wednesdays the fact that Terry's car is in the driveway is no reason to believe she took the bus. The fact that she car-

²³Other views also seem to have this result. For example, the view of Marcus (2012).

²⁴For some literature on basic actions, see Enc (2003), Hornsby (2005), Lavin (FC).

²⁵Even if you disagree about some of the more fine grained details of the previous few paragraphs, it is easy enough to see how to modify my analyses in small ways in order to suit your tastes. For example, if you think that beliefs can be produced by the reasons for which they are held, then it is easy to see how to modify my analysis to reflect this (similarly for the case of action).

²⁶Of course, this is not to say that there aren't some facts that are necessarily reasons for certain things. It's plausible that some facts are (to use terminology I'm about to introduce in the main text) reasons that are not undercuttable. I of course think you can act and believe for those reasons. I'm thus committed to thinking that one can have dispositions that can never be manifested (*viz.*, the disposition to revise if the reason is undercut). Fortunately, I think everyone has to think that there are dispositions that can't possibly be manifested. For a plethora of examples, see Jenkins & Nolan (FC).

pools on Wednesday is, to use terminology originating in Pollock (1974), an undercutting defeater of the reason once provided by the fact that Terry's car is in the driveway. It makes it the case that there is no reason provided by the fact that Terry's car is in the driveway.

With this phenomenon in mind, we can get a grip on what the relevant dispositions need to be. In order for some normative reason p to sustain a belief of A 's that q , it needs to be that A is disposed to revise q if p ceases to be a reason to believe q .²⁷ If A has this disposition, then her belief that q is intuitively sensitive to the fact p is a reason to believe q . This is what we want.

Now on to action. The sentiment is the same even if the details are different. The details are different because actions are events that are produced rather than states that are sustained. Actions are caused or produced by normative reasons when they are done for normative reasons. That is, what it is for an action ϕ to be done in virtue of the fact that r is a normative reason to ϕ is for ϕ to be caused by the fact that r is a normative reason to ϕ . Just like in the case of belief, this just pushes the issue back. But I also have a view of what it is for ϕ to be caused by the fact that r is a normative reason. It's also a dispositional account.²⁸

When one ϕ s, one manifests some disposition, which is just to say that there was some state of the agent that led to ϕ -ing when combined with some manifestation conditions. When our acts are done for normative reasons, they are the result of a certain kind of disposition manifesting. These dispositions are sensitive to normative reasons *qua* normative reasons. When A ϕ s for the normative reason r , one's ϕ -ing is the manifestation of a disposition to ϕ when r is a reason to ϕ . Thus, when A ϕ s for some normative reason r , it's (usually) true of A that if r hadn't been a normative reason to ϕ , A wouldn't have ϕ -ed.²⁹ If A 's ϕ -ing is a manifestation of such a disposition, then intuitively A is sensitive to the fact that r is a normative reason. This is precisely what we want.

One objection to the way I've set things up holds that I've just reintroduced the deviancy problem by ultimately analyzing what's going on in terms of sustaining and production. This is not right, though. What I've done is *isolated* where the problem lies. The solution to the deviancy problem is found by figuring out which feature of the fact doing the causal work must do the causing in order for it to be causation of the right kind. I think I have given an adequate account of this in dispositional terms.

The fully fleshed out view of what it is to believe for a normative reason r is given by Normative-Belief and the fully fleshed out view of what it is to act for a normative reason r is given by Normative-Act:

Normative–Belief: What it is for A to believe p for some normative reason r is for A to be disposed to revise p if r ceases to be a normative reason to believe p .

Normative–Action: What it is for A to ϕ for a normative reason r is for A 's ϕ -ing to be the manifestation of a disposition to ϕ when the fact that constitutes r is a normative reason to ϕ .

²⁷Remember that we are limiting ourselves to the reasons that are possessed. Because of this, I assume that *ceasing to possess* is a kind of undercutting defeat. The fact that constitutes the reason needn't cease being a reason when you cease to possess it, but it does cease to be a reason *you possess*. One needs to be sensitive to this in the same way one needs to be sensitive to traditional undercutting defeaters in order to believe (and act) for normative reasons.

²⁸I'm deeply indebted to Wedgwood (2006, 2007) here.

²⁹Two things to note about this. First, it's not meant to imply a counterfactual account of dispositions. I think all such accounts are false. Second, despite the falsity of counterfactual accounts, the counterfactuals are usually true when one has the dispositions. This makes the counterfactuals a good (but fallible) guide to whether one has the relevant dispositions.

4.4.1 Why This is, Alas, Not Enough to Get All that We Want

I think the preceding accounts are correct accounts of what it is to act and believe for normative reasons. As we saw above, one traditional motivation for giving such accounts was to figure out when a belief or action is justified, rational, blameworthy, or creditworthy. The common thought is that it's not enough to be justified, rational, or creditworthy for an act or belief to be supported by normative reasons. One also had to perform the act or hold the belief for the normative reasons.

This is the distinction between *ex ante* rationality and *ex post* rationality introduced above. ϕ -ing is *ex ante* rational (or justified or praiseworthy) just in case ϕ -ing is supported by sufficient normative reasons. One's ϕ -ing is *ex post* rational just in case ϕ -ing is *ex ante* rational and one ϕ -s for the reasons that make ϕ -ing *ex ante* rational. Giving a successful account of *ex post* rationality is the main motivation for almost all discussions of believing for normative reasons found in the epistemology literature.

The preceding accounts aren't enough to give us a successful account of *ex post* rationality. The problem stems from the fact that when ϕ -ing is *ex ante* rational, it is supported by *sufficiently strong* normative reasons. Not all acts and beliefs supported by normative reasons are rational. Only the ones that are supported by sufficiently strong normative reasons are.

This will lead to yet another batch of deviancy counterexamples—cases where one acts or believes for some normative reason r , but doesn't act or believe for a sufficient normative reason. To see this, let's modify Affirming the Consequent a bit. Suppose that Sam knows that Terry's car is in the driveway and she believes Terry took the bus to work. Moreover, she is disposed to revise her belief that Terry took the bus if the fact that Terry's car is in the driveway ceases to be a reason to believe Terry took the bus—e.g., if it's Wednesday. Thus, Sam believes Terry took the bus for the normative reason provided by the fact that Terry's car is in the driveway. However, it doesn't follow that she believes Terry took the bus for the *sufficient* normative reason provided by the fact that Terry's car is in the driveway. Home Sick shows this:

Home Sick

In the actual world, the fact that Terry's car is in the driveway is a sufficient reason to believe that she took the bus. However, consider the world where it isn't a sufficient reason even though it's still a reason. Suppose Sam calls Terry's office to inquire about dinner. Terry's secretary tells Sam that Terry is home sick (and has been all day). This defeats the reason provided by the fact that Terry's car is in the driveway. However, suppose that Sam, despite being disposed to revise her belief if the reason ceased being a reason, isn't disposed to revise when she finds out this new information, and thus does nothing.

In Home Sick, Sam does in fact believe for a normative reason that happens to be sufficient. This is because she believes for a normative reason and, in the actual world, that reason is sufficient. However, it doesn't seem like she believes for a sufficiently strong normative reason. This is because she isn't sensitive to the fact that it's a sufficiently strong normative reason, as is shown by her behavior in the world where the reason in question stops being sufficiently strong. Thus, believing for a normative reason that happens to be sufficiently strong isn't enough for believing for a sufficiently strong normative reason.

Just to be clear, in Home Sick Sam is sensitive to undercutting defeaters—this is what it is to believe for a normative reason. She is, in other words, sensitive to whether the fact in question is a reason. But she's not sensitive to the fact that the normative reason for which she believes is sufficient. We can see this by noting that she is not disposed to revise her belief when the reason is outweighed. It seems like you need to be sensitive to both kinds of defeat in order to believe for a sufficient normative reason.

The form the proper response should take is exactly the same as our response to the deviancy problems from §3.3. The problem with the view that what it is to act for a sufficient normative reason is to act for a normative reason that happens to be sufficient is the same problem we saw for the account that holds that what it is to act for a normative reason is to act for a reason that happens to be normative—they are both composite accounts. The proper way to solve the deviancy problems in both cases is to move to a prime account.

Just as in the previous section, I think the correct account holds that when A ϕ s for a sufficiently strong normative reason r , there is a tight connection between the fact that A ϕ s and the fact that r is a sufficiently strong normative reason to ϕ . On my view, the relation is the causal *in virtue of* relation. I think Sufficient Normative is the correct analysis of acting for sufficiently strong normative reason:

Sufficient Normative: What it is for A to ϕ for some sufficient normative reason r is for A to ϕ in virtue of the fact that r is a sufficient normative reason to ϕ .

I spell out this *in virtue of* relation dispositionally, just as before. The only difference being that the disposition involved with acting for a sufficient normative reason is not only sensitive to whether the consideration in question is a reason—i.e., to whether it has been undercut—but also to whether the reason is a sufficient reason—i.e., it is sensitive to whether it has been outweighed by stronger reasons on the other side. Thus, when one acts or believes for sufficient normative reasons, one is sensitive to both kinds of defeaters, which is intuitively what we want from an agent who acts and believes rationally or justifiably.

The difference between states like beliefs and actions applies here, as well. States are sustained whereas actions are produced. Thus, the relevant dispositions will be different depending on whether it is a state or an action. Rather than going through all the details of this again, I will just state my analyses of what it is to believe for sufficient normative reasons and act for sufficient normative reasons.

Sufficient Normative–Belief: What it is for A to believe p for some sufficient normative reason r is for A to be disposed to revise p if r ceases to be a sufficient normative reason to believe p .

Sufficient Normative–Action: What it is for A to ϕ for a sufficient normative reason r is for A 's ϕ -ing to be the manifestation of a disposition to ϕ when the fact that constitutes r is a sufficient normative reason to ϕ .

On my view, you have *ex post* rational or justified beliefs only when you meet the conditions of Sufficient Normative–Belief and your actions are *ex post* rational or justified only when you meet the conditions of Sufficient Normative–Action.³⁰ These views have at least two things going for

³⁰There are serious objections to this account of *ex post* rationality. One serious objection has to do with forgotten evidence. I offer a response to this problem in chapter 5, along with responses to several other problems with the account.

them, at this point (I will point out several more virtues in the final section). First, they capture the extremely natural thought that when you act or believe for sufficient normative reasons, your acts and beliefs are sensitive to defeaters of those reasons. Second, they solve the deviancy problems in precisely the way my more general diagnosis of what's going on in the deviancy examples suggests.³¹

4.5 What it is to Act and Believe for Reasons

In §4.3.3 we saw that acting and believing for normative reasons is a different topic from acting and believing for reasons—one is not parasitic on the other. In the last section I offered my analyses of what it is believe for a normative reason and what it is act for a normative reason. In this section I'm going to explicate my analyses of what it is to act for a reason and believe for a reason. I'll argue that these analyses solve the deviancy problems found in §4.3.2. First, let's remind ourselves of those cases. I gave two deviancy counterexamples for causal accounts of acting for reasons, the first from Chisolm (1966) and the second from Davidson (1980):

Dead Uncle

Dan has decided to kill his uncle Fred because he is set to inherit a vast amount of money from his uncle. The thought of inheriting the money excites him so much that he loses control of his car and kills a pedestrian. The pedestrian happens to be his uncle. He inherits the fortune.

Dead Climber

Dan and Fred are climbing an arduous mountain together. Dan decides that he will only survive the climb if he rids himself of the burden Fred brings to the project. He thus believes that he must get rid of Fred in order to survive. This thought so terrifies him, though, that it causes him to loosen his grip on the rope that is supporting Fred. Fred falls to his death. Dan survives.

In both cases there is some consideration that causes some event even though it doesn't seem like the event caused is an action done for that consideration. The trick, of course, is to explain the difference between these cases and cases where Dan does perform the actions for the relevant reasons.

I gave one case of a deviant causal chain in the case of belief from Pollock & Cruz (1999):

Late for Class

Joe is late for class. This causes him to quicken his pace. This causes him to slip, which causes him to fall on his back and look skyward. This causes him to believe there are birds in the tree.

The structure is the same. There is some consideration that causes Joe's belief that there are birds in the tree, but intuitively that consideration is not a reason for which Joe believes. The trick is to

³¹Those who are mainly interested in normative matters could now skip to §4.7 without loss. Readers wanting to read a shorter version of the chapter are also encouraged to do this.

explain the difference between cases like Late for Class and a case where Joe does hold the belief for the relevant consideration.³²

It's trickier to see exactly what's going wrong in these cases than in the deviancy cases for causal accounts of ϕ -ing for normative reasons. The main reason for this, I think, is that in the cases of acting for normative reasons, there is an independent property the facts have—the property of being a normative reason—that one needs to stand in a certain relation to in order to act or believe for it. The normative reason is always going to be there, by hypothesis. This gives us something to work with. Now, of course, there is also some independent entity—the consideration—in these more basic cases. But the analogous move in these cases—to say that one must act in virtue of the fact that the consideration is a consideration—is unilluminating. Something else must be going on.³³

The key, I think, is that in non-deviant cases the agent is acting and believing in virtue of the fact that the consideration has a certain status in one's psychology. The status, I claim, is that of *being treated as a normative reason to hold the belief or perform the action*.^{34, 35} When you ϕ for a reason r , you ϕ in virtue of the fact that you treat r as a normative reason to ϕ .

The last chapter provided an analysis of what it is to treat something as a reason. I don't think I need to rely on all the details of that analysis here. I will instead point to two of the core features of treating.

The first feature is that one must believe the propositions one treats as normative reasons. Of course, this isn't to say that we aren't *disposed* to treat lots of things we don't believe as normative reasons. We conditionally treat lots of things as normative reasons, conditioned on them obtaining (the manifestation conditions of these disposition are just us believing they obtain).

The second, more substantial, condition is that when we treat p as a normative reason to ϕ , we are disposed to think p is positively relevant to the question of whether to ϕ when we're reasoning about the merits of ϕ -ing under certain conditions. The conditions are, roughly, what Gibbard (1990) calls *normative discussion*. Normative discussions are unconstrained in certain ways. For example, in normative discussion social pressures against seeing certain things as salient are lifted. Moreover, in normative discussions we are trying to discover *all* of the merits of ϕ -ing. Thus, there are not the usual pragmatic pressures to only reason about the considerations that are especially

³²It's harder to imagine the good case with this example. Here's one: Joe's class is held during lunch, when a lot of the university members feed the birds. The birds are thus usually not in the trees during this time because they are on the ground to get the food. Joe notices he's late for class (and lunch is thus over). He infers from this that the birds are back in the trees.

³³Note that I'm going start using 'consideration' more than 'reason' so as to not confuse the things for which we act when we merely act for reasons with the things for which we act when we act for normative reasons. I might occasionally use 'reason.' When you see 'reason' unadorned by 'normative,' I am talking about considerations.

³⁴Schossler (2007, 2012) takes the notion of treating as a reason as central to his account of intentional action. He doesn't use the notion to solve the deviancy problem, though. On his view, the deviancy problem is a separate issue. I think this makes the view border on circularity. Moreover, his solution to the deviancy problem appeals to the types of dispositions I think must hold to act and believe for normative reasons. I think this will not explain all the cases of acting for reasons. In short, I think his view ends up extensionally inadequate because he neglects to notice that acting and believing for reasons is a different topic than acting and believing for normative reasons.

³⁵Many others have thought treating as a normative reason plays an important role in the correct account of believing and acting for normative reasons. See, e.g., Milligan (1974), Velleman (2000), Broome (1997), Schroeder (2007). Setiya (2007b) argues against the claim that treating as a normative reason is necessary, but still appeals to treating as a reason. As far as I know, no one has used the notion the way I do.

weighty.³⁶ When you treat p as a normative reason to ϕ , you are disposed to think that p is relevant to the question 'why ϕ ?' in normative discussion. Of course, it should be noted that for many of the things we treat as normative reasons, we are also disposed to think they are relevant to reasoning about ϕ -ing in ordinary circumstances.

The million dollar question at this point is what this *thinking p is relevant* comes down to.³⁷ On the most natural interpretation, it involves a belief with a normative content (or something close to that)—e.g., that p counts in favor of ϕ -ing, or p is a reason to ϕ . Some theorists are attracted to this (e.g., Scanlon (1998), Milligan (1974), Broome (1997)). According to some other theorists, it amounts to a completely separate—perhaps *sui generis*—psychological state. Marcus (2012) prefers to describe the state as representing *that p is to be believed as a consequence of the fact that q is to be believed*.³⁸ According to a third group of views, the state one enters into is just a belief that p , but one that has extra phenomenal character—that of counting in favor of ϕ -ing (see Dorsch (FC)). On this view, what it's like to believe that p 'feels' differently when you treat p as a reason to ϕ than when you don't.

As we saw in the last two chapters, I think the first view is misguided. This is because it over-intellectualizes. It doesn't seem like we need to have or be disposed to have beliefs with explicit normative contents in order to treat things as normative reasons. This would bar those that lack enough of the concept of a normative reason to have beliefs with explicit normative contents as being able to treat things as normative reasons. Since I think it's clear that small children and many animals do treat things as normative reasons, this seems like a decisive objection.

I'm not sure what to make of the second view. It all depends on what the new state is like. I think that Marcus (2012)'s version of the view overintellectualizes in several ways and is no real improvement over the first type of view. However, it's possible to have this type of view without overintellectualizing. My worry with this, though, is that it feels like we're just stipulating into existence a state with the right properties. Perhaps treating is non-reducible to other states, but the burden for showing this is quite heavy.

For these reasons, I think the third view is quite attractive. It at least has the right formal features. That is, I think it avoids overintellectualization while also providing plausible explanations of the core features of treating. Although I don't want to commit myself to the view here, I think it is the the best option of the three. Nevertheless, despite this philosophical puzzle about some features of treating, I don't think anyone should be skeptical that there is such a phenomenon. We are all intimately familiar with the phenomenon (even when we reason about whether such a phenomenon exists!).

I think it's no surprise that the notion of treating would play a pivotal role in the correct account of acting and believing for reasons. Preliminary evidence for this is the fact that the considerations we see as salient when reasoning about what to do and believe are, if things go well, the considerations we believe and act for.

There is a composite way to incorporate the notion of treating into an analysis of ϕ -ing for a reason. According to this view, what it is to ϕ for some reason r is for r to cause one to ϕ and

³⁶See Schroeder (2007) for an excellent cataloguing of the effect such pragmatic features have on our normal discourse about reasons.

³⁷This is going beyond the discussion in the last chapter. In the terms of the last chapter, this is a deeper investigation into what the Normative Condition on treating is like.

³⁸He argues that this can't just be a doxastic attitude we already posit and thus that it is a new state.

for one to treat r as a normative reason to ϕ . Unsurprisingly, this view leads to further deviancy problems, as is shown by the following example (inspired by (Schroeder, 2007, chapter 8)).

Virus

Joe likes to dance and thus is disposed to treat the fact that there will be dancing at Liz's party tonight as a reason to go to Liz's house (when he believes that Liz is throwing a party with dancing). However, he also has a bizarre virus that compulsively causes him to go to Liz's house whenever he believes there is a party at Liz's house. He comes to believe that there is a party at Liz's house (and that there will be dancing at said party). The virus causes him to compulsively head to Liz's house. The fact that he treats the dancing as a reason to go plays no role in the etiology of his behaviour.

Joe meets the two conditions laid down by the composite account. He is caused to go to Liz's house by a consideration that he happens to treat as a normative reason to go to Liz's house. However, he intuitively doesn't act for that reason—perhaps he doesn't act at all. So the composite version of the view must be false.

No need to worry, though, for there is a prime version of the view that doesn't have this problem. The form of the view is by now familiar. What needs to happen for Joe to act for the reason is that he needs to go to Liz's house in virtue of the fact that he treats the dancing as a normative reason to go to Liz's house. This isn't true of Joe in Virus. The fact that he treats the consideration as a normative reason is idle.

This nicely highlights the basic idea behind my general account of what it is to ϕ for a consideration r .

Consideration: What it is for A to ϕ for some consideration r is for A to ϕ in virtue of the fact that A treats r as a normative reason to ϕ .

In order for this to be fully satisfying, an account of the *in virtue of* relation must be given. Once again, I think it should be fleshed out dispositionally. And once again, the differences between states like beliefs and events like actions make a difference to what the relevant dispositions are. When you ϕ in virtue of the fact that you treat r as a normative reason to ϕ , your action is the result of a disposition that manifests when you treat r as a normative reason to ϕ . When you believe q in virtue of the fact that you treat p as a normative reason to believe q , you are disposed to revise q if you cease treating p as a normative reason to believe q . This is just to say that Consideration–Belief and Consideration–Action are true:

Consideration–Belief: What it is for A to believe q for a consideration p is for A to be disposed to revise the belief that q if A ceases to treat p as a normative reason to believe q .

Consideration–Action: What it is for A to ϕ for a consideration r is for A 's ϕ -ing to be the manifestation of a disposition to ϕ when A treats r as a normative reason to ϕ .

These accounts solve the deviancy problems. Dan does not hit the pedestrian in virtue of the fact that he treats the fact that he will inherit a large sum from his uncle if his uncle dies as a normative reason to kill the pedestrian. Likewise, Dan does not loosen his grip in virtue of the fact that he

treats the fact that he will survive only if his partner dies as a normative reason to loosen his grip. Moreover, Joe does not believe there are birds in the trees in virtue of the fact that he treats the consideration that he is late for class as a normative reason to believe there are birds in the trees. The dispositions that manifest when Dan behaves in those ways and when Joe forms that belief have nothing to do with treating the considerations that cause the behavior and belief as normative reasons. In Dan's cases, it's his excitement and terror, respectively, and in Joe's case it's his disposition to treat the contents of visual perceptions as reasons.³⁹

4.5.1 Again, This is Not Enough

At the end of the last section we saw that we couldn't capture everything we want to capture with a theory of ϕ -ing for normative reasons without an analysis of what it is to ϕ for sufficient normative reasons. Moreover, we saw that ϕ -ing for sufficient normative reasons isn't the same as ϕ -ing for normative reasons that are sufficient. More needs to happen in order to ϕ for sufficient normative reasons.

Something similar happens when it comes to ϕ -ing for reasons. To see this, compare Jack to Jill. Jack treats the consideration that smoking is harmful to his health as a normative reason to not smoke. Moreover, he doesn't smoke for that reason. Not only this, but Jack treats this as a sufficient normative reason not to smoke. This means, at first pass, that Jack is disposed to see that consideration as ending deliberation under circumstances of normative discussion. Jill, on the other hand, treats the fact that smoking harms her health as a normative reason not to smoke and she doesn't smoke for that reason. However, she doesn't treat that consideration as a sufficient normative reason. She treats it as a normative reason, but not as one that is weighty enough to end deliberation.

Something is going wrong with Jill that isn't going wrong with Jack. Jill acts for a reason she doesn't treat as weighty enough to settle deliberation. Jack, on the other hand, does. Other things being equal, it doesn't seem like you should move from being in Jill's position to acting and it does seem like you should move from being in Jack's position to acting. So it seems like we need an analysis of what Jack is doing, for at this point we just have an analysis of what Jill is doing.

What Jack is doing, I think, is refraining from smoking in virtue of the fact that he treats the fact that smoking is harmful to his health as a sufficient normative reason not to smoke. Just like before, this *in virtue of* relation is spelled out dispositionally. The reason why Jack acts in virtue of the fact that he treats certain considerations as sufficient normative reasons is that Jack's act is the manifestation of a disposition that manifests when Jack treats that consideration as a sufficient normative reason to perform that act.

This leads to Sufficient Consideration–Action and Sufficient Consideration–Belief:

Sufficient Consideration–Belief: What it is for A to believe p for a sufficient consideration r is for A to be disposed to revise p if A ceases to treat r as a sufficient normative reason to believe p .

³⁹A key difference between Dan's cases and Joe's is that in Dan's cases there doesn't seem to be an action at all, and hence there is no action done for a reason. In Joe's case, by contrast, there is a belief that is presumably held for reasons.

Sufficient Consideration–Action: What it is for A to ϕ for a sufficient consideration r is for A 's ϕ -ing to be the manifestation of a disposition to ϕ when A treats r as a sufficient normative reason to ϕ .⁴⁰

This concludes the part of the chapter that explicates my analyses and argues they can solve deviancy problems. I think that I've provided a general diagnosis of what leads to deviancy problems and argued for analyses that avoid the relevant features. Thus, I think I've provided analyses that have the virtue of solving the deviancy problems while at the same time being independently attractive. In the final part of the chapter, I'll spell out three ways in which these analyses are even more attractive, by showing that they can help us make progress on several other debates across philosophy. Before getting to that, though, I should explain the relationship(s) between acting and believing for normative reasons and acting and believing for reasons.

4.6 The Connection Between Acting for Reasons and Acting for Normative Reasons

As we've seen, I deny that there is any parasitic relation between ϕ -ing for reasons and ϕ -ing for normative reasons. This entails that you stand in a different relation to normative reasons when you ϕ for them than you do for considerations when you ϕ for them. I think this is forced upon us by cases like Affirming the Consequent. The only way to be sensitive to normative reasons *qua* normative reasons is to be sensitive to the fact that they are normative reasons. But one needn't be sensitive in this way to ϕ for reasons. So it can't be the same relation.

Despite this, it sure *seems* like it's the same relation. It's plausible that, from the first person perspective, the relation seems to be the same. In other words, when you act for reasons, it seems like you act for normative reasons. And when you act for normative reasons, it seems like you act for reasons. Since I deny that they are the same relation, I must explain why these things seem the way they do. I'll take them one at a time.

4.6.1 The Transparency of Acting for Reasons

We'll start with what I take to be the most important piece of data, which I'll call the Transparency of Acting for Reasons:

Transparency of Acting for Reasons: When one ϕ s for a consideration p , it seems as if one is ϕ -ing for a normative reason p .⁴¹

Views that hold that acting for normative reasons involves acting for reasons have an easy time explaining the Transparency of Acting for Reasons. Acting for reasons is transparent because you

⁴⁰Both of these analyses are prime. It's interesting to note that in this case, unlike in the others, a composite account is not at all tempting. A composite account holds that acting for a sufficient consideration is a matter of acting for a consideration that happens to be sufficient, where the fact that you act for the consideration and the fact that it is sufficient are not related. If the two conditions were independent, then we would have a new batch of deviancy cases. I'll leave it as an exercise for the reader to construct a case of this kind.

⁴¹Although this claim is intuitively plausible and endorsed by many, it should be noted that it is controversial. Setiya (2007b) offers purported counterexamples.

stand in the same relation to the considerations on which you act in both cases. My view can't give this explanation.

Regardless, my view still has an elegant explanation of the transparency of acting for reasons. Recall that in order to ϕ for a consideration p , on my view, you have to treat p as a normative reason to ϕ . This is why acting for reasons is transparent. It's transparent because acting for reasons essentially involves treating those considerations as normative reasons to act the way you do.

Indeed, I think that no matter what you think about the relationship between acting for reasons and acting for normative reasons, you still need to explain the transparency of acting for reasons in this way. That is to say, I think that my account of acting for reasons is correct, even if acting for normative reasons is somehow parasitic on acting for reasons.

4.6.2 The Transparency of Acting for Normative Reasons

Not only does it seem like you act for normative reasons when you act for reasons, it seems like you act for reasons when you act for normative reasons. In other words, Transparency of Acting for Normative Reasons is true:

Transparency of Acting for Normative Reasons: When one ϕ s for a normative reason p , it seems as if one is ϕ -ing for a consideration p .

Unlike in the last case, it's tempting to think that it's necessarily true that whenever you ϕ for a normative reason p , you are ϕ -ing for a consideration p . I think we should give in to this temptation. Thus, I think that it seems like we ϕ for reasons when we ϕ for normative reasons because we actually do. Unfortunately, nothing I have said so far guarantees this. I thus have to supplement the account to show this. Fortunately, I think the account can be supplemented by independently plausible assumptions that entail that whenever one ϕ -s for a normative reason r , one ϕ -s for a consideration r .

There are three important claims. First, in order to ϕ for a normative reason r , one needs to possess r as a reason to ϕ . Second, in order to possess r as a reason to ϕ , one must treat r as a reason to ϕ . Third and finally, not only does one need to treat r as a reason to ϕ , one must treat r as a reason to ϕ in virtue of the fact that r is a reason to ϕ . I'll briefly defend each claim here.⁴²

It is very plausible that in order to ϕ for a normative reason r , one needs to possess r as a reason.⁴³ This is because it is plausible that in order for r to be the normative reason for which you act or believe, r must be within your epistemic ken. If you are ignorant of r , then r cannot be the normative reason for which you act or believe. Part of what it is to possess r as a normative reason to ϕ is for r to be in your epistemic ken. The claim that possession is necessary in order to be able to ϕ for r explains why ignorance precludes one from being able to ϕ for r .

However, as we saw in both chapter 2 and chapter 3, standing in some privileged epistemic relation to r isn't sufficient for possessing r as a normative reason to ϕ . We often know facts that are reasons to ϕ but fail to possess them as reasons to ϕ because we fail to see the normative connection. To adapt the case from earlier chapters, suppose that there is a scientific consensus that some

⁴²For a defense of the first claim see chapter 7 and for a defense of the second claim see chapter 3.

⁴³The claim is often endorsed (e.g., by Turri (2009, 2010)). Moreover, it is often a background assumption (e.g., in Schroeder (2007, 2009b), Hornsby (2008), Hyman (2006), Gibbons (2001)).

bacteria found in your favorite cheese is harmless. You thus don't treat the presence of the bacteria as a reason to refrain from eating the cheese. As it turns out, the scientists are wrong. Given certain background conditions, the bacteria is harmful. Moreover, the background conditions are met in your case. Thus, the fact that the cheese contains that bacteria is in fact a reason not to eat the cheese. Intuitively, though, you don't possess that fact as a (normative) reason not to eat the cheese. This is because, I claim, you don't treat that fact as a normative reason not to eat the cheese. This provides good reason to think that in order to possess r as a normative reason to ϕ , one must treat r as a normative reason to ϕ .

Finally, a brief defense of the claim that in order to possess r as a normative reason to ϕ , one must treat r as a normative reason to ϕ in virtue of the fact that r is a normative reason to ϕ . In the context of this chapter, this claim shouldn't be surprising. A major lesson from cases like Affirming the Consequent is that one can treat some fact as a normative reason to ϕ in many different ways. Some of those ways are seriously disconnected from the actual normative facts. It is quite plausible, I think, that when one treats some fact as a normative reason to ϕ in these ways, one doesn't possess that fact as a normative reason. One needs to track the normative facts in order to possess the reason.

If one needs to treat r as a normative reason in order to possess r as a normative reason and one needs to possess r in order to ϕ for r , then it follows that when one ϕ -s for a normative reason r , one will be sensitive to whether one treats r as a normative reason. For example, if one believes q for some normative reason p , one's belief that q will be sensitive to whether one possesses p . It follows that one will be sensitive to whether one treats p as a normative reason to believe q . For if one ceases to treat p as a normative reason to believe q , one will cease possessing p . Since one believes for a normative reason, one is disposed to revise the belief that q when one ceases to possess the normative reason p . But this sensitivity—a sensitivity to whether one treats p as a normative reason—just is what it is to believe for that reason. Thus, when one believes for a normative reason, one also believes for a reason. This just is Transparency of Acting for Normative Reasons.

4.7 Further Upshots

I think there are several upshots of the views defended here, over and above the fact that they can solve the deviancy problems. I'll start with two debates in epistemology before moving on to two debates in practical philosophy.

4.7.1 The Relationship Between *Ex Post* and *Ex Ante* Justification

A common thought is that we can understand *ex post* rationality in terms of *ex ante* rationality plus believing for reasons.⁴⁴ That is to say, what it is for A 's belief that p to be *ex post* rational is for the proposition believed to be *ex ante* rational—i.e., for there to be some set of reasons S that sufficiently supports p —and for A to hold her belief that p for the members of S .

