

© 2014

Kurt Ludwig Sylvan

**ALL RIGHTS RESERVED**

ON THE NORMATIVITY OF EPISTEMIC  
RATIONALITY

BY KURT LUDWIG SYLVAN

A dissertation submitted to the  
Graduate School—New Brunswick  
Rutgers, The State University of New Jersey  
in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy  
Graduate Program in Philosophy

Written under the direction of

Ernest Sosa

and approved by

---

---

---

---

---

New Brunswick, New Jersey

May, 2014

## ABSTRACT OF THE DISSERTATION

### On the Normativity of Epistemic Rationality

by Kurt Ludwig Sylvan

Dissertation Director: Ernest Sosa

Many epistemologists equate the rational and the justified. Those who disagree have done little to explain the difference, leading their opponents to suspect that the distinction is an *ad hoc* one designed to block counterexamples. The first aim of this dissertation—pursued in the first three chapters—is to improve this situation by providing a detailed, independently motivated account of the distinction. The account is unusual in being inspired by no particular theoretical tradition in epistemology, but rather by ideas in the meta-ethical literature on reasons and rationality. The account is also unusual in proposing that the distinction between rationality and justification can be derived from a reasons-based account of justification. Historically, this is a striking claim. In epistemology, reasons-based accounts of justification are standardly treated as paradigmatically internalist accounts, but this dissertation argues that we should believe the reverse: given the best views about reasons—again drawn from meta-ethics—we should expect reasons-based accounts of justification to be strongly externalist.

The first half of the dissertation might leave one wondering why rationality matters from the epistemic point of view. The second aim of the dissertation is to answer this question. The final two chapters argue (1) that we can only explain why rationality matters from the epistemic point of view if we reject the nearly universal assumption that all derivative epistemic value is instrumental value, and (2) that there are powerful

reasons to reject this assumption, since it is the true origin of the so-called *swamping problem*. It is then argued that if we reject the instrumentalist assumption, we can get a *truth-oriented* account of epistemic value that provides a unified explanation of how rationality, justification, and knowledge matter from the epistemic point of view. This result is unprecedented: while some epistemologists find room for both internalist and externalist species of epistemic value, virtually all assume that these properties cannot have a common evaluative ground, and especially not a common *truth-oriented* one. The concluding moral of the dissertation is that once we properly understand what it is to be truth-oriented, we see that this common assumption is mistaken.

## Acknowledgements

"I never thought I'd need so many people...."

—David Bowie ("Five Years")

That line resonates with me as I put some rushed finishing touches on this collection of papers—as does the sad refrain: "Five years! That's all we've got!" It is hard to believe that the five (and a half, at least) greatest years of my life are coming to an end. Certain people—a lot of them—are responsible for the greatness of these years. I'm lucky not only to have had their support and company, but to have the promise of interacting with many of them for years to come. It's high time to acknowledge them.

But before I acknowledge these many people, I'd like to thank the ones who made it possible for me to get as far as graduate school in the first place. First, I thank my mother—not only for single-handedly supporting me and my education but also for never seriously questioning my goal to become a philosopher. I announced that goal to her when I was fifteen, after stumbling on philosophy while reading a biography of the singer of my epigraph. She was unfazed. I can only recall one moment—I was taking my first symbolic logic class—when she urged me to become a lawyer. This was doubtless because I was annoyingly overcompensating for my previous ignorance of analytic philosophy by barraging her with arguments.

Among (living) professional philosophers, my oldest debts are to my undergraduate advisors Evan Fales, Richard Fumerton, and Sarah Buss. To Evan and Richard, I owe my original interest in metaphysics and epistemology; to Sarah, my original interest in the philosophy of normativity. Evan was the best kind of advisor—unbelievably generous with his time, supportive, but also vigorously demanding. With him I did numerous formative independent studies in metaphysics, the philosophy of science, and epistemology. Although I didn't end up doing much metaphysics or philosophy of science

as a graduate student and my epistemological views have drifted far from the classical foundations on which he raised me, I fully expect to return to the issues I studied with him in my career. Richard and Sarah, in turn, rank among the most impressive philosophical interlocutors I've ever encountered and they both prepared me well for the kinds of conversations I now have daily. They both continue to serve as ultimate models of serious philosophy, each possessing a nearly unrivaled combination of depth, analytical rigor, dialectical ingenuity, and undying concern for the most central—well, the *real*—issues in philosophy. Having now seen much of the philosophical world, I remain convinced that I was brought up in one of its best locales. I have met very few philosophers nearly as great as Evan, Richard and Sarah. The only equals I know well are on my committee (or almost on it, in Parfit's case).

At Rutgers and the greater New York metropolitan area, I am indebted to seemingly countless philosophers. I'll attempt to fix on the most central influences in chronological order, starting with faculty and then turning to fellow graduate students.

In my first semester, I had the luck of taking a proseminar with Alvin Goldman that focused on epistemology and the philosophy of mind. It was here that I received my first direct taste of Rutgers epistemology at its finest. Alvin's impact on my development as a philosopher is tremendous and hard to measure—in addition to being partly responsible for my turn to externalism, he has provided many of the best comments I've received on my work, taught me more than anyone about how to write philosophy well, encouraged and stimulated numerous projects of mine (including my first publication) while pushing me with gracious patience to improve my style, attended my reading groups, presented in them, and stayed on my committee even when it seemed that my thesis might end up being 400+ pages long. It is an immense honor to have had such support from one of the world's greatest epistemologists, and I'm constantly thankful for it.

In my second semester, I had the further luck of taking a seminar with Ruth Chang and Derek Parfit that was, I am sure, the most influential class I've ever taken. If I hadn't taken this class, it is certain that I would have written a completely different and surely less interesting dissertation. Besides this, it was taught by two of most personally significant philosophers I've known at Rutgers. I'm not sure I would have

survived graduate school without the encouragement of Ruth and Derek. Ruth, I don't know how you do it, but keep doing it for your students: your comments have always been among the best and sternest I've received, and yet you've also given me more confidence as a philosopher than anyone by being so supportive and kind to me, making me feel like my work is actually worth something and to be taken seriously. Thank you also for doing independent studies with me that led me to be aware of much of the axiological literature that inspired Chapters 4 and 5 (and disabused me of the notion that value = utility), and giving me so much of the stylistic advice after my prospectus defense that led to the dissertation's current form.

Derek was a huge inspiration—so much so that I almost called the dissertation *On What Epistemically Matters*. He is committed to philosophy with an otherworldly level of passion and rigor, his investment in his projects has a feeling of life-or-death urgency, and yet, one-on-one, he's proven the humblest philosopher I've ever met, always treating me as at least an equal and often—remarkably—as a superior. Whenever I think I'm becoming a good philosopher, I think of Parfit and realize how much harder I must work and how much humbler I must be.

Now comes my second year, where I met the crowning influence: Ernie Sosa, than whom no better advisor could be conceived. With their distinctively exploratory dialectical style and boundless riches of points and counterpoints, Ernie's seminars were the most exciting I've ever taken, and many of my best ideas were stimulated by being a member of these seminars. After one seminar I was hooked, and by becoming his student I got the bonus of experiencing the same kind of excitement in the dissertation group that Ernie has run for decades, producing many of the best epistemologists currently around. To be part of this tradition is a huge honor, but greater still is the honor of having spent years corresponding and talking with one of the world's greatest epistemologists, kindest persons, and most encouraging mentors. There is no chance I could have thrived as I have without Ernie; to him I dedicate the dissertation.

After that year, two influences at the faculty level stand out. I was thrilled when Susanna Schellenberg was hired at the end of my second year. I wish I'd arrived at

Rutgers later, so that I could have formally taken her classes. But while our philosophical interactions only began in my third year, her impact has been huge. Besides frequently providing me with philosophically incisive comments and lots of encouragement, Susanna more than anyone has helped me to adapt to my transition from being a graduate student to being an actual professional philosopher. I am particularly grateful for the advice she's given me about the job market, about writing for journals and submitting to journals, and more generally about marketing my work.