On the common picture, then, *ex post* rationality can be understood just in terms of *ex ante* rationality and believing for (normative) reasons. Two of the more forceful challenges to this traditional view, it turns out, try to create problems for it by appealing to cases with the same structure

⁴⁴In epistemology, the focus is usually on justification rather than rationality. I will follow the majority of epistemologists in thinking these are the same. If they aren't, you can read this sub-section as being just about justification.

as Affirming the Consequent. These are the arguments of Millar (1991) and Turri (2010). Millar argues in light of the cases that we need some other condition above *ex ante* rationality and believing for the reasons that confer *ex ante* rationality, while Turri argues that we need to analyze *ex ante* rationality in terms of *ex post* rationality. Both Millar's and Turri's arguments share a common premise. Namely, that believing for a normative reason is a matter of believing for a reason that happens to be normative—both assume that the correct view of believing for normative reasons is a composite one.

Turri appeals to two cases.⁴⁵ The second case he uses is very close to Affirming the Consequent. We'll call it Spurs Win (see pg. 317 of Turri (2010)):

Spurs Win

Both Mr. Ponens and Mr. F.A. Lacy know that the Spurs will win if they play the Pistons and know that the Spurs will play the Pistons. They are both thus propositionally [i.e., *ex ante*] justified in believing that the Spurs will win. They also both infer that the Spurs will win from those two propositions (and only those two propositions). However, Mr. Ponens uses *modus ponens*, whereas Mr. F.A. Lacy uses *modus profusus*: For any p, q , and r , $(p \wedge q) \rightarrow r$.⁴⁶

Obviously Mr. F.A. Lacy is not justified. Moreover, it is obvious that the reasons for which he believes the Spurs will win is that they will play the Pistons and that if they play the Pistons, they will win. Finally, it is obvious that those facts constitute sufficiently strong reasons to believe the Spurs will win. But this doesn't mean the common view is false. My view predicts that Mr. F.A. Lacy's token belief is irrational and unjustified. This is because, according to my view, Mr. F.A. Lacy doesn't believe for sufficiently strong normative reasons—he merely believes for reasons that happen to be normative (and sufficient). He is just like Sam in Affirming the Consequent. If this is right, then cases like Spurs Win are no threat to the traditional view of *ex post* rationality. The key is understanding what acting and believing for reasons actually amounts to.

Millar's and Turri's arguments get traction against the traditional analysis of *ex post* rationality only if the cases in question truly are cases of believing for normative reasons. Since they're not, they pose no challenge to the claim that *ex post* rationality is analyzed in terms of *ex ante* rationality and acting and believing for normative reasons. Thus, a virtue of my view is that it allows us to hold onto the common analysis of *ex post* rationality.

4.7.2 Speckled Hens and the Epistemology of Perception

It's plausible that experience provides us with reasons. When I see that the table is red, I gain a reason to believe that the table is red. Moreover, it seems like I can gain *ex post* rational beliefs when I experience the table as red and believe for the reason provided. An objection to these thoughts runs through the fact that experience can represent very specific things. Moreover, sometimes it seems like beliefs that those specific things obtain can be caused by experiences that represent those specific things *even though* it's intuitively clear that those beliefs are not *ex post* rationalized by reasons provided by the experiences.

⁴⁵I won't mention Millar's cases because he unfortunately never gives any. He only refers to cases with this structure (see (Millar, 1991, pgs. 57-64). Turri's argument was anticipated by Swain (1988).

⁴⁶I'm following Turri here by making this a conditional. I assume he meant to make this an inference rule. Otherwise Mr. F.A. Lacy is just using *modus ponens* from a particularly dumb conditional premise.

The case of the speckled hen brings out the problem. Suppose you have a visual experience of a speckled hen. This visual experience represents the hen's 106 spots. It also causes you, suppose, to form the belief that the hen has 106 spots. According to a naïve causal view about believing for normative reasons, this is sufficient for your belief to be held for the reason provided by the visual experience. So it seems like the view that holds that having an *ex post* rational belief is a matter of believing for a sufficient normative reason predicts that this belief should be rational. But it doesn't always seem to be *ex post* rational.

This is just another deviancy case. Your belief might be caused in some way by the experience, but it won't be sustained by the right feature of the experience, which is the fact that it is a sufficient reason to hold the belief. Since your perceptual apparatus lacks certain discriminatory abilities, it's plausible you aren't disposed to revise the belief if the reason in question ceases to be sufficient. For example, suppose that your initial perception was non-veridical—the hen has another spot that was hidden from view at first. Suppose the hen changes positions and reveals the 107th spot. The reason provided by the first experience is now defeated. But it's plausible that you won't revise your belief, given your lack of discriminatory abilities. This is strong evidence that you don't actually believe for a sufficient normative reason, and hence my view doesn't predict that you are *ex post* rational. Thus, my view can explain what's going wrong in the speckled hen case. This is a virtue.⁴⁷

4.7.3 The Causal Efficacy of the Normative

It has been highly contentious since at least Harman (1984) whether normative properties—if they exist at all—are *causally efficacious*. Although there is some dispute as to what it would take for them to be causally efficacious, most seem to think that they would need to play an essential role in some correct causal explanation in order to be causally efficacious.

As Wedgwood (2007) has pointed out, views about acting and believing for normative reasons like the one defended here (and the one defended by Wedgwood himself) do seem to vindicate the causal efficacy of the normative. Let me explain.

Suppose I know p and if p , then q and I infer q from these two known facts. Suppose also that I believe q for the sufficient normative reasons constituted by those two facts. Thus, I believe q in virtue of the fact that the set of reasons containing just those two reasons is sufficient. This means that a certain property of that set—the property of being sufficiently strong—sustains my belief. Explanations of how a state is sustained are causal explanations. Moreover, the property of being sufficiently strong is a normative property. Thus, it looks like we have a case where a correct explanation of how my belief is sustained essentially appeals to a normative fact.

If this is right, then of course practical reasons will also be causally efficacious. Suppose I ϕ for some sufficiently strong reason r . Thus, my ϕ -ing was the result of the manifestation of a disposition to ϕ when r is a sufficient normative reason to ϕ . That means that some property of r —viz., that it constitutes a sufficient reason to ϕ —causally explains why I ϕ -ed. That property is a normative one. Thus, it seems like there is a correct explanation of why I ϕ -ed that makes essential reference to a normative property.

The most common strategy for resisting the causal efficacy of the normative runs through the idea that the normative supervenes on the descriptive. This means that for any normative property P , there will be some descriptive fact D that realizes P . This means that for any correct explanation

⁴⁷For more on this, see Lord & Sylvan (MS), which was partially inspired by Sylvan (MSa) and Sylvan (MSc).

that appeals to *P*, there will be some other correct explanation that only appeals to *D*. Following Harman (1984), many have taken this to show that the normative is causally idle. We can just explain everything in terms of the descriptive properties that the normative ones supervene on.

I think the proper way to respond to this argument is to deny that there can only be one correct explanation.⁴⁸ There is some correct explanation that only refers to the descriptive properties that the normative ones supervene on. Fair enough. But that doesn't show that the causal explanations that appeal to the normative aren't also correct. The rest of this chapter provides evidence that they are correct.

In fact, I think the rest of the chapter provides evidence that they must be correct if there is any acting or believing for normative reasons at all. Recall that I argued that in order to solve the deviancy problem, one must pick out the property that must do the causing in order for it to be causation of the right kind. This property, I've argued, must be the property of being a normative reason. The explanation of the action or attitude that appeals to the descriptive properties that the normative properties supervene on won't always explain acting for normative reasons. This is because those descriptive properties need not constitute the normative properties. Thus, appealing to the normative is essential to fully explaining our acts and attitudes when we perform them or hold them for normative reasons.

One might think this argument shows too much.⁴⁹ More specifically, one might think this argument, if successful, would also vindicate the causal efficacy of the arithmetical. For there are all sorts of true counterfactuals that have arithmetical concepts in their antecedents. For example, suppose my kitchen backsplash contains 17 rectangular tiles. It might be true that if my backsplash were to have 18 rectangular tiles, then I would be able to make a square. One might think this true counterfactual is analogous to counterfactuals I appeal to in my account of believing for normative reasons—e.g., usually when one ϕ s for a normative reason *r*, it will be true that if *r* weren't a normative reason, one wouldn't ϕ .⁵⁰ If the claims were analogous, then one might think my argument in this section could be extended to show that the arithmetical is causally efficacious. Since it's (supposedly) not, this would be a major problem for the argument.

Fortunately, the two cases aren't analogous. The important point is that my claim about the normative is stronger than the claim that the various counterfactuals are true. It is possible for the various counterfactuals to be true even if the normative is causally inefficacious. The truth of the counterfactuals follows from the causal efficacy of the normative, but it could also follow from claims that are compatible with the causal inefficacy of the normative. Similar things can be said about the arithmetical. The counterfactuals follow, in our case, from claims that are compatible with the causal inefficacy of the arithmetical. Namely, they follow from claims about *the tiles*. No purely arithmetical fact is shown to be causally efficacious by the truth of the counterfactual. So my argument doesn't overgenerate in this way.

Still, even if it doesn't overgenerate in that way, it might overgenerate in a more obvious way.⁵¹ Let's take an example involving arithmetic more analogous to paradigm cases of acting for nor-

⁴⁸This is following (Wedgwood, 2007, chapter 8). See that for much more discussion and defense of this type of account of the causal efficacy of the normative.

⁴⁹Thanks to Gideon Rosen for raising this objection and discussing it with me.

⁵⁰It's important to remember that I don't think that the claim that one ϕ s for a normative reason *r* is made true by the counterfactual. It's made true by one's dispositions (which aren't made true by counterfactuals, either). I'll just ignore this important feature of my view for the sake of argument.

⁵¹Thanks to Michael Smith for raising this objection and discussing it with me.

mative reasons. Suppose Michael is taking an arithmetic exam. Question number 28 asks what $212+432$ is. Michael comes to the conclusion that the answer is 644. He then writes '644' in the appropriate place. One might think that if my argument above is sound then, in this case, the fact that the $212+432=644$ causally explains Michael's writing down '644' in the appropriate spot. This would vindicate the causal efficacy of the arithmetical.

I think the right thing to say about this case is that it is not analogous to a case of acting for normative reasons. This is because *just is* a case of acting for normative reasons. The arithmetical fact has a normative property—viz., being a reason to write '644' in the appropriate spot—that is causally efficacious in this case. Does this vindicate the causal efficacy of the arithmetical? In a sense it does. Arithmetical facts can have properties—normative ones—that are causally efficacious. But any kind of fact can have those properties as well. So it's not as if there is something special about the fact that, in this case, an arithmetical fact happens to have the property. In any case, I am not bothered at all by this result. This seems like the correct amount of generation.

4.8 Conclusion

This chapter has had several aims. The main aim was to provide analyses of believing and acting for reasons, believing and acting for normative reasons, and believing and acting for sufficient normative reasons that avoid deviancy problems. I also argued that one has to provide separate analyses of each phenomenon in order to avoid all deviancy problems. While each analysis was close to the others, none were reducible to the others.

My final aim was to show that the views defended here have important payoffs in other parts of epistemology and metaethics. They allow us to hold onto the traditional view of the relationship between *ex ante* and *ex post* rationality, they allow us to solve the problem of the speckled hen, and they allow us to defend the causal efficacy of the normative.

Part II

Foundationalism, Deception, and the Importance of Being Rational

Chapter 5

Foundations for the Foundational

5.1 Introduction

SUPPOSE you rationally believe some proposition p . Here's a question for an epistemologist: Why are you rational in believing p ? And here's a natural answer: Because you derived p from q , which you rationally believe. But why are you rational in believing q ? Suppose the epistemologist answers this question like she answered the first. You're rational in believing q because you derived it from r , which you rationally believe. We'll then just ask the same question about r . If she keeps giving the same type of answer, we'll continue to ask the same question. Suppose this goes on *ad infinitum*. This exercise will be tedious. But that isn't the half of it. The epistemologist would have a very bad theory of rationality if her theory required there to be an infinite chain of rational beliefs for any particular belief to be rational.

Enter foundationalism. Foundationalism's main thesis is that there are some beliefs that are *foundationally rational*—i.e., their rationality is not dependent upon you being rational in holding any other beliefs. The foundational beliefs stop the regress. If foundationalism is true, then eventually we'll get to a belief that will allow the epistemologist to give a new type of answer. She'll be able to say of that belief that it is rational in some way other than being derived from another rational belief. A nice trick if you can pull it off.

This chapter isn't about why we should be foundationalists. I think we should, but that won't matter for our concerns here. This chapter is about what form the best foundationalist theory should take. More specifically, it's about whether foundationalists should think the foundational beliefs are both supported by sufficient reasons and based on sufficient reasons. This issue has mostly been neglected. There have been, however, several foundationalists who have argued that foundational beliefs needn't be based on reasons. I think these theorists are wrong. This chapter will thus be a defense of the claim that the best foundationalist theory holds that the foundational beliefs are based on reasons. In short, the foundational beliefs have foundations.

The plan is this. In the next two sections we'll get clearer on the regress, what it's supposed to show, and how to respond to it. Then, in §5.4, I will present a positive argument for the claim that all foundational beliefs are supported by sufficient reasons. §§5.5-5.6 will investigate arguments against the claim that the foundational beliefs are based on reasons. I'll argue that these arguments founder because they presuppose implausible views of reasons or of what it is to base a belief on a reason (or both). Once we hold the correct view about these notions, it will quite plausible that the foundational beliefs are based on sufficient reasons.

5.2 The Regress(es), with Possible Replies

Let's start by reminding ourselves of the distinction between propositional and doxastic rationality.¹ We'll do it by example. Suppose both Jack and Jill read the National Weather Service's forecast for today. It says that the high temperature today is going to be between 50 and 60 degrees. This is an excellent reason to believe that the high temperature is going to be between 50 and 60 degrees. Suppose Jack believes it will be between 50 and 60 degrees today because of the NWS's forecast. Jill also believes it will be between 50 and 60 degrees. However, despite the fact that she read the weather report, she doesn't hold the belief because the report says it will be between 50 and 60 degrees. Instead, she believes it will be because she infers it from her belief that it will 57 degrees. She believes it will be 57 degrees because it's the 19th of the month, and she always thinks that the day's temperature can be determined by multiplying the day of the month by 3.

In this case, both Jack and Jill possess excellent evidence that the temperature will be between 50 and 60 degrees—this is because they both read the NWS's forecast. The proposition *the high temperature will be between 50 and 60 degrees* is thus rational for both of them. However, they don't both rationally believe that the high temperature will be between 50 and 60 degrees. Only Jack's belief is rational. This is because Jack is the only one that believes the propositionally rational proposition because of the excellent reasons both he and Jill possess. Jill believes for a bad reason, and thus doesn't have a doxastically rational belief.

With this in mind, now recall how I described our regress above. I said that the problem looming for our epistemologist was the need to posit an infinite chain of beliefs derived from other rational beliefs. This certainly is an unattractive posit. Notice, though, that there are actually two different problematic parts to the posit. On the one hand, there is a claim about propositional rationality. If the epistemologist continues to give the same type of answer, then she has to posit an infinite number of propositionally rational propositions in order to explain how I'm propositionally rational in believing p . On the other hand, there is a claim about doxastic rationality and how you can get it. If the epistemologist keeps giving the same answer, then she has to posit an infinite number of inferences (or, as I put it in the introduction, derivations) to explain why I'm doxastically rational in believing p . In short, there are two regresses, and it's worth explicitly separating them.

The first regress, then, is about propositional rationality. According to a naïve view, in order to be propositionally rational in believing any proposition p , one must be propositionally rational in believing some other proposition q , which one can rationally infer p from. This view gives rise to at least one troubling regress. It seems to follow immediately from this view that in order to be propositionally rational in believing any proposition p , you have to be propositionally rational in believing an infinite number of other propositions. Since this view posits this infinite set of propositionally rational propositions, let's call it Propositional Infitism. And let's call this regress problem The Foundations Problem. The Foundations Problem strikes most as a fatal flaw of propositional infinitism.

I think it's important to ask why it is that The Foundations Problem is so damning for propositional infinitism. Often, I think, philosophers talk as if The Foundations Problem is problematic in ways in which it isn't. This happens most often when they run it together with another problem

¹In this chapter I am going to use the propositional/doxastic terminology rather than the *ex ante/ex post* terminology. This is mostly to make it easier to name some of the views we'll be discussing. I also don't need the level of generality afforded by the *ex ante/ex post* terminology because I will only be talking about beliefs.

that crops up for views that give rise to a second regress. Consider the view that holds that in order to have a doxastically rational belief that p , you have to derive p from some other doxastically rational belief in some other proposition, q . Call this view Doxastic Infinitism. Doxastic Infinitism also gives rise to a troubling regress. This is a regress of *inferences*. If doxastic infinitism is true, then in order to be doxastically rational in believing anything at all, one must make *an infinite number of inferences*. Let's call this The Processing Problem for doxastic infinitism.²

Propositional infinitism is not doxastic infinitism, and The Foundations Problem is not The Processing Problem. Propositional infinitism is a theory about propositional rationality. It's not the case that propositional infinitism entails doxastic infinitism, nor is it the case that doxastic infinitism entails propositional infinitism. That is to say, one could be a propositional infinitist but not a doxastic infinitist—this is the default position amongst propositional infinitists—³ and one could be a doxastic infinitist without being a propositional infinitist. The latter view is a strange one. On that view, p can be propositionally rational for A in a way that is not dependent on there being propositional rationality for an infinite number of propositions (this is the denial of propositional infinitism) even though, in order to have a doxastically rational belief one has to perform an infinite number of inferences (this is doxastic infinitism). Although this view is very strange, it is not incoherent.

What's most important right now is that propositional infinitism doesn't entail doxastic infinitism. Thus, propositional infinitism doesn't have The Processing Problem. So what is bad about The Foundations Problem? The first problem is that it is implausible on its face that *every* propositionally rational proposition requires an infinite number of propositionally rational propositions. The classic response to this is to point out that we do seem to have propositional rationality to believe infinite sets of propositions—for example, that $2 > 1$, that $3 > 1$, that $4 > 1$, et cetera. This might be true, but it does little to explain why it is that in order to be propositionally rational in believing that I have an itch on my left leg I must be propositionally rational in believing an infinite number of other propositions. What would those other propositions even be? It's hard enough to find a candidate for the *first one*—i.e., the proposition that my propositional rationality that I have an itch depends upon—let alone an infinite number.⁴

If we reject propositional infinitism, then we're left with two options. Either we can be propositional foundationalists or we can be propositional coherentists. Propositional coherentists maintain that in order for p to be propositionally rational for A , p must cohere in certain ways with A 's other beliefs. The most important feature of propositional coherentism for our purposes is that the propositional coherentist believes that all propositionally rational propositions for A are at least partially rational in virtue of the other propositions that are propositionally rational for

²Just to give two examples of philosophers running these things together, Fumerton (2010) speaks of The Processing Problem as *the* problem that motivates foundationalism. And Armstrong (1973) consistently casts the debate in terms of inferences. This leads him into drawn out discussions of what counts as an inference. While these discussions are interesting, I don't think they are particularly germane to the regress that should motivate us to be foundationalists. Even when they don't run the problems together in harmful ways, hardly anyone explicitly notes that these are distinct problems. (Pryor, 2005, §3) is an admirable exception.

³See especially Klein (1998).

⁴The best candidate for the first one is something like *it appears that I have an itch*. This in itself has problems. For one thing, this seems to run into skeptical worries (see, e.g., (Williamson, 2000, ch. 8)). But even if those problems could be addressed, what's the propositionally rational proposition that I could infer *it appears that I have an itch* from? This is where it gets puzzling. This type of argument for foundationalism is given by Pryor (2005).

A. Propositional foundationalists deny this; they think that there are some propositions that are propositionally rational independently of the any other propositionally rational propositions.

Propositional foundationalists, then, think that there is some propositional rationality that is not dependent upon there being any other propositional rationality. When your propositional rationality for p is like this, we'll say that it is *immediate*. The immediately rational propositions stop the regress and serve as the foundation for all of the *mediated rationality*—all of the propositional rationality that is dependent on there being other propositional rationality.

It's again important to stress that the propositional foundationalist doesn't immediately have any solution to The Processing Problem. Doxastic infinitism is compatible with propositional foundationalism. It might be that there are some propositions that are immediately rational *even though* in order to be doxastically rational in believing those propositions, you must complete an infinite number of inferences. This is obviously an odd view, and it is one that no propositional foundationalist I know of has ever accepted. Instead, propositional foundationalists usually hold that the immediately rational beliefs are *uninferentially doxastically rational* as well. A belief is uninferentially doxastically rational just in case it is doxastically rational in a way other than via inference from some other doxastically rational belief. Let's call those who think that there are uninferentially doxastically rational beliefs doxastic foundationalists. Most propositional foundationalists speak as if propositional foundationalism is doxastic foundationalism. This is not true. The two views aren't even about the same thing.

Unfortunately, a discussion about terminology is needed here. Often in the debate about foundationalism, theorists have used the term 'non-inferential' to pick out the rational beliefs, whatever they're like, that solve The Foundations Problem. I think that when you set up the dialectic in the way that I have, this is bizarre. For one thing, it's important to be specific about which foundationalism we're talking about. If we're talking about propositional foundationalism, then it's quite unclear why we'd want to call the propositionally rational beliefs that solve The Foundations Problem 'non-inferentially' rational. Neither the problem nor its solution have much at all to do with inferences.⁵

The oddity of this terminology is made even more obvious when one considers early views about what the 'non-inferential' beliefs are like. Followers of Descartes have often thought that the 'non-inferentially' rational beliefs are the beliefs that are infallible.⁶ Others have thought the 'non-inferentially' rational beliefs are the ones you are certain of. It's quite odd to think that those sets of beliefs would be appropriately picked out by the term 'non-inferential.' It's much better, I think, to talk about immediate rationality when we are talking about the fight between the propositional foundationalist, the propositional coherentist, and the propositional infinitist.

If I could have my way, I would use the term 'non-inferential' to pick out the special kind of basing that the doxastic foundationalist believes in. Alas, I don't have the power to reverse 50 years of poor technical usage of the term 'non-inferential.' So I have to drop that term. Instead, I use the uglier 'uninferential.' To repeat, a doxastically rational belief that p is uninferentially rational just

⁵Of course, inference does play some role in propositional infinitism. This is because the propositional infinitist thinks that in order for p to be propositionally rational, there must be an infinite number of propositionally rational propositions that one *could* infer p from. They stress that one merely needs to be able to do this, not that one has to do it. And, of course, the propositional foundationalist denies this is necessary, so her view really doesn't have anything to do with inference.

⁶See, for example, Price (1950).

in case it is doxastically rational in some other way than via inference from another doxastically rational belief.

The upshot of this discussion is that there are two regresses and two kinds of foundationalist views. Each foundationalist view offers a solution to one of the regresses, but only one. Henceforth we will be interested in both propositional foundationalism and doxastic foundationalism.

5.3 Foundationalism, Reasons, and Basing

In the introduction I said that the claim that we are most interested in is whether foundationalists should think that all rational beliefs are based on sufficient reasons. Now we can see that this claim can be broken into two parts corresponding to two different foundationalist views. That is, this can be broken down into two questions, one for the propositional foundationalist and one for the doxastic foundationalist.

Question for Propositional Foundationalist

Do all propositionally rational propositions have to be supported by sufficient reasons?

Question for Doxastic Foundationalist

Do all doxastically rational beliefs have to be based on sufficient reasons?

We'll assume that the propositional foundationalist thinks that all mediated propositionally rational propositions have to be supported by sufficient reasons. The sufficient reasons that support the mediated rational propositions are just the reasons that mediate one's rationality. The contentious claim is whether the *immediately* rational propositions have to be supported by sufficient reasons.

Similarly, we can assume that the doxastic foundationalist thinks that all of the inferentially rational beliefs are based on sufficient reasons. The sufficient reasons are just the propositions or mental states that one infers the inferentially rational beliefs from. The contentious claim for the doxastic foundationalist is whether all of the *uninferentially* rational beliefs have to be based on sufficient reasons. Thus, we can refine our two questions.

Refined Question for Propositional Foundationalist (the Propositional Question)

Do all *immediately* propositionally rational propositions have to be supported by sufficient reasons?

Refined Question for Doxastic Foundationalist (the Doxastic Question)

Do all *uninferentially* doxastically rational beliefs have to be based on sufficient reasons?

I think that the answer to both questions is Yes. Herein I will refer to the refined question for the propositional foundationalist as the Propositional Question and refer to the refined question for the doxastic foundationalist as the Doxastic Question. In the next section I will lay out an argument for concluding that Yes is the correct answer to the Propositional Question. It will be clear which premise opponents of the argument must deny. The bulk of the chapter will be dedicated

to showing that extant challenges to the crucial premise founder. In so doing I will begin to develop positive views that vindicate the crucial premise. I will then argue that once you answer Yes to the Propositional Question, it's very plausible you should answer Yes to the Doxastic Question. Finally, I will consider an objection to answering Yes to the Doxastic Question even if you answer Yes to the Propositional Question.

5.4 The Argument from the Nature of Defeat

Here's my argument for why we should answer Yes to the Propositional Question—that is, why foundationalists should think that all immediately rational beliefs are supported by sufficient reasons.

The Argument from the Nature of Defeat

- (1) If p is not supported by sufficient reasons, then the set of reasons to believe p is defeated.
- (2) If the set of reasons to believe p is defeated, then p is not propositionally rational.
- (C) Thus, if p is not supported by sufficient reasons, then p is not propositionally rational.

(2) is just a truism about the relationship between defeat and propositional rationality. Indeed, it's tempting to think that (2) is an analytic truth. Given this, (1) is what my opponents will have to deny. Here is a quick argument for (1).

The Argument from the Nature of Defeaters

- (1) If p is not supported by sufficient reasons, then there are decisive reasons not to believe p .
- (2) If there are decisive reasons not to believe p , then set of reasons to believe p is defeated.
- (C) If p is not supported by sufficient reasons, then the set of reasons to believe p is defeated.

Again, I take (2) in this argument to be a truism about the relationship between reasons and defeat. Indeed, I think it's a truism that what it is to be a (full) defeater of the reasons to believe p is to be a decisive reason to either believe $\neg p$ or withhold judgment concerning p .⁷ So my opponents will want to deny (1) in this argument as well.

I take it that my opponent will agree that both of the first premises are true when it comes to mediately rational propositions. The controversy is over whether they are true for some of the immediately rational propositions. In order for my opponent to be successful, she must drive a wedge between lacking sufficient reasons for believing p and having decisive reasons not to believe p .

Extant theorists have pursued a common strategy in trying to drive this wedge. According to this strategy, the relevant wedge can be driven home because *there aren't enough reasons to go around*. More specifically, in cases of immediate propositional rationality, there just aren't enough reasons

⁷What it is to be a partial defeater is to just be a reason either to believe $\neg p$ or to withhold judgement.

around to sufficiently support the propositions in question. Thus, cases of immediate propositional rationality will be cases where there aren't sufficient reasons supporting p , but there aren't decisive reasons not to believe p —i.e., p is not defeated.

There are two ways of pursuing this strategy. The first is to argue that it follows from the nature of reasons that there won't be enough reasons to go around in cases of immediate propositional rationality. The second—more common—strategy is to appeal to cases where it seems as if there is immediate propositional rationality but it is unclear whether there are sufficient reasons to support the propositions that are rational.

In the next section we will look into some of the details of these arguments. I will argue that both ways of pursuing this strategy can be resisted.

5.5 Arguments in Favor of Answering No to the Propositional Question

By answering No to the Propositional Question, the propositional foundationalist commits herself to the view that there are some immediately propositionally rational propositions that are not supported by sufficient reasons.

As I said before, there are at least two strategies for showing what needs to be shown. The first is to argue that the nature of reasons is such that there aren't always going to be sufficient reasons in the case of immediate propositional rationality. The second strategy is to appeal directly to intuitions about cases. Let's look at each strategy in turn.

5.5.1 From Mentalism about Reasons to a No Answer: Lyons's Argument

(Lyons, 2009, ch. 3, §4) wants to show that only rational beliefs can be reasons, and thus that only rational beliefs can provide reason based rationality for other beliefs. Lyons also assumes that propositional foundationalism is true. That is, he assumes that there are some propositionally rational propositions whose rationality is independent of propositional rationality for any other propositions. When you combine his claim about reason-based rationality—that it is only provided by rational beliefs—with his propositional foundationalism, you get the result that some propositionally rational propositions are rational in some way other than being supported by sufficient reasons. Lyons thinks that these beliefs—the basic beliefs—are rational because they are the outputs of certain reliable cognitive faculties. He thus holds a disjunctive theory of propositional rationality. Propositionally rational propositions are rational either because they are supported by one's other rational beliefs or because they are the output of certain reliable cognitive faculties.

Lyons starts by asking why it is that rational beliefs can provide reasons-based rationality for other beliefs.⁸ The basic idea of his view is that when I am rational in believing p and if p, q , then any reasons that justify p will justify q .⁹ This means that when you gain a rational belief that q by inferring it from a rational belief that p , your belief that q isn't rationalized by p ; rather, it's rationalized by what rationalizes p . If we assume that this is how it works for all reason-based rationality, then

⁸He uses different terminology. Instead of saying 'reasons-based rationality', he says 'evidential justification.' He makes it clear on p. 23 that he takes 'reasons' and 'evidence' to pick out the same thing.

⁹He intends for this to only be a *ceteris paribus* claim. We needn't be concerned with that here.

we have an explanation for why irrational beliefs don't justify.¹⁰ Suppose I infer q from an irrational belief that p . Since all rational inferential beliefs are rationalized by the reasons that rationalize the premise of the inference and by hypothesis my belief that p isn't supported by sufficient reasons, we get the result that my belief that q is not supported by sufficient reasons. And this is so for any p . Thus, given the assumption that all inferentially rational beliefs are rationalized by the reasons that justify the premises, we get an explanation for why irrational beliefs cannot transmit rationality. Analogous arguments show why nondoxastic states like desires, hopes, etc. can't transmit rationality.

Lyons explains the upshot of this argument by writing, 'Analogous reasoning explains why desires, fears, bare conceptions, and the like cannot evidentially justify beliefs. It is not so much because they are not *beliefs* but because they are not *justified*. Of course, since only beliefs can be epistemically justified, it follows that only beliefs can evidentially justify beliefs' (75). Again, if only beliefs can provide reason-based rationality and we assume foundationalism, it follows that there are some rational beliefs that aren't rationalized by reasons. Thus, if foundationalism is true and Lyons's argument is sound, then we must answer No to the Propositional Question.

It's important to stress that Lyons sees his argument for the claim that only beliefs can provide reasons-based rationality as an argument to the best explanation. The anchor of his argument is that his view provides the best explanation for why irrational beliefs (and nondoxastic states) can't justify other beliefs. This claim about the best explanation is what I will challenge.

Lyons crucially assumes that reasons are *mental states*. His entire discussion is centered on the question of whether reasons are just beliefs, or whether reasons can also be certain kinds of experiential states. *Given the assumption that reasons are mental states*, it does seem mysterious how it is that all rational beliefs are supported by sufficient reasons. This is so for many reasons. Lyons's argument picks up on one of these.

In fact, I think that a mentalist view of reasons faces much larger problems than the one Lyons points to. The problems can be stated as a trilemma. It seems like we are bound to get caught up in either some type of threatening regress, or we will have to think that some nondoxastic states are reasons, or we'll have to think that some beliefs are literally self-rationalizing. The regress option arises for views that hold that only beliefs are reasons, and this is what motivates us to adopt foundationalism in the first place; the obvious problem with allowing nondoxastic states to be reasons is that it's hard to see exactly which ones to allow;¹¹ and it's hard to see any plausible defense of the claim that some beliefs are literally self-rationalizing (more on this below).

So much the worse for the assumption that reasons are mental states! I think that the problems for that view provide a good motivation for the view that reasons are propositions.¹² To show this, let's first look at how the propositional view, when supplemented with plausible assumptions, can offer an alternative explanation of why irrational beliefs don't transmit rationality. We'll then look

¹⁰It's slightly odd that he thinks he can assume this. But since I think his argument fails anyway, I won't harp on why this might not be a good assumption.

¹¹This is just the most obvious problem. The bulk of (Lyons, 2009, ch. 3) convincingly argues against views of this kind.

¹²It's also good motivation for thinking reasons are facts. I am inclined towards this view if we treat facts as true propositions. However, I also think that the problems for the mental state view is good motivation for a less restrictive propositional view, one which allows both true and false propositions to be reasons. Schroeder (2008, 2009b, 2011) argues for a view like this.

at how the propositional view can handle the mental state view's trilemma. This greatly mitigates Lyons's inference to the best explanation.

Why You Can't Get Rational Beliefs from Irrational Beliefs via Inference

The propositional view is compatible with the central thought behind Lyons's view, which is that when rationality transmits via inference from a rational belief that p to a rational belief that q , what rationalizes q is not p , but instead what rationalizes p .¹³ Let's suppose that this is true. If it is, then the propositionalist (the defender of the propositional view) can offer a very similar explanation to Lyons' of why rationality transmits via inference. To wit, if you're rational in believing p and you correctly infer q from p , then the reasons that rationalize your belief that p will also rationalize your belief that q .

Of course, the main question is what explanation can the propositional view give for why rationality cannot be transmitted via inference from irrational beliefs. At first blush, the propositional view does have a problem with this. After all, it seems like the propositionalist ought to say that one can irrationally believe p and p can be a sufficient reason to believe q . In cases like this, why is it that a belief inferred from p is irrational? After all, p is a sufficient reason to believe q .

The beginning of a proper response is to recall that the reasons that are important to whether you are rational are the reasons you *possess*. It's plausible that in order to possess a reason, you have to be at least rational in believing that reason. If this is right and only the reasons you possess count towards whether you're rational, it follows that when you irrationally believe p , p doesn't contribute towards whether you're rational in believing the propositions that it supports. Thus, when you infer q from p , your evidential position towards q is not enhanced by p .¹⁴

'But wait,' one might say, 'we all know that you might have a doxastically irrational belief that p even though p is propositionally rational. In those cases, you will possess reasons to believe q —viz., the reasons that confer propositional rationality on p . Yet we still shouldn't think that you are rational in believing q .' This objection is well taken. However, this objection only gets traction if the above explanation is supposed to be of why *doxastic* rationality doesn't transmit via inference from a doxastically irrational belief. But the above explanation is not of this; instead, the above explanation is of why *propositional* rationality doesn't transmit via inference from a propositionally irrational belief. And if you aren't propositionally rational in believing p , then you won't possess sufficient reason to believe p . Thus, it won't be the case that you possess sufficient reasons to believe q .

Lyons is also seeking to explain why propositional rationality doesn't transmit via inference from a propositionally irrational belief. This is clear because he assumes that *it follows* from the fact that one is irrational in believing p that there are no rationalizers for p . As we've seen, the former

¹³Just because the propositional view is compatible with this, it doesn't mean we should accept it. Indeed, I don't think we should accept this because p is going to be, independently of one's epistemic relation to it, a reason to believe q .

¹⁴Recall that this is similar to the argument I gave against Low Bar views of possession in chapter 3. It's also important here to remember that Schroeder provides another explanation for why rationality isn't transmitted via inference from irrational beliefs that (1) is compatible with answering Yes to the Propositional Question and (2) allows irrational beliefs to be reasons (see Schroeder (2011)). According to Schroeder, the contents of irrational beliefs can be reasons, but that doesn't mean they transmit rationality because the reasons are guaranteed to be defeated. For his argument for this explanation, see (Schroeder, 2011, §3). For more argument against this, see chapter 3.

claim doesn't follow from the latter claim if we are talking about doxastic rationality (and Lyons certainly recognizes this).¹⁵

The propositional view, when supplemented with a claim about which reasons contribute towards rationality, can offer an explanation of why rationality doesn't transmit via inference from irrational beliefs. Lyons never considers this view. Given the *prima facie* plausibility of this view, I take it the force Lyons's inference to the best explanation is severely mitigated. It will become even more mitigated after I show how the propositional view can avoid the trilemma for mentalist views of reasons, which I will now do.

Avoiding the Trilemma

Let's take each horn of the trilemma in turn. The first horn is the vicious regress. If you thought all rational beliefs must be supported by sufficient reasons and you thought that all reasons were beliefs, then for each rational belief you'd need another belief. And if you thought there couldn't be loops in the chain—i.e., if you aren't a coherentist—then you'd have to posit an infinite regress of beliefs to serve as reasons. Since the propositionalist doesn't think that reasons are ever beliefs (let alone always beliefs), she doesn't even get started down the road to the regress.

In response to the first horn most accept that some nondoxastic states can be reasons too.¹⁶ The nondoxastic states *par excellence* are nondoxastic perceptual states. The thought is that those states are reasons that support perceptual beliefs but aren't in need of any support themselves since they are nondoxastic. Thus, they can stop the regress.

I cannot offer a complete argument against these views. Matters become enormously complicated very quickly.¹⁷ However, I do think the basic idea behind one argument from Lyons (2009) is quite compelling.¹⁸

The beginning of Lyons's argument is that there is a spectrum of possible states one could appeal to in order to stop the regress. On one end of the spectrum there are states with very unsophisticated contents—let's call these mere sensations. On the other end of the spectrum there are states that have sophisticated contents—let's call these percepts. Indeed, percepts have the same contents as the perceptual beliefs. The basic idea of Lyons's argument is that the sensations cannot do the requisite epistemological work and the percepts *constitute* the beliefs in question.

Lyons has several complicated arguments for the first horn of the dilemma that I won't go into here. One simple argument is anchored in the claim that sensations can't be reasons because they can't stand in the right kind of probabilistic relations with beliefs/the contents of beliefs. Indeed, it's even worse than this: *The contents* of the sensations can't even stand in these relations because

¹⁵What is the explanation of why doxastic rationality doesn't transmit via inference from an doxastically irrational yet propositionally rational proposition? It's because in these cases you will base your belief that *p* on the wrong reasons, and if you infer *q* from *p*, then *q* will also be partially based on those wrong reasons. Thus, your belief that *q* will not be properly based, and thus will not be doxastically rational.

¹⁶Here's a short and far from comprehensive list of those who make this move: Alston (1988), Bonjour (2002), Brewer (1999), Feldman & Conee (1985), Fumerton (2010), Huemer (2001), Pollock (1986), and Pryor (2000, 2005).

¹⁷By way of example, recall my initial objection to this move above. Namely, how are we supposed to choose which nondoxastic states count as reasons? It doesn't seem like we should let all of them be reasons, so we need some way of deciding which ones are and which ones aren't. It becomes quite difficult to see how to do this once it's pointed out that the states themselves can't stand in probabilistic relations with the contents of one's beliefs. Without even this much, it's hard to see how to determine under which conditions nondoxastic states can be reasons.

¹⁸He makes essentially the same argument in Lyons (2008).

they aren't propositional.¹⁹ Maybe there's some other way they can relate to beliefs that explains how they could be reasons for beliefs. But the burden is squarely on the shoulders of the sensations-as-reasons theorist to tell us how.²⁰

The second horn holds that the percept can't serve as an independent reason for the belief because in the cases in question the percept constitutes the belief. There's a widely discussed argument against this claim. The argument's main premise is that it's possible to have a percept without having a belief. The conclusion drawn from this claim is that percepts are never beliefs. The premise is certainly true. If I were to dip my pen into a clear glass full of water, I would have a percept that the pen is crooked even though I wouldn't believe that the pen is crooked. However, the argument has a hidden premise; namely, that if any percepts are beliefs, all percepts are beliefs.

It's not clear why we should think this hidden premise is true. After all, the most plausible accounts of the ontology of beliefs hold that playing a certain functional role is at least a necessary condition for something being a belief. And it's not clear why we shouldn't think that percepts sometimes play the functional roles of belief. Moreover, it seems like the percepts are representational in the right kind of way. Being representational in a certain way is also often thought to be a hallmark of belief. Since percepts seem to be representational in the right kind of way, representationalists about belief don't seem to have any *prima facie* reason to reject the claim that percepts are sometimes beliefs.

One obvious objection to the thought that percepts are sometimes beliefs is that it is clear that the percepts often lack all the *modal properties* of the beliefs.²¹ Suppose I see that there is a banana on the table and believe that there is a banana on the table. It is possible (and very likely) that this belief has different modal properties than the percept. If I were to look away from the table, the percept would cease to exist, but my belief would not. This is some evidence that the percept is not the same as the belief.

This is why I was careful not to say that the percept and the belief are *identical*. This is not my view.²² My view is that sometimes percepts *constitute* beliefs. In the banana case, there is just one state when you are having the percept. That state is the percept state and the belief state. But the belief is constituted by a different state when you lose the percept. This shows that the belief isn't identical to the percept, but it doesn't show that the belief isn't constituted by the percept when the percept is around.

Once again, this isn't completely decisive, but it does push the burden on the other side.²³ It's up to the other side to either defend the hidden premise or provide another argument for why percepts are never beliefs.²⁴

If neither is forthcoming, then the percept-as-reasons theorist will have to think that the reason provided by one's perceptual experience is the same state as the belief that it supports. This will

¹⁹Often times in contexts like this 'content' is short for 'propositional content.' I'm obviously not using it that way. The contents of the sensations are just the data—whatever it turns out to be—that is the non-agential relata in the sensation relation. I take no stand exactly what the non-propositional contents of the sensations are. (Compare: Many think that pain states have contents that are non-propositional.)

²⁰For some discussion of this and some very hand-wavy suggestions at how to meet this burden, see Pryor (2005).

²¹Thanks to Tom Kelly for raising this objection.

²²I think it is unclear whether this is Lyons's view. Sometimes he seems to suggest that it is.

²³Lyons does a lot more work than I've done sharpening each horn. I think that his discussion is close to decisive.

²⁴In Lord (MSe) I provide further arguments for the claim that some percepts are beliefs. Lyons sees this as a key move in showing that some rational beliefs aren't supported by sufficient reasons. In Lord (MSe) I show that holding that some percepts are beliefs significantly helps my kind of view get out of problems.

push them to the third horn, which is the view that the basic beliefs are literally self-rationalizing.²⁵ That is, they will have to think that the perceptual belief is rationalized by a reason constituted by the perceptual belief itself. This view is very *prima facie* implausible. In fact, it's hard to see why a view like this earns the right to answer Yes to the Propositional Question at all. We might as well take it that it's just a brute fact that one is rational in such cases.

The propositionalist can split these last two horns. She needn't be concerned with sensations. For she should instead say that the *content* of the percept is the reason. Moreover, she should hold that *even if* the percept is sometimes the perceptual belief, that belief is not *self*-rationalizing. It's rationalized by the content, which is (at least in the good case) an antecedently excellent reason to have the belief.²⁶

It's natural to wonder at this point why, if at all, the propositionalist thinks the processes that produce foundational beliefs—perception being the paradigm—are particularly important epistemically. It's clear why the sensations-as-reasons and the percepts-as-reasons theorists think they are important. It's because they think that some of the output states of those processes are reasons. Moreover, they think that the states that are the reasons aren't themselves in need of rational support. Thus, those states not only are reasons, but they play a particularly important theoretical role by solving The Foundations Problem. The propositionalist, on the other hand, doesn't think this. She thinks that the contents of the output beliefs are the reasons. And in most cases those contents will have nothing to do with having the perception. So it seems like perception doesn't play a big role in the propositionalist's story.

In response to this, the propositionalist should highlight the fact that perception plays a key role in *giving us reasons*.²⁷ That is to say, we come to possess reasons via perception. It's not that the states themselves are reasons; rather, it's that by having the perceptual experiences we come to possess independently existing reasons. Furthermore, the propositional foundationalist who holds this view should insist that (at least some of) our perceptual beliefs are immediately propositionally rational. The rough story of why this is is that via perception we gain a decisive reason to hold that belief, the gaining of which is not dependent upon what we are already rational in believing.

The upshot of this discussion is that the propositionalist view has some serious advantages over the mentalist view. Not only can the propositionalist provide an explanation of why rationality doesn't transmit via inference from irrational beliefs, it also avoids a trilemma for mentalist views. Given the virtues of the propositional view, Lyons cannot simply assume mentalism. Without mentalism, Lyons's argument doesn't work. So much the worse for Lyons's argument.

²⁵There are obviously other ways to be pushed into thinking that the basic beliefs are self-justifying. (Pollock & Cruz, 1999, pps. 84-6) argue that any foundationalist view that holds that all rational beliefs are rational by other beliefs have to think that some beliefs are self-justifying.

²⁶Obviously you can't think that the content is always a good reason (or a reason at all) if you think that reasons are factive. I do think reasons are factive, and thus I have to explain what's going on in the bad cases. The next chapter is devoted to this. However, it's important to stress that *at this point in this chapter's dialectic*, I am neutral between factive and non-factive propositional views. Non-factive views can say that the content is a good reason in both the good and bad cases. The advantages one accrues from this form the overriding motivation for Schroeder (2008, 2011)'s acceptance of a non-factive view.

²⁷This point is made by, amongst others, Williamson (2000, 2009), Schnee (MS) and Sylvan (MSa).

5.5.2 Dispatches from the Counterexample Wars

The main reason why explicit arguments are rarely given against answering Yes to the Propositional Question is that many epistemologists think there are counterexamples to the claim that all propositionally rational beliefs are supported by sufficient reasons. Now that I've given a rough articulation of the kind of propositionalist view I favor, we should look at some of the purported counterexamples to see how the propositionalist view fares. I don't think the counterexamples are decisive either way. I think the debate will need to be adjudicated elsewhere.

Before we consider what I take to be the most pressing examples, let me mention and put to the side a certain batch of examples. Many mention—but never fully spell out—cases of *a priori* rationality and introspective rationality as counterexamples.²⁸ I have two reasons for not considering these cases. First, I think that an analogous story to the story told about perception above and in the next chapter can be applied to those cases as well.²⁹ I thus don't find these potential examples very moving. The second reason I want to put them to the side is that I've never seen someone provide a fully spelled out case to evaluate. Usually the domains in question are simply referred to. This is dubious for two reasons. First, it makes it very hard to even have intuitions about whether these domains provide counterexamples to the claim we're interested in. Second, it allows my opponents to avoid offering up any substantive views about the domains in question. Since they inevitably make suppressed assumptions about the domains, it's very hard to track the claims their arguments depend upon.

With that said, let's turn to the cases that I think are a serious challenge. These are 'forgotten evidence' memory cases.³⁰ Suppose that, when I was 16, I read in my world history book that Archduke Ferdinand was killed in June 1914. Suppose I come to believe that he was killed in June 1914. Given normal background conditions, I come to know that he was killed in June 1914. Now suppose I don't have an occurrent judgement that he was killed in June 1914 until I'm trying to come up with a counterexample 10 years later. By the time the belief becomes occurrent again, I have completely forgotten where I learned that he died in June 1914.

In this case, it's plausible that my belief is rational. However, it's not clear what reasons I have for holding it. It certainly doesn't seem like I have the reason for which I originally formed the belief. I don't remember where or when I read this information. And it seems like the reason for which I originally formed the belief was something along the lines of *My world history textbook says that AF was killed in June 1914*. Since I don't know that anymore, I don't have the reason for which I originally formed the belief. So if my belief now is supported by a sufficient reason, it's going to have to be a reason other than the one for which I originally formed the belief.

I'm not so sure we have enough information to know what to think about this case, despite my opponents thinking otherwise. For one thing, it's not clear to me that in all cases where you forget your original reasons and there is no evidence against your belief, you are rational. Think about it in terms of credences. I find it very implausible that you shouldn't lower your credence upon realizing that you don't remember where you read about AF's death. When I'm in these situations, I imme-

²⁸See, for example, Plantinga (1993), Grundman (2009), and Sosa (2009).

²⁹Let me note that I *don't* think this commits me to a strong rationalist view of the *a priori* in the sense that I don't think this commits me to thinking that there is some independent faculty of 'rational intuition' that is analogous to our perceptual faculties (for theorists that do hold this view, see Huemer (2008), Bealer (2002), Bengson (MS), and Chudnoff (FC)). I'm much more sympathetic to reductionist views of our faculties of intuition. For more, see Williamson (2007).

³⁰See Plantinga (1993), Grundman (2009), and Goldman (2009a).

diately become less confident about the proposition in question upon realizing I don't remember the source of the information. And I don't think I'm irrational in doing that. Indeed, I think I would be irrational if I didn't do that.

This doesn't immediately settle the question of whether it's possible to be rational in believing that AF was killed in June 1914, in these cases. It just shows that the fact that you forgot the original reason has a bearing on how confident you should be. This is a mark against any view that thinks rationality floats free from the reasons you possess (at least in these cases).

I think that you won't be rational in these cases if you lack some other knowledge though. That is, rationality here doesn't just depend on the preservation of a belief that was once based on sufficient reasons that hasn't been defeated by counterevidence. For example, the fact that it memorially seems to me that AF died in June 1914 is a reason to believe this. And it often is one of the reasons I base my later judgement on. This reason can't explain why my stored belief is rational after I forget the other reasons, if it is. For you don't have the memorial seeming until you bring it out of storage. Seemings are conscious; you thus don't have them about beliefs in storage. I'll thus only have this reason after I bring the belief out of storage. So while this reason can help justify the later judgement, it won't explain why the stored beliefs are rational once I forget the original reasons.

I think that when the beliefs are in storage, they are based on reasons provided by more general things that I know. For example, the fact that the sources of information I trust are generally reliable. If I didn't know that the sources of information I trust are reliable, then I probably wouldn't be rational. This isn't because this lack of knowledge is itself a defeater; it's just that I would lack a reason to trust my memory.³¹ In forgotten evidence cases where you are rational, this is often the sufficient reason you have for the belief. In my own case, if I hold onto the belief despite forgetting my initial reasons, it is almost always for this reason.³²

That isn't the only potential reason, of course. I also know that my sources of information in high school were generally reliable. And I know that my teachers were epistemically trustworthy. I could go on. The point is simply that I often know lots of general things that are reasons that support stored beliefs, and support the claim that my memory is trustworthy.

I anticipate at least two objections to this line. First, one might agree that all of those facts are reasons and that when I do know them, they support the beliefs in question. However, the objector continues, *those* reasons are reasons you presumably possessed when you initially formed your belief, and those reasons were likely partial bases of your belief. Thus, you can't appeal to them because, *ex hypothesi*, you have forgotten those reasons.

I would be happy to bite this bullet, if only I thought there was any bullet to bite. It's fine for the objector to stipulate both that those reasons were partial bases of my original belief and that I've since forgotten them. But then it seems highly implausible to me that you are rational at all in these

³¹Of course, if you knew that the sources of information you usually trust when forming beliefs are generally unreliable, then that would be a defeater. But you can fail to know that and fail to know that the sources of information you trust in forming beliefs are generally reliable.

³²Of course, you might not always *know* that the sources of information you trust when forming beliefs are generally reliable, for a variety of reasons. For example, in a demon world you won't know this because it won't be true. In those worlds, you will know that the sources seem to be reliable. And that will be a good reason to keep the belief (they will also seem to be reliable in the good worlds too, and that fact will be a reason in those worlds, too). In worlds where reliable sources are available but you usually stick to the unreliable ones, you'll possess the reason provided by the fact that it seems to you that they are reliable, but given normal background conditions that reason will be defeated, at least after some amount of time.

cases. Imagine you were to realize that you'd forgotten *all* of those reasons. Surely your confidence in the proposition would plunge, and rightly so.

The second objection goes the other way. According to this line, you can bootstrap your way into being rational in believing anything simply by forming the judgement and letting it go into storage.³³ This is almost true. It also has to be, of course, that you don't possess any reasons that defeat the proposition believed. Meeting this condition for any old ridiculous thing is going to be harder than the objector thinks. Perhaps highly deluded people will be able to pull this off, or people who are in very impoverished epistemic positions, but those people will be few and far between.

'Fine,' the objector might continue, 'let's just focus on those poor fools. Isn't it implausible that they can bootstrap into rational beliefs?' I don't think so. It's a sad fact that if you're in a bad enough epistemic position you'll be rational in believing all manner of things. Indeed, I think the set of propositions that no one could ever be rational in believing is very small, perhaps empty. The important point, though, is that this isn't going to be easy. It would be bad if it were easy to bootstrap. But it's implausible to think that it will be. In order to think that it'd be easy, I think you'd have to implausibly narrow the set of reasons we normally have that defeat crazy things.

To illustrate, suppose Jack comes to irrationally believe at time t that the moon is made out of blue cheese. That belief goes into storage and Jack forgets about why he originally formed the belief (at t_1 , say). Now if the objection is going to be convincing, it has to be that it is relatively easy for this belief to be rational after t_1 . But this is implausible. For if Jack is at all normal, he will possess a lot of reasons that call into doubt this belief. He will, for example, know a bit about the general makeup of matter in space and he will know a bit about the kind of stuff cheese is made of. This will provide compelling reason to be skeptical about his belief. Of course, it's not necessarily the case that he knows about this. And so it's possible for his belief to become rational once it goes into storage. But this is unlikely enough, I claim, to be worrisome.

Let me recap the lessons I think should be gleaned from this. First, I think we should be skeptical of alternative views that hold that the rationality of our beliefs in these cases floats entirely free from the reasons we possess. This is because I think it's very plausible that when you do forget those original reasons, you should lower your confidence. I don't see how alternative views could explain this. Second, I think there are lots of reasons one can appeal to in order to explain why we are sometimes rational in these types of cases. Thus, I don't think these types of examples are close to being decisive. They might be if certain controversial things are true. But a good deal of work must be done to show those things to be true. Giving the examples isn't enough. This debate will not be settled by these types of examples.

5.6 From Propositional Foundationalism to Doxastic Foundationalism

We saw in §5.2 that Propositional Foundationalism and Doxastic Foundationalism are, strictly speaking, independent of each other. Either could be true while the other false. That said, some combinations are more natural than others. In particular, it would be odd if Propositional Foundationalism is true even though Doxastic Foundationalism is false. The two views most naturally go together. That is, it is natural to think that all doxastically rational beliefs are based on sufficient rea-

³³This type of objection is raised to a similar view in Huemer (1999)

sons if you think that all propositionally rational beliefs are supported by sufficient reasons. Thus, it is very natural to answer Yes to the Doxastic Question if you answer Yes to the Propositional Question.

It would be odd to answer Yes to the Propositional Question and No to the Doxastic Question because it is odd to think that that all propositionally rational propositions are supported by sufficient reasons even though sometimes one's doxastically rational beliefs can be explanatorily divorced from those reasons. It is much more natural to put those reasons to work in one's theory of doxastic rationality. After all, everyone will think that the reasons that support the premises of one's inferences will play a role in explaining why one's inferentially doxastically rational beliefs are rational. Once you think that the immediately propositionally rational propositions are *also* supported by sufficient reasons, it is odd to think that the uninferentially doxastically rational beliefs can float free of those reasons.

If this is right, then the best way to argue that some uninferentially doxastically rational beliefs are not based on sufficient reasons is to argue that the immediately propositionally rational propositions aren't always supported by sufficient reasons. That is, the best way to answer No to the Doxastic Question is by answering No to the Propositional Question.

Since I've already argued that Yes is the correct answer to the Propositional Question, I've thus already met the biggest challenge to the thought that all doxastically rational beliefs are based on sufficient reasons. However, at least one worry remains. The basic idea behind the worry is that even if we grant that all of the immediately propositionally rational propositions are supported by sufficient reasons, there is no plausible story about how we can base our uninferentially doxastically rational beliefs on those reasons. This threatens a Yes answer to the Doxastic Question even if Yes is the correct answer to the Propositional Question.

I think that the worry can be overcome. Once it is, I think the most natural answer to the Doxastic Question is Yes. In other words, we should think all doxastically rational beliefs are based on sufficient reasons.

5.7 The Duplicate Reasons Problem

Recall the propositionalist's story about why the immediately rational perceptual beliefs are supported by sufficient reasons. Her story was that—at least in the good case—the sufficient reason supporting one's belief that p is p .³⁴ After all, it's hard to find a better reason to believe p than p (at least in the good case)! Suppose you agree with this. Still, you might insist that even if p always sufficiently supports your immediately rational belief that p , p can't be your *reason for* believing p . That is, it can't be the reason that you base your belief on if your belief is to count as doxastically rational. That would be viciously circular reasoning, and viciously circular reasoning can't produce doxastic rationality.

³⁴Objection: The propositionalist can't give an adequate story of the bad case—i.e., the non-veridical case. Reply: There are two possible replies for the propositionalist *in this dialectic*. First, she could hold that false propositions can sometimes be reasons, and thus hold that your reasons in the good case are the same as they are in the bad case (see Schroeder (2011, 2008)). Second, she could hold that the reasons are different in the two cases, but one is still rational in the both cases. Schroeder's main (only?) argument for the first option turns on it giving a better story of the bad case than the second view. The next chapter will be a sustained argument for the second option, and thus against Schroeder. What's particularly important point here is that *as far as this chapter goes* I am neutral between these options. So the bad cases are not necessarily a problem for the propositionism I defend here.

This type of objection is leveled by Anthony Brueckner in two recent papers (Brueckner (2005, 2009)). Following Schnee (MS), I'll call this the duplicate reasons problem. It threatens a Yes answer to the Doxastic Question even if you answer Yes to the Propositional Question.

Brueckner starts out by arguing that in order for p to be a reason you possess, you must at least believe p . I argued against this in chapter 3. But for our purposes here we can assume that this is true. If it is, though, then it seems as if basing a belief that p on p will always be a viciously circular piece of reasoning. Here's (Brueckner, 2005, p. 441):

So now we have arrived at the following position: my belief of C [my cup is red] is rational in virtue of my belief of the evidential proposition that my cup is red. That is to say, my belief of the proposition that my cup is red is rational in virtue of my belief of the proposition that my cup is red!

Given our purposes, this is slightly ambiguous. On one reading, this is an objection to a certain account of propositional rationality. The objection would be that it can't be that one's belief that p is propositionally rational in virtue of one's belief that p . That would be, it seems, literal self-rationalization. This objection is not germane here. This is so for two reasons.

First, the best version of the objection turns on mentalism about reasons. It turns on, that is, the claim that the reason that is justifying the belief is *constituted* by the belief itself. That is literal self-rationalization. The propositionalist, of course, doesn't face this objection. For the propositionalist doesn't think the belief is justifying itself. She thinks the belief is rationalized by its content. Moreover, she should think that the content is an excellent reason to have the belief no matter one's relationship to the content.

The second—more obvious—reason why this is not germane here is that we aren't interested anymore in objections to the view that all propositionally rational propositions are supported by sufficient reasons. We've moved past that issue.³⁵ We're focusing now on objections to the claim that all uninferentially doxastically rational beliefs are based on sufficient reasons.

There is a germane reading of the objection. According to this reading, the propositionalist is committed to thinking that in duplicate reasons cases, one's belief is derived from that very belief. To put the point in terms of reasoning, in duplicate reason cases one reasons from the content of one's belief to that very belief! This seems obviously bad. You can't become rational in believing p via a derivation from p !

A first pass at a response points out that the cases we're interested in are the *uninferential* cases. Surely it's illegitimate to *infer* p from p . Rationality can't transmit via inference from a belief that p to a belief that p . For if your belief that p was already rational, then the inference is superfluous, and if it wasn't already rational, then there's no rationality to transmit. This is right as far as it goes. But this doesn't yet threaten the view we are concerned with, which is the view that all of uninferentially doxastically rational beliefs are based on sufficient reasons.

I think that there is a large grain of truth in this first pass response. However, I don't think it is fully adequate. This is because the objector doesn't have to insist that the type of reasoning is only illicit when it's inferential reasoning. There is a natural sense in which one is reasoning when one forms uninferential perceptual beliefs. Or, perhaps more perspicuously, there is a sense in

³⁵Of course, if Brueckner's objection required some new maneuvering on my part, the discussion in the last section would be incomplete. But, as I showed in the last paragraph, I already have all the needed resources to respond to the objection construed as an objection to my view of propositional rationality.

which one's perceptual capacities are engaging in reasoning when one forms uninferential perceptual beliefs. To use our earlier terminology, one's perceptual capacities process the sensations one receives through the sense modalities in order to output the more fine-grained percepts. This, I think, should count as reasoning in some broad sense.

Given this, we can refine the objection. The refined version runs as follows. Suppose that in order to possess a reason, you must believe it. And suppose that your reason to believe p in the uninferential perceptual case is p . Thus, you believe p —and this next phrase is important—in *virtue of* the fact that you believe p . But that is illicitly circular reasoning. So you can't properly base your belief that p on p in the uninferential case.

I think this objection fails, roughly, because it equivocates on *in virtue of*. On all views in the family of views I'm defending here, in order for r to be a reason you possess, it's necessary to stand in some epistemic relation to r . I've been assuming with Brueckner that the epistemic relation is belief. Thus, you possess reasons in virtue of believing them. But that doesn't mean that, when some belief of yours that p is rationalized by r , that belief is rational *in virtue of* the fact that you believe it. The belief is rational in virtue of the fact it's based on a sufficient reason. And, importantly, that reason isn't constituted by the belief.

This gets us part of the way, but there is still some more work. For I haven't explained yet why the type of processing in the uninferential case yields rational beliefs. That is, I haven't explained why one can be doxastically rational in believing p because of p in the uninferential case but not in the inferential case. I think I already hinted at an explanation above when I explained why the reasoning is illicit in the inferential case. When inference goes well, rationality gets *transmitted*. The point of inference is to transmit rationality from antecedent beliefs to new beliefs. But this *isn't* how uninferential belief forming processes work, and for good reason. The uninferential processes put us in contact with the world.

Think of it this way. In the inferential cases, we are trying to expand our rational beliefs merely through reasoning with the beliefs we already have. Given this, it's clear why p won't be rational via an inference from p . If p is rational, it will be rational for some other reason. Paradigm uninferential reasoning isn't like this.³⁶ In the paradigm cases, we expand our rational beliefs through direct contact with the world. That is, the inputs to the uninferential processes are facts about the world, whereas the inputs to our inferential reasoning processes are other beliefs. This difference explains why (at least some of) the uninferential beliefs are immediately rational. They are immediately rational because the world provides the inputs.

So here's the picture. The reason why basing your belief that p on p yields doxastic rationality in the uninferential case is that p is itself the input to the process that yields the belief that p . This is a large part of why it's *uninferential*. As long as your belief that p is sustained by p once you form it, your belief is based on p . And p , we've agreed, is about as good a reason as you can get for believing p . In contrast, inferring p from p is illicit because it's either superfluous or there won't be any rationality for the inference to transmit.

This gives us both an explanation of why duplicate reasons are okay in the non-inferential case, and an explanation of why they aren't in the inferential case. This, it seems to me, is all we need to solve the duplicate reasons problem.

³⁶As I'll argue in the next chapter, there are cases of uninferential basing that aren't like the paradigm.

5.8 Summary and Conclusion

This chapter has been about what the best version of *foundationalism* is. Upon reflection, there are two independent foundationalisms of interest. According to Propositional Foundationalism, there are some immediately propositionally rational propositions, which is to say that the rationality of some propositionally rational propositions is not dependent upon any other propositionally rational propositions. According to Doxastic Foundationalism, there are some uninferentially doxastically rational beliefs, which is to say that the rationality of some doxastically rational beliefs are not inferred from other doxastically rational beliefs.