Jonathan Dancy has played a different sort of role in helping me to adapt to the transition to being a real professional. When we met at the USC/UCLA conference and Jonathan praised the paper I presented, I felt like I'd finally "made it". Finally, I thought, I've managed to write something that one of my heroes sincerely likes. Of course, I've been lucky to work with many of my heroes, but as an outsider Jonathan had no obligation to try to take my work seriously. But he did—and even better, he joined my committee, has given me excellent comments, and will, I hope, be a contact for years to come. Jonathan's substantive influence on me as a philosopher is also huge, but most of that influence came well before we met. I can't think of anyone who has written so much with which I agree. Indeed, when Jonathan disagrees with something I've written, I feel like it's my conscience talking, and I instantly want to change my mind. (But wait: the author of "The Logical Conscience" wouldn't say I therefore ought rationally to do so. So now what?)

I also owe a lot to an enormous number of fellow students at Rutgers, Princeton, NYU and beyond. Errol Lord has been my closest interlocutor for the past four years; many of the ideas in this dissertation were developed through discussions with him, and the ideas about the nature of believing for normative reasons in Chapter 1 are really joint ideas which you'll find in our paper "Prime Time (for the Basing Relation)". On the ideas in this dissertation, I've received many important comments from Greg Antill, Marcello Antosh, Rima Basu, Bob Beddor, David Black, Thomas Blanchard, Rodrigo Borges, Tim Campbell, Marilie Coetsee, Lindsay Crawford, Kate Devitt, Megan Feeney, Will Fleisher, Amy Floweree, Daniel Fogal, Georgi Gardiner, Simon Goldstein, Beth Henzel, Stephanie Leary, Katy Meadows, Ricardo Mena, Lisa Miracchi, Shyam



Nair, Kate Nolfi, Pamela Robinson, Blake Roeber, Daniel Rubio, Justin Sharber, Erica Shumener, Peter van Elswyk, Tobias Wilsch, and audiences at the Georgetown Reasons and Reasoning Conference, the Northwestern-Notre Dame Epistemology Conference, the University of Edinburgh, and the USC/UCLA Graduate Conference. I can't limit the list to graduate students, however, as I've also received crucial comments from Matt Benton, Dan Greco, Joachim Horvath, Jonathan Ichikawa, Clayton Littlejohn, Heather Logue, Conor McHugh, Carlotta Pavese, Ted Poston, Andrew Sepielli, Daniel Star, Mike Titelbaum, Lee Walters, Jonathan Way, Daniel Whiting, Eric Wiland, and audiences at the St. Louis Annual Conference on Reasons and Rationality and the University of Southampton.

Finally and foremost, I thank my lovely and awesome fiancée Kate for tirelessly supporting me, providing me with intellectual companionship and self-confidence, weeding out some of my lesser ideas with her tremendous wit, making sure I took breaks, proof-reading the entire dissertation before I sent it out to letter writers, providing supreme stylistic advice, making me explain philosophy accessibly (and helping me see that what I write actually has accessible interest), helping calm my dissertation-induced nerves, and putting up with other neurotic academic behavior. I also thank our cat Seneca, whose distinguished ears assessed the cadences of the sentences to follow as I read them aloud, and who also caused me to take breaks with his not-so-Stoic miaows for food.

# Table of Contents

<b>Abstract</b> . . . . .	ii
<b>Acknowledgements</b> . . . . .	iv
<b>Introduction</b> . . . . .	1
<b>1. Reasons and the Metaphysics of Epistemology</b> . . . . .	6
1.1. The Reasonable Picture . . . . .	6
1.2. Reasons: Some Ontology and Ideology . . . . .	12
1.2.1. Normative Reasons and the Relation of Possession . . . . .	12
1.2.2. The Ontology of Normative Reasons . . . . .	14
1.2.3. Motivating Reasons . . . . .	16
1.3. Are Reasons Necessary for Justified Belief? . . . . .	19
1.3.1. Examples against the Necessity of Reasons: Obvious Math . . . . .	19
1.3.2. The Slip . . . . .	21
1.3.3. Possession: An Achievement . . . . .	22
1.3.4. The Achievement in More Detail: Two Kinds of Competence . . . . .	24
Access as Apt Seeming . . . . .	24
Sensitivity as Apt Treating . . . . .	25
Non-Inferential Cases Again . . . . .	27
1.3.5. The Extended Account . . . . .	28
1.3.6. Further Applications: Lyons on Perceptual Justification . . . . .	29
1.3.7. Taking Stock . . . . .	31
1.4. Is Belief for Sufficient Reasons Enough? . . . . .	32
1.4.1. The Insufficiency Argument . . . . .	32
1.4.2. Believing Something Because You Have Good Reasons . . . . .	33

1.4.3.	Other Reasons to Reject RP’s Account of Doxastic Justification?	35
1.5.	Two Levels of Explanation . . . . .	36
1.5.1.	Examples of the Two Levels: Buck-Passing and Hedonism . . . . .	37
1.5.2.	Reliabilism : RP :: Hedonism : Buck-Passing . . . . .	38
1.5.3.	Against the Deflationary Response . . . . .	39
1.6.	Conclusion . . . . .	39
<b>2.</b>	<b>What Apparent Reasons Appear to Be . . . . .</b>	<b>45</b>
2.1.	Reasons and Apparent Reasons . . . . .	45
2.2.	Apparent Reasons <i>De Dicto</i> and <i>De Re</i> . . . . .	47
2.2.1.	<i>De Dicto</i> Views . . . . .	47
2.2.2.	Understanding the Overintellectualization Objection . . . . .	48
2.2.3.	Other Concepts? . . . . .	52
2.2.4.	The <i>De Re</i> View . . . . .	53
2.3.	The Failure of <i>De Re</i> Views . . . . .	54
2.3.1.	The Problem of Unapparent Reasons . . . . .	54
Defeat? . . . . .		56
Why Not Demand the Enablers To Be Apparent? . . . . .		58
Morals . . . . .		60
2.3.2.	The Problem of Wouldn’t-Be Reasons . . . . .	61
2.4.	Competence, Appearance, and Treating . . . . .	64
2.4.1.	A Key Ingredient: Treating . . . . .	64
2.4.2.	Another Key Ingredient: Competence . . . . .	65
2.4.3.	CAT and the Problem of Wouldn’t-Be Reasons . . . . .	67
2.4.4.	Further Analysis of Key Ingredients . . . . .	70
More on Treating and Attraction . . . . .		70
More on Competence . . . . .		72
2.4.5.	Taking Stock . . . . .	75
2.5.	But Don’t Envelopes Slice All the Theories? . . . . .	76

2.6.	Concluding Remarks . . . . .	79
<b>3.</b>	<b>Rationality and Justification: Reasons to Divorce?</b> . . . . .	<b>83</b>
3.1.	Indifference . . . . .	83
3.1.1.	A History of Indifference . . . . .	83
3.1.2.	Why Indifference Is Striking . . . . .	85
3.1.3.	Why Resisting Indifference Matters . . . . .	87
3.1.4.	Antecedents and the Need for Detail . . . . .	88
3.1.5.	The Plan . . . . .	90
3.2.	Objective Reasons, Possession, and Apparent Reasons . . . . .	90
3.2.1.	Objective Reasons and the Early Distinction . . . . .	90
3.2.2.	Apparent Reasons vs. Possessed Objective Reasons . . . . .	92
3.2.3.	Taking Stock: A Spectrum of Reason-Based Evaluations . . . . .	95
3.3.	A Deeper Rationale for the Distinction . . . . .	96
3.3.1.	Apparent Reasons: Two Views . . . . .	96
3.3.2.	Treating, Competence and Appearance: A Better View . . . . .	99
3.3.3.	Grounding the Distinction: Competence and Performance . . . . .	100
3.3.4.	Possessing an Objective Reason to $\phi$ . . . . .	101
3.3.5.	The R/J Distinction: A New Instance of an Old Pattern . . . . .	103
3.4.	Extending the Distinctions to Epistemology . . . . .	104
3.4.1.	Objective Epistemic Reasons: Unpossessed and Possessed . . . . .	104
3.4.2.	Apparent Epistemic Reasons . . . . .	107
3.4.3.	Rationality vs. Justification in Epistemology . . . . .	108
3.4.4.	Further Divergences: Structural Rationality . . . . .	109
3.4.5.	Comparisons with Other Bifurcating Proposals . . . . .	110
3.5.	Implications for Epistemology . . . . .	111
3.5.1.	How the Internalism/Externalism Debate Shouldn't Proceed . . . . .	111
3.5.2.	A Challenge for Internalists . . . . .	113
3.5.3.	Challenging Concessive Externalist Views . . . . .	114