We were interested in two questions. First, should the Propositional Foundationalist think that all propositionally rational propositions are supported by sufficient reasons? And second, should the Doxastic Foundationalist think that all doxastically rational propositions are based on sufficient reasons? Since both type of foundationalist will answer Yes to these question in cases of mediate and inferential rationality, the fight is really over the immediate and uninferential cases. I'll briefly summarize the dialectic in each case.

Propositional Foundationalism

My opponent's basic strategy is to claim that in the cases of immediate propositional rationality, there just aren't enough reasons around to sufficiently support the immediately propositionally rational propositions. Thus, those will be cases where there aren't sufficient reasons supporting p , but there aren't decisive reasons not to believe p —i.e., p is not defeated. This was the basic idea behind both Lyons's argument and the counterexamples. On my opponent's picture, then, at least in some cases of immediate rationality, there aren't any reasons to sufficiently support the proposition in question. At least in the case of Lyons, this leads him to a disjunctive view of propositional rationality.

Hopefully we are now in a position to see that I've already sketched a view that, if true, goes a long way towards rebutting my opponent. This is because I've already sketched a view that closes the gap in the immediately rational cases. The first step in sketching this view was rejecting Lyons's mentalism about reasons. Lyons, recall, thinks that only beliefs can be reasons. This is a more restrictive mentalism than most. Most mentalists also think that perceptual states of some sort or other can also be reasons. Lyons thinks that it's implausible to think that perceptual states of any kind can be reasons. He thinks his arguments for this claim—that it's implausible that any perceptual states are reasons—is key to showing that not all immediately rational beliefs are supported by sufficient reasons. Here's one way to reconstruct his argument:

Lyons's Argument from Elimination for Beliefs-Mentalism³⁷

- (1) Reasons are presentational mental states.
- (2) The presentational mental states are all either perceptual states or beliefs.
- (3) Perceptual states aren't reasons.
- (C) Thus, only beliefs are reasons.

³⁷Recall that Lyons also provides an argument to the best explanation for Beliefs-Mentalism. In §5.4.1 I showed that Beliefs-Mentalism didn't provide the best explanation of the relevant data.

Unlike most, I'm sympathetic to his argument for (3). The premise I think we should reject is (1). We shouldn't think reasons are mental states; we should think they are propositions. Rather than rehearsing the various arguments for the propositional view again, I will instead explain how this view helps with the defense of the two first premises.

Once you hold the propositional view, it's easy to see what the relevant reasons are in the immediately rational cases. In the good case, the most obvious reason will be p itself. p is about as good a reason as you can get for believing p . So if you gain p when you see that p , then you've gained a great reason to believe p by seeing p . Propositionalists will split when it comes to the bad case. If they hold a factive view of the relevant epistemic reasons, then they will have to hold that if one is rational at all in the bad case, the reasons that propositionally justify p in the bad case are different than the reasons that propositionally justify p in the good case. Propositionalism *per se* doesn't commit you to this, though.³⁸ If you hold a non-factive view of the relevant epistemic reasons, then you can hold that one has the same reasons in both the good and bad cases.

No doubt the choice between a factive and non-factive account will be momentous. At the end of the day, I prefer the factive version. Adequately explaining the bad case is *the* hardest problem for the factive theorist. For this reason, I reserve the entire next chapter to this problem. In this chapter the most important point is that you can have all the advantages of a propositional view without having this problem.

The second type of argument designed to drive home the relevant wedge was purely case based. That is, the supposed counterexamples were supposed to be theory neutral. The cases were supposed to be ones where everyone agrees that there are propositionally rational beliefs that aren't supported by sufficient reasons. I argued that the cases were not plausible enough to adjudicate the dispute. There are plenty of facts that are intuitively reasons that I can appeal to in order to account for the cases. Because of this, I don't think the dispute can be settled by the cases.

This is how I see the dialectic, at this point. The main motivation for alternative views was the claim that there aren't enough reasons to explain all of the propositional rationality that there is. This, in turn, motivates a disjunctive theory of a certain kind. According to the disjunctive theory, in many cases it will follow from the fact that there aren't sufficient reasons for p that p is defeated. But not all the time. In some cases, no such thing will follow. In these special cases—which will (most likely) be the cases of immediate rationality—we have to see whether some other conditions are met to know if p is rational (e.g., the reliabilist thinks that meeting some reliability constraint is the relevant condition).

Given that there is a plausible theory that doesn't require this type of disjunctivism, I think the disjunctivism is unmotivated. The reasons view is much simpler and more unified. This is even clearer when you keep in mind that you can give an important place to reliability in a reasons framework. You can do this by incorporating reliability into your analysis of reasons, basing, or both (a la Comesana (2010), Bedke (2010), Goldman (2011), Sylvan (MSb)). I say we stick with reasons. They can provide the basis for a simple, unified theory that can give principled explanations of all the cases.

Doxastic Foundationalism

As we saw, the best way to argue that not all doxastically rational beliefs are based on sufficient reasons is to argue that not all propositionally rational beliefs are supported by sufficient reasons.

³⁸Existence proof: Schroeder (2009b, 2011).

Once you answer Yes to the Propositional Question, it is very natural to answer Yes to the Doxastic Question.

That said, we considered a recent challenge to the thought that the uninferentially doxastically rational beliefs can be based on sufficient reasons, even if they are always supported by sufficient reasons. The main idea behind the challenge was that one cannot properly base a belief that p on p . All such basing would be objectionably circular.

I argued that this is too hasty. There are two parts. First, it was happily conceded that one can never base a belief that p on p via inference. There is an independent explanation of this, though. Namely, inference transmits rationality. So if one is already rational in believing p , no rationality will be transmitted via an inference. And if one is not rational in believing p , then no rationality can be transmitted via the inference.

The second part of the response was to insist that in the uninferential cases one can base one's belief that p on p . Uninferential (immediate) basing is unlike inferential basing in that it doesn't seek to transmit rationality. New rationality is created. This is achieved when one is sensitive to the reason in a certain way. In the terms of the view advocated in the last chapter, this is achieved when one's belief that p is sensitive to the right normative features of p . There is nothing objectionably circular about this type of uninferential basing.

Concluding Thoughts

As we've seen, the epistemology of immediate and uninferential rationality provides one of the main challenges from epistemology to my claim that rationality is a matter of correctly responding to the normative reasons you possess. Epistemologists challenge both the claim that all rational beliefs are supported by sufficient reasons and the claim that in order for a belief to be rational it has to be that one correctly responded to the reasons one possesses in holding that belief.

As we've also seen, I think that these challenges fail. Thus, I think that the epistemology of immediately and uninferentially rational beliefs does not topple my view. Moreover, I think the view that has emerged here is attractive. It tells a unified and plausible story about how rationality works across the board.

Chapter 6

Defeating the Externalist's Demons

True propositions can make a false proposition probable, as when someone is skillfully framed for a crime of which she is innocent. If perceptual evidence in the case of illusions consists of true propositions, what are they? The obvious answer is: the proposition that things appear to be that way. (Williamson, 2000, pg. 198)

6.1 Introduction

EXTERNALIST theories of rationality hold that the rational status of an attitude or action supervenes on more than just the non-factive internal states—the states one can be in even if the content of the state is false—of the agent. Internalist theories hold that the rational status of an attitude or action does supervene on the non-factive internal states of the agent. At least since Cohen & Lehrer (1983) and Cohen (1984), it has widely been thought that the deceived pose a special problem for externalist theories. To see this, compare The Good Case to The Bad Case:

The Good Case

Sam is in a favorable environment and has reliable belief forming capacities. She sees that the table is red, hears the baby crying, and learns interesting truths from her history textbook.

The Bad Case

Pam is Sam's non-factive internal state duplicate. So when Sam perceives the table is red, Pam has a phenomenologically indistinguishable visual perception; when Sam has an auditory perception of a baby crying, Pam has a phenomenologically indistinguishable auditory perception; and when Sam makes a certain inference because of what her history textbook says, Pam makes a phenomenologically indistinguishable inference. The only difference between Sam and Pam is that Pam is radically deceived by an evil demon. Nearly all of her beliefs about the external world are false.

It's very natural to think that both Sam and Pam are fully rational. It's easy for an internalist to explain this. This is because internalists think that Sam and Pam have all the same rationalizers.

Since Sam and Pam are non-factive internal state duplicates, they are in identical situations when it comes to rationality. Thus, given that Sam is rational to degree n , Pam is rational to degree n .

It's not easy to see how an externalist can explain the natural thought that Sam and Pam are rational to the same degree. This is because the externalist is committed to saying that Sam and Pam don't have the same rationalizers. They are in radically different situations when it comes to the external factors relevant in externalist theories of rationality. Take the externalist view Cohen (1984) was arguing against—viz., a naïve form of reliabilism. According to this view, a belief that p is justified only if it is produced by a reliable process. By hypothesis, Pam's beliefs are not formed by a reliable process. Thus, the reliabilist is committed to saying that she is not rational. So some externalist views—e.g., naïve reliabilism—clearly cannot vindicate the thought that Sam and Pam are rational to the same degree.¹ This has come to be known as The New Evil Demon problem for externalism.

A trend amongst externalists has been to bite the bullet and then do damage control. The basic idea behind this type of response is that two separate properties are conflated when one states the New Evil Demon problem. The externalist view is correct about the more demanding state and the internalist view is correct for the less demanding one. To illustrate, take Goldman (1988)'s reply to the problem. According to Goldman, his reliabilist view is correct when it comes to what he calls strong justification. Pam is not strongly justified. However, Pam is weakly justified, or rational. Thus some internalist view is true when it comes to rationality or weak justification. The New Evil Demon problem only seems like a problem because the internalist has conflated strong and weak justification. Or so the response goes.²

While I am somewhat sympathetic with this type of response (see Lord (MSc)), I am going to defend a different type of response here. I am going to argue that it's not necessary for the externalist to bite the bullet—or at least not any real bullets. I'll show that at least two externalist views can explain why Sam and Pam are both rational and why it's plausible to say that, in some sense, they are both rational to the same degree.

The structure will be as follows. In the next section I'll set out the type of externalist view I will be concerned to defend. I will then explicate the particular version of the New Evil Demon problem that plagues views of this type. I will then show how two views of the type I favor can solve the problem. My work won't be done, yet, because there is a related problem—The New New Evil Demon problem—that still plagues the type of view I'm defending. I'll introduce that problem in §6.4. I'll then offer a solution to The New New Evil Demon Problem. Along the way I'll also show how my solutions to the Evil Demon problems can also solve the problem raised by so-called Knowledge from Falsehood cases (see Warfield (2005), Fitelson (2010), Arnold (2011)) and a puzzle about non-veridical perception raised by Schroeder (2011).

6.2 Preliminaries and a Statement of the Problem

The New Evil Demon problem is a problem for any view that is committed to Different Rationalizers:

¹That's not to say there aren't reliabilist replies to the problem. See Comesana (2002) for a nice explication of perhaps the most popular reliabilist response.

²Others who pursue this type of response—although they don't all think the properties being conflated are strong and weak justification—include Littlejohn (2009, 2012), Weatherson (2008), Bach (1985), Engel (1992).

Different Rationalizers: Sam and Pam do not share the same rationalizers.

Lots of views will hold Different Rationalizers. I am interested in one type of view. I'm interested in the view that holds both that what rationalizes beliefs are epistemic reasons and that Sam and Pam have different epistemic reasons. That is, I'm interested in views that hold Reasons Rationalize and Different Reasons:

Reasons Rationalize: Epistemic reasons are the rationalizers of beliefs.

Different Reasons: Sam and Pam don't have the same reasons.

Again there are a variety of views that hold Different Reasons. The most natural—and the one I'm interested in—holds that Different Reasons is true because Reasons Factivity is true:

Reasons Factivity: p can be a reason to ϕ only if p is true.

If Reasons Factivity is true, then there are many more reasons for Sam to believe what she believes than there are for Pam to believe what she believes. This is because Sam believes a lot of truths that constitute reasons. While Pam believes those claims, they aren't reasons for Pam because they are false.

It's important to keep in mind, of course, that it is implausible that *all* epistemic reasons are rationalizers for everyone *all the time*. It's plausible that only reasons one *possesses* can rationalize one's beliefs. I won't take a stand as to what that relation is like, yet, but it obviously follows from Reasons Factivity that the epistemic relation must be factive.³

Cohen (1984)'s original argument against naïve reliabilism shows that the naïve reliabilist can't account for the fact that Sam is rational *at all*. This is because a necessary condition for a belief to be rational is that the belief be formed by a reliable process. Since Sam's belief isn't formed by a reliable process, it isn't rational at all. Thus, the *very first step* one must make in solving The New Evil Demon problem is vindicating Same Status:

Same Status: Sam and Pam are both rational.

Vindicating Same Status is just the first step. This is because many have a stronger intuition. Many think that not only are both Sam and Pam rational, they are *equally* rational. That is to say, one must also account for Same Degree:

Same Degree: Sam and Pam are equally rational.

³If you carve up the epistemic relations in one way, this won't be true. It is true that whenever you have p as a reason by standing in some epistemic relation r with p , p will be true. But you might think that the *epistemic* relation is, e.g., justification. It will follow from Reasons Factivity that only your true justified beliefs will be reasons you have. But since justification isn't factive, the epistemic relation won't be factive. I prefer to think of the relation in this case as being *true justified belief*. That relation is factive. I don't think it matters too much how finely you individuate the relations. I prefer my way because it more easily allows us to discriminate against two views, one of which is factive (true justified belief), the other of which is not factive (justified belief). Thanks to Tom Kelly for discussion here.

It's often unclear what it would take for Sam and Pam to be equally rational. Usually it seems that objectors mean that Sam and Pam have the same rationalizers. If this is what people mean by equally rational, then obviously those who accept Reasons Factivity (and Reasons Rationalize) will have to deny that Sam and Pam are equally rational. I am going to question below whether this is the proper way to think about a group of agents being equally rational. But even if one takes a harder line and bites the bullet with respect to Same Degree, one needs to explain away the widespread intuition that Sam and Pam are equally rational.

The final preliminary is that, at this point, we are only talking about *ex ante* rationality. That is to say, we are only talking about what it takes for a belief to be rationalizable. Plausibly, in order for a particular belief to be rational—to be *ex post* rational—rather than being merely rationalizable, the belief has to be suitably connected to the rationalizers. Wishful thinkers might sometimes hold beliefs that are rationalizable even though their token beliefs aren't rational because they aren't suitably related to the reasons. I take the New Evil Demon problem to be a problem about *ex ante* rationality. This is because it arises for views that hold that Sam and Pam have different rationalizers. Rationalizers are direct inputs into one's theory of *ex ante* rationality. They will also be highly relevant to one's theory of *ex post* rationality, but only derivatively through one's view of *ex ante* rationality.⁴

The New Evil Demon problem applied to the views I'm interested in here looks like this: Same Status is surely true. Yet it's not clear why if Reasons Rationalize and Reasons Factivity are both true. For if they are, then Sam and Pam differ dramatically in the epistemic reasons that they possess. Sam has a lot more reasons to hold her beliefs than Pam does. Indeed, it's not clear that Pam has *any* reasons to believe what she believes. So it's mysterious how the defender of Reasons Factivity can account for Same Status. Moreover, even if she could, it's especially mysterious how she could vindicate Same Degree. After all, according to her, Sam has a lot more reasons to believe what she does than Pam. Moreover, Sam's reasons seem to be much better than Pam's. So it seems like she must be more rational than Pam, and thus it seems the friend of Reasons Factivity cannot account for Same Degree.

With this in hand, it's interesting to note that two more specific problems raised for Reasons Factivity are (close to) special cases of The New Evil Demon problem. The first are so-called *knowledge from falsehood* cases. Here's one purported case from Warfield (2005):⁵

Reliable Watch

Fritz has a 7 pm meeting and extreme confidence in the accuracy of his fancy watch. Having lost track of the time and wanting to arrive on time for the meeting, he looks carefully at his watch. He reasons: 'It is exactly 2:58 pm; therefore I am not late for my 7 pm meeting'. He knows the conclusion, but as it happens it's exactly 2:56 pm, not 2:58 pm.

It's very plausible that Fritz knows that he's not late for his 7:00 meeting. Moreover, it's intuitive that a false proposition—viz., that it's exactly 2:58 pm—plays an essential rationalizing role in his

⁴Some disagree about what the problem's focus is (see especially Silins (2005)). Don't worry, I'll address the *ex post* version below—this is what I call the New New Evil Demon Problem. I think the New New Evil Demon problem is a much bigger problem than the New Evil Demon problem.

⁵I've modified the wording slightly. These cases are also discussed in Fitelson (2010), Arnold (2011).

obtaining the knowledge that he's not late. If this is right, then Reasons Factivity is in trouble. Call this the Knowledge from Falsehood problem.

It's reasonable to expect, I think, that the solution to the Knowledge from Falsehood problem will be just like a solution to the New Evil Demon problem. For in solving both problems one has to pick out some possessed reasons that rationalize the relevant beliefs. In The Bad Case we need to find some facts that constitute reasons that rationalize Pam's beliefs, and in Reliable Watch we need to find some facts that constitute reasons that rationalize Fritz's belief.

The second problem comes from Schroeder (2011). Schroeder's problem has to do with how it is that non-veridical perceptual beliefs are rationalized. Those who accept Reasons Factivity have a nice story to tell about veridical perceptions—the perceptual beliefs are rationalized by the contents of the perceptions.⁶ This obviously won't do for non-veridical perceptions since the contents of those perceptions are false.

There is an obvious way in which this is an argument against Reasons Factivity. Schroeder also argues—and this is what he's primarily interested in showing—that this is a strong reason to deny High Bar views of the *possession* relation. Recall from chapter 3 that High Bar views hold that in order to possess a reason, you have to stand in a positive epistemic relation with that reason. Many High Bar views commit one to Reasons Factivity, but not all do. So, for example, the view that in order to possess *p* you have to know *p* entails Reasons Factivity, but the view that in order to possess *p* you have to justifiably believe *p* does not (at least given plausible assumptions about justification). So Schroeder's argument is more than just an argument against Reasons Factivity.

Again, I think we should expect a solution to the New Evil Demon problem to provide (part of) a solution to Schroeder's problem.⁷ The key to a solution is to show that there are some other facts that constitute possessed reasons that rationalize one's beliefs in non-veridical perceptual cases. As we'll see, Schroeder ends up arguing that in the case of non-veridical, basic, perceptual justification, it's not plausible to appeal to other facts to account for beliefs rationalized by non-veridical perceptions. This is a version of the New New Evil Demon problem. This will be discussed at length below. For now the important point is that *at the very least* the friend of Reasons Factivity needs to appeal to other facts to explain why some beliefs are rationalized by non-veridical perceptions.

The upshot is this. The New Evil Demon problem poses a grave risk to views that hold Reasons Factivity (plus Reasons Rationalize). Claims like Same Status and Same Degree are very plausible, and one needs to account for them if one is going to successfully defend Reasons Factivity. Moreover, there are at least two related problems that seem to be open to the same type of solution as the New Evil Demon problem. If one can solve the New Evil Demon problem, then one can solve the related problems too. So there is a large payoff awaiting the theorist who can provide a compelling solution to the problem. In the rest of the chapter I aim to collect this payoff.

6.3 A Solution to the New Evil Demon Problem

The core of the New Evil Demon problem is that it's mysterious how the externalist can account for why both Sam and Pam possess sufficient reasons to hold their beliefs. As we saw above, many

⁶This is the view I developed in the last chapter. See also Williamson (2000).

⁷It won't solve all of it because Schroeder's problem is an amalgam of the New Evil Demon problem and what I'll call the New New Evil Demon problem. We thus won't have a full solution to Schroeder's problem until we solve the New New Evil Demon problem.

also have the intuition that Sam and Pam have equally strong reasons to believe what they do. It's particularly mysterious how the externalist can account for that intuition. In order to solve the version of the problem that plagues the type of view I'm interested in, one has to first explain what could rationalize the beliefs of the systematically deceived.

At this point we need to be more precise about what the possession relation comes to. This is because a necessary condition on some fact being a rationalizer is that it is possessed by the agent in question—only the reasons that are possessed can rationalize beliefs. On the views I'm interested in—views that hold Reasons Factivity—the possession relation is factive. I'll assume for most of the chapter that the possession relation is knowledge—I'll dub this view $P \rightarrow K$. This is mostly for convenience's sake. This is likely the view readers are most familiar with thanks to Williamson (2000). More importantly, the main points I'll make don't turn on exactly which factive relation turns out to be the possession relation. That said, in the conclusion I will point out one problem with $P \rightarrow K$ and show how a related view—the view that the possession relation is being in a position to know—solves this problem. For an extended argument that being in a position to know is the relevant relation, see chapter 3.

Accounting for Same Status

Let's start with how to account for Same Status—the datum that Sam and Pam are both rational. The trick is to find some facts in the bad case that are sufficient reasons to believe that the relevant agents possess. The most prominent defender of $P \rightarrow K$, Timothy Williamson, has explicitly addressed this challenge. He writes,

In unfavorable circumstances, one fails to gain perceptual knowledge, perhaps because things are not the way they appear to be. One does not know that things are that way, and $E=K$ [i.e., $P \rightarrow K$] excludes the proposition that they are as evidence. Nevertheless, one still has perceptual evidence, even if the propositions it supports are false. True propositions can make a false proposition probable, as when someone is skillfully framed for a crime of which she is innocent. If perceptual evidence in the case of illusions consists of true propositions, what are they? The obvious answer is: the proposition that things appear to be that way (Williamson, 2000, pg. 198).

Given normal background conditions, the fact that it appears to you that p is a reason to believe p . One's perceptual faculties wouldn't be very trustworthy if this were false.⁸ Williamson thus thinks that the reasons that rationalize the beliefs of the deceived are facts about how the world appears to them.

Note that if this is going to work for all the cases, it has to be that *nearly every time* we perceive p , we *know* that it appears that p . In other words, if this is going to work, Always Knows has to be true:

Always Knows: *Ceteris paribus*, every time A perceives p , A knows that it appears that p .

⁸Note that I'm *not* making a claim about reliability. I'm making a claim about epistemic trustworthiness. It might be that perceptions are trustworthy but not reliable (although they won't ever be trustworthy when one knows they're unreliable). The claim in the text thus doesn't conflict with claims made by some internalists about justification (see, e.g., Silins (2005), Pryor (2000)).

Always Knows is *prima facie* implausible. After all, in some cases one simply doesn't have time to go on to form the extra belief that it appears that *p*. And sometimes it would simply be a waste of time to go on to form the extra belief that it appears that *p*.

In the sentences immediately following the passage already quoted, Williamson addresses this challenge. He writes

Of course, unless one has reason to suspect that circumstances are unfavorable, one may not consider the cautious proposition that things appear to be that way; one may consider only the unqualified proposition that they really are that way. But it does not follow that one does not know that things appear to be that way, for one knows many propositions without considering them. When one is walking, one normally knows that one is walking, without considering the proposition. Knowing is a state, not an activity. In that sense, one can know without considering that things appear to be some way (*ibid.*).

This passage is a bit opaque, but I think that with a little unpacking we can see how the P→K theorist has a somewhat plausible response to the challenge. I take it that Williamson agrees with the thought that it would be a waste of time to always form a certain kind of belief about how things appear to you upon having a perception. This would require you to *consider* whether things appear to you in a certain way. He argues here, though, that one needn't consider something in order to know it—one usually knows one is walking when one is walking without considering whether one is walking. I'm not entirely sure how Williamson wants to spell this out. The way I think of it is that you have tacit knowledge when you know things you've never considered. This doesn't seem crazy. It does seem like, when I'm walking, I have tacit knowledge that I'm walking.⁹

Even granting that the facts about how things appear to the deceived are possessed reasons, one isn't home free yet. For we need the facts about how things appear to the deceived to be *sufficiently weighty* to make the deceived's beliefs rational. Let's focus on a particular belief of Pam's. Suppose both Sam and Pam have a visual perception that represents what looks to be a red table. Sam sees an actual red table, Pam does not. However, it appears to Pam that there is a red table. It's plausible that this fact is a reason to believe *the table is red*. Moreover, it's not crazy to think that Pam knows that *it appears that the table is red*.

I think it's also quite plausible that this reason is sufficiently strong. This is because what's relevant to whether Pam's belief is rational is the weight of the reasons Pam possesses. Given the reasons Pam possesses, it's plausible that the reason provided by the appearance fact is sufficiently weighty to rationalize Pam's belief. In order for the reason not to be sufficiently weighty, she would have to possess some defeater. This would have to be some strong evidence against the claim that the table is red or it would have to be some kind of undercutting defeater—some fact that severs the indicatory connection between her appearance and the table being red. Now certainly there are some facts that would be defeaters of either kind *if Pam came to possess them*. The rub, though, is that she doesn't *possess* any of those reasons. The world seems to her the same way the world seems to Sam. Thus, it seems plausible that the reason provided by the fact that the table appears red is sufficiently strong for Pam.

⁹Unfortunately, not nearly enough work has been done on tacit knowledge. Most work suggests that it will be quite hard to successfully analyze tacit knowledge as a state we're actually in when we have tacit knowledge. See especially (Lycan, 1988, ch. 3).

Crucially, I think that it follows from the nature of non-factive internal state duplicates that Sam possesses a sufficient reason to believe p just in case Pam does. This is because if Sam possesses some reason r to believe p , then Pam will have some corresponding non-factive state that allows her to know that it appears that r . And if the reason provided by Pam's appearance gets defeated, then Sam will have some internal state that is also a defeater of r . For example, if it appears to Pam that r' and r' defeats r , then it will also appear to Sam that r' . And this will be true for any r and r' .

We should reply in the same way to the Knowledge from Falsehood problem. There are other possessed sufficient reasons for Fritz to believe he won't be late. For example, Fritz knows that it appears to be 2:58 pm. He also knows that his extremely reliable watch says that it is 2:58 pm. Because of the reasons provided by these pieces of knowledge, he also knows that it's earlier than 4 pm, earlier than 5 pm, earlier than 6 pm etc. Those are all sufficient reasons Fritz possesses to believe that he won't be late.¹⁰

We can also say similar things to solve the problem that Schroeder poses. Indeed, we needn't go past the first Williamson quote to have a (partial) solution to Schroeder's problem. The sufficient possessed reason for one to believe as one does in paradigm cases of non-veridical perception is the fact that things appear to be thus and so. The appearance fact is what is doing the rationalizing.

My explanation of Same Status is thus quite simple. There are nearby truths that are sufficient reasons to believe the relevant propositions. Moreover, those in the bad cases do possess those reasons. Thus, their beliefs are *ex ante* rational—supported by possessed sufficient reasons.¹¹

Accounting for Same Degree

If what I said in the last subsection is true, then both Sam and Pam are in a position to have rational beliefs. Both of their beliefs are well supported by the reasons they possess. As we saw above, this is weaker than the conclusion many want to draw about characters like Sam and Pam. Many want to say that not only are Sam and Pam both rational, they are *equally rational*.

Before we start, let me say that I am not interested in showing that all of Sam and Pam's beliefs are *necessarily* rational to the same degree. I don't think that's true. But this is just part and parcel of being the type of externalist I am.¹² What I aim to show is that, given a very natural way of understanding the cases, there is an important sense in which Sam and Pam are rational to the same degree. Moreover, I am going to argue that the ways in which my view predicts they are not rational to the same degree are virtues of my theory, not bugs.

With that said, it's not always clear what is meant by 'equally rational.' Often times, it seems to me, this is synonymous with 'don't have the same rationalizers.' If that is what is meant by 'equally rational', then of course the externalist has to deny that Sam and Pam are equally rational! One might find externalism implausible on its face, but mounting an objection to the view in this way is dialectically unhelpful.

I think the best way of thinking of a set of reasons making a belief more rational is by thinking that a belief, b_1 , is more rational than a belief, b_2 , just in case the set of possessed reasons support-

¹⁰Perhaps not surprisingly, this is the most popular solution offered in the literature on knowledge from falsehood (see Coffman (2008), Arnold (2011), Fitelson (2010)).

¹¹I imagine some readers are jumping up and down about the fact that I've said nothing about how those in the bad cases are suitably connected to the nearby truths to count as *ex post* rational. Fear not, I will deal with that issue in a bit.

¹²Indeed, I think it's very plausible that this is part and parcel of being an externalist at all.

ing b_1 is *weightier* than the set of possessed reasons supporting b_2 . At the very least, this is a natural starting point. The weight of reasons comes in degrees. The phenomenon we're investigating—the relative rationality of beliefs—also comes in degrees. It's thus natural to think that the relative rationality of beliefs tracks the weight of the possessed reasons supporting those beliefs.

One might be a bit surprised that I've initially committed myself to this view of what it is for one belief to be more rational than another. After all, at first blush it seems very plausible to think that Sam must be more rational than Pam if externalism is correct because the reasons Sam possesses are, in an obvious way, better than the reasons Pam possesses. After all, Sam possesses the reason constituted by the fact that the table is red, whereas Pam only possesses the reason constituted by the fact that it appears as if the table is red. A natural way to spell out the way in which Sam's reason is better than Pam's is in terms of weight—it's better because it's weightier.

While I agree that this is a natural thought, I think that it is mostly irrelevant to the question at hand. This is because it confuses what we can call a reason's *atomic* weight with a set of reasons' *relative* weight. Atomic weight—if it exists—is the weight that a reason has independently of how it interacts with the other reasons in a particular situation.¹³ Relative weight is the weight that sets of reasons have after they are weighed against reasons for alternatives in any given situation. Sam's reason is *necessarily* weightier only in the atomic sense. We can't know what a set of reason's relative weight is until we know what the other reasons are.

The rub, though, is that the relevant weighing for assessing how rational a particular belief is is the relative weighing. So we aren't going to learn that Sam's belief is more rational than Pam's just by learning that her reason has greater atomic weight. We need to know what other reasons there are.

Notice, though, that the natural way to fill in the details of The Good Case and The Bad Case is to assume that not only does Sam have better reasons for the beliefs she holds, she also has better reasons not to hold the beliefs she holds. So while Sam *knows* that she sometimes makes perceptual errors, Pam only knows that it appears to her that she sometimes makes perceptual errors. Sam might know that Tom said the store was to the left (even though she believes it to be to the right; she knows Tom is terrible at directions and has other reasons to believe it to be to the right), but it only appears to Pam that Tom said the store was to the left, etc.

If this is how we fill in the case, then the reasons Sam has for not believing what she does will have greater atomic weight than the reasons Pam has for not believing what she does. Indeed, I boldly conjecture that the difference in relative weight between Sam's and Pam's reasons for believing what they do is *identical* to the difference in relative weight between Sam's and Pam's reasons for not believing what they do in fact believe. If this is right, then the so-called gap between the reasons for believing and reasons against believing will be the same. This just is a way of measuring relative weight. If this is right, then there is a natural sense in which the relative weight of Sam's reasons is the same as the relative weight of Pam's—the ratio between the reasons for and reasons against will be the same for both Sam and Pam.

Of course, we're not forced to fill in the cases this way. It's possible for Sam to just have the reasons Pam does against what she believes even though she has more reasons to believe what she does believe. In that filling out of the case, the relative weight of Pam's reasons for believing what she does will be greater than the relative weight of Sam's reasons for believing what she does. Two things to note about this. First, it's not damning because it's so obviously true. Second, this is *not*

¹³I am going to assume for now that there are atomic weights. I will question this below.

the most natural way of filling out the case. The most natural way of filling out the case is that where Sam is well-informed about the way the external world is, Pam is misinformed. This goes for reasons against as much as it goes for reasons for. Thus, I don't think those who endorse Same Degree are thinking of the case in this way. They are thinking of it in a way that suggests that the relative weight of Sam's reasons is the same as the relative weight of Pam's.

Thus, I think that there is a true reading of Same Degree. Sam and Pam are equally rational because the relative weights of the sets of reasons they possess for the beliefs they hold are the same. Moreover, I think that this is the most important reading of Same Degree. This is because what we should be interested in determining the relative rationality of a group of agents is the relative weights of the sets of reasons they possess.

I've run the explanation by appealing to atomic weights. It is important to note that this isn't necessary. As long as there is some way to measure the strength of all the reasons together—i.e., as long as there's a way to measure the strength of the reasons all things considered—then we will have what my explanation needs. And anyone who thinks that we can build up strict notions like obligation, justification, and rationality out of reasons and their weight is committed to thinking that there is some relative notion.

On the other hand, not everyone is committed to thinking that there is some atomic notion. Indeed, in the epistemic case it is a bit mysterious what the atomic weights could be. This is because the orthodox way of measuring the strength of some evidence is by using conditional probability. We can measure the strength of a reason r to believe p by measuring how likely p is given r and some background information. Usually, the background information will just be the rest of one's evidence. So on this view, there really isn't anything like atomic weights. This is because in order to measure the strength of any particular reason r , you have to use the other evidence. So each measure is relative in an important way.

There are two important conclusions to draw from this. First, it seems like the anti-externalist thought we started out with—viz., the thought that Sam's reasons must be weightier because her reason is weightier—tacitly assumes a notion of atomic weight that is mysterious. Second, my explanation doesn't need to assume that there are mysterious atomic weights.

There is another objection to my idea in the ballpark. According to this line, Sam's reasons must be weightier because Pam has more reasons to believe what she believes than Pam does. It appears to Sam that the table is red just like it appears to Pam that the table is red. So Sam not only has *the table is red* as a reason to believe it's red, she also has the appearance fact as a reason to believe it's red. So she has two reasons to Pam's one. It is natural to think that if A has reasons r_1 and r_2 to ϕ and B just has r_1 , then A is guaranteed to have weightier reasons to ϕ than B .

This is natural, but also resistable. To see this, consider my belief that $2+2=4$. My belief that $2+2=4$ is maximally rational. Moreover, my maximal rationality is *overdetermined* in an important way. That is to say, there are several proper subsets of my set of possessed reasons to believe that $2+2=4$ that are all strong enough on their own to make me maximally rational in believing $2+2=4$. Consider two of the subsets, S_1 and S_2 . By hypothesis, each is weighty enough to maximally rationalize—so each is as weighty as they get. Now consider the union of S_1 and S_2 . Imagine I came to possess the set comprised of the union, call it S_3 . Plausibly, possessing S_3 is in some way *better* than merely possessing S_1 or S_2 . But it can't be that my belief that $2+2=4$ is more rational given S_3 than in it is given merely S_1 or S_2 . My belief is maximally rational no matter which set I possess. So what's better about possessing S_3 ?

I think the answer is that the rational status of my belief is more modally robust when I possess S_3 . One has to go further out into modal space to get to a world where the set of my possessed reasons to believe $2+2=4$ fails to rationalize a belief in $2+2=4$. It's *not* that S_3 is weightier in any helpful sense of the word. Rather, it's just as weighty as S_1 and S_2 . It's rather that its weightiness is overdetermined.

I think that something similar is happening with Sam and Pam (although I also think there is a feature of the $2+2=4$ case that isn't present in the Sam and Pam cases). It's not that the set of reasons Sam has is weightier in any helpful sense of the word. It's rather that her rationality is overdetermined. She could fail to have one of the reasons she does have and still be rational.

A promising objection to this application of the overdetermination thought is that it might seem to commit me to the idea that Pam's (and Sam's) belief that the table is red is maximally rational. After all, the fact that there are multiple proper subsets of my evidence for $2+2=4$ that maximally rationalize is what started us down the overdetermination path to begin with. Because of this, it's plausible to think, at this point, that the proper application of the overdetermination lesson is only in cases where one's reasons maximally rationalize.

This isn't right. It's just that the cases of maximal rationalization are the easiest to glob onto. I think what is overdetermined in these cases is the *sufficiency* of the (relative) weight of reasons to believe $2+2=4$. When I possess S_3 , the sufficiency of my reasons to believe $2+2=4$ is more overdetermined than when I possess just S_1 or S_2 . Similarly for Sam. The sufficiency of the set of her reasons to believe the table is red is more overdetermined than the set of Pam's reasons to believe the table is red.¹⁴

The important point is that it doesn't follow that the relative weight of Sam's reasons is greater than the relative weight of Pam's just because Sam has two reasons to Pam's one. Sam's reasons are more overdetermined than Pam's, but the relative weight is the same. It's just that Sam's reasons has that weight further out in modal space than Pam's.

I'll close by noting that if something like this is correct, then the externalist has a ready answer to a seemingly forceful complaint made by Silins (2005). Silins writes,

As far as propositional [i.e., *ex ante*] justification is concerned, the evidential externalist should allow that internal twins can have different degrees of justification for their beliefs. Suppose an externalist said that, even though internal twins have different evidence, they are equally justified in their beliefs. Perhaps, given that one is not certain of what evidence one has in the good case, the differences between the evidence one has in each case end up washing out. On this sort of externalist position, however, evidential differences between internal twins are epistemically idle. Evidential differences here would not even explain differences in knowledge—it seems that one can have knowledge in the good case even if one has the same evidence in the bad case. However, I take it that evidential differences between internal twins are interesting and important only if those differences can generate differences in propositional justification between twins. So I take it that, if Evidential Externalism is true, then internal twins can fail to have the same degree of propositional justification (386).

Silins seems to think that if Sam's reasons for believing the table is red are better than Pam's, then it must be because they are weightier. Moreover, he thinks that if the externalist denies that Sam's rea-

¹⁴The difference between Sam and me, I think, is that S_3 overdetermines the sufficiency much more than Sam's reasons do.

sons are weightier, then the difference between Sam and Pam is epistemically idle. I think neither of these claims are true. Sam's reasons might be more robust than Pam's even though Pam's reasons are just as weighty as Sam's. Moreover, the robustness of Sam's reasons is not epistemically idle. She is in a better position than Pam precisely because the weight of her reasons is overdetermined.

There are thus two parts to my accounting for Same Degree. Most importantly, I think that there is a natural way in which Sam and Pam are equally rational—i.e., there is a reading of Same Degree that is true. Sam and Pam are equally rational because the relative weight of the set of reasons Sam has for believing what she believes is equal to the relative weight of the set of reasons Pam has for believing what she does. If you posit atomic weights, then there is at least one sense in which Sam's reasons are weightier than Pam's. The reasons Sam has for believing as she does have greater atomic weight than the reasons Pam has for believing what she does—but, then again, the reasons Sam has for not believing as she does also have greater atomic weight than the reasons Pam has for not believing as she does. Despite these (purported) facts about the atomic weights, atomic weights are not the weights we should be most interested in. We should be most interested in the relative weighing. This is because a belief that p is (*ex ante*) rational only if the weight of the reasons possessed to believe p is at least as weighty as the set of reasons to disbelieve p and the set of reasons to withhold belief in p . This is a relative weighing. Moreover, it's quite plausible that the reasons Sam has to believe as she does have the same relative weight as the relative weight of the set of reasons Pam has to believe as she does.

The second part of my story is about what it is that makes Sam's position better than Pam's. Some have argued that to concede that there is something better about Sam's situation is to concede Same Degree. This is true in one sense. If you think *what it is* for Sam to be more rational than Pam is for Sam to be in a better position than Pam, then obviously I deny Same Degree. But I don't think this is the right way to see things. For I think the best way of thinking of Same Degree is in terms of the relative weights of Sam and Pam's reasons. Those are the same, and thus I think we should think of them as equally rational. On the other hand, Sam is in a better position. This is because her reasons are more modally robust. The reasons she has to believe as she does would retain their relative weight further out into modal space than Pam's reasons would. This is why Sam is in a better position. This is compatible with thinking Same Degree is true.

6.4 But Wait, There's More: The New New Evil Demon Problem

The new evil demon problem, as I've characterized it, is about *ex ante* rationality. It is thus primarily a search for suitable rationalizers of the beliefs of those in the bad cases. Moreover, as we saw, the Knowledge from Falsehood problem and Schroeder's non-veridical perception problem were also largely about what rationalizes subjects in less than ideal circumstances. I argued above that there are plausible solutions to these problems.

Unfortunately for me, the New Evil Demon problem is not the only, nor is it the hardest, problem for externalists posed by the deceived. This is because we are not only interested in *ex ante* rationality, we are also interested in *ex post* rationality. It's not just that Pam is *ex ante* rational—i.e., it's not just that the things she believes are sufficiently supported by possessed reasons; her beliefs themselves are rational, at least if Sam's are.

The thing that fills the gap between *ex ante* rationality and *ex post* rationality is *basing*. In order for a token belief to be *ex post* rational, there must be the right connection between the belief and

the rationalizers of the belief—on my view, the reasons that sufficiently support the belief. Thus, if Pam is *ex post* rational, it must be that her beliefs are connected to the reasons that sufficiently support her belief. The same goes for the likes of Fritz and for those who have rational non-veridical perceptual beliefs.

This problem is much harder than the original problem. This is because it's completely mysterious how to account for the full range of cases, at least given the orthodox understanding of how large the range of cases is (foreshadowing!). On the orthodox understanding, there are two interesting types of cases. The first type of basing relationship is the *inferential* relationship. Many of Sam's and Pam's beliefs will be inferential in their origins. For example, both Sam and Pam might infer certain historical beliefs from other beliefs they have about what their history textbooks say (or, in Pam's case, what appears to be a history textbook). Focus on a particular belief: the belief that Fermat's last theorem was proved by Andrew Wiles.

Sam infers that Fermat's last theorem was proved by Andrew Wiles from her knowledge that her textbook says that he proved Fermat's last theorem. Since her textbook is reputable, this seems like a sufficient reason to believe that Wiles proved Fermat's last theorem. Moreover, provided that her inference was normal, it's very plausible that her belief is *ex post* rational partially because of the inference. By inferring the belief from a rational belief in a sufficient reason, she comes to base her belief on the sufficient reason.

The etiology of Pam's belief is very similar to the etiology of Sam's. She also infers that Wiles proved Fermat's last theorem from a belief that her history textbook says so. However, there is no history textbook in front of Pam and in her world Wiles didn't prove Fermat's theorem. Thus, given Reasons Factivity, Pam's belief cannot be rationalized by the reasons that rationalize Sam's beliefs since those aren't reasons for Pam. This means that Pam's inference can't be doing the epistemic work that Sam's inference is. Sam's inference is what's connecting her to the reasons in the right kind of way to make her belief *ex post* rational. Pam's inference must not be doing this since the content of the belief is not a reason to have the belief. Above I argued that the reason that sufficiently supports Pam's belief is something like *it appears that a reputable history textbook is reporting that Andrew Wiles proved Fermat's last theorem*. It's very implausible that Pam *infers* her belief from *that* claim. Thus, if she is going to be suitably connected to the sufficiently strong possessed reason, it looks like the connection can't be inferential.

Note that this problem also crops up in Fritz's knowledge from falsehood case. Fritz infers that he won't be late to his 7 pm meeting from his belief that it is exactly 2:58 pm. Given Reasons Factivity, *it's exactly 2:58 pm* cannot be the sufficiently strong possessed reason that rationalizes Fritz's belief. Thus, Fritz's inference must not be what's putting Fritz in suitable contact with the rationalizer of his belief. So, if he is *ex post* rational—which he must be if he has knowledge—then the process by which he becomes suitably related to his sufficient reasons must not be inferential.¹⁵

The second type of case is the *non-inferential* case. It's very plausible that some of Sam's and Pam's beliefs have non-inferential origins. The paradigm non-inferential beliefs are perceptual beliefs. Let's take a particular perceptual belief of Sam's: the belief that the table in front of her is red. Given that this is a normal perceptual belief, it's plausible that Sam and Pam don't infer this belief from some other beliefs. The etiology is rather different. The belief is rather caused by their perceptual systems themselves—the etiology is sub-inferential in some sense. In Sam's case, this perceptual system puts her into appropriate contact with a sufficient reason to believe that the table

¹⁵This is pressed forcefully in the knowledge from falsehood literature by Fitelson (2010), Arnold (2011).

is red. It's controversial which reason that is; for simplicity's sake, I'll stick to the simplest proposal and the proposal I like best: the reason is the fact that the table is red.¹⁶ Even though Sam doesn't infer the table is red from another belief, she is still suitably connected to her reason because of the etiology.

The etiology of Pam's belief is very similar to the etiology of Sam's. Pam's belief, just like Sam's, is caused by her perceptual system—it's sub-inferential in the same way. The only difference is that Sam's belief is in some sense caused by the redness of the table, whereas Pam's isn't. This is a crucial difference, according to those who believe Reasons Factivity. For it means that only Sam's belief can be rationalized by *the table is red*, despite the similar etiologies. Above I argued that the sufficient reason Pam possesses to believe the table is red is rather the fact that it appears to her that there is a red table. But it's plausible that her belief isn't connected to *that* reason in the way Sam's belief is connected to the fact that the table is red—it's not paradigmatically non-inferential. But, again, it's not inferential either. So Pam's belief, if it is going to be *ex post* rational, has to be connected to her sufficient possessed reasons in some other, mysterious, way.

The problem is magnified if we make a widely held assumption about the nature of non-inferentially rational beliefs. They are often thought to be *foundational*. Foundational beliefs are rational beliefs whose rational status doesn't depend on one having any other rational beliefs. One doesn't need other rational beliefs in order to have rational non-inferential beliefs. This is what makes non-inferential beliefs foundational beliefs.¹⁷

It's natural to think that both Sam's and Pam's beliefs about the table are foundational. Sam's certainly seems foundational. She doesn't need to rationally believe anything else in order to rationally believe that the table is red. After all, her perceptual system puts her into contact with the fact that the table is red. This is enough to rationalize her belief. But this can't be what's going on with Pam if Reasons Factivity is true and my story in the last section is correct. For according to that story, in order for Pam to be rational, she must know that the table appears red to her. Thus, Pam does need to have another rational belief—the belief that it appears that the table is red—in order to be rational. Thus, her belief isn't foundational. On the other hand, it's not inferential either, which is the paradigm type of non-foundational rational belief.

It's interesting to note that Schroeder presents a more nebulous version of this very problem. He writes,

That means that this view predicts that internal duplicates with duplicate pasts in very similar environments can nevertheless differ not just in whether they know, but in whether their belief is justified [i.e., rational], if for the very first time one sees something green in front of her while the other experiences a vivid hallucination as of something green in front of her, and both form the perceptual belief that there is something green in front of them...So unfortunately the high bar on having evidence is not a bar that can be met in the full range of foundational perceptual cases. At most it can be met in the case of veridical perceptual experiences, but it would be best to avoid need-

¹⁶Other externalist candidates include the fact that Sam sees a red table (see, e.g., McDowell (1994), Neta & Pritchard (2007), Pritchard (2012)).

¹⁷I am running roughshod over my careful parsing of the terminology in the last chapter. In the terminology of the last chapter, the foundational beliefs are both immediate and uninferential. However, I am going to use uninferential in a slightly different way in this chapter.

ing to distinguish between the justification provided by veridical and by non-veridical perceptual experiences (Schroeder, 2011, pg. 217).

The reason why I think this is nebulous is that it runs together the New Evil Demon problem and the New New Evil Demon problem. I don't think Schroeder ever shows that, according to externalists of a certain kind, internal duplicates will differ when it comes to *ex ante* rationality (or justification). But it's plausible that he shows that externalists of a certain kind can't think that those in the bad cases will have rational beliefs that are *foundational*. If you think that all the foundational beliefs are non-inferential, you think that there are only inferentially or non-inferentially based rational beliefs, and you think that those in the bad case don't have inferentially rational perceptual beliefs, then it follows that internal duplicates will differ when it comes to *ex post* rationality.

The New New Evil Demon problem, then, is a puzzle about basing. On the one hand, those committed to Reasons Factivity have trouble accounting for some seemingly inferentially rational beliefs that are based on falsehoods. If Reasons Factivity is true, then it can't be that one's beliefs are rationalized by falsehoods. So they can't be rationalized by the propositions they are inferred from. But then one must tell some story about how it is that there is a suitable connection between one's beliefs and actual possessed reasons. It doesn't seem like the connection can be inferential. Moreover, it doesn't seem like it can be non-inferential in the usual sense of non-inferential. This is because it's clear that the beliefs aren't foundational beliefs. Indeed, according to the story I told above, in order for Pam to have an *ex post* rational belief that Andrew Wiles proved Fermat's last theorem, she has to have a *ex post* rational belief that it appears to her that a reputable history textbook says Andrew Wiles proved Fermat's last theorem. So it can't be a foundational belief and thus it doesn't seem like it can be non-inferential. Thus, it can't be inferential or non-inferential.

On the other hand, those committed to Reasons Factivity have trouble accounting for perceptual beliefs caused by non-veridical perceptions. While there is a nice story to tell about The Good Case—perceptual beliefs in The Good Case are based on their contents—this story doesn't carry over to The Bad Case. Pam's belief that the table is red can't be based on *the table is red* because that isn't a reason for Pam. However, Pam's belief clearly isn't inferential. Moreover, according to the story above, her belief can't be foundational either. This is because in order for her belief that the table is red to be rational, she has to rationally believe that it appears to her that there is a red table. If it's not foundational, then it doesn't seem like it can be non-inferential. So it's not inferential or non-inferential.

6.5 Solving The New New Evil Demon Problem

The main idea behind my solution to the New New Evil Demon problem is that there is another type of basing besides paradigmatic inferential basing and paradigmatic non-inferential basing. I'll call this type of basing *uninferential basing*.¹⁸ Here are some examples of uninferential basing:

Lottery

¹⁸This type of basing has been bumped up against at points in the literature, although no one has picked it out as relevantly different than inferential and non-inferential basing. See especially Kelly (2002), Evans (FC), Harman (1973).

I learn that Richard won the lottery. I then infer from this knowledge that Richard is rich. Suppose, though, that I later come to learn that when Richard won the lottery he was in deep debt. He had so much debt, in fact, that winning the lottery didn't make him rich; it only allowed him to get out of debt. I also learn, however, that he has subsequently become rich because of some good investment decisions. This knowledge becomes the basis of my belief that he is rich *and* my knowledge that he won the lottery ceases to be the basis of my belief that he is rich.

Red Widgets¹⁹

You are touring a widget factory. You visually perceive a red widget and come to believe that it's a red widget. The factory owner comes up to you and says 'That's a fine red widget, isn't it? I just finished painting it.' He then goes on to tell you that you have to be careful about which widgets are red because they light the place with red lights. You look up and, sure enough, the factory floor is lit by red lights. The fact that the widget looks red ceases to be the basis of your belief but you still believe the widget is red because of the owner's testimony—that becomes the basis for your belief.

In both of these cases, it's very plausible that the basis of your belief changes over time even though you never drop the belief. In Lottery, you continuously believe that Richard is rich even though the basis of that belief changes over time. For a while you believe this because Richard won the lottery. You later come to find out that that reason wasn't sufficient to rationalize believing that he was rich because of his debt. However, you don't drop the belief because you also find out that Richard is rich through smart investments. This fact becomes the basis of your belief. Something similar happens in Red Widgets. You continuously believe the widget is red despite the fact that the basis for this belief changes. At first you believe it because it looks red. However, the factory owner alerts you to a defeater of this reason—the red lights. Before defeating your original reason, though, the factory owner provides another sufficient reason through his testimony. This testimonial reason becomes your belief's new basis.

Despite the fact that these are clear cases of basing, I don't think they have the hallmarks of either inferential or non-inferential basing. They are obviously not non-inferential in the way we've been understanding non-inferential. This is because the beliefs in question are obviously not foundational. You only rationally believe Richard is rich because you rationally believe he got rich through smart investments and you rationally believe it's a red widget only because you rationally believe the owner confirmed it was red. Since the output beliefs aren't foundational, the basing isn't non-inferential.

On the other hand, it's not obviously inferential, either. This is because paradigm cases of inference are *creative*. New beliefs come into existence because of the inference. Inference is a deliberative activity that results in new beliefs. This is not happening in these cases. In both cases your output belief stays the same. What changes is how those beliefs relate to other information you possess.²⁰ In this way, the basing isn't paradigmatically inferential.

¹⁹This comes from Evans (FC) and is a variation of a case from Pollock (1974).

²⁰On some conceptions of inference, these cases might be inferential. Wedgwood (FCa), e.g., holds that as long as you 'reaffirm' a belief in a certain type of way, you count as inferring the belief. I would be fine if this turned out to be inferential. My point here is just that it isn't paradigmatically inferential.

It will be important in what follows to have some idea of why it is that the basis of your beliefs change in these cases. Although I have a worked out theory of exactly what is going on, I don't think we need to rely on this theory to get the basic idea about what's happening in these cases.²¹ The basic idea is that you no longer base your belief that Richard is rich on the fact that he won the lottery when you find out about the debt because your belief that he's rich ceases to be *sensitive* to your belief that he won the lottery. On the other hand, when you find out about his successful investing, your belief that he is rich becomes sensitive to your belief that he became rich after successful investing. It's hard to know exactly what this sensitivity comes to (see chapters 4 and 5 for what I think it comes to), but I think a necessary condition is (roughly) that you are disposed to drop your belief that he's rich when either you stop believing he became rich through successful investments or that belief is defeated.²² It seems like you do have this disposition when it comes to your belief that he became rich through investment, but don't have this disposition when it comes to the belief that he won the lottery. This is the explanation for why the former is the basis of your belief and the latter isn't.

For all I've said, this is only correct when it comes to uninferential basing. Perhaps what it is generally to base a belief that q on p is not to be sensitive in this way. I doubt it. I think this type of theory generalizes nicely, but I don't want to be wedded to that here.

I am going to use this third type of basing to solve the New New Evil Demon problem. We'll take the two types of cases one at a time, starting with the inferential case.

6.5.1 The Inferential Case

We discussed two inferential cases. The primary case has to do with Sam and Pam's belief that Andrew Wiles proved Fermat's last theorem. The problem for Reasons Factivity is that the only reason that Pam plausibly possesses to rationalize her belief is constituted by the fact that it appears to Pam that a reputable history book says that Andrew Wiles proved Fermat's last theorem. But this isn't what Pam infers her belief from—she infers it from the claim that a reputable history book says AW proved Fermat's last theorem. So it doesn't seem like her basing is inferential. But it also isn't non-inferential since her belief that AW proved Fermat's last theorem isn't foundational—it depends on her rationally believing that it appears to her that a reputable history book says that AW proved Fermat's last theorem. Thus, it's not clear how her belief could be based on the reason provided by the appearance fact.

My answer, obviously enough at this point, is that her belief is uninferentially based on the reason provided by the appearance fact. It's plausible that her belief that AW proved Fermat's theorem is sensitive to her belief that it appears that a reputable history textbook says that he did. She is disposed to give up the former belief if she loses the latter. For example, if she found out that a demon was tricking her in this case, she would likely drop her belief that AW proved Fermat's last

²¹My view about basing—or believing for normative reasons—is defended in chapter 4. I utilize it in chapter 5 as well.

²²This actually isn't my view about the relevant necessary condition for being sensitive to the *objective reason* provided by the fact that he become rich through successful investment. My view is that you have to be disposed to drop the belief that he's rich if the reason constituted by that fact ceases to be a sufficiently strong possessed reason. This can happen either if you lose the reason entirely or if the reason becomes defeated. Due to the fact that that view is very controversial, I'll stick to the less controversial view in the text, which spells out the disposition just in terms of ceasing to believe.

theorem because her reason would then be defeated. Or if she gained reason to believe the particular history book wasn't reputable, then she would drop her belief about who proved Fermat's last theorem. More extremely, it seems like she would drop the belief that AW proved Fermat's last theorem if she stopped believing that it appeared to her that the book says he did. Since the appearance facts provide sufficient reason to believe that AW proved Fermat's last theorem and Pam bases her belief on that reason, her belief is *ex post* rational.

Not only is all of this plausibly true for Pam, but it also is for Sam. It's plausible that Sam's belief is also sensitive to the appearances. It's very plausible that she is disposed to drop her belief if she loses confidence in the trustworthiness of the appearances. Of course, it's also plausible that she is sensitive in the right kind of way to the way things actually are, unlike Pam. But this is just to say that not only is her *ex ante* rationality overdetermined, so is her *ex post* rationality. She bases her belief on more than one sufficiently strong reason. This seems like the right result, and it's a virtue of a view that embraces uninferential basing that it can explain how it is that Sam bases her belief on the reason provided by the appearance fact.

The same thing is going on, I conjecture, in Fritz's knowledge from falsehood case. His belief that he won't be late for the 7 pm meeting is sensitive to the reasons that actually rationalize his belief—e.g., the fact that his extremely reliable watch says it's 2:58 pm. It's plausible that he's disposed to drop his belief if he either stopped believing that or if that was defeated in some way. Of course, this doesn't change the fact that Fritz's belief is causally related to his belief that it's exactly 2:58 pm. His inference causally explains why he started believing what he does. But part of what explains (at least partially) why he doesn't drop the belief is the fact that his belief is sensitive to the appearances. Since the appearance facts rationalize his belief and his belief is based on those reasons, his belief is *ex post* rational.

6.5.2 The Non-Inferential Case

The non-inferential cases we focused on are all perceptual cases. Sam sees a red table; Pam has a phenomenologically indistinguishable perception and comes to believe that the table is red. The problem for Reasons Factivity is that while it's possible for Sam to base her belief on the fact that the table is red, it's not possible for Pam to base her belief on that reason since it's not a reason for Pam. My story about Pam's *ex ante* rationality is that her belief is *ex ante* rational because the fact that it appears as if there is a red table provides a sufficiently strong possessed reason to believe the table is red. The problem is that Pam's belief isn't inferred from that reason, nor is it non-inferentially based on that reason. Thus, it's not clear how Pam's belief can be based on the possessed sufficient reason.

Once again, my answer is that her belief is uninferentially based on the reason provided by the appearance fact. That is, her belief that the table is red is sensitive to whether the content of her belief that it appears as if there is a red table is a sufficient reason to believe that it is a red table. Other things equal, if she stopped believing it appeared as if there is a red table, she would stop believing there is a red table. Moreover, other things equal, if the reason provided by the appearance fact was defeated, she would stop believing there is a red a table. If this is right, then her belief is suitably connected to the sufficient possessed reason. Thus, she is *ex post* rational.

Recall that Schroeder claimed that 'it would be best to avoid needing to distinguish between the justification provided by veridical and by non-veridical perceptual experiences' (16). Given the contours of the present discussion, this is multiply ambiguous (which is not to say that it was

ambiguous in Schroeder's context). Schroeder might just mean that it would be best to not have to posit that Sam and Pam have different rationalizers. I take it he doesn't mean this, though, since it's obvious that externalists have to claim that Pam's belief, if rationalized at all, will be rationalized by different rationalizers than Sam's. More importantly, externalists think of this part of their view as a *virtue*. Williamson (2000), e.g., thinks that making this move is a key component of responding to a certain type of skeptic.

A more plausible interpretation is that Schroeder thinks we shouldn't distinguish between the way in which Sam's and Pam's beliefs are based. If we think that the only types of basing are inferential and non-inferential, then the problem is acute. It's very implausible in the perceptual cases (and the inferential cases, for that matter) that Pam either inferentially bases her belief *or* non-inferentially bases her belief on the appearance fact. Nevertheless, I've argued, she is still sufficiently sensitive to those facts to count as basing her belief on those facts. This sensitivity doesn't arise because her beliefs about the appearances play a direct causal role in the *formation* of the belief. The perceptual process is what directly causes the production of the belief. Nonetheless, the belief itself is sensitive to the appearance facts in the right kind of way.

Of course, we've already noted that Pam's perceptual beliefs won't be foundational. They are only rational because she holds rational beliefs about the appearances. Schroeder seems to think that this is a cost. I think the externalist should be happy to embrace this result. This is because the externalist should think that, at least for those of us in favorable epistemic worlds, our foundational beliefs about the external world are rationalized by the external facts themselves. This is one side of the coin that gives us an advantage over the skeptic. Moreover, despite the fact that Pam's belief about the color of the table isn't foundational, it's plausible that her belief about the appearances is foundational. It is a type of foundational self-knowledge. Thus, at least in this case, the chain of non-foundational beliefs only goes back one link before getting to a foundational belief.²³

6.6 Conclusion (With A Reason to Reject $P \rightarrow K$)

The discussion so far has assumed that the possession relation is knowledge. It's important to point out that there is at least one other view of possession that can accept my solutions to the Evil Demon problems. This is the view that the possession relation is being in a position to know— $P \rightarrow PTK$. This is because one is automatically in a position to know what one knows. Thus, if $P \rightarrow PTK$ is true, then one has all the reasons $P \rightarrow K$ predicts one has. Moreover, nothing about $P \rightarrow PTK$ prevents the possibility that Pam's beliefs are uninferentially based on the reason she has in virtue of knowing that it appears that *p*.

It is not my aim in this chapter to fully adjudicate between these two views. However, it is worth pointing out that $P \rightarrow PTK$ has an easier time solving the New Evil Demon problem.

$P \rightarrow PTK$ has an easier time solving the New Evil Demon problem because it can do so using weaker assumptions than $P \rightarrow K$. Most importantly, the $P \rightarrow PTK$ theorist needn't be committed to Always Knows:

Always Knows: *Ceteris paribus*, every time *A* perceives *p*, *A* knows that it appears that *p*.

²³Surely things are a bit more complicated in some of the inferential cases, but I think a foundational belief won't ever be too far from the non-foundational belief in question.

The $P \rightarrow K$ theorist needs Always Knows in order to secure the claim that in normal cases one possesses the reason provided by the fact that it appears that p when one has the perceptual experience that p . However, according to the $P \rightarrow PTK$ theorist, one needn't actually know p in order to possess it. One merely needs to be in a position to know p . Thus, in order to solve the New Evil Demon problem the $P \rightarrow PTK$ theorist just needs Always PTK:

Always PTK: *Ceteris paribus*, every time A perceives p , A is in a position to know that it appears that p .

Always PTK is much more intuitively plausible than Always Knows. It is not at all obvious that every time one perceives p , one knows that it appears that p . However, it is very plausible that every time one perceives p , one is in a position to know that it appears that p . Thus, $P \rightarrow PTK$ has at least this advantage over $P \rightarrow K$. Whether there is conclusive reason to accept $P \rightarrow PTK$ over $P \rightarrow K$ is a topic for another occasion.

My main aim in this chapter was to give an existence proof of a plausible externalist view of rationality that solves two problems posed by those in less than ideal environments. I've argued that there are at least two plausible externalist views that predict that the deceived can not only be rational, but can also be rational to the same degree as their undeceived counterparts. This doesn't mean, however, that those of us in favorable environments don't have epistemic advantages, but it does show that one can be externalist and respect the rationality of the radically deceived.

Chapter 7

What You're Rationally Required to Do and What You Ought to Do (Are the Same Thing!)

"Be reasonable" is a categorical imperative. (Gibbons, 2010, pg. 352)

The question of permissibility is a question that can be asked by a deliberating agent, and one that a normal agent can be expected to be able to answer. The answer to this question is not just a matter of what is in fact the case (whether anyone could know it or not). But at the same time, permissibility is not merely a matter of what a particular agent believes the facts to be. It also depends on what is reasonable for the agent to believe in the situation, what it is reasonable for the agent to do to check those beliefs, and whether the agent has done those things. (Scanlon, 2008, pgs. 51-2)

7.1 Introduction

IT has been notoriously difficult to vindicate the natural thought that moral considerations necessarily bear on what one ought to do. That is to say, it has been difficult to explain why it is that the facts that determine what one is morally required to do necessarily affect what one ought to do, full stop. The requirements of morality thus face a sort of skeptical challenge. The challenge is to explain the *deontic significance* of morality. In order to meet the challenge, one must explain why we always have most reason (or any reason) to do what morality requires.