3.5.4.	Light Shed on Puzzling Cases . . . . .	115
3.6.	Concluding Remarks . . . . .	118
<b>4.</b>	<b>Veritism without Instrumentalism . . . . .</b>	<b>122</b>
4.1.	Veritism, Instrumentalism, and Fundamental Epistemic Value . . . . .	122
4.2.	Instrumentalism and the Swamping Problem . . . . .	128
4.2.1.	The Old Problem (and Why It Remains Important) . . . . .	129
4.2.2.	The Deeper Problem . . . . .	131
4.2.3.	Why Other Forms of Derivative Value Help . . . . .	132
4.2.4.	A General Problem for Instrumentalists . . . . .	134
4.3.	Veritism without Instrumentalism . . . . .	136
4.3.1.	Other Ways of Valuing in General . . . . .	137
4.3.2.	Ways to Place Value on Accuracy in Thought . . . . .	138
Commitment . . . . .	139	
Respect . . . . .	140	
Compliance . . . . .	141	
4.3.3.	Deriving Epistemic Value via Hurka's Principle . . . . .	141
4.3.4.	Comparative Propriety and Rankings of Value . . . . .	143
4.3.5.	Advantages . . . . .	144
4.4.	Five Objections Answered . . . . .	145
4.4.1.	Too Demanding? . . . . .	146
4.4.2.	Doesn't the View Then Collapse into Familiar Views? . . . . .	147
4.4.3.	What About Other Epistemic Values? . . . . .	148
4.4.4.	Trivial Truths and the Valuing of Accuracy . . . . .	150
4.4.5.	Does Hurka's Principle Really Generalize? . . . . .	151
4.5.	On the Competition . . . . .	152
4.5.1.	What About Reliabilist Virtue Epistemology? . . . . .	153
4.5.2.	First Horn . . . . .	154
4.5.3.	Second Horn . . . . .	156

4.5.4. A Middle Way Between Reliabilism and Responsibilism . . . . .	157
4.6. Conclusion . . . . .	158
4.7. Appendix A: Must Reliabilists Be Instrumentalists? . . . . .	159
4.8. Appendix B: Does Type Instrumentalism Make a Difference? . . . . .	161
<b>5. Rationality and Fundamental Epistemic Value . . . . .</b>	<b>164</b>
5.1. A Puzzle about the Value of Rational Belief . . . . .	164
5.1.1. Disclaimer . . . . .	168
5.2. Instrumentalism $\wedge$ Derivativeness $\rightarrow$ $\neg$ Retention of Worth . . . . .	169
5.2.1. The Irrelevance of Subjective Instrumentalism . . . . .	170
5.2.2. The Failure of Counterfactual Instrumentalism . . . . .	171
5.2.3. The Failure of Ingredient Instrumentalism . . . . .	174
5.3. What Should We Conclude? . . . . .	176
5.4. Reminder: Independent Reasons to Reject Instrumentalism . . . . .	177
5.5. The Value of Rational Belief Depuzzled . . . . .	179
5.6. Implications for Epistemology . . . . .	182
<b>Bibliography . . . . .</b>	<b>186</b>

## Introduction

### The Goals

The dissertation has two major goals. The first is to defend a novel version of the distinction between rationality and justification in epistemology that profits from recent theorizing about rationality and reasons outside of epistemology. In meta-ethics, it has become common to separate the question of whether it would be rational for one to  $\phi$  from the question of whether one possesses objectively good reasons to  $\phi$ . Rationality only requires us to heed the *apparent reasons* we have and the pressures of *coherence* exerted by our attitudes, and both can fail to coincide with objectively good reasons. My distinction is inspired by this picture. While justification does not require one to heed the objective reasons given by all the facts, it does require one to possess some objectively good reasons that are not outweighed by other objectively good reasons that one possesses. Rationality does not require so much: indeed, one can be rationally required to comply with reasons that aren't objectively good at all.

The second major goal is to explain how epistemic rationality could have any necessary significance from the epistemic point of view, given the pressures to divorce rationality and justification. The parallel distinction in meta-ethics has led many to think that it is an open question whether rationality necessarily has any normative significance. Of course, it is intuitive that rationality has this significance. But one craves an explanation. This is so even if we are merely interested in the significance of epistemic rationality from the epistemic point of view: for as I argue, epistemic rationality does not necessarily reliably promote fundamental epistemic values like accurate belief and has no necessary connection to objectively good epistemic reasons.

To explain the necessary significance of rationality from the epistemic point of view,

I argue that we must reject widely assumed views about epistemic value and normativity. We must deny that all derivative epistemic value is merely instrumental epistemic value and deny that our epistemic obligations are best understood in consequentialist terms. To undermine these views, I show that they are the source of long-standing problems. Indeed, I show we cannot even explain the epistemic value of properties that *are* instrumental to truth given a purely instrumentalist model of epistemic value. If we reject that model, we can get a unified account of the epistemic value of rationality, justification, and knowledge and solve some long-standing problems. Rationality matters in the same way that justification and knowledge matter: epistemologists have just been wrong *across the board* about what it is to matter epistemically.

A striking feature of my view is that it preserves the idea that accurate belief is the fundamental epistemic value but unpacks this idea in a non-consequentialist way. In a crude slogan: accuracy is fundamentally to be respected as the norm of belief, not to be produced as a goal. The way to respect the norm of accuracy just is to form beliefs in the ways characteristic of epistemic rationality. So, there is a truth-oriented explanation of why rational belief matters, though it is not a consequentialist one. A similar kind of explanation, I show, can and should be given of the epistemic value of justified belief and knowledge. In this way, we can secure truth-oriented unity in the theory of epistemic value even though rationality doesn't guarantee reliability.

Epistemologists have so far been unable to grasp this unity because they have made some optional assumptions about how derivative epistemic value is grounded in fundamental epistemic value. While the epistemic value of rationality looks puzzling from an instrumentalist angle, that fact is evidence that this angle is mistaken. Rationality is not the exception but the rule: we cannot fully understand the value of any epistemically significant ways of believing from a purely instrumentalist angle.

## **Chapter Summaries**

The dissertation has two parts that fulfill my two goals. The first consists of three chapters, and the second consists of two.



In Chapter 1 (“Reasons and the Metaphysics of Epistemology”), I defend a reasons-based account of justification that profits from ideas and distinctions from theorizing about reasons outside of epistemology. I argue that externalists in epistemology can and should hold a reasons-based account, and that the best general views about reasons (which I take from meta-ethics) strongly support an externalist—indeed reliabilist—epistemology. This chapter is crucial for the analogies with meta-ethics that underpin my justification/rationality distinction. But the chapter has independent interest: among other things, it shows that it is a mistake to think that the distinction between internalist and reliabilist theories of justification corresponds to the distinction between reasons-based and non-reasons-based theories.

In Chapter 2 (“What Apparent Reasons Appear to Be”), I lay the foundations for the other half of my distinction between justification and rationality. Many meta-ethicists understand rationality as a matter of correctly responding to apparent normative reasons, where apparent normative reasons are not necessarily real normative reasons. I endorse this general idea, but think that the accounts of apparent reasons that meta-ethicists have so far offered are flawed, and defend my own account. On my account, apparent normative reasons are apparent facts that we are competently disposed to treat like objective normative reasons, where the notion of competence is analyzed indirectly in terms of objective reasons and a competence/performance distinction is drawn. While my account allows that apparent reasons can easily fail to be real reasons, it implies that the two are likely to coincide when the conditions for the exercise of competence are normal. I defend this as a virtue of my account.