Despite this skepticism about morality, other types of requirements have traditionally been seen to be on steadier deontic footing. A paradigm case is *rationality*. Rationality, it is often thought, is clearly deontically significant. The considerations that bear on what's rational seem to necessarily affect what one ought to do, full stop. In fact, something stronger seems to be true. Namely, that you always ought to be rational. It does not seem like much of a challenge at all to explain why we always have most reason (and thus some reason) to do what rationality requires. It is a striking fact that many of the classic positions—both skeptical and non-skeptical—concerning the deontic significance of morality assume the deontic significance of rationality.¹ This is evidence enough that rationality is usually seen to be on steady deontic ground.

¹This breaks down roughly into Humean skeptics and Kantian rationalists. On the skeptical side, see (just to name two) Foot (1972), Harman (2000). On the rationalist side, see Smith (1994), Korsgaard (1996).

Despite this (or perhaps because of it), it has become increasingly popular in the literature on practical reason to deny that you should be rational.² And given some assumptions about rationality and about what determines what you should do, this actually seems plausible. Moreover, those assumptions about rationality and about what determines what you should do are thought to be supported by powerful arguments. Because of these rather surprising facts, skepticism about the deontic significance of rationality is alive and kicking.³ Back to this in a moment.

Another corner of normative philosophy is concerned with what determines what one ought to do, full stop. This debate is about what the best theory is of what we might call the *ought of deliberation*. This is the 'ought' that figures in the central deliberative question: What ought I do?⁴ It's commonly assumed that the answer to the central deliberative question is the thing that you ought to do, full stop—it *settles (correct) deliberation*. What you ought to do and believe full stop is, we'll say, what you ought to do and believe. The debate I'm interested in here is about which facts *determine* what you ought to do (and believe). The relevant question is this: Do all the facts determine what you ought to do or do only some of them?

Consider an example. Jack's mother is in the hospital. She needs an operation in order to survive past this week. Her insurance won't pay. Jack, being a fledgling Art Historian/dealer, doesn't have the money. It looks like his mother is going to die. She will, however, be extremely comforted by Jack's presence in her final days. She lives in California; Jack lives in New York. Jack needs to decide whether to go see her. As it happens, a pawn shop owner in Queens has just unknowingly (and legitimately) bought a rare Picasso. He's selling it at a fraction of the price it's worth. If Jack were to buy it, he would be able to use it as collateral for a loan that would pay for his mother's surgery. The rub, of course, is that he has no idea that this pawn shop even exists, much less that such a deal is to be had there. Interesting question: What ought Jack do with his day? Go to California or go to Queens?⁵

If you think that what you ought to do is determined by all the facts, then you are committed to thinking that the place for Jack to go is Queens. If you don't like that answer, then you should think that what you ought to do is determined by only some of the facts.⁶ The trick is to explain which facts are relevant.

There is a tight connection between the correct resolutions to these two debates—i.e., the debate about whether you ought to be rational and the debate about what determines what you ought

²I will follow the standard practice of treating 'should' as a synonym for 'ought.'

³This skepticism is usually described as skepticism about the 'normativity' of rationality. I think this is unfortunate terminology because it is not specific enough. I know of no one that denies that rationality is *evaluatively* relevant—e.g., it's relevant to whether one is a bad agent or reasoner. Indeed, Kolodny (2005)—the *locus classicus* of skepticism about the 'normativity' of rationality—explicitly considers whether rationality is merely evaluative rather than 'normative.' This is confusing, I think. This is because evaluative properties are presumably normative. What Kolodny and others have worried about is the more specific property of being deontically significant. Thanks to Kurt Sylvan for discussion about this.

⁴In setting things up this way I'm following (among others) Broome (MS), Kiesewetter (2011), Ross (2012), Kolodny & MacFarlane (2010).

⁵Some think that there isn't really any good question here. This is because, claim these theorists, all we need to do is distinguish between a subjective and objective sense of 'ought'. Once we do this, we can see that Jack objectively ought to go to Queens and subjectively ought to go to California. I disagree that this is enough to solve the problem, as do many others. I spell out why I disagree in §7.4.1.

⁶Of course, the last two sentences are shorthand. There are all sorts of ways all the facts could determine what one ought to do. According to some of those ways, it might not be true that you ought to go to Queens. But on all the most plausible accounts that take into consideration all of the facts, you ought to go to Queens.

to do. I will argue in this chapter for my theory of rationality—which I'll here dub the *Reasons Responsiveness* account—by showing that it vindicates the deontic significance of rationality. According to the Reasons Responsiveness account, rationality consists in correctly responding to the normative reasons you possess. If this view is correct, I will argue, then what you ought to do *just is* what you are rationally required to do. In order to show this, I will argue for a view about what determines what you ought to do. I will argue that there are independent reasons to think that what determines what you ought to do are the reasons you possess. The basic idea behind my main argument for this is that the reasons that obligate have to be potentially action guiding in a certain sense—we have to be able to act *for* the reasons that obligate. I will argue that in order to be able to act for a reason in the appropriate way, one has to possess the reason. It is a few short steps from here to the claim that the reasons one possesses determine what one ought to do.

The chapter has the following structure. I begin by canvassing the recent literature in metaethics about rationality and its deontic significance. The upshot will be that the two most discussed views have serious problems vindicating the deontic significance of rationality. This will initially motivate the Reasons Responsiveness account, which immediately entails that we always possess reasons to be rational. Unfortunately, it doesn't immediately follow from the Reasons Responsiveness view that one always ought to be rational. In order to settle this question, we must investigate the second debate. So in the second half of the chapter we will focus on this. I will argue that it is independently very plausible that what you ought to do is determined by the reasons you possess. Given that the reasons you possess determine what you are rationally required to do, according to the Reasons Responsiveness view, the view vindicates the full deontic significance of rationality. This is a powerful reason to accept the view.

7.2 Why Be Rational?

In the middle of the last decade, a wave of philosophers started to question the deontic significance of rationality.⁷ Unfortunately, what exactly is being questioned is itself a matter of debate. Because of this, I'll start by explicating what I take to be the core of the challenge. I will then introduce the main views discussed in the literature and show how they have a hard time meeting the challenge. This will be by way of motivating my positive account, which will be introduced in §7.3 (everyone should read §7.2.1).

7.2.1 Preliminaries

The easiest way into the question being asked is to think about the most obvious way in which rationality could turn out to be deontically significant. Rationality would obviously be deontically significant if you *should* be rational. This, of course, is not simply to say that you *rationally should* be rational—i.e., that you are required by rationality to do the things rationality requires of you. It's certainly not open to dispute whether you rationally should be rational. Rather, the question is whether you should, full stop, be rational.

⁷See Kolodny (2005, 2007a, 2008a,b), Broome (2005a,b, 2008a, MS), Raz (2005a,b), Southwood (2008), Way (2010a,b).

One way to think about the question is to imagine that there are different systems of requirements—e.g., morality, prudence, etiquette, and rationality.⁸ Some of these requirements are deontically significant in particular kinds of ways. Some of them, for example, might be such that if you are required by that system to ϕ , you thereby ought full stop to ϕ . Some of them might not be this strict. Some of them might just be such that if you are required by that system to ϕ , there is thereby a reason for you to ϕ . The question has been whether rationality is a system that is either such that if it requires you to ϕ , then you thereby ought to ϕ or such that if it requires you to ϕ , then you thereby have a reason to ϕ .

Despite the fact that the literature has focused on these rather narrow questions, I don't think this is the best way to see things. These questions are *too narrow*. This is because they assume that rationality is deontically significant only if *the fact that rationality requires ϕ -ing* grounds either a full stop requirement to ϕ or a reason to ϕ . I agree that these are two ways that rationality could turn out to be deontically significant. But we can also ask the broader questions of whether it's true that we always ought to do what rationality requires or whether it's true that we always have a reason to do what rationality requires. It might be true that we always ought to do what rationality requires or have reason to do what rationality requires *even if* those facts aren't grounded in the fact that rationality requires one to do the thing in question. Let's focus, then, on the broader questions. To ease discussion, let's say that rationality is **Strongly (Deontically) Significant** if it turns out that we always ought to do what rationality requires. And let's say that rationality is **Weakly (Deontically) Significant** if it turns out that we always have reason to do what rationality requires.

7.2.2 Coherence, Rationality, Deontic Significance

Now that we have some understanding of what the relevant questions are, we need to know something about what rationality requires. In the next section I will explicate what I think rationality requires. This isn't of primary importance now. What we need to know now is how rationality is conceived in the recent literature about the deontic status of rationality. As we saw in chapter 2, the literature has focused on two different views. Both views are motivated by the thought that rationality has a tight connection with being coherent in certain ways. Common examples are means-end coherence, belief consistency, and following your conscience. Both views think that you are rational when you are means-end coherent, and that you are irrational when you are means-end incoherent.

The views diverge when it comes to *why* you are rational when you are means-end coherent and irrational when you are means-end incoherent. That is to say, the views diverge when it comes to which requirement you comply with when you are means-end coherent and which requirement you violate when you are means-end incoherent. According to the first view, you are irrational when you fail to intend to ψ when you intend to ϕ and believe that in order to ϕ you must intend to ψ because Means-End N is true:

Means-End N: If you intend to ϕ and believe that in order to ϕ you must intend to ψ , then you are rationally required to intend to ψ .

⁸This is how the question is normally asked in the literature. This is largely because this is the way John Broome thinks of the issues. See for example Broome (2007b, MS). Ultimately, I think that this way of thinking of things is misleading, but it is a good way to initially glob onto what's going on.

Since 'rationally required' just scopes over the consequent of Means-End N, we'll call it a narrow-scope requirement. We'll call those that accept narrow-scope requirements narrow-scopers. Narrow-scopers think that you are irrational when means-end incoherent because you lack an intention rationality requires you to have.

Usually the narrow-scope view is introduced merely as a foil for what has come to be the dominant view (more on why it has become the foil below). This view holds that you are irrational when you are means-end incoherent because you violate Means-End W:

Means-End W: You are rationally required to [intend to ψ if you intend to ϕ and believe that you must intend to ψ in order to ϕ .]

'Rationally required' takes wide-scope over the conditional in Means-End W. Thus, it is a wide-scope requirement and those that accept it are wide-scopers. Wide-scopers think you are irrational when means-end incoherent because a conditional you are required to make true is false.

The literature has assumed that we can discover what the requirements of rationality are in a way that's independent from asking whether they are deontically significant. This is because it has been assumed that rationality is paradigmatically a matter of coherence, and that we can home in on questions about coherence without having to talk about the deontic significance of the resulting requirements.

The wide-scope view, being the dominant view, has been the main target of the skeptical arguments. Its failure to account for the deontic significance of rationality is all the more surprising because it is largely motivated by the thought that only it can account for why you should be rational.⁹ The argument that motivates the wide-scope view goes something like this:¹⁰

- (1) Either the wide-scope view is true or the narrow-scope view is true.
- (2) If the narrow-scope view is true, then it's not the case that you always should be rational.
- (3) You always should be rational.
- (4) Thus, the narrow-scope view is false.
- (C) Therefore, the wide-scope view is true.

Let's first see why (2) is supposed to be true. The thought is that if (2) were false, then rationality would give rise to objectionable bootstrapping. To see this, suppose John intends to kill Becky and believes that in order to kill Becky he must stab her. If Means-End N is true, then it thereby follows that rationality requires John to intend to stab Becky. This is in itself implausible. But it gets even worse if you think that rationality is deontically significant. If it's strongly significant, then it would follow that John ought full stop to intend to stab Becky. This is certainly false. If it's weakly significant, then it would follow that John has a reason to intend to stab Becky. This isn't as obviously false as the last claim, but it still strikes nearly everyone as very implausible.¹¹

⁹The thought goes back at least until the 1970s/early 1980s (see, e.g., Hill (1973), Greenspan (1975), Darwall (1983)). It became widely thought about and accepted because of the work of John Broome, especially Broome (1999).

¹⁰Cf. Schroeder (2004, 2005a).

¹¹Schroeder (2004, 2005b) argued that the narrow-scope requirements are deontically significant in the weaker sense. He has since given this up (see Schroeder (2009b) for his retraction of the earlier view).

Means-End W doesn't have these problems. This is because John can comply with Means-End W without intending to stab Becky. He can comply by giving up his end or his means-end belief. Thus, it doesn't follow from Means-End W that John is rationally required to intend to stab Becky when he intends to kill her and believes that he must stab her in order to kill her. Thus, there is no bootstrapping even if rationality is strongly or weakly significant.¹²

Despite the wide-scopers' (apparent) ability to avoid bootstrapping, it has become increasingly unclear whether the wide-scope requirements are deontically significant in either the strong or the weak senses. Before seeing why, it's very important to add one more clarification about the question being asked.¹³

Suppose I believe I am writing and disbelieve I am writing. It's plausible that I'm irrational. The wide-scooper thinks I'm irrational because I'm rationally required to [not disbelieve I am writing if I believe I am writing]. When asking whether this requirement is strongly significant, we are asking whether I ought to not-disbelieve-I'm-writing-or-not-believe-I'm-writing. We are not merely asking whether I ought to be in some state other than the incoherent one. It's plausible that it's never permissible to hold contradictory beliefs. In this case, either my evidence will sufficiently support the claim that I'm writing or it will sufficiently support the claim that I'm not or neither claim will be sufficiently supported. My evidence will never sufficiently support both claims at the same time. So whenever I both believe and disbelieve I'm writing right now, it will be the case that I ought not have (at least) one of those beliefs. In this case, it's the belief that I am not writing right now.

So we can stipulate that I ought not disbelieve that I'm writing right now. And it's true that if I comply with that requirement, I will be coherent. It is very important to note that the literature in question assumes that *this is not to say* that I ought to comply with the wide-scope requirement.¹⁴ In other words, we are assuming there is a gap between being required to do something that guarantees I'll be coherent and being required to be coherent. It should be noted that it's quite plausible that it doesn't immediately follow from the fact that you are required to do something that guarantees that *p* obtains that you are required to see to it that *p* obtains. To take a famous example, if I comply with my requirement to post the letter, I will guarantee it is true that *I post the letter or I burn the letter*. It is counterintuitive that it *follows* that I'm required to post the letter or burn the letter.¹⁵ So it goes for the coherence requirements as well. It doesn't follow from the fact that I ought to have some attitude that will guarantee coherence that I ought to be coherent.¹⁶

¹²Some argue that the wide-scope view does have a bootstrapping problem in cases where it's impossible to comply in certain ways. See, for example, Greenspan (1975), Setiya (2007a), Schroeder (2009b).

¹³The following clarification only applies to the wide-scope view. The narrow-scope view doesn't think you are required to be coherent in the way that the wide-scope view does. That is, the narrow-scope view doesn't hold that the thing you are required to do is be coherent. The narrow-scope view holds that given that you hold certain attitudes, you are required to hold some other attitudes. It turns out—by design—that whenever you hold the attitudes you are required to hold, you will be coherent in the right ways. This allows the narrow-scooper to avoid this type of objection, but it also seems to lead directly to the argument against the narrow-scope view sketched above.

¹⁴Unfortunately participants in the literature have been very bad at making this clear. In fact, in many cases it appears as if participants don't even notice that there needs to be a gap between being required to do something that guarantees you'll be coherent and being required to be coherent in order for the question to be interesting at all. The only two places I know of where this is explicitly mentioned (besides Lord (MSb) and §2.5.1 of chapter 2) is (Kolodny, 2007a, n. 17) and Schroeder (MSa).

¹⁵This, of course, is Ross' Paradox. See Ross (1944). For a nice discussion in this context, see Broome (2007b).

¹⁶Another way to make the point is to stipulate that not-disbelieving-I'm-writing-or-not-believing-I'm-writing is a disjunctive action. If you thought this, then the wide-scope requirements require you to perform that disjunctive

Similar points apply to vindicating the weak significance of the wide-scope requirements. We can stipulate that I have reason to not disbelieve that I'm writing right now. And thus I have reason to do something that guarantees I'll be coherent. This is not yet to say that I have reason to be coherent. I have a reason to not disbelieve I'm writing. This is not necessarily a reason to not-disbelieve-I'm-writing-or-not-believe-I'm-writing. When we're asking whether the wide-scope requirements are weakly significant, we are asking whether there is always reasons to be like that, not merely whether there are always reasons to be in states that guarantee that you'll be like that.¹⁷

Once this is made clear, I think the problem becomes acute. What would make it the case that we ought to always comply with the wide-scope requirements? Remember it *can't* merely be that we are always required to do things that guarantee we'll be coherent. In order to show that the wide-scope requirements are strongly significant, we'd have to show that there are reasons over and above the reasons we have for individual attitudes that make it the case that we ought to be coherent. But it's completely mysterious what reasons there could be to always make it the case that I ought to be coherent.

To bring out the mystery, let's look at the three strategies most often pursued in the literature.¹⁸ The first strategy is that you always instrumentally ought to be coherent. According to this strategy, it's always the case that by being coherent you are doing something else that you ought to do. Although it's very plausible that you are always doing something else you ought to do when you are coherent *in certain ways*, it's not plausible that you are always doing something else you ought to do by being coherent in any way. To use our earlier example, it's stipulated that I am doing something else I ought to do when I fail to disbelieve I am writing right now. But it seems very implausible that I am doing something else I ought to be doing if I instead drop my belief I am writing right now and continue to disbelieve I am writing right now. This is a way to do what the wide-scope requirement requires. Thus, it's not generally true that I do something else I ought to do by complying with the wide-scope requirement.

The second strategy claims that you ought to comply with the wide-scope requirements because doing so is constitutive of some activity you're engaged in. Candidate activities include believing, intending, and reasoning.¹⁹ The underlying idea is that if you are a creature whose attitudes never cohere in the ways mandated by the wide-scope requirements, then you are not a creature with beliefs/intentions or a creature who reasons. While this might be true, it simply does not follow from this that you always ought to comply with the wide-scope requirements. This is because, *inter alia*, it doesn't follow from this that you cease being an creature with beliefs, intentions or a creature who reasons *every time* you are incoherent. Perhaps a requirement to *sometimes* be coherent follows from these supposed agential facts, but this is a far cry from a full vindication of the wide-scope requirements.²⁰

action. They are strongly significant if you ought to perform that action and weakly significant if you have reason to perform that action. But that action is different than the actions that are the disjuncts of the disjunctive action. Thus, it doesn't follow you ought to perform the disjunctive action just because you always ought to perform one of the actions that are the disjuncts.

¹⁷Again, it's important to note that it's very plausible that there is such a gap. If there wasn't, then my reason to post the letter would thereby be a reason to post the letter or burn it.

¹⁸Both Broome and Kolodny pursue these strategies at length in the work cited in n. 2. What follows is mostly just a rehashing of what they say. See also Way (2010a,b).

¹⁹See Bratman (2009a,b), Buss (MS) for a defense of the first two, and see Hussain (MS) for a defense of the reasoning account.

²⁰Even this weaker claim is open to a form of the agency, shmagency objection a la Enoch (2006).

The final strategy holds that coherence is intrinsically very good. So good that you always ought to be coherent. This is very implausible. It is not that implausible to think that coherence is an intrinsic good. But it does seem very implausible that it is so good that its goodness always makes it the case that you ought to be coherent. This is what would need to obtain for coherence's intrinsic goodness to vindicate the strong significance of the wide-scope requirements.

None of the three strategies seem promising when it comes to vindicating the strong significance of the wide-scope requirements. Do any of them hold promise for vindicating rationality's weak deontic significance? I think the same type of arguments can be given against the analogous strategies for vindicating the weak significance of the wide-scope requirements. The instrumental strategy holds that you have a reason to be coherent because by being coherent you will always be doing something else you have reason to do. But it seems possible that you can comply with the wide-scope requirements in ways that are such that you aren't doing something else you have reason to do by complying in those ways. The constitutivist strategy holds that some agential facts ground a reason to be coherent. Again, it's not clear which agential facts could do this. Some agential facts might ground reasons to sometimes be coherent, but this is a far cry from always grounding reasons to be coherent. And the intrinsic goodness strategy holds that the intrinsic goodness of coherence grounds a reason to be coherent.

While the intrinsic goodness strategy is the most plausible of the three when it comes to weak significance, it is still somewhat implausible that coherence is intrinsically good. What is *intrinsically* good about it, exactly? Moreover, even if coherence did turn out to be intrinsically good, it's not plausible that the reason grounded by such goodness would be very strong, very often. It's cold comfort for those of us who think rationality is very important to find out that we have weak reasons to be rational because coherence is an intrinsic good! These three strategies seem the most *prima facie* plausible. Since none of them seem promising, it's plausible to think that the wide-scope requirements are neither weakly nor strongly deontically significant.

A recap is in order. We started off by clarifying two ways in which rationality could be deontically significant. It is *strongly significant* if it is such that we always ought to do what rationality requires. It is *weakly significant* if it is such that we always have reason to do what rationality requires. The next step was to investigate what rationality requires. Those in the literature start with the assumption that there is a tight connection between rationality and coherence. There are two main views about what the connection is. The narrow-scope view was rejected because if it were true then it's very plausible that rationality is not deontically significant in either the strong or the weak ways. This left the wide-scope view. Although the wide-scope view is traditionally motivated by the thought that it alone can account for the deontic significance of rationality, it turns out to be mysterious how the wide-scope view can account for the deontic significance of rationality. What we're left with is skepticism about the deontic significance of rationality.

7.3 A Better Way: Reasons Responsiveness

Recall that we started our discussion of rationality in the previous section by assuming that there is a tight connection between rationality and coherence. This led us to the narrow- and wide-scope views, which led us to skepticism about the deontic significance of rationality. That is not a good place to be. I think the first thing we should do in response to these problems is explore the possibil-

ity that rationality doesn't have the type of connection to coherence that narrow- and wide-scopers assume.

Instead, I think we should hold Reasons Responsiveness.

Reasons Responsiveness: Rationality consists in correctly responding to the objective reasons one possesses.²¹

To reiterate, the objective reasons are the reasons constituted by facts. And the objective reasons you possess are the objective reasons you stand in some privileged epistemic relation with.²²

I should be explicit (and reiterate a point important for the last chapter) from the outset that I think that reasons come cheaper than most ethicists. To take two stark examples, I think that testimonial facts can be objective reasons and I think that existential facts about what there is reason to do and what one ought to do can also be objective reasons. So, if a reliable advisor tells you that you really shouldn't go into the next room but doesn't tell you why, he still gives you a reason not to go into the next room. In the simplest case, this will just be the fact that you shouldn't go into the next room. When the reliable advisor happens to say something false, it will be the fact that a reliable advisor said you shouldn't go into the next room. Many deny that these types of facts can be objective normative reasons.²³

Why do I think these facts can be objective reasons? Because they bear the earmarks of objective reasons. First, they intuitively count in favor of some actions and attitudes. You should treat what the advisor says as a reason in deliberation. Indeed, you should treat it as a very strong reason. Because of this, it's plausible they can justify actions and attitudes. Moreover, they can also be the reasons *for which* we act or hold attitudes.²⁴

It will be important for the overall plausibility of my view that objective reasons come cheap. For one thing, this allows me to account for the rationality of those who have false but rational beliefs. The most extreme version of this is the radically deceived characters that were the subject of the last chapter. I won't rehash those details here. It's worth keeping in mind in this chapter, however, that I think reasons come cheap.

It's also important to recall a main lesson from chapter 2. Namely, that Reasons Responsiveness doesn't commit you to thinking that you can be rationally incoherent. Indeed, as I argued in chapter 2, I think it's quite plausible that you are rational only if you are coherent. This is because it's quite plausible that when you are incoherent, you aren't correctly responding to all the reasons

²¹Broome (2007a, MS) considers whether rationality consists in correctly responding to reasons. He argues that it doesn't. As I show in chapter 2, his arguments don't have any traction against certain views that appeal to the reasons you possess.

²²I won't defend here any particular view about what it takes to possess a reason. I think that all of the plausible views will be able to explain what's going on in the cases that will interest us. The three most plausible views, I think, are the views that in order to possess some reason *r*, you have to know *r*, justifiably believe *r*, or be in a position to know *r*. In chapter 3 I defend the last option.

²³See, for example, Broome (2008b), McKeever & Ridge (2012), McNaughton & Rawling (2011).

²⁴Furthermore, I think that stock objections to the idea that they can be objective reasons are misguided. Some objections involve misguided worries about doublecounting (see Schroeder (2009a) and Väyrynen (2006) for reasons why these objections are misguided). Others appeal to intuitions about right-making, holding that these facts can't intuitively be right-makers (see, e.g., Broome (2008b), McKeever & Ridge (2012), McNaughton & Rawling (2011)). This thought cuts both ways, however, for it does seem like these facts can justify actions and attitudes. Given this, it's unclear why we shouldn't think that these reasons are counterexamples to whatever theory of right-making objectors have in mind. Because of this, it is very unlikely these intuitions will adjudicate the dispute.

you possess. To use our earlier example, it's plausible that there will never be a single time where I possess sufficient reasons to believe that I am writing and possess sufficient reasons to disbelieve that I am writing. I will always possess decisive reasons to not have at least one of those beliefs. In our example, it's disbelieving that I am writing. I possess decisive reasons to not have that belief. Thus, when I hold it, I am irrational.

Note that it follows immediately from Reasons Responsiveness that rationality is weakly deontically significant. This is because in order to be rationally required to ϕ , I must possess objective reasons to ϕ . Thus, there will always be objective reasons to ϕ when rationality requires me to ϕ . Securing the weak deontic significance of rationality is thus very easy for this view.

Nevertheless, it's far from clear that this view can vindicate the strong deontic significance of rationality. Even though the view entails that there will always be reasons to do what rationality requires, it's far from clear that these reasons will always be weighty enough to ground an obligation. Moreover, according to a very popular view about what one ought to do, it's obvious that sometimes the reasons one possesses won't be sufficiently strong to make it the case that one ought to do what's rationally required.

According to this view, which we'll call *objectivism*, what you ought to do is determined by all the objective reasons. You ought to ϕ , according to the objectivist, only if the balance of all the reasons decisively supports ϕ -ing. If objectivism is true, then Reasons Responsiveness cannot vindicate the strong deontic significance of rationality. This is clear by reflecting on cases. Here's the case from §7.1:

Sick Mother

Jack's mother is in the hospital. She needs an operation in order to survive past this week. Her insurance won't pay. Jack, being a fledgling Art Historian/dealer, doesn't have the money. It looks like his mother is going to die. She will, however, be extremely comforted by Jack's presence in her final days. She lives in California; Jack lives in New York. Jack needs to decide whether to go see her. As it happens, a pawn shop owner in Queens has just unknowingly (and legitimately) bought a rare Picasso. He's selling it at a fraction of the price it's worth. If Jack were to buy it, he would be able to use it as collateral for a loan that would pay for his mother's surgery. The rub, of course, is that he has no idea that this pawn shop even exists, much less that such a deal is to be had there.

In Sick Mother, Jack possesses good reasons to go to California—his mother will be greatly comforted by his presence. However, there is also a very good reason for him to go to Queens—viz., the fact there is a cheap Picasso to be had. The objectivist says that, in this case, Jack ought to go to Queens even though it seems like it would be irrational for him to go to Queens. It would be irrational because the reasons he possesses strongly support going to California.

This is a serious challenge. In order to see if it can be met, we must look at what the best theory is of what we ought to do. If objectivism is the best theory, then Reasons Responsiveness cannot vindicate the strong deontic significance of rationality. This would be a big blow to those of us who think rationality is deontically significant. In fact, I don't much care about a mere vindication of the weak significance of rationality. I will only be fully satisfied if rationality turns out to be strongly deontically significant.

The rest of this chapter will aim to meet the challenge. I don't think that objectivism is the best theory of what we ought to do. Instead, I think that the best theory holds that what we ought to do is what is decisively supported by the reasons we possess. I think this view can be well motivated independently of any debate about rationality. This, I think, is as it should be if rationality is really strongly significant.

7.4 Ignorance and Obligation

In this section I will argue for Possessed Reasons:

Possessed Reasons: What you ought to do is determined by the reasons you possess.²⁵

I will first discuss purported counterexamples to objectivism. I will argue that these cases are in fact counterexamples. Moreover, I will argue that Possessed Reasons provides a very plausible explanation of what's going on in these cases. The second argument runs through the claim that if the members of some set of reasons S make it the case that you ought to ϕ , then you can ϕ because of the members of S . I'll argue that you can only ϕ because of the members of S if you possess them. If sound, this argument will establish that a necessary condition for some set of reasons S making it the case that you ought to ϕ is that you possess the members of S . I will end this section by arguing that once you accept this necessity claim, the relevant sufficiency claim will follow from very plausible assumptions.

7.4.1 Ignorance and Envelopes

It is obviously true that we are almost always ignorant of the full effects our actions will have. Take this example discussed by Thomson (1990) and Scanlon (2001, 2008).

Day's End

Jack always comes home at 9:00 P.M., and the first thing he does is flip the light switch in his hallway. He did so this evening. His flipping the switch caused a circuit to close.

²⁵One of the major objections to anti-objectivist views appeals to bystanders with more information. It seems as if those with more information can have true thoughts about what one ought to do that isn't what one ought to do relative to what one knows. My goals in this chapter do not include answering this objection. My main goal is to provide a new argument for Possessed Reasons. However, I think that this problem can be solved. My solution has two parts. The first part is to show that Possessed Reasons is compatible with the thought that deliberation aims at what's best (or what's supported by all the reasons). Possessed Reasons is compatible with this claim because correct pursuit of one's aims might be constrained in certain ways. Here are two examples. Our obligations might be constrained by our physiological abilities even though we aim to do what's best. Our epistemic obligations might be constrained by, e.g., the evidence, even though the aim of epistemic deliberation is to believe the truth. In other words, it doesn't follow from the claim that the aim of epistemic deliberation is to believe the truth that we always ought to believe the truth. The second part of the solution appeals to the semantics of 'ought' in English. 'Ought' is a flexible word insofar as we can relativize our 'ought' claims to different bodies of information. This means that advisors with more information can have 'ought' thoughts about our obligations that are relativized to their information. This is what they're doing, I claim. It's right that this means they aren't having thoughts about our obligations *per se*. This is unsurprising, though, given that the aim is to do what's best. So even though they aren't thinking about our deliberative obligations, they are having thoughts that are relevant to the deliberative project. See Lord (MSa) for more.

By virtue of an extraordinary series of coincidences, unpredictable in advance by anybody, the circuit's being closed caused a release of electricity (a small lightning flash) in his neighbor's house next door. Unluckily, his neighbor was in its path and was therefore badly burned.²⁶

Many have thought, and I am one of these people, that cases like Day's End strongly support the rejection of objectivism. After all, the objectivist thinks that Jack ought not flip the switch in his hallway. This is because the fact that it will lead to his neighbor being badly burned is an objective reason not to. Relative to all the reasons there are, this reason is decisive. But since flipping the switch will lead to his neighbor being burned only because of 'an extraordinary series of coincidences' that he has no way of knowing about, it seems quite plausible that, to use Scanlon's words, 'If it is true that [Jack] ought not to have flipped the switch, this is true only in a sense of 'ought not' that seems to me to lack the moral content that the idea of permissibility has. Both [Jack] and [his neighbor] may wish, after the fact, that [Jack] had not flipped the switch, but in doing so [Jack] did not act impermissibly' (Scanlon, 2008, pg. 48).

Unfortunately for foes of objectivism, there is a standard objectivist reply to this line of reasoning.²⁷ The core insight is that we must cleave apart the deontic facts—facts about what ought to be done—from the hypological facts—facts about blame and praise. Once we do this, the response goes, we can see that arguments against objectivism like Scanlon's try to draw conclusions about the deontic from conclusions about the hypological. They get their bite only if we assume that the fact that Jack is not blameworthy for flipping the switch entails that he acted permissibly. Once we give up the idea that permissibility lines up with blamelessness, the argument falls apart. Jack might be blameless, but he still does what he ought not do. Indeed, it seems like there is a good explanation of why Jack is blameless despite doing wrong in this case. Jack is blameless because it is rational for him to believe that flipping the switch is permitted by all the reasons. But his belief is *false*. This is what explains why he still doing something wrong.

I agree that we must separate the deontic facts from the hypological facts.²⁸ Moreover, I agree that cases like Day's End lose their initial dialectical bite once this standard move is made.²⁹ Luckily, though, there are other well known cases where the standard move is less than effective. Three Envelopes is a case with this structure:³⁰

Three Envelopes

²⁶This is almost exactly how Scanlon (2008) presents this case, following (Thomson, 1990, pg. 229). I've changed the pronouns.

²⁷This move goes back to Moore (1912). Thomson herself replies this way in Thomson (1986, 1990, 2008). See also Graham (2010), Sylvan (2014).

²⁸Many other anti-objectivists do this as well. See especially Scanlon (2008) and Zimmerman (2008).

²⁹This is not to say that I think the objectivist is right about these cases, nor do I think that these cases lack bite. In the next subsection I will provide an argument that, I think, provides an elegant explanation for why cases like Day's End actually do have bite against the objectivist. My claim here is just that making the standard move robs these cases of unprincipled intuitive bite. The trick for the anti-objectivist is to give a principled explanation for why objectivism is wrong in these cases. I will try to do that in the next subsection.

³⁰This particular case initially comes from Ross (2006). It is discussed further in Schroeder (2009b), Ross (2012), Ross & Schroeder (FC). Ross's case was inspired by a case from Broome (MS), which was inspired by a case in Parfit (2011), which was inspired by a case in Regan (1980). This case and/or cases with the same structure are discussed, among other places, in Schroeder (2009b), Wedgwood (2007), Kearns & Star (2009), Kolodny & MacFarlane (2010), Jackson (1991), Zimmerman (2008), Kiesewetter (2011), Graham (2010).

Suppose Margaret is given the choice to pick one of three envelopes placed in front of her. Margaret is informed that the third envelope contains \$900. She is also informed that either the first envelope or the second envelope contains \$1000, and that whichever envelope doesn't have the \$1000 in it is empty. So, given her evidence, there is a .5 chance that the first envelope contains \$1000 and a .5 chance that the second envelope contains \$1000.

Intuitively, Margaret ought to choose the third envelope—i.e., the one she knows has \$900 in it. Moreover, there is an important difference between Three Envelopes and Day's End. Namely, in Three Envelopes Margaret *knows* that choosing envelope three is not the best option. That is, she knows that it is not the option decisively supported by all the reasons. The option best supported by all the reasons is choosing the envelope with \$1000 in it. Despite the fact that Margaret knows that choosing envelope three is the second best option, it still seems like she ought to choose it.

This difference between Day's End and Three Envelopes makes the standard objectivist move much less plausible when it comes to Three Envelopes. For in Day's End, it's rational for you to think that flipping the switch is permitted by the balance of all the reasons. It's not rational to think that choosing envelope three is permitted by the balance of all the reasons in Three Envelopes. In fact, Margaret *knows* that it's not. Still, it seems like Margaret ought to choose envelope three. This is important because the standard move seems plausible in Day's End only because you have a rational yet false belief about what the balance of all the reasons supports. The rationality of this belief helps explain why you are blameless, and the falsity explains why what you did was actually impermissible. Three Envelopes doesn't have this structure. Margaret seems blameless all right, but she lacks the false belief. The fact that it still seems like she should choose the second best option despite the lack of the false belief strongly suggests that she is blameless because choosing the third envelope is what she ought to do.

Before moving on, let me mention a common reaction to Three Envelopes in order to set it aside. Many ethicists' first reaction is to draw a distinction between what you objectively ought to do and what you subjectively ought to do.³¹ They then use this distinction to explain intuitions. First, you have the intuition that what you ought to do is choose the envelope with \$1000 in it because that's what you objectively ought to do. Second, you have the intuition that what you ought to do is choose envelope three because that's what you subjectively ought to do. This is supposed to resolve the puzzle.

The problem is that it doesn't resolve the puzzle.³² At best it changes the subject. That is, at best we find out that there is some sense in which Margaret ought to choose envelope three. It turns out that there are independent reasons to believe in such a sense, and the explanation it gives for why, in some sense, Margaret ought to choose envelope three is a principled one. But this is not the question we were asking. We were asking whether Margaret ought to choose envelope three. We were asking whether Yes would be the correct answer to a question Margaret might ask herself, viz., Ought I choose envelope three?

Here's an objectivist nicely putting the point:

Even if there are two types of moral obligation, the question remains: in which are we interested when doing moral theory about action? That is, in which sense of 'moral

³¹This move is made in response to this case in at least Schroeder (2009b), Smith (2008), Ross (2006).

³²Many—both objectivists and anti-objectivists—have pointed this out. See, e.g., Jackson (1991), Graham (2010), Kiesewetter (2011), Kolodny & MacFarlane (2010), Bjornsson & Finlay (2010).

obligation' are we interested when we debate whether consequentialism or some other theory provides the correct account of our moral obligations with respect to action? Even if there are multiple senses of 'moral obligation' we still need an answer to this question (Graham, 2010, pg. 95).³³

And here's an anti-objectivist nicely putting the point.

Others again try to solve the puzzle by distinguishing different senses of "ought"... I am willing to concede that it might be useful to speak of what an agent ought to do relative to certain considerations, and that different qualified notions of "ought" might be important in their own right. Nevertheless, I believe that there is a substantial question at issue between objectivists and perspectivists [anti-objectivists] when it comes to what might be called the "overall ought" of practical deliberation. This is the concept involved in the deliberative question, "What ought I to do?" (or "What should I do?") and deliberative conclusions of the form, "I ought to ϕ ". Practical conclusions of this sort are supposed to guide rational decision-making and action directly. In other words, the "ought" at issue is the one that is appealed to in the common idea that it is irrational, or akratic, not to intend what one believes one ought to do. Now, in order to make a rational decision guided by a belief that one ought to do something, one needs a univocal concept of "ought" that figures in such beliefs. It is perfectly consistent to believe, "I ought to ϕ , relative to X," and, "I ought not to ϕ , relative to Y," but one cannot rationally intend both to ϕ and not to ϕ . There must be one sense of "ought," the belief in which is the relevant one for decision-making. We need to be able to judge, "I ought to ϕ , full stop" (Kiesewetter, 2011, pg. 2).

The important point is that mere appeal to the distinction between subjective and objective oughts doesn't answer our question because it doesn't specify which obligation provides an answer to the central deliberative question. If Margaret is conceptually sophisticated enough, she is in a position to know that she subjectively ought to choose envelope three and that she objectively ought to choose envelope one or envelope two. But even if she did know those facts, she could meaningfully wonder which requirement she ought to satisfy. She thus doesn't seem to learn what she ought to do, full stop, by learning what she subjectively and objectively ought to do. There is another question that hasn't been answered yet—viz., What ought Margaret do?

I hasten to add that there are at least two ways to use the objective/subjective distinction to answer our question. The first way is to insist that what you ought to do is just what you subjectively ought to do. The second way is to insist that what you ought to do is what you objectively ought to do. I obviously don't think that what you ought to do is what you objectively ought to do. But for all I've said, it might be that what you ought to do is what you subjectively ought to do. In fact, it might be that the theory I give about what you ought to do is the correct theory about what you subjectively ought to do.³⁴ Because of this, that view is somewhat orthogonal to the current

³³It's worth pointing out that Graham is only concerned with moral requirements and explicitly tries to distance himself from the heart of the debate we're interested in. Nevertheless, I think the debate he's interested in is very close to the one I'm interested in, certainly close enough for the passage quoted to be relevant.

³⁴It's not always clear what role the subjective ought is supposed to play, but it's commonly thought that it's supposed to have a very tight connection with blame—sometimes it's suggested that you subjectively ought to ϕ

dialectic, whereas the view that what you ought to do is what you objectively ought to do is central to the current dialectic.

Many think that cases like Three Envelopes are counterexamples to objectivism. I agree. Surprisingly, though, hardly anyone has fleshed out anti-objectivist views in any detail.³⁵ Usually theorists are happy to appeal to 'the evidence' and leave it at that. Given the huge variety of views about evidence found in epistemology, the range of evidential views is huge. Blanket appeals to claims with such a large variety of potential interpretations is generally unfortunate. Moreover, I think that it makes it quite a bit easier for objectivists to respond to the cases.

The question, then, is this: Why is it that Margaret ought to choose envelope three? My answer is that it is because that is what the reasons she possesses decisively supports.³⁶ Which reasons does she possess and which ones does she lack? She obviously possesses the reason constituted by the fact that there is \$900 in envelope three, and she obviously lacks the reason provided by the fact that there is \$1000 in whichever envelope the \$1000 is in. This isn't enough, though, because it's very intuitive that she possesses some reason to choose envelope one and some reason to choose envelope two. I think the reason to take envelope one is the fact that there is a .5 chance that there is \$1000 in envelope one and the reason there is to choose envelope two is a .5 chance that there is \$1000 in envelope two.³⁷ Those facts are objective reasons to choose envelopes one and two respectively. Moreover, they are reasons Margaret possesses. Nevertheless, neither is weighty enough to make choosing those envelopes sufficiently supported by the reasons that are possessed. Thus, the reason to choose envelope three is the only reason possessed that is sufficient, and it is thus decisive.³⁸

Day's End is slightly different. In that case, it's just that you fail to possess a strong reason to not flip the switch—viz., the fact that doing so will lead to your neighbor being badly burned. You possess all of the normal reasons one has to flip on the switch—it allows you to see what you're doing,

only if you'd be blameworthy for not ϕ -ing. If this is true of the subjective ought, then I doubt the view I'm arguing for is the best view of the subjective ought. Indeed, I doubt there is any *deontic* concept that has that tight of a connection to blame.

³⁵An admirable exception is Zimmerman (2008). My main problem with Zimmerman is methodological. I think that we should try to fit the view within a reasons framework. I think this is so for two reasons. First, it's plausible that reasons are normatively fundamental. Second, my view is pitched at a higher level of abstraction than Zimmerman's. This allows me to be neutral about the consequentialism/deontology debate. This is because we can see that debate as one about the structure of reasons and their weights. So both sides can accept my claims because they are at a higher level of abstraction. Zimmerman's view, on the other hand, is committed to denying non-consequentialist views.

³⁶As I said above, I'm not going to defend here any particular view of what it is to possess a reason. See chapter 3 for a defense of such a view. The most intuitive views hold that in order to possess a reason r , you have to have some type of epistemic access to r —e.g., you have to know r , justifiably believe r , or be in a position to know r .

³⁷A note about what the chances are here. I am thinking that they are conditional probabilities. They are the probabilities that the \$1000 is in the respective envelopes given what Margaret knows (this is what Williamson (2000) calls evidential probability). I want to be neutral about whether objective chances provide objective reasons (I suspect they do if they exist) or whether subjective chances provide objective reasons (I suspect they don't, at least not always).

³⁸Schroeder (2009b) gives a similar explanation. My explanation differs from Schroeder's in key ways. For example, Schroeder thinks that, at least in the relevant sense, Margaret ought to pick envelope three *just so long as* she believes as she does in the original case. Thus, even if she was in a position to know that her beliefs were false it would be the case that she ought to choose envelope three. My explanation doesn't have this result because the relevant facts have to actually be facts (i.e., objective reasons) in order to make it the case that she ought to choose envelope three.

it makes it easier to navigate etc. Possessed Reasons holds that the fact that by flipping the switch you will badly burn your neighbor doesn't affect what you ought to do because you don't possess that reason. You don't possess that reason because you don't stand in the appropriate epistemic relation with it. Since the reasons you do possess sufficiently support flipping the switch, it's permissible to flip the switch. This is a very natural explanation of both why the fact that flipping will lead to your neighbor being burned doesn't affect what you ought to do and why it's permissible to flip the switch.

There are a few things to notice about the explanations given by Possessed Reasons. First, only objective reasons determine what you are permitted and forbidden from doing. This is a virtue because it allows easy assimilation into a broader picture of normativity. According to this view, what it is to be normative is to be analyzed in terms of objective reasons. Since Possessed Reasons analyzes the property of being what you ought to do in terms of objective reasons, it fits in with this broader picture.³⁹

The second thing to notice is just how natural and illuminating these explanations are. Appealing to possessed reasons gives one a very plausible explanation for why (at least non-culpably) unknown reasons don't affect what you ought to do. It's intuitively plausible that the possessed/unpossessed reasons distinction is poised to do the kind of work needed in these cases.

7.4.2 Ignorance and Acting for the Right Reasons

I take the last subsection to have demonstrated that Possessed Reasons does a nice job explaining some of the facts on the ground. In this subsection I'll offer a more abstract argument for Possessed Reasons. I will start by arguing that a necessary condition for some set of reasons S to make it the case that you ought to ϕ is that you possess the members of S . This is not enough to establish Possessed Reasons. However, if it is true, then objectivism is false. Moreover, once you accept this necessity condition, it follows from plausible assumptions that a sufficient condition for being obligated to ϕ is that one possesses a set of decisive reasons to ϕ . Once we have both the necessity and sufficiency claims, we'll be much closer to establishing Possessed Reasons.

We'll start with an argument for the following claim: If the members of some set of reasons S make it the case that you ought to ϕ , then you possess the members of S . The argument's plausibility will be enhanced by some build up.

Upon reflection, it seems very plausible that we aren't always required to bring about the best state of affairs. For example, consider Dunk for Money:

Dunk for Money

Mark Cuban, being the eccentric billionaire that he is, decides to have a raffle. The person whose name is chosen gets a shot at a \$10 million prize. In order to win the prize, one has to do a 360° dunk on an NBA regulation sized hoop. Just for fun, Sam

³⁹It's worth mentioning that it also does this without having to appeal to any type of modal analysis. Schroeder (2009b), for example, tries to explain cases like Three Envelopes in a framework compatible with the thought that reasons are basic by appealing to the notion of a subjective reason. Subjective reasons are analyzed in terms of objective reasons, but modally—some proposition p is a subjective reason to ϕ just in case p is an objective reason to ϕ if true. I think this will inevitably lead to intractable conditional fallacy worries. It's a big virtue that my view doesn't have those problems.

enters the raffle. She is the lucky winner. Unfortunately for Sam, she is only 4'11" tall. Because of this, she lacks the ability to dunk on an NBA regulation sized hoop.

Question: Ought Sam dunk the ball? Pretty clearly not, I think. Why is that? It would, after all, be much better for Sam to dunk rather than not. Despite this, the answer is obvious. Sam *can't* dunk. She lacks the ability to *perform the relevant action*. Of course she can try to dunk, and presumably she ought to try to dunk. But that's not to say that she ought to dunk. It's not true that she ought to dunk because she lacks the ability to dunk.

Since Sam can't perform the relevant action at all in Dunk for Money, she also can't dunk *for* the fact that by dunking she'll win \$10 million. When you perform actions for reasons, your acts are explained in a certain way by those reasons. Consider some examples. I put on my coat this morning because it was cold outside. I believe the heat is on in my office because it's warm in my office and cold outside. I'll go home tonight at 5 because dinner is at 6. These explanations are not merely causal (although I think they are causal). They also have a normative element to them. The explanans are not merely the causes of the actions and beliefs, they are also the justifiers of them. It follows from the fact that Sam can't perform the relevant action in Dunk for Money that she can't dunk in order to win \$10 million.

Of course, I don't think Dunk for Money is a counterexample to objectivism. Objectivists can get out of the problem by invoking some version of 'ought' implies 'can'. I think a deeper explanation is that the fact that Mark Cuban will give Sam \$10 million if she dunks is not a reason *for Sam* to dunk. Thus, that fact won't be relevant to the objectivist theory that holds that what Sam ought to do is a function of all of the reasons (for Sam).

Although Dunk for Money is not a counterexample to objectivism, it's helpful build up to my argument. Consider Day's End again. In Day's End you face the choice between flipping on the hall light or not. There are good reasons for you to switch it on and a good reason for you to not switch it on. Moreover, in Day's End, unlike in Dunk for Money, it seems like you can perform the relevant action—i.e., you have the ability to refrain from flipping on the switch. Not flipping the switch can be an action of yours. So there both is a reason for you not to flip it on and you can perform that action. Thus, the objectivist thinks that you ought to refrain from flipping the switch. The reason why is that it will badly burn your neighbor. That fact is a decisive reason not to flip, according to the objectivist.

There's a catch, though. Although you can refrain from flipping the switch, it *doesn't* seem like you can do so *because* flipping the switch will badly burn your neighbor. It doesn't seem like that could be your reason for acting in the right kind of way. For one thing, you don't believe that flipping the switch will harm your neighbor. You've never even considered that, and rationally so. Given this, there is no way that you could refrain from flipping the switch because flipping it will harm your neighbor.

I think something even stronger is true. I think that even if you did irrationally believe that flipping the switch would badly burn your neighbor and decided to bumble around in the dark, you wouldn't have refrained because flipping the switch would badly burn your neighbor. You'd be acting like an idiot if you did that. We might be able to give some psychological explanation of what you did that cited your belief that flipping the switch would badly burn your neighbor, but we couldn't say you did it because flipping the switch would badly burn your neighbor. This is because you have no idea that flipping the switch will burn your neighbor.

To make this more plausible, compare Delusional Andy to Surprised Andy:

Delusional Andy

Andy knows that his wife has always been an extremely loyal person. He also knows that he has no reason to think that she is cheating on him. Despite this knowledge, he does believe that she is cheating on him. He thus moves out and files for divorce. In fact, his wife is cheating on him.

Surprised Andy

Andy knows that his wife has always been an extremely loyal person. However, much to his surprise, he learns that she is cheating on him—her best friend tells him, he finds some love letters, and he catches his wife with her lover. He thus moves out and files for divorce.

In both Delusional Andy and Surprised Andy, Andy reasons from a belief that his wife is cheating on him to an intention (and subsequent action) to move out and file for divorce. Despite this similarity, it's plausible that only Surprised Andy moves out and files for divorce because his wife is cheating on him. Delusional Andy is, well, delusional. He has no reason to think that she is cheating on him. He gets completely lucky. Because of this, it doesn't seem like he files for divorce because she's cheating on him.⁴⁰

Notice that you are a lot like Delusional Andy when you more or less arbitrarily believe that flipping the switch will badly burn your neighbor. It turns out that you are right, but you are just lucky. It's not like that reason is *guiding you*. You have about as little contact with that reason as one can have while believing the proposition that constitutes it. (We'll return to Delusional Andy and Surprised Andy momentarily.)

I think the fact that you can't refrain from flipping the switch because it will badly burn your neighbor *explains* why it's not the case that you ought to refrain. In other words, I think that Acting On is plausible:

Acting On: If the reasons in some set S make it the case you ought to ϕ , then you can ϕ because of the members of S .

It's important to point out that if you knew that flipping the switch would burn your neighbor, then you could refrain for that reason. It's just that given how you are when you walk in the door, you aren't in a position to refrain for that reason. Thus, I think we should understand the relevant sense of 'can' as relativized to one's epistemic position. A stab at the relevant necessary condition: One can ϕ because of r only if one's epistemic position doesn't bar one from ϕ -ing because of r .⁴¹ It's plausible that a sufficient condition for one's epistemic position barring one from ϕ -ing because of r is one's epistemic position barring one from knowing r . This might be necessary, too. It all depends on the correct analysis of ϕ -ing because of r . I don't need to give such an analysis, though.⁴² All I need here is this sufficient condition.

⁴⁰For similar arguments for a similar conclusion, see Hyman (2006), Hornsby (2008), Gibbons (2001), Marcus (2012). The most visible detractor is Jonathan Dancy (see, e.g., Dancy (2000)). I think that Dancy is really talking about what it is to act for reasons, not what it is to act for normative reasons. We are interested in the latter here. He's right about acting for reasons, but not about acting for normative reasons. See chapter 4 for more on this.

⁴¹There are obviously other necessary conditions—e.g., one has to have certain physiological abilities.

⁴²See chapter 4 for an analysis.

What's the difference between the cases where you can ϕ because of some reason r and cases where you can't? My answer is that the cases where you can ϕ because of r are always cases where you *possess* r . You don't possess r when you are (non-culpably) ignorant of r , and thus can't ϕ because of r when you are (non-culpably) ignorant of r .⁴³ When you combine this with Acting On, you get an argument for the claim that possessing the members of a set of reasons S is a necessary condition for those reasons making it the case that you ought to ϕ .

- (1) If the reasons in some set S make it the case you ought to ϕ , then you can ϕ because of the members of S (Acting On).
- (2) If you can ϕ because of the members of S , then you possess the members of S .
- (C) If the reasons in some set S make it the case you ought to ϕ , then you possess the members of S .

If this argument is sound, then objectivism is false—Day's End is a counterexample. So objectivists need to deny one of the premises. In fact, objectivists need to resist what I said about (1). Objectivists have two choices when faced with (1). They can either deny it or they can argue, *pace* what I argued above, that you can act for the reasons that require you to ϕ when you're ignorant of those reasons. I think both options are implausible. In order to see this, it's important to consider the connection between ϕ -ing because of the reasons that require you to ϕ and your action being creditworthy.

Suppose I am required to buy Anne a hat for her birthday because it would make her happy. And suppose I do buy her a hat. These conditions are not sufficient for my action to be creditworthy. In order for my action to be creditworthy, there must be the right kind of connection between the reasons that obligate the action and my performance of the action. To illustrate, my act would be creditworthy if I were to buy her a hat because it would make her happy. On the other hand, if I bought her the hat only because it would make *me* happy (it would cover her up her hair, which I find distasteful), it doesn't seem like my act would be creditworthy. It would only be an accident that the following conjunction is true: I bought the hat and the hat makes Anne happy. If it's merely an accident that the act I actually perform is the particular act I ought to perform, then it doesn't seem like I deserve credit for it.

This fact is partly explained, I think, by the fact that when it's an accident that the act one performs is the act one ought to perform, one doesn't act for the right reasons—i.e., the reasons that make it the case you ought to do that thing.⁴⁴ It's plausible that there is some tight connection between acting for the right reasons and being creditworthy. One way to flesh this out is Link:

⁴³Not all ignorance is created equal. If we hold that A is ignorant about p just in case A fails to believe p , then I think it's implausible that you always fail to possess the reasons you are ignorant of. This is why, in chapter 3, I argue for a *non-holding* account of possessing reasons. Non-holding accounts claim that you can possess reasons you don't believe. However, I think you are culpable for the reasons you possess that you are ignorant of. Hence the qualification in the text.

⁴⁴This type of view is developed in the case of actions in Wedgwood (2006, 2007), Schossler (2007). I further develop this thought in chapter 4.

Link: When A is required to ϕ by the members of some set of reasons S , A is creditworthy for ϕ -ing just in case A ϕ s because of the members of some subset of S that are sufficiently strong to require A to ϕ .^{45,46}

When you combine Link with the claim that I don't act for the right reasons when I buy Anne the hat because it will make me happy, you get a principled explanation of the fact that I'm not creditworthy.

With Link in hand, let's return to the objectivist. To repeat, the first premise isn't a problem for the objectivist if it is either false or if, my arguments before notwithstanding, it's possible to act for the right reasons when you are ignorant of them. Let's examine each possibility in turn, starting with the latter.

One type of objectivist will insist that when you are required to ϕ because of some set of reasons S , you can always ϕ because of the members of S . There are two different kinds of cases that are relevant. Day's End is an instance of the first type of case. In that type of case, you not only are not in a position to know the relevant fact, you don't even believe it. I take it that it's extremely plausible that, in this type of case, you can't ϕ because of the reason. You can't refrain from flipping the switch because it will badly burn your neighbor if you don't even believe that your neighbor will be badly burned if you flip the switch. For reasons that will become clear, I don't think this yet decisively tells against this first strategy.

The second kind of case is like Delusional Andy (and the extension of Day's End where you irrationally believe flipping the switch will burn your neighbor). In that case, you do believe the relevant things, though irrationally. This first objectivist reply has serious hope of sticking if it's possible to do the right thing for the right reason in these cases. This is because it's always possible in the relevant sense of possibility to go from the first type of case to the second. Recall that something is possible in the relevant sense only if nothing about your epistemic position bars you from doing that thing. Although one's epistemic position does bar one from knowing the relevant fact in both kinds of case, nothing about one's epistemic position in either type of case bars one from believing the relevant proposition.

Unfortunately for the objectivist, it's simply not plausible that you can do the right thing for the right reason in the second type of case. We can see this by reflecting on Link. It's very plausible that Delusional Andy is *not* creditworthy for his moving out and filing for divorce, whereas Surprised Andy is. Similarly, it seems like you are not creditworthy in the extension of Day's End where you refrain from flipping the switch because you irrationally believe that flipping the switch will badly burn your neighbor. It might be that you end up doing what's best. But you just got lucky, and hence are not creditworthy. It follows from Link that Delusional Andy doesn't move out for the right reason, nor do you refrain from flipping the switch for the right reason. If this is right, then believing the relevant fact (and acting in light of that belief) is not sufficient for acting for the right reason. Thus this first strategy fails.

The second route the objectivist can take is to deny (1). I think Link sheds light on the plausibility of this, as well. If you deny (1), then you are committed to thinking that there are cases where

⁴⁵Markovits (2010, FC) defends a very similar view to this about moral worth. If you replace 'credit worthy' with 'moral worth', then you get something very close to Markovits's view. It also resembles, of course, an old Kantian idea.

⁴⁶The bit about subsets of S is necessary because some obligations are overdetermined. That is to say, in some cases there are multiple proper subsets of the set of all the reasons to ϕ that require one to ϕ . At least in some of these cases, it seems one can be creditworthy in ϕ -ing just so long as they act for the reasons in one of those subsets.

(i) the members of some set of reasons S make it the case that you ought to ϕ and (ii) it's impossible to ϕ because of the members of S . It follows from Link that there are cases where the members of S make it the case that you ought to ϕ and it's impossible for you to ϕ and be creditworthy for ϕ -ing. I think this is deeply implausible.

It's implausible because it robs our full stop obligations of a certain kind of *action guidingness*. If we cannot perform the actions we are obligated to perform in a way that deserves credit, then those obligations are not action guiding. To put it another way, if our full stop obligations are action guiding, then our actions can be guided by the facts that determine our full stop obligations. The paradigm of being guided by some fact is acting because of that fact. But we've just seen that in denying (1), the objectivist is denying that we can always act because of the right-makers of the acts that we are obligated to perform. Thus, there is an obvious way in which the objectivist has to deny that our obligations are always action guiding. This seems very implausible.

The point can be put less abstractly. Consider Day's End again. If objectivism is true and (1) is false, then there is simply no way you can do what you ought for the right reason. In other words, the only way you can do what you ought to do is by doing something idiotic. There's no way for the right-makers of your act to get any legitimate grip on you. But, we're supposed to believe, they require you stumble around in the dark all the same. This is dubious.

Thus, it looks like (1) is enough to topple the objectivist. This isn't enough for me, however. This is because I not only want to defeat the objectivist, I want to establish Possessed Reasons. So we need to consider (2), as well. (2) holds that possessing r is a necessary condition for ϕ -ing because of r .

I think the most plausible route for denying (2) is to hold that in order to possess r , you have to stand in some positive epistemic relationship to r —e.g., you have to know r —but one needn't stand in this relation to r to ϕ because of r . The most natural version of this view holds that you merely have to believe r in order to ϕ because of r .⁴⁷ Fortunately for us, we needn't dwell too much on this proposal, for we have already seen that it is implausible. It's implausible because it's implausible that merely believing r (and acting in light of this belief) is sufficient for ϕ -ing because of r . If this were true, then it would follow from Link that Delusional Andy is creditworthy for moving out. But he's not.⁴⁸ Thus, it seems like you have to stand in some positive justificatory relationship to r in order to ϕ because of r . If you also have to stand in some positive justificatory relationship to r to possess r , it's plausible that in order to ϕ because of r , you must possess r .

Even if this is right, I've only established that a necessary condition for a set of reasons S to make it the case that you ought to ϕ is that you possess the members of S . This is significant in the dialectic with the objectivist, for it entails that objectivism is false. But it's not a full vindication of Possessed Reasons. In the rest of the section I'll argue for two claims that get us closer to a vindication of Possessed Reasons.

The first claim is a strengthening of the necessity condition. It has to be not only that you possess the members of S , it has to be that the members of S are *decisive reasons* to ϕ . If they are decisive, then they (i) aren't defeated by other reasons you possess—i.e., the members of S will be weightier than the reasons you possess not to ϕ —and (ii) they are much weightier than the reasons

⁴⁷It's possible to hold that in order to possess r you have to know r , but in order to ϕ because of r you just have to justifiably believe r . I can't see any motivation for this, though. It's more plausible to think that there is a high bar on possession but a low one on acting for a reason.

⁴⁸This argument works just as well against views that accept (2) but hold that there is a low bar on having (e.g., Schroeder (2008, 2011)). Thus, this argument supplements my arguments against those views in chapter 3.

you possess for any alternative action. Thus, I think we should replace the conclusion of the above argument with Necessary:

Necessary: If the members of some set of reasons S make it the case that you ought to ϕ , then the members of S are decisive reasons to ϕ that you possess.

I don't think this strengthened necessity condition needs much more defense. It's very plausible that a set of reasons needs to decisively support ϕ -ing in order to make it the case you ought to ϕ . If the members of S are defeated or are not much weightier than the alternatives, then it's independently plausible that they don't make it the case that you ought to ϕ .⁴⁹

We'd be much closer to a vindication of Possessed Reasons if we were to show that Sufficient is true:

Sufficient: If the members of S are decisive reasons to ϕ that you possess, then the members of S make it the case that you ought to ϕ .

If Necessary is true and Sufficient is false, then there is some condition over and above possessing decisive reasons to ϕ that must be met in order for it to be the case that you ought to ϕ . Obviously there are a lot of possible conditions that one could propose. I don't have a good enough imagination to think of all of them (nor do I have the space to consider them all). But it's worth considering two of them.

The first might be some kind of ability condition. It might be that 'ought' implies 'can' but that 'possessing decisive reasons' doesn't imply 'can.' I agree (obviously!) that there is some ability condition. Not only do I think that you must be able to ϕ in order to be required to ϕ , I think you must be able to ϕ because of the right-makers in order to be required to ϕ . But whatever plausibility these claims have when it comes to what you ought to do seems to apply equally well to what you possess decisive reasons to do. We can see this by considering examples. Take the basketball example discussed earlier. It's true that it's not the case that Sam ought to dunk the basketball, but it's equally plausible that she doesn't possess reasons to do that. A plausible generic claim in the vicinity is something like: If r is a reason for A to ϕ , then it's (metaphysically) possible for A to ϕ because of r .

A second proposal is some type of hypological condition. For example, perhaps it's true that if you ought to ϕ , then you are blameworthy for not ϕ -ing, but it's not the case that you are always blameworthy for not doing what you possess decisive reasons to do. If this were true, then there would be cases where you possess decisive reasons to ϕ but wouldn't be blameworthy if you didn't ϕ and thus it wouldn't be the case that you ought to ϕ . One problem with this proposal is that the hypological condition on 'ought' is false. It's just not true that you are always blameworthy for not

⁴⁹Things are complicated when the reasons you possess in favor of one option are slightly weightier than the reasons you possess for the other options. In some cases, it will seem like you ought to take the option that is most supported. For example, if you have the choice between taking a \$1000 prize or a \$1002 prize and everything else is equal, you ought to take the \$1002 prize. However, it doesn't seem like you are obliged to take the option most supported in every case. For example, in some cases of mere permissibility—cases where one has sufficient reasons to perform several actions and is not required to perform any particular action—it seems like adding some lightweight reasons doesn't oblige you to take the option that you add the reasons to. So, if it's merely permissible for me to take job A and merely permissible to take job B , I won't come to be obliged to take job A simply because job A offers me an extra \$50 per year. But surely the \$50 provides a reason to take job A .

doing what you ought. There are simple recipes for counterexamples. One is to take a case where your pet theory of obligation says you ought to ϕ at time t . At t_1 , bombard your attention with useless data. You can set the case up so that it is very hard to do what you ought (and thus most of the time you won't), but nevertheless the conditions that obligate you to ϕ are still in place. In this type of case, you will seem blameless for not ϕ -ing because of how hard ϕ -ing becomes, even though you will still be obligated to ϕ .