In Chapter 3 (“Rationality and Justification: Reasons to Divorce?”), I defend my distinction between rationality and justification. I also explain its implications for issues like the possibility of higher-order defeat, the normative impact of perceptual seemings with “checkered” etiologies, and several internalism/externalism debates. On my view, rationality has a structural and a substantive side, and both are separable from justification. Structural rationality consists in complying with the pressures of coherence exerted by our attitudes, which are *hypothetical* in the sense that one can rationally

escape these pressures by dropping the relevant attitudes. Substantive rationality consists in heeding the apparent reasons, where these are understood in the way suggested in Chapter 2. Unlike rational  $\phi$ -ing, justified  $\phi$ -ing requires one to possess an objective reason to  $\phi$  that is not outweighed by other objective reasons one possesses.

In Chapter 4 (“Veritism without Instrumentalism”), I defend a picture of epistemic value that proves crucial for my second goal. I argue that friends of the view that true belief is the fundamental epistemic value (“Veritism”) should reject the idea that all non-fundamental epistemic value is merely a species of instrumental epistemic value (“Instrumentalism”). Instrumentalism, I argue, is responsible for a generalized version of the swamping problem, and Veritists can avoid it iff they reject Instrumentalism.

To be concrete, I sketch an alternative version of Veritism that profits from a form of value derivation that I take from Thomas Hurka’s work on virtue. For a toy illustration outside of epistemology, consider beauty and its proper appreciation. Appreciating beauty is good *because* beauty is good. But that ‘because’ signals no instrumental explanation: appreciating beauty does not reliably cause more beauty to exist, nor is it the product of anything beauty-conducive. How is this type of value derivation relevant to epistemology? On my picture, there are certain *ways of placing value on accuracy in thought* that are manifested by rational belief, justified belief, and knowledge. These forms of belief are epistemically good because they manifest ways of placing value on accuracy in thought, and the ways of placing value on accuracy in thought are, in turn, epistemically good because accuracy is epistemically good. But like the ‘because’ in the earlier example, these are not purely instrumental ‘because’s. This makes a crucial difference, I argue, because this form of value derivation is immune from swamping by the presence of the more fundamental value (unlike instrumental derivation).

In Chapter 5 (“Rationality and Fundamental Epistemic Value”), I deepen the objection to the instrumentalist model of epistemic value by arguing that there is no satisfactory instrumentalist explanation of the worth of epistemic rationality or rational beliefs. Instrumentalists must deny that epistemic rationality necessarily has any value from the epistemic point of view. Extending the arguments from Chapter 4, I argue

that this is really a symptom of a more general problem, and that instrumentalists cannot even explain the epistemic value of forms of belief (e.g., justified belief) that *are* instrumentally connected to truth! Using the account developed in Chapter 4, I then show that we can explain why rationality has necessary rather than merely contingent epistemic value if we reject the instrumentalist model. I conclude with some broader morals about what is really at stake in the dispute between reliabilists and internalists in epistemology: the dispute does not, as is almost universally believed, boil down to a dispute about the place of truth in epistemology. Internalists can also secure a truth connection —just not an instrumental one.

## Chapter 1

### Reasons and the Metaphysics of Epistemology

#### Overview

On one familiar tradition in epistemology, reasons figure centrally in the analysis of epistemic justification. On the most ambitious version of this traditional view, justified belief is nothing over and above belief for all-things-considered good epistemic reasons. In recent years, this reasons-based view has met with considerable opposition. Some opponents deny that reasons are necessary for justification. Others deny that belief for all-things-considered good reasons is sufficient for justified belief. In this paper, I answer these lines of opposition to reasons-based views, though in an untraditional way. I show how insights from recent meta-ethical literature on reasons can secure reasons-based views from the opposition, albeit in a novel externalist form. I also show that a common division of labor in this literature eliminates the apparent rivalry between reasons-based views and reliabilist views, revealing them to operate at different levels of meta-normative analysis. Along the way, I uncover two neglected ways in which reliabilist ideas can underpin the traditional view—viz., in analyzing (i) the relation of *possession* that one must bear to reasons to get justified, and (ii) the relation of *believing something because one has a good reason*, which matters for doxastic justification. These observations, I argue, resolve other worries.

#### 1.1 The Reasonable Picture

What role should reasons play in the analysis of epistemic justification? According to a traditional view that I will call the *Reasonable Picture* ("RP" for short), they should play a prominent role. On the most ambitious version of this traditional view, justified

belief is nothing over and above belief for all-things-considered good reasons.<sup>1</sup> In recent years, this reasons-based view has encountered considerable opposition, largely from a sect of externalists I'll label the *Resistance*. Some members of the Resistance insist that reasons are unnecessary for justified belief. Others maintain that belief for all-things-considered good reasons is insufficient for justified belief. If they were right, reasons would do little to illuminate the essence of epistemic justification.

But I think they are wrong, and the aim of this paper is to offer a new line of response to their doubts. The Resistance, I'll argue, crucially overlooks some attractive views about reasons that have become common in new literature on reasons and rationality outside of epistemology. They also crucially neglect the fact that friends of RP can understand both the *possession* of good reasons and the type of *basing* that is necessary for doxastic justification as *achievements* that essentially manifest reliable cognitive ability.<sup>2</sup> If friends of RP invoke reliable cognitive abilities in these ways, they can undermine key sources of opposition to their view. Admittedly, some members of the Resistance might take this reply to show that reasons are *epiphenomenal* in the analysis of justification: reliable cognitive abilities are what *really* matter, they might say. But that reply is compelling, I will argue, only if we conflate two levels of meta-normative

---

<sup>1</sup>Throughout this paper, I will take all-things-considered good reasons for belief to be good reasons that are not defeated by reasons for other doxastic attitudes (disbelief and suspension), in either the undercutting or the rebutting way. Following Schroeder (2012) and (Ms), I take reasons to suspend judgment on whether P to be undercutting defeaters for any reasons that might have counted in favor of believing that P. As a result, one cannot, as I am using the term, have all-things-considered good reasons to believe that P if one isn't permitted to judge that P; for if one isn't permitted to judge that P, one has a conclusive reason to suspend judgment. For this reason, I will also use the term "all-things-considered good reason" interchangeably with "sufficient reason".

<sup>2</sup>Here and throughout the paper, I intend the underlying suggestion to be neutral with respect to different versions of reliabilism. I don't mean to place overwhelming weight on an *ability*-based formulation, though I do find it attractive. The point to be defended is the more abstract one that there are unanticipated places for reliabilist themes to enter into a reasons-based picture of justification without any sacrifice or compromise—specifically, in the analysis of *possession* and the analysis of what it is to believe something *because one has a good reason*. I am perfectly open to stating everything in the classic *process*-based reliabilist framework of Goldman (1979), and use his formulations at various points. And at certain points, I will also allude to reliable *faculties* (as they appear, say, in the theories of Sosa (1991) or Plantinga (1993)) and reliable *indicators* (as they appear, say, in Alston (1988)). But I would hope that the point made by this paper is fully general, and that I'm revealing new places where reliabilist ideas from any particular reliabilist tradition can figure into a reasons-based account of justification. Of course, that is a substantive claim and one try to argue that my view is only really capable of accommodating themes from *certain* reliabilist traditions, to the (misguided) exclusion of others. But in thinking things through so far, I've discovered no reason to think that the arguments in what follows are limited to any particular form of reliabilism.

analysis. Once separated, we see that reasons-first epistemology and reliabilism are not only consistent, but conflict no more than the buck-passing account of value conflicts with hedonism.

Before defending these claims, I want to say a bit more about why I am interested in RP and why it is a significant doctrine. Here I will be departing from the concerns that have often moved epistemologists to embrace RP, and following a new wave of writers in the meta-ethical literature on reasons and rationality. So, while RP is a traditional view, I propose to reframe it and defend it in an untraditional way—so much the better, I will maintain.

*The Significance of RP: A New Take.* As I see it, RP is best viewed as a contribution to a larger project that has become popular in recent literature on the nature of normativity—viz., the “Reasons First” project. The aim of this project is to simplify the internal structure of the normative, by reducing all normative facts to facts about reasons. The project got started in meta-ethics with the work of people like T. M. Scanlon and Derek Parfit.<sup>3</sup> It has made its way into epistemology *via* the efforts of Mark Schroeder.<sup>4</sup> The project in epistemology consists in trying to reduce all epistemically normative facts to facts about good epistemic reasons. This project is consistent, of course, with the idea that normativity can be naturalized. Indeed, it is an orthogonal project: one could think that reasons are the most fundamental units of normativity while being a naturalist or a nonnaturalist. And in meta-ethics at least, both sides are represented: Reasons-Firsters include hardcore naturalists like Schroeder and hardcore nonnaturalists like Parfit.

Although the project is neutral on that issue, it remains significant. Often we are interested in considering the internal structure of a domain and asking: “What are the basic elements of the domain? How can other elements be reduced to them?” Some major issues in the philosophy of mind have this form. Consider discussions about whether phenomenal states reduce to intentional states, or about whether intentionality is the mark of the mental; and consider narrower questions, like whether intentions

---

<sup>3</sup>See Scanlon (1998) and Parfit (2001, 2011).

<sup>4</sup>See Schroeder (2007, 2008, 2011).

reduce to belief-desire pairs. The question of whether reasons ground justification is interesting in the same way that these questions are interesting.

Because many epistemologists have been interested in how the epistemic is grounded in the non-epistemic, they might take RP to be a trivial doctrine. But this would be a mistake. It would be like thinking that the view that intentions reduce to belief-desire pairs is trivial because it does not help us to understand how the mental reduces to the non-mental. Questions of fundamentality within domains are real questions worthy of focus in their own right.

So, following the meta-ethicists, we should distinguish two questions:

(Q1) What are the most basic units of normativity and how are other normative facts grounded in facts about these units?

(Q2) How is the normative grounded in the non-normative?

We can clearly address the question “What is justified belief?” in two compatible ways in line with this distinction.<sup>5</sup> If we address that question as part of the larger project of answering Q1, RP makes an important contribution. Q1 attracts great interest among a wide array of theorists interested in the nature of normativity. Many of these theorists are not epistemologists. But epistemologists are missing out.

So, RP is a significant doctrine. Nevertheless, one might be left wondering why RP is supposed to be attractive. Why think that reasons are prior to justification rather than *vice versa*? Here is one major reason. Being justified, like being permitted, is in the first instance an *overall* or *all-things-considered* status. Reasons, on the other hand, are *contributory*.<sup>6</sup> They are *pro tanto* items with weights that can be balanced. If the contributory is generally more fundamental than the overall, reasons will be more fundamental than justification. Of course, this assumes that reasons modally co-travel with justification. But here I think common sense supports RP: it is a truism that good

---

<sup>5</sup>As I will note again toward the end of the paper, Goldman (1979) was explicitly aware of this distinction from the very first page. But the distinction has come to be neglected in the epistemological literature, and this has caused theorists to see rivalries where there needn't be any.

<sup>6</sup>The overall vs. contributory contrast and terminology comes from Dancy (2004).

reasons are necessary for justified belief. To be sure, the Resistance disagrees and I will defend the truism from their objections. But the truism has an intuitive upper hand.

Does the argument work, though? Well, epistemologists have introduced the term ‘*prima facie* justification’ to refer to something contributory, and this might seem to undermine the argument. Indeed, that term is now entrenched. But that term has always looked stipulative, like W. D. Ross’s term ‘*prima facie* duty’. At least Ross (1930: 20) was honest:

[T]he phrase ‘*prima facie* duty’ must be apologized for, since it suggests that what we are speaking of is a certain kind of duty, whereas it is in fact not a kind of duty, but something related in a certain way to duty. Strictly speaking, what we want is not a phrase in which ‘duty’ is qualified by an adjective, but a separate noun.

The separate noun that slipped Ross’s mind was ‘reason’. Similar thoughts extend to ‘*prima facie* justification’. The referent of this expression plausibly *just is* a normative reason or set of reasons. To agree that *prima facie* justification is more basic than justification proper is then to agree that reasons are more basic than justification proper.

*The Neutrality of RP.* Why bother explaining why I think RP is significant? Partly because I want to distance myself from common ways of understanding its significance. RP is often associated with other views in epistemology, but I think that RP can and should be separated from these views. Indeed, once we see that RP is neutral on many important issues in epistemology, one might naturally wonder why we should care about it. Hence the foregoing context.

On what familiar epistemological issues is RP neutral? First of all, RP is consistent with foundationalism. Foundationalists only claim that there is a set of non-inferentially justified beliefs from which all of our other beliefs derive their justification. We should not assume that all good epistemic reasons are inferential justifiers. Some reasons might be immediate justifiers (e.g., experience-given reasons). A foundationalist can embrace RP if she is willing to claim that all non-inferentially justified beliefs are justified by reasons of this immediate sort. Certain examples of non-inferentially justified beliefs



raise worries for RP, as we will see. But the minimal foundationalist thesis is consistent with RP.

Secondly, RP is not an essentially internalist thesis for four reasons:

- i. One can understand the *ontology* of good epistemic reasons in externalist ways. One can, for instance, view reasons as facts and allow these to include extra-mental facts.<sup>7</sup>
- ii. One can understand the *quality* of epistemic reasons in externalist ways. One can, for instance, hold that what makes a consideration R a good epistemic reason to believe P is that R-type considerations are reliable indicators of the truth of P-type propositions.<sup>8</sup>
- iii. One can understand the relation of *possession* that we must bear to reasons to be poised to receive justification from them in externalist ways. One might, for instance, hold that factive mental states like *seeing that R* are ways to possess R as a reason.<sup>9</sup>
- iv. One can understand what it takes to believe that P *because one has good reasons* in externalist ways. One might hold that one can only believe that P *because one has good reasons* to believe P if one manifests certain reliable cognitive abilities.<sup>10</sup>

Finally, RP is not an essentially evidentialist thesis for two reasons. On the one hand, RP speaks only of good epistemic reasons. It is an open question whether all good epistemic reasons are evidence. While the clearest examples are pieces of evidence, there might be other sorts of good epistemic reasons (indeed, I suspect this is true).<sup>11</sup>

---

<sup>7</sup>Cf. Dancy (2000).

<sup>8</sup>Cf. Alston (1988).

<sup>9</sup>Cf. McDowell (1995) and Pritchard (2012).

<sup>10</sup>Reliabilists (e.g., Goldman (2012)) have proposed accounts of *proper basing* in this spirit. But the claim being made here is a stronger one, concerning what I take to be the ordinary notion of believing something because one has a good reason rather than the (arguably) technical notion of proper basing.

<sup>11</sup>See Schroeder (2012) for arguments.

On the other hand, evidentialism is often associated with mentalism about evidence and internalism about justification. Conee and Feldman are clear that mentalism is a further thesis and that the core evidentialist thesis doesn't entail internalism. But many assume otherwise. This seems to me as mistaken as assuming that consequentialism is essentially hedonistic just because the fountainhead consequentialists held a hedonistic version. But I don't want to quibble: what matters is that RP abstracts away from Conee and Feldman's views about the ontology and quality of good reasons just like the minimal consequentialist thesis abstracts away from Bentham's views about good consequences.

*The Plan.* With these clarifications out of the way, here is the plan. Because I am defending RP with the help of insights from the meta-ethical literature, I will begin in §2 by introducing some common distinctions and points from that literature that will matter throughout the paper. In §3, I turn to address some arguments against the *necessity* of reasons for justification (propositional or doxastic). In §4, I address some arguments against the thought that believing something because one has good reasons is *sufficient* for doxastic justification. I address the reply that reasons are *epiphenomenal* in the analysis of justification in §5. I draw things to a close in §6, explaining how my positive view differs from recent syntheses of evidentialism and reliabilism, and answering a few remaining objections to RP.