Furthermore, even if you think that you are always blameworthy for not doing what you ought, it's hard to see why you would want to deny that you are always blameworthy for not doing what you possess decisive reasons to do. The most plausible reason I can think of is that what you ought to do is transparent in some way but what you possess decisive reason to do isn't (perhaps because being decisive isn't transparent). You could think this, but I'm not sure why you would. If what you ought to do is transparent, then it seems like it will always be transparent what the facts support, but this is just another way of saying that it will be transparent what the reasons you possess decisively support.

It doesn't seem like these two proposals are plausible, and thus it seems like having decisive reasons to ϕ is sufficient for it being the case that you ought to ϕ . This gives us Necessary & Sufficient:

Necessary & Sufficient: The members of S are decisive reasons to ϕ that you possess iff the members of S make it the case you ought to ϕ .

Necessary & Sufficient is, I think, just a precisification of Possessed Reasons. It tells you how your obligations are a function of the reasons you possess.

7.5 Summary of Results (Or: Why You Ought to Be Rational)

It was established in §7.3.3 that if the Reasons Responsiveness view is true, then rationality is weakly deontically significant. This is because there is always reason to do what you possess reason to do. So if you always possess reasons to do what rationality requires, then there will always be reasons to do what rationality requires. This only goes so far, though. It would still be a considerable win for skeptics about the deontic significance of rationality if rationality turns out to only be weakly significant. In order to fully vindicate the deontic significance of rationality, one needs to show that it's strongly significant—i.e., that you ought to be rational.

The last section aimed to show this indirectly. The main result of the last section was that what you ought to do is determined by the reasons that you possess. The plausibility of this claim is independent of any considerations having to do with rationality—this is why it was indirect.

When you combine §7.3 with the results of §7.4, you see that Reasons Responsiveness can vindicate the strong deontic significance of rationality. If what you are rationally required to do is what you possess decisive reasons to do and what you ought to do is what you possess decisive reasons to do, then the requirements of rationality *just are* the requirements you ought to comply with. It is a serious virtue of my view that it can give such a plausible account of why rationality is strongly deontically significant. This is especially weighty since it's completely mysterious how rival views can even account for the weak deontic significance of rationality, let alone the strong deontic significance. Given that it is a truism that we should be rational, the fact that only my view can vindicate the strong deontic significance of rationality is a strong reason to accept my view—one which you now possess.

Bibliography

- Alston, W. (1988). An internalist externalism. *Synthese*, 74, 265--283. [cited on pg.]
- Armstrong, D. (1973). *Belief, Truth, and Knowledge*. Cambridge University Press. [cited on pg.]
- Arnold, A. (2011). Some evidence is false. *Australasian Journal of Philosophy*, 91(1), 165--172. [cited on pg.]
- Arpaly, N. (2006). *Meaning, Merit, and Human Bondage*. Princeton University Press. [cited on pg.]
- Arpaly, N. & Schroeder, T. (2012). Deliberation and acting for reasons. *The Philosophical Review*, 121(2), 209--239. [cited on pg.]
- Audi, R. (1993). *The Structure of Justification*. Cambridge University Press. [cited on pg.]
- Bach, K. (1985). A rationale for reliabilism. *The Monist*, 68(2), 246--263. [cited on pg.]
- Ballantyne, N. & Coffman, E. (2011). Uniqueness, evidence, and rationality. *Philosophers' Imprint*. [cited on pg.]
- Bealer, G. (2002). Modal epistemology and the rationalist renaissance. In T. Gendler & J. Hawthorne (Eds.), *Conceivability and Possibility*. Oxford University Press. [cited on pg.]
- Beddor, B. (MS). Evidentialism, circularity, and grounding. Manuscript, Rutgers University. [cited on pg.]
- Bedke, M. (2010). Developmental process reliabilism: On justification, defeat, and evidence. *Erkenntnis*, 73(1), 1--17. [cited on pg.]
- Bengson, J. (MS). The intellectual given. Manuscript, University of Wisconsin. [cited on pg.]
- Bjornsson, G. & Finlay, S. (2010). Metaethical contextualism defended. *Ethics*, 121(1), 7--36. [cited on pg.]
- Boghossian, P. (FC). The nature of inference. *Philosophical Studies*. [cited on pg.]
- BonJour, L. (2002). Toward a defense of empirical foundationalism. In M. DePaul (Ed.), *Resurrecting Old-Fashioned Foundationalism*. Rowman and Littlefield. [cited on pg.]
- Bratman, M. (2009a). Intention, belief, practical, theoretical. In S. Robertson (Ed.), *Spheres of Reason*. Oxford University Press. [cited on pg.]

- Bratman, M. (2009b). Intention, practical rationality, and self-governance. *Ethics*, 119(3), 411--443. [cited on pg.]
- Brewer, B. (1999). *Perception and Reason*. Oxford University Press. [cited on pg.]
- Broome, J. (1997). Reasons and motivation. *Proceedings of the Aristotelian Society*, 71(1), 131--146. [cited on pg.]
- Broome, J. (1999). Normative requirements. *Ratio*, 12, 389--419. [cited on pg.]
- Broome, J. (2001). Are intentions reasons? and how should we cope with incommensurable values? In C. Morris & A. Ripstein (Eds.), *Practical Rationality and Preference: Essays for David Gauthier*. Cambridge University Press. [cited on pg.]
- Broome, J. (2005a). Does rationality give us reasons? *Philosophical Issues*, 15, 321--337. [cited on pg.]
- Broome, J. (2005b). Have we reason to do as rationality requires?—a comment on Raz. *Journal of Ethics & Social Philosophy*, 1. [cited on pg.]
- Broome, J. (2007a). Does rationality consist in responding correctly to reasons? *Journal of Moral Philosophy*, 4, 349--374. [cited on pg.]
- Broome, J. (2007b). Requirements. In J. J. Toni Rønnow-Rasmussen, Björn Petersson & D. Egonsson (Eds.), *Homage à Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*. Online Tribute (url: <http://www.fil.lu.se/hommageawlodek/>). [cited on pg.]
- Broome, J. (2007c). Wide or narrow scope? *Mind*, 116(462), 359--370. [cited on pg.]
- Broome, J. (2008a). Is rationality normative? *Disputatio*, 11. [cited on pg.]
- Broome, J. (2008b). Reply to Southwood, Kearns and Star, and Cullity. *Ethics*, 119(1), 96--108. [cited on pg.]
- Broome, J. (FC). Comments on Boghossian. *Philosophical Studies*. [cited on pg.]
- Broome, J. (MS). *Rationality Through Reasoning*. Manuscript, Oxford University. [cited on pg.]
- Brueckner, A. (2002). Williamson on the primeness of knowing. *Analysis*, 62(3), 197--202. [cited on pg.]
- Brueckner, A. (2005). Knowledge, evidence, and skepticism according to Williamson. *Philosophy and Phenomenological Research*, LXX(2), 436--443. [cited on pg.]
- Brueckner, A. (2009). E=K and perceptual knowledge. In P. Greenough & D. Pritchard (Eds.), *Williamson on Knowledge*. Oxford University Press. [cited on pg.]
- Brunero, J. (2007). Are intentions reasons? *Pacific Philosophical Quarterly*, 88(4), 424--444. [cited on pg.]

- Brunero, J. (2010). The scope of rational requirements. *The Philosophical Quarterly*, 60(238), 28--49. [cited on pg.]
- Brunero, J. (2012). Instrumental rationality, symmetry, and scope. *Philosophical Studies*, 157(1), 125--140. [cited on pg.]
- Buss, S. (MS). Norms of rationality and the superficial unity of the mind. Manuscript, University of Michigan. [cited on pg.]
- Chang, R. (1997). Introduction. In R. Chang (Ed.), *Incommensurability, Incomparability, and Practical Reason*. Harvard University Press. [cited on pg.]
- Chisolm, R. (1966). Freedom and action. In *Freedom and Determinism*. Random House. [cited on pg.]
- Chudnoff, E. (FC). What intuitions are like. *Philosophy and Phenomenological Research*. [cited on pg.]
- Coffman, E. (2008). Warrant without truth? *Synthese*, 162(2), 173--194. [cited on pg.]
- Cohen, S. (1984). Justification and truth. *Philosophical Studies*, 46(3), 279--295. [cited on pg.]
- Cohen, S. & Lehrer, K. (1983). Justification, truth, and knowledge. *Synthese*, 55(2), 191--207. [cited on pg.]
- Comesana, J. (2002). The diabol and the demon. *Philosophical Studies*, 110, 249--266. [cited on pg.]
- Comesana, J. (2010). Evidentialist reliabilism. *Noûs*, 44(4), 571--600. [cited on pg.]
- Dancy, J. (2000). *Practical Reality*. Oxford University Press. [cited on pg.]
- Dancy, J. (2004a). *Ethics without Principles*. Oxford University Press. [cited on pg.]
- Dancy, J. (2004b). Two ways of explaining actions. *Royal Institute of Philosophy Supplement*. [cited on pg.]
- Dancy, J. (2009). Reasons and rationality. In S. Robertson (Ed.), *Spheres of Reason*. Oxford University Press. [cited on pg.]
- Darwall, S. (1983). *Impartial Reason*. Cornell University Press. [cited on pg.]
- Davidson, D. (1980). *Essays on Actions and Events*. Oxford University Press. [cited on pg.]
- Dorsch, F. (FC). The phenomenal presence of reasons. In F. D. Fiona Macpherson, Martine Nida-Rümelin (Ed.), *Phenomenal Presence*. Oxford University Press. [cited on pg.]
- Enc, B. (2003). *How We Act: Causes, Reasons, and Intentions*. Oxford University Press. [cited on pg.]

- Engel, M. (1992). Personal and doxastic justification in epistemology. *Philosophical Studies*, 67(2), 133--150. [cited on pg.]
- Enoch, D. (2006). Agency, shmagency: Why normativity won't come from what is constitutive of agency. *The Philosophical Review*, 115(2), 169--198. [cited on pg.]
- Evans, I. (FC). The problem of the basing relation. *Synthese*. [cited on pg.]
- Feldman, R. (1988). Having evidence. In *Philosophical Analysis*. Kluwer Publishing Co. [cited on pg.]
- Feldman, R. (2003). *Epistemology*. Prentice-Hall. [cited on pg.]
- Feldman, R. (2007). Reasonable religious disagreement. In L. Antony (Ed.), *Philosophers without Gods*. Oxford University Press. [cited on pg.]
- Feldman, R. & Conee, E. (1985). Evidentialism. *Philosophical Studies*, 48(1), 15--34. [cited on pg.]
- Fitelson, B. (2010). Strengthening the case for knowledge from falsehood. *Analysis*, 70(4), 666--669. [cited on pg.]
- Foot, P. (1972). Morality as a system of hypothetical imperatives. *The Philosophical Review*, 81(3), 305--316. [cited on pg.]
- Foot, P. (1997). Recantation 1994 of "Morality as a System of Hypothetical Imperatives". In A. G. Stephen Darwall & P. Railton (Eds.), *Moral Discourse and Practice*. Oxford University Press. [cited on pg.]
- Fumerton, R. (2010). Foundationalist theories of epistemic justification. *Stanford Encyclopedia of Philosophy*. [cited on pg.]
- Gibbard, A. (1990). *Wise Choices, Apt Feelings*. Oxford University Press. [cited on pg.]
- Gibbons, J. (2001). Knowledge in action. *Philosophy and Phenomenological Research*, 62(3), 579--600. [cited on pg.]
- Gibbons, J. (2006). Access externalism. *Mind*, 115(457), 19--39. [cited on pg.]
- Gibbons, J. (2010). Things that make things reasonable. *Philosophy and Phenomenological Research*, 81(2), 335--361. [cited on pg.]
- Ginet, C. (1990). Justification. *Journal of Philosophical Research*, 15(93), 93--107. [cited on pg.]
- Ginet, C. (2008). In defense of a non-causal account of reasons explanations. *Journal of Ethics*, 12(3/4), 229--237. [cited on pg.]
- Goldman, A. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and Knowledge*. D. Reidel. [cited on pg.]

- Goldman, A. (1988). Strong and weak justification. *Philosophical Perspectives*, 2, 51--69. [cited on pg.]
- Goldman, A. (2009a). Internalism, externalism, and the architecture of justification. *The Journal of Philosophy*, 106(6), 309--338. [cited on pg.]
- Goldman, A. (2009b). Williamson on knowledge and evidence. In D. Pritchard (Ed.), *Williamson on Knowledge*. Oxford University Press. [cited on pg.]
- Goldman, A. (2011). Toward a synthesis of reliabilism and evidentialism? or: Evidentialism's problems, reliabilism's rescue package. In T. Dougherty (Ed.), *Evidentialism and its Discontents*. Oxford University Press. [cited on pg.]
- Goldman, A. (2012). *Reliabilism and Contemporary Epistemology*. Oxford University Press. [cited on pg.]
- Graham, P. A. (2010). In defense of objectivism about moral obligation. *Ethics*, 121(1), 88--115. [cited on pg.]
- Greenspan, P. (1975). Conditional oughts and hypothetical imperatives. *The Journal of Philosophy*, 72(10), 259--276. [cited on pg.]
- Grundman, T. (2009). Reliabilism and the problem of defeaters. *Grazer Philosophische Studien*, 79(1), 65--76. [cited on pg.]
- Hare, C. (2009). Perfectly balanced interests. *Philosophical Perspectives*, 23(1), 165--176. [cited on pg.]
- Hare, C. (2010). Take the sugar. *Analysis*, 70(2), 237--247. [cited on pg.]
- Harman, G. (1973). *Thought*. Princeton University Press. [cited on pg.]
- Harman, G. (1984). Logic and reasoning. *Synthese*, 60(1), 107--127. [cited on pg.]
- Harman, G. (1986). *Change in View*. MIT Press. [cited on pg.]
- Harman, G. (2000). Is there a single true morality? In *Explaining Value and Other Essays in Moral Philosophy*. Oxford University Press. [cited on pg.]
- Hill, T. (1973). The hypothetical imperative. *The Philosophical Review*, 82(4), 429--450. [cited on pg.]
- Hookway, C. (1999). Epistemic norms and theoretical deliberation. *Ratio*, 12(4), 380--397. [cited on pg.]
- Hornsby, J. (2005). Semantic knowledge and practical knowledge. *Proceedings of the Aristotelian Society*, 79(1), 107--130. [cited on pg.]
- Hornsby, J. (2007). Knowledge in action. In *Action in Context*. De Gruyter. [cited on pg.]

- Hornsby, J. (2008). A disjunctive conception of acting for reasons. In A. Haddock & F. Macpherson (Eds.), *Disjunctivism: Perception, Action, Knowledge*. Oxford University Press. [cited on pg.]
- Huemer, M. (1999). The problem of memory knowledge. *Pacific Philosophical Quarterly*, 80(4), 346--357. [cited on pg.]
- Huemer, M. (2001). *Skepticism and the Veil of Perception*. Rowman & Littlefield. [cited on pg.]
- Huemer, M. (2008). Revisionary intuitionism. *Social Philosophy & Policy*, 25(1), 368--392. [cited on pg.]
- Hussain, N. (MS). The requirements of rationality. Manuscript, Stanford University. [cited on pg.]
- Hyman, J. (2006). Knowledge and evidence. *Mind*, 115(460), 891--916. [cited on pg.]
- Jackson, F. (1991). Deontic-theoretic consequentialism and the nearest-dearest objection. *Ethics*, 101(3), 461--482. [cited on pg.]
- Jenkins, C. & Nolan, D. (FC). Disposition impossible. *Noûs*. [cited on pg.]
- Kearns, S. & Star, D. (2009). Reasons as evidence. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics, v. 4*. Oxford University Press. [cited on pg.]
- Kelly, T. (2002). The rationality of belief and other propositional attitudes. *Philosophical Studies*, 110(2), 163--196. [cited on pg.]
- Kelly, T. (2010). Peer disagreement and higher-order evidence. In R. Feldman & T. Warfield (Eds.), *Disagreement*. Oxford University Press. [cited on pg.]
- Kiesewetter, B. (2011). "Ought" and the perspective of the agent. *Journal of Ethics & Social Philosophy*. [cited on pg.]
- Klein, P. (1998). Foundationalism and the infinite regress of reasons. *Philosophy and Phenomenological Research*, 58(4), 919--925. [cited on pg.]
- Kolodny, N. (2005). Why be rational? *Mind*, 114(455), 509--563. [cited on pg.]
- Kolodny, N. (2007a). How does coherence matter? *Proceedings of the Aristotelian Society*, 107, 229--263. [cited on pg.]
- Kolodny, N. (2007b). State or process requirements? *Mind*, 116(462), 371--385. [cited on pg.]
- Kolodny, N. (2008a). The myth of practical consistency. *European Journal of Philosophy*, 16(3), 366--402. [cited on pg.]
- Kolodny, N. (2008b). Why be disposed to be coherent? *Ethics*, 118(3), 437--463. [cited on pg.]
- Kolodny, N. (MS). Instrumental transmission. Manuscript, University of California. [cited on pg.]

- Kolodny, N. & MacFarlane, J. (2010). Ifs and oughts. *The Journal of Philosophy*, 107(3), 115--143. [cited on pg.]
- Korsgaard, C. (1996). *The Sources of Normativity*. Harvard University Press. [cited on pg.]
- Lavin, D. (FC). Must there be basic action? *Noûs*. [cited on pg.]
- Lehrer, K. & Paxson, T. (1969). Knowledge: Undefeated justified true belief. *The Journal of Philosophy*, 66(8), 225--237. [cited on pg.]
- Littlejohn, C. (2009). The externalist's demon. *Canadian Journal of Philosophy*, 39(3), 399--434. [cited on pg.]
- Littlejohn, C. (2012). *Justification and the Truth-Connection*. Cambridge University Press. [cited on pg.]
- Lord, E. (2008). Dancy on acting for the right reason. *Journal of Ethics & Social Philosophy*. [cited on pg.]
- Lord, E. (2010). Having reasons and the factoring account. *Philosophical Studies*, 149(3), 283--296. [cited on pg.]
- Lord, E. (MSa). Acting for the right reasons, abilities, and obligation. Manuscript, Princeton University. [cited on pg.]
- Lord, E. (MSb). The explanatory problem for cognitivism about practical reason. Manuscript, Princeton University. [cited on pg.]
- Lord, E. (MSc). How to be rational without being justified. Manuscript, Princeton University. [cited on pg.]
- Lord, E. (MSd). Intentions, intensification, and mere permissibility. Manuscript, Princeton University. [cited on pg.]
- Lord, E. (MSe). Seeing is believing. Manuscript, Princeton University. [cited on pg.]
- Lord, E. & Sylvan, K. (MS). Prime time (for the basing relation). Manuscript, Princeton University and Rutgers University. [cited on pg.]
- Lycan, W. G. (1988). *Judgement and Justification*. Cambridge University Press. [cited on pg.]
- Lyons, J. (2008). Evidence, experience, and externalism. *Australasian Journal of Philosophy*, 86(3), 461--479. [cited on pg.]
- Lyons, J. (2009). *Perception and Basic Beliefs*. Oxford University Press. [cited on pg.]
- Marcus, E. (2012). *Rational Causation*. Harvard University Press. [cited on pg.]
- Markovits, J. (2010). Acting for the right reasons. *The Philosophical Review*, 119(2), 201--242. [cited on pg.]

- Markovits, J. (FC). Saints, heroes, sages, and villains. *Philosophical Studies*. [cited on pg.]
- McDowell, J. (1994). *Mind and World*. Harvard University Press. [cited on pg.]
- McKeever, S. & Ridge, M. (2012). Elusive reasons. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, volume 7. Oxford University Press. [cited on pg.]
- McNaughton, D. & Rawling, P. (2011). The making/evidential reason distinction. *Analysis*, 71(1), 100--102. [cited on pg.]
- Mele, A. (2007). Reasonology and false beliefs. *Philosophical Papers*, 36(1), 91--118. [cited on pg.]
- Millar, A. (1991). *Reasons and Experience*. Oxford University Press. [cited on pg.]
- Millar, A. (2004). *Understanding People*. Oxford University Press. [cited on pg.]
- Miller, C. (2008). Motivation in agents. *Noûs*, 42(2), 222--266. [cited on pg.]
- Milligan, D. (1974). Reasons as explanations. *Mind*, 88(330), 180--193. [cited on pg.]
- Moore, G. E. (1912). *Ethics*. Oxford University Press. [cited on pg.]
- Neta, R. (2008). What evidence do you have? *British Journal for the Philosophy of Science*, 59(1), 89--119. [cited on pg.]
- Neta, R. & Pritchard, D. (2007). McDowell and the new evil genius. *Philosophy and Phenomenological Research*, 74(2), 381--396. [cited on pg.]
- Parfit, D. (1984). *Reasons and Persons*. Oxford University Press. [cited on pg.]
- Parfit, D. (1997). Reasons and motivation. *Aristotelian Society Supplementary Volume*, 71(1), 99--130. [cited on pg.]
- Parfit, D. (2011). *On What Matters*. Oxford University Press. [cited on pg.]
- Paul, L. (2000). Aspect causation. *The Journal of Philosophy*, 97(4), 235--256. [cited on pg.]
- Plantinga, A. (1993). *Warrant and Proper Function*. Oxford University Press. [cited on pg.]
- Pollock, J. (1974). *Knowledge and Justification*. Princeton University Press. [cited on pg.]
- Pollock, J. (1986). *Contemporary Theories of Knowledge*. Rowman & Littlefield. [cited on pg.]
- Pollock, J. & Cruz, J. (1999). *Contemporary Theories of Knowledge*. Rowman & Littlefield. [cited on pg.]
- Price, H. (1950). *Perception*. Methuen. [cited on pg.]
- Pritchard, D. (2012). *Epistemological Disjunctivism*. Oxford University Press. [cited on pg.]
- Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*, 34(4), 517--549. [cited on pg.]

- Pryor, J. (2005). There is immediate justification. In *Contemporary Debates in Epistemology*. Blackwell. [cited on pg.]
- Raz, J. (2005a). Instrumental rationality: A reprise. *Journal of Ethics & Social Philosophy*, 1. [cited on pg.]
- Raz, J. (2005b). The myth of instrumental rationality. *Journal of Ethics & Social Philosophy*, 1. [cited on pg.]
- Regan, D. (1980). *Utilitarianism and Cooperation*. Oxford University Press. [cited on pg.]
- Reisner, A. (2009). Unifying the requirements of rationality. *Philosophical Explorations*, 12(2), 243--260. [cited on pg.]
- Rosen, G. (2001). Nominalism, naturalism, epistemic relativism. *Nous*, 35, 69--91. [cited on pg.]
- Ross, A. (1944). Imperatives and logic. *Philosophy of Science*, 11(1), 30--46. [cited on pg.]
- Ross, J. (2006). *Acceptance and Practical Reason*. PhD thesis, Rutgers University. [cited on pg.]
- Ross, J. (2012). Rationality, normativity, and commitment. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, volume 7. Oxford University Press. [cited on pg.]
- Ross, J. & Schroeder, M. (FC). Reversibility or disagreement. *Mind*. [cited on pg.]
- Scanlon, T. (1998). *What We Owe to Each Other*. Harvard University Press. [cited on pg.]
- Scanlon, T. (2001). Thomson on self-defense. In A. Byrne, R. Stalnaker, & R. Wedgwood (Eds.), *Fact and Value: Essays on Ethics and Metaphysics for Judith Jarvis Thomson*. MIT Press. [cited on pg.]
- Scanlon, T. (2004). Reasons: A puzzling duality? In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*. Oxford University Press. [cited on pg.]
- Scanlon, T. (2007). Structural irrationality. In *Common Minds: Themes from the Philosophy of Phillip Pettit*. Oxford University Press. [cited on pg.]
- Scanlon, T. (2008). *Moral Dimensions: Permissibility, Meaning, Blame*. Harvard University Press. [cited on pg.]
- Scanlon, T. (MS). *Being Realistic about Reasons*. Manuscript, Harvard University. [cited on pg.]
- Schnee, I. (MS). Duplicate reasons. Unpublished Manuscript, University of Western Kentucky. [cited on pg.]
- Schossler, M. (2007). Basic deviance reconsidered. *Analysis*, 67(295), 186--194. [cited on pg.]
- Schossler, M. (2012). Taking something as a reason for action. *Philosophical Perspectives*, 41(2), 267--304. [cited on pg.]

- Schroeder, M. (2004). The scope of instrumental reason. *Philosophical Perspectives*, 18, 337--364. [cited on pg.]
- Schroeder, M. (2005a). The hypothetical imperative? *Australasian Journal of Philosophy*, 83(3), 357--372. [cited on pg.]
- Schroeder, M. (2005b). Instrumental mythology. *Journal of Ethics & Social Philosophy*, 1. [cited on pg.]
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford University Press. [cited on pg.]
- Schroeder, M. (2008). Having reasons. *Philosophical Studies*, 139(1), 57--71. [cited on pg.]
- Schroeder, M. (2009a). Buck passers' negative thesis. *Philosophical Explorations*. [cited on pg.]
- Schroeder, M. (2009b). Means-end coherence, stringency, and subjective reasons. *Philosophical Studies*, 143(2), 223--248. [cited on pg.]
- Schroeder, M. (2011). What does it take to 'have' a reason? In A. Reisner & A. Steglich-Petersen (Eds.), *Reasons for Belief*. Cambridge University Press. [cited on pg.]
- Schroeder, M. (FC). Stakes, withholding, and pragmatic encroachment on knowledge. *Philosophical Studies*. [cited on pg.]
- Schroeder, M. (MSa). Hypothetical imperatives: Scope and jurisdiction. Manuscript, University of Southern California. [cited on pg.]
- Schroeder, M. (MSb). Knowledge is belief for sufficient (objective and subjective) reason. Manuscript, University of Southern California. [cited on pg.]
- Setiya, K. (2007a). Cognitivism about instrumental reason. *Ethics*, 117(4), 649--673. [cited on pg.]
- Setiya, K. (2007b). *Reasons Without Rationalism*. Princeton University Press. [cited on pg.]
- Silins, N. (2005). Deception and evidence. *Philosophical Perspectives*, 19(1), 375--404. [cited on pg.]
- Skorupski, J. (2011). *The Domain of Reasons*. Oxford University Press. [cited on pg.]
- Smith, M. (1994). *The Moral Problem*. Blackwell Publishing. [cited on pg.]
- Smith, M. (2008). Consequentialism and the nearest-dearest objection. In I. Ravenscroft (Ed.), *Minds, Ethics, and Conditionals: Themes from the Philosophy of Frank Jackson*. Clarendon Press. [cited on pg.]
- Smithies, D. (2006). *Rationality and the Subject's Point of View*. PhD thesis, NYU. [cited on pg.]
- Sosa, E. (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge, vol. II*. Oxford University Press. [cited on pg.]

- Southwood, N. (2008). Vindicating the normativity of rationality. *Ethics*, 118(1), 9--30. [cited on pg.]
- Sutton, J. (2007). *Without Justification*. MIT Press. [cited on pg.]
- Swain, M. (1981). *Reasons and Knowledge*. Cornell University Press. [cited on pg.]
- Swain, M. (1988). Alston's internalistic externalism. *Philosophical Perspectives*, (2), 461--473. [cited on pg.]
- Sylvan, K. (2014). *On the Normativity of Epistemic Rationality*. PhD thesis, Rutgers University. [cited on pg.]
- Sylvan, K. (MSa). Doing epistemology with reasons. Manuscript, Rutgers University. [cited on pg.]
- Sylvan, K. (MSb). Reasons and the metaphysics of justification. Manuscript, Rutgers University. [cited on pg.]
- Sylvan, K. (MSc). Subtle reasons, coarse motivations, and epistemologists' blindspots. Manuscript, Rutgers University. [cited on pg.]
- Thomson, J. J. (1986). Imposing risks. In *Rights, Restitution, and Risk*. Harvard University Press. [cited on pg.]
- Thomson, J. J. (1990). *The Realm of Rights*. Harvard University Press. [cited on pg.]
- Thomson, J. J. (2008). *Normativity*. Open Court Press. [cited on pg.]
- Turri, J. (2009). The ontology of epistemic reasons. *Nouûs*, 43(3), 490--512. [cited on pg.]
- Turri, J. (2010). On the relationship between propositional and doxastic justification. *Philosophy and Phenomenological Research*, 80(2), 312--326. [cited on pg.]
- Turri, J. (2011). Believing for a reason. *Erkenntnis*, 74(3), 383--397. [cited on pg.]
- Unger, P. (1979). *Ignorance*. Oxford University Press. [cited on pg.]
- Väyrynen, P. (2006). Moral generalism: Enjoy in moderation. *Ethics*, 116, 707--741. [cited on pg.]
- Velleman, D. (2000). *The Possibility of Practical Reason*. Oxford University Press. [cited on pg.]
- Vogelstein, E. (FC). Subjective reasons. *Ethical Theory and Moral Practice*. [cited on pg.]
- Warfield, T. (2005). Knowledge from falsehood. *Philosophical Perspectives*, 19(1), 405--416. [cited on pg.]
- Way, J. (2009). Two accounts of the normativity of rationality. *Journal of Ethics & Social Philosophy*. [cited on pg.]

- Way, J. (2010a). Defending the wide-scope account of instrumental reason. *Philosophical Studies*, 147(2), 213--233. [cited on pg.]
- Way, J. (2010b). The normativity of rationality. *Philosophy Compass*, 5(12), 1057--1068. [cited on pg.]
- Way, J. (2011). The symmetry of rational requirements. *Philosophical Studies*, 155(2). [cited on pg.]
- Way, J. (2012). Explaining the instrumental principle. *Australasian Journal of Philosophy*, 90(3), 487--506. [cited on pg.]
- Way, J. (FC). Intentions, akrasia, and mere permissibility. *Organon F*. [cited on pg.]
- Weatherson, B. (2008). Deontology and descartes' demon. *The Journal of Philosophy*, 105(9), 540--569. [cited on pg.]
- Weatherson, B. (MS). Knowledge, bets, and interests. Manuscript, University of Michigan. [cited on pg.]
- Wedgwood, R. (2002). Practical reason and desire. *Australasian Journal of Philosophy*, 80(3), 345--358. [cited on pg.]
- Wedgwood, R. (2003). Choosing rationally and choosing correctly. In S. Stroud & C. Tappolet (Eds.), *Weakness of Will and Practical Irrationality*. Oxford University Press. [cited on pg.]
- Wedgwood, R. (2006). The normative force of reasoning. *Noûs*, 40(4), 660--686. [cited on pg.]
- Wedgwood, R. (2007). *The Nature of Normativity*. Oxford University Press. [cited on pg.]
- Wedgwood, R. (FCa). Justified inference. *Synthese*. [cited on pg.]
- Wedgwood, R. (FCb). The right thing to believe. In *The Aim of Belief*. Oxford University Press. [cited on pg.]
- White, R. (2005). Epistemic permissiveness. *Philosophical Perspectives*, 19(1), 445--459. [cited on pg.]
- Williams, B. (1981). Internal and external reasons. In *Moral Luck*. Cambridge University Press. [cited on pg.]
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press. [cited on pg.]
- Williamson, T. (2007). *The Philosophy of Philosophy*. Blackwell. [cited on pg.]
- Williamson, T. (2009). Replies to critics. In P. Greenough & D. Pritchard (Eds.), *Williamson on Knowledge*. Oxford University Press. [cited on pg.]
- Wright, C. (FC). Comments on Boghossian's "The Nature of Inference". *Philosophical Studies*. [cited on pg.]
- Zimmerman, M. J. (2008). *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge University Press. [cited on pg.]