## 1.2 Reasons: Some Ontology and Ideology

### 1.2.1 Normative Reasons and the Relation of Possession

In the broader literature on reasons and rationality, it is widely appreciated that the concept of a reason can be used in importantly different ways.<sup>12</sup> The most crucial use for the purpose of assessing RP is the normative use. Normative reasons are considerations

---

<sup>12</sup>I do not mean this to be incompatible with Dancy (2000: 2)'s claim that there are not really different *kinds* of reasons but rather different questions that the concept of a reason can be used to answer. I agree with Dancy that in good cases, motivating and normative reasons can coincide, which I take to be the primary reason why he opposes the thought that there are different kinds of reasons.

that count in favor of actions and attitudes. In both ethics and epistemology, normative reasons divide into two categories. Some are not essentially possessed by anyone. Call these *objective normative reasons*. Undiscovered evidence E for some hypothesis H is an example of an objective epistemic reason. Objective normative reasons *alone* justify no attitudes or acts—at least not in the way that interests epistemologists. Only normative reasons that we *possess* can justify in that way. Call these reasons *possessed normative reasons*.

Some view possessed normative reasons as a species of objective reasons: namely, the ones to which we bear the relation of possession. Others reject this view and see the reasons that bear on justification as belonging to an irreducibly distinct category.<sup>13</sup> I favor the first view, and I would distinguish between possessed objective reasons and merely apparent normative reasons. But nothing in this paper will hang on this.

The nature of possession matters greatly for the status of RP. Indeed, I will try to show that some of the main concerns of the Resistance can be resolved by gaining a better understanding of this relation. Of course, there is pressure in the other direction. One might suspect that the relation of possession must be understood in epistemically normative terms, and on this basis reject RP. After all, one might find it implausible that merely believing that P can enable one to possess P as a good reason to believe other things.<sup>14</sup> If so, one might worry that possession must be understood in terms of justified belief. If this worry were apt, RP would be circular.<sup>15</sup>

But circularity is not inevitable. One could appeal to normative concepts more primitive than the concept of justification—e.g., ones drawn from virtue epistemology, such as *competence*. Although this view suggests that there is a more basic kind of epistemic normativity that enables reasons to do their justifying work, this conclusion alone does not harm RP. Defenders of RP can ground justification in possessed reasons and then ground possessed reasons in facts about competence.

---

<sup>13</sup>See Schroeder (2008) for the second view. See Lord (2010) for a defense of the first view.

<sup>14</sup>Schroeder (2011) defends the low bar view and argues that its seemingly implausible consequences can be explained away by a defeat strategy. I will mention some doubts about this in the next section.

<sup>15</sup>For a circularity worry in this spirit, see Goldman (2011).

Another possibility is to appeal to non-normative factors to ground the possession relation. We can do this without lowering the bar on possession to mere belief. For example, we might directly appeal to broadly reliabilist ideas to explain the conditions under which certain mental states enable us to possess normative reasons. I do this in §3, and this is one crucial move that I will be making to address the Resistance. There are two important conditions on possession that I analyze reliabilistically: *access* and *sensitivity to the favoring relation* between the reason and the relevant doxastic attitude.<sup>16</sup> Access gets analyzed in terms of *apt seemings* and sensitivity in terms of *apt treatings* of considerations like objective reasons with certain favoring profiles. Aptness in both cases is defined in terms of manifestations of certain cognitive abilities.

We will get clearer on the details in §3. For now, note that we are not forced to analyze possession in normative terms even if we dislike low bars on possession. The account I develop is not the only option. Another option would be to pursue a McDowellian view on which factive but non-doxastic mental states like seeing that P, remembering that P, intuiting that P, etc., are the pathways to possession.<sup>17</sup> One could also adopt a Williamsonian view on which knowledge is a mental state and possession gets analyzed in terms of knowledge. Ambitious knowledge-firsters will certainly claim to avoid circularity.

### 1.2.2 The Ontology of Normative Reasons

Another important issue for the status of RP is the ontology of normative reasons. In contemporary meta-ethics, it is common to deny that normative reasons are mental states. The preferred view is that objective normative reasons are facts and possessed normative reasons are facts or apparent facts to which we stand in partly mental relations. I will call this view *factualism*, but for simplicity count as factualist a view on which possessed normative reasons can be merely apparent facts (like Schroeder (2008)'s

---

<sup>16</sup>By 'sensitivity', I don't mean anything that can be captured by a Nozickian counterfactual. As we will see in §3, I mean something *dispositional*. Dispositions are not analyzable in terms of counterfactuals.

<sup>17</sup>See McDowell (2002) for arguments that seeing that P is non-doxastic and can precede knowledge. Bernecker (2007) makes the same points about remembering that P.

view).

The facts that constitute normative reasons are often extra-mental, though facts about one's mental life can be normative reasons. The reasons that support introspective beliefs, for example, are mental facts. And although the facts we perceive can constitute reasons for perceptual beliefs, they are not the *only* reasons for perceptual beliefs. In a happy world, perceptual appearances are reliable indicators of reality. In such a world, the fact that things perceptually appear some way is an objective reason to believe things are that way.

I favor factualism. It is, I believe, crucial for defending RP. Epistemologists often spurn factualism—but for unconvincing reasons. Let's correct some misconceptions.

Factualism about normative reasons does not make justification *factive*.<sup>18</sup> This is so even if one requires possessed normative reasons to be real rather than possibly merely apparent facts. One might make a competent inductive inference from nothing but facts and arrive at falsehood. A strong factualist view only requires the ultimate *inputs* to justified inferences to be facts, not the *outputs*. Since the view does not entail that mental facts cannot be good reasons, it also allows for false justified perceptual beliefs.

Factualism by itself has no bad implications about demon worlds even if it takes the strong form. It might have bad implications if one adds further commitments—e.g., if one holds that F is a good reason to believe P only if F-type facts are reliable indicators of P-type facts or if one thinks that *only* extra-mental facts can be normative reasons. But charitably defined, factualism entails neither claim. Indeed, no one has ever believed the second claim, which makes obvious nonsense of introspective justification. Anyway, remember that some proponents of the ontology at issue allow possessed normative reasons to include merely apparent facts. These people face no *prima facie* worries here.

Another objection to the factualist ontology is that normative reasons that are merely “out there” are irrelevant to justification. The response is (i) to agree that normative reasons must be possessed to justify and (ii) to agree that mental states *help*

---

<sup>18</sup>Justification for believing P is *factive* in the sense I have in mind here when it entails that P.

*us to possess* (or “*provide*”) these reasons, but (iii) to deny that mental states *are* the reasons. Williamson (2000: 197) put it well: “Experiences provide evidence; they do not consist of propositions. So much is obvious. But to provide something is not to consist of it.”

We can view mental states as things that help us to possess and reason with normative reasons, while viewing the normative reasons that are possessed as facts or apparent facts. This picture secures the advantages of factualism while allowing us to honor the role of mental states in providing reasons and in good reasoning.

### 1.2.3 Motivating Reasons

The concept of a reason can be used to pick out some phenomena that are not essentially normative. *Explanatory reasons* provide one example. An explanatory reason for someone’s thinking or acting in some way is a factor that explains why this person thinks or acts in this way. Not all the factors that explain why we act or think are the reasons *for which* we think or act: the fact that someone is depressed might explain why this person tends to draw certain pessimistic conclusions, but it would be unusual for someone to have as *their* reason for drawing these pessimistic conclusions the fact that they are depressed.

This brings out a second use of the concept of a reason that is not essentially normative. This second use will matter more for our discussion of RP. *Motivating reasons* are the reasons that help to explain why people think or act by being the reasons *for which* they think or act. Motivating reasons are not essentially good reasons: we often act and think for bad ones. In using the term ‘motivating’ in the epistemic domain, I don’t assume doxastic voluntarism. A motivating epistemic reason is just a reason for which someone thinks. Still, to save words and avoid misunderstanding, I will often use the term ‘rationale’.

Many meta-ethicists view rationales as apparent facts or propositions rather than mental states.<sup>19</sup> This view receives support from linguistic phenomenology. We often

---

<sup>19</sup>See Dancy (2000) for the apparent facts view; see Lord (2008) for a defense of the propositional view.

say things like: “Her rationale for being an atheist is that there is unnecessary suffering” or “His rationale for going to the party is that there is dancing.” The view also receives support from the phenomenology of motivation: we rarely navel-gaze at our mental states when reasoning. Another source of support is the need to understand how we could act for good reasons. In ethics, good reasons routinely include facts about the world. If we can act for these good reasons, it arguably must be possible for them to *be* our motivating reasons.<sup>20</sup>

I favor this non-statist view, though I am neutral about which version is correct. Sometimes I will talk *as if* the propositional version is correct but only for linguistic ease.

Resistance to this view tends to rest on misunderstandings or uncharitable interpretations. Clearly, mental states matter for motivation. But the non-statist view does not say otherwise. By being in certain mental states, certain considerations become our rationales. But our rationales are not themselves mental states: mental states enable different considerations to be our rationales. This is not to deny that our rationales can be considerations about our mental states. That is possible, especially in epistemology.

Sometimes it might look circular to think of one’s rationale for belief as a proposition or apparent fact. But this rests on a misunderstanding. Notice that it is not circular to have as one’s reason for believing that one is in pain the sheer fact that one is in pain. That fact is a great reason to believe that one is in pain—none better! Circularity is a property of *inferences*, and not all rationales are *inferential*. Sometimes rationales are direct ones on which our beliefs are based non-inferentially. When we are directly acquainted with certain facts, for example, they can serve as direct rationales. This happens in the pain case. There is no circularity there, since there is no inference.<sup>21</sup>

---

<sup>20</sup>See Dancy (2000) for this argument. Lord (2008) and Turri (2009) offer some responses.

<sup>21</sup>This answers one of Turri (2009)’s objections to non-statism. Turri also raises a problem from reasons for withholding. But Turri does not notice some options. He considers a case where Nevil withholds on whether there is unnecessary suffering in the world, believes that God exists only if there is not, and so withholds on whether God exists. What is Nevil’s reason for the second instance of withholding on the non-statist view? Turri thinks the non-statist must say it includes <that there is unnecessary suffering in the world>. This makes Nevil look irrational. But non-statists are not forced to say this. The obvious thing to say is that Nevil’s ultimate rationale for withholding is <that the evidence he possesses for thinking that there is unnecessary suffering in the world is insufficient>. This is just another proposition or apparent fact. If Nevil is rational, this fact is surely the one that

Rationales matter for doxastic justification. To get doxastically justified, it is not enough that one possesses good reasons: one must also believe what one believes *because* one has these good reasons. For that to happen, rationales must enter the picture. How so?

Virtually everyone assumes that the following is the whole story:

*(Coincident Rationale)* S believes that P because S has a good epistemic reason to believe that P iff (i) S has a rationale *r* for believing that P and (ii) *r* happens to correspond to a good epistemic reason that there is for S to believe that P.

It is not the whole story. Believing something because you have a good reason is an *achievement*. If Coincident Rationale were the whole story, believing something because you have a good reason would be no achievement. You could pull it off by incompetent luck, by accidentally landing on a rationale that happens to coincide with a good reason. But it isn't so easy. To believe something *because you have a good reason*, the fact that your reason *is a good reason* must help to explain why you believe as you do. For that fact to play that explanatory role, it must trigger certain reliable cognitive abilities of yours.<sup>22</sup>

As we will see, the Resistance crucially assumes Coincident Rationale and uses examples of incompetent basing against RP's account of doxastic justification. But their examples of incompetent basing undermine Coincident Rationale, not RP. I defend this verdict in §4 and a similar verdict about possession in §3.

---

motivates Nevil's first instance of withholding! There is no reason why a non-statist cannot appeal to this fact unless the view is defined uncharitably.

<sup>22</sup>Again, reliabilists such as Goldman (2012) have recommended similar accounts of what they call *proper basing*. But the claim being made here is a stronger claim that, if true, would show that RP does not need to be modified at all. If one could only appeal to proper basing, RP could not be vindicated in its pure form. But I'm claiming that it can be vindicated in that form, so long as the defender of RP has the correct account of believing something because one has a good reason. Or, equivalently, I'm claiming that there is no difference between believing something because you have a good reason and proper basing—which I take to be an interesting claim that reveals a new avenue for friends of RP.



### 1.3 Are Reasons Necessary for Justified Belief?

Let's turn to consider the arguments of the Resistance, beginning with some arguments against the necessity of reasons for epistemic justification (propositional and doxastic).

The arguments have a common aim. They try to show that certain examples of non-inferentially justified belief cannot be understood as RP suggests. In reply, I will make the following points. Rather than showing that we can have justified beliefs without good reasons, these examples reinforce (i) the virtues of a factualist ontology of reasons and (ii) the need for a better view about the relation of possession, which I will provide.

Following Sosa (2007), I will introduce this worry with basic *a priori* beliefs. After sketching my reply, I will offer a theory about the relation of possession that gets the right verdict in non-inferential cases, and generalize the view to inferential cases. Having done this, I'll show how the same points address Lyons (2009a)'s version of the worry.

#### 1.3.1 Examples against the Necessity of Reasons: Obvious Math

Consider simple arithmetic beliefs, like your belief that the proposition  $\langle 1 + 1 = 2 \rangle$  is true.<sup>23</sup> Plausibly, your belief is non-inferentially justified.

What reason justifies your belief? Some believe that *intuition* can provide reasons for such *a priori* beliefs. But what reason does intuition *provide* in our case, given the distinction between providing and being a reason? Just the introspective fact that one has the intuition? If that fact is the candidate, members of the Resistance will raise two objections.

To see the first objection, note that members of the Resistance are often reductionists about intuitions and other seemings, taking them to consist in epistemically evaluable states. Ernest Sosa, for example, argues that intuitions are just a species of *attractions to assent*.<sup>24</sup> The mere fact that one is attracted assent to  $\langle P \rangle$  is not a serious reason to believe  $\langle P \rangle$ . People with no mathematical competence might find "obvious" many

---

<sup>23</sup>Only because the difference will now prove crucial, I will use expressions like ' $\langle P \rangle$ ' for propositions and expressions like 'the (apparent) fact that P' for their (apparent) truth-makers.

<sup>24</sup>See Sosa (2007: Ch.3).

abstruse mathematical truths that they aren't justified in believing. They may also be self-deceived or introspectively incompetent, unable to notice the first-order incompetence that explains why they are unjustified. By a sheer fluke, they may land on only truths. If so, they will *possess* no defeaters. By hypothesis, they lack the second-order competence to see their first-order incompetence. But they are also lucky, landing by chance on truths.<sup>25</sup>

Might one appeal to unpossessed objective defeaters to explain away the intuitions? Not plausibly. There can be unpossessed objective defeaters for perfectly justified beliefs (e.g., in many Gettier cases), but our cases are not remotely like these cases. It is plausible in our cases that doxastic justification is lacking, and implausible that it is lacking in standard cases involving unpossessed defeating evidence.

There is a second objection: viewing the relevant reasons as introspective facts about our intuitions is an overintellectualization. If the apparent fact that P is one's rationale for believing that Q, one plausibly needs the concepts necessary to grasp the proposition  $\langle P \rangle$ . Yet young children can justifiably believe that  $1 + 1 = 2$  without having the concept of an intuition or being in a position to engage in any second-order reflection.

It is true that children may believe that  $1 + 1 = 2$  *because* it strikes them as intuitively obvious. But that 'because' does not signal that this introspective fact is their rationale. By comparison, if one infers  $\langle P \rangle$  from  $\langle Q \rangle$ , one will believe  $\langle P \rangle$  because one believes  $\langle Q \rangle$ . But one's rationale for believing  $\langle P \rangle$  is  $\langle Q \rangle$ , not  $\langle I \text{ believe that } Q \rangle$ . Belief in  $\langle Q \rangle$  simply enables  $\langle Q \rangle$  to serve as one's rationale for believing  $\langle P \rangle$ . Similar points apply to the intuition that P. This is clear with children, who lack the ability to exploit the proposition that they intuit that P as a rationale.

So, the Resistance may doubt on other grounds that the justifying reason is *that one has the intuition*. Yet, the Resistance will ask, if this isn't the reason, what is it? Is it just the sheer fact that  $1 + 1 = 2$  which makes true the proposition  $\langle 1 + 1 = 2 \rangle$ ?

---

<sup>25</sup>This is a problem for Schroeder (2011)'s defense of a low bar on possessing reasons. Schroeder is right that our intuitions about whether someone "has no reason" are unreliable. He is also right that an appeal to defeat can often explain away these intuitions. But the strategy is harder to defend in extended cases involving both first-order incompetence and higher-order competence.

Surely, they will insist, that isn't right. One can be attracted to treat any fact as given if one is sufficiently incompetent. And wouldn't it be circular to have as one's reason for believing  $\langle 1 + 1 = 2 \rangle$  the sheer fact that  $1 + 1 = 2$ ? Discarding the idea, the Resistance spurns RP.

### 1.3.2 The Slip

This conclusion is hasty. First, recall that we can sometimes be justified in believing  $\langle P \rangle$  on the basis of the clear fact that makes  $\langle P \rangle$  true. One can be justified in believing that one is in pain on the basis of the clear fact that one is in pain. There is no circularity here. It would be circular if one tried to *infer*  $\langle I \text{ am in pain} \rangle$  from  $\langle I \text{ am in pain} \rangle$ . But our case is non-inferential. The proper response to the relevant reason is not an inferential response. Since circularity is a property of inferences, it is wrong to dismiss as "circular" this case of believing that P on the basis of the fact that makes P true.

The proper conclusion to draw is that we can sometimes possess the sheer fact that P as a direct license to believe the proposition that this sheer fact makes true.<sup>26</sup> So, why not apply this model in the basic intuitive case, and say that in virtue of *reliable* intuition (or a reliable intuitive process), one can *possess* the sheer fact that  $1 + 1 = 2$  as a reason to believe what this fact makes true?<sup>27</sup>

I cannot see why not. Sure, not just any old intuition can put us into contact with facts so as to enable these facts to serve as direct licenses to believe what they make true. One can incompetently find anything intuitive. But this reveals *nothing special about the case of intuition*. Similar points apply even to introspection. Consider:

(FRED) Fred contemplates his visual field and the number of floaters in it. He has weak subitizing ability and cannot reliably discriminate more than 4 floaters without long focus. Yet he is self-deceived, fancying himself like Rain Man. Faced

---

<sup>26</sup>Parallel remarks hold for apparent facts, which are apparent truth-makers.

<sup>27</sup>Again, I'm hoping the core point here to be neutral between the process reliabilist themes of Goldman (1979), the faculty reliabilist themes of Sosa (1991), the virtue reliabilist themes of Greco (1999) and Sosa (2007), and other reliabilist themes. I don't see a strong reason to think that the point couldn't equally well be hosted by any of these traditions.

now with 16 floaters in his visual field, he finds himself attracted to assent to the correct answer by incompetent luck, and then judges that there are 16.

Fred's belief is unjustified—both propositionally and doxastically. Why? Because he lacks certain discriminatory abilities.<sup>28</sup> So, not even all conscious mental facts can serve as non-inferential licenses without the help of ability. Yet the “no reasons needed” conclusion is implausible here. Imagine Sally, a more discriminating subject who glimpses the fact that there are 16 floaters in her visual field and then judges that there are 16. That fact is clearly relevant to why her judgment is justified. The key is just that Sally, unlike Fred, possesses that fact as a non-inferential reason *in virtue of* her greater abilities.

We can say the same things in the intuitive case. One can possess intuitively obvious facts as direct licenses to believe what they make true *when* the underlying intuitions manifest cognitive ability. Reasons are hardly unnecessary. In happy introspective and intuitive cases, we have the best reason we could possibly have *by* exercising cognitive ability!

### 1.3.3 Possession: An Achievement

What we learn is that possessing normative reasons is itself an achievement. It is the achievement of getting reasons in one's grip for justifying use in virtue of certain cognitive abilities. This recommends a bi-level picture that grounds non-inferential justification in possessed non-inferential reasons and the possession of such reasons in manifestations of cognitive ability by the attractions to assent that constitute intuitive and other seemings.<sup>29</sup> I will make some revisions, but we can start with:

---

<sup>28</sup>Goldman (1976), of course, invoked such abilities in his classic relevant alternatives account of perceptual knowledge. Here I'm suggesting that the usefulness of the notion extends to one's account of justification *by way of* one's account of reason-possession. What's original here is not, of course, the appeal to discriminatory abilities period in an account of justification (which one can trace back to Goldman, and find in current forms of virtue epistemology such as Greco (2009)'s view), but rather the thought that these abilities should play an essential role *in our account of possession*. The original claim is that there is already a natural place within a reasons-based framework—namely, the crucial notion of *possession*—where we should invoke these reliabilist themes and *thereby* bring them into our account of justification.

<sup>29</sup>I will assume with Sosa (2007) that *attractions to assent* are the things to which seemings reduce. But all that matters is that we have some reductionist picture of seemings that entails that seemings

(L1) S is non-inferentially propositionally justified in believing  $\langle P \rangle$  iff S possesses some (apparent) fact F as a non-inferential reason to believe  $\langle P \rangle$  and no defeating reasons, and S is non-inferentially doxastically justified in believing  $\langle P \rangle$  iff S believes  $\langle P \rangle$  because S has that good non-inferential reason and no defeating reasons.

(L2) S possesses F as a good non-inferential reason to believe  $\langle P \rangle$  iff

- (i) F is a good epistemic reason to believe  $\langle P \rangle$ ,
- (ii) F makes it non-doxastically seem to S that  $\langle P \rangle$ , and by no inferential route,
- (iii) the attraction to assent that constitutes the seeming in (ii) manifests a disposition of S's to be attracted to assent non-inferentially to  $\langle P \rangle$ -type propositions upon considering them only if they are true.

I'll call dispositions like the one in (iii) 'competences', but this is a useful term of art for something non-normative—though it is something epistemically *good*, like true belief.

To avoid misunderstandings, mark two facts. First, competences can be fallible. A competence to  $\phi$  does not give one the ability to infallibly  $\phi$ . This doesn't mean that the competence is a competence to do something weaker than  $\phi$ -ing. It just means that it is a less than infallible competence to  $\phi$ . So, if one has the disposition in (iii), it does not follow that one will infallibly be attracted to assent only to truths.

Second, there is a *competence/performance* distinction. Competences are dispositions. Like other dispositions, there is a set of favorable conditions for their manifestation. Unfavorable conditions do not remove competence. An archer in a hurricane isn't deprived of her archery skill. Competences can even be displayed—though not strictly manifested—in abnormal conditions, like cars in a display room. So, while competences are reliable in favorable conditions, they are not reliable *period*.<sup>30</sup>

---

are epistemically evaluable states. So, the reader should not think I place too much weight on the "attractions" ideology.

<sup>30</sup>See Sosa (2010) for discussions of the metaphysics of competence, and Sosa (1992, 1993) for classic discussions of how the second point addresses the new evil demon problem.

The bi-level picture of non-inferential justification is general, applying to perceptual and introspective cases as well as intuitive cases. It leaves open how exactly we are to understand cases of *false* non-inferentially justified belief. Here one could view F as a real fact but not the fact that makes  $\langle P \rangle$  true. If so, one ends up with a Williamsonian model on which the relevant reasons differ in good and bad cases. But one could instead view F as a merely apparent fact and hold the weaker version of factualism.

### 1.3.4 The Achievement in More Detail: Two Kinds of Competence

We do not yet have a general account of possessing a good reason R to believe  $\langle P \rangle$ . We want an account that applies to inferential cases. To extend the account to these cases, we must distinguish two achievements: the achievement of *access* to reasons and the achievement of *sensitivity to the favoring relations* between reasons and beliefs. Both are generally necessary, but they are easier to separate in inferential cases.

Consider an inferential case to fix ideas. Suppose  $\langle P \rangle$  follows deductively from  $\langle Q \rangle$  and  $\langle R \rangle$ . The existence of this deductive relation ensures that the truth of  $\langle Q \rangle$  and  $\langle R \rangle$  is an objective reason to believe  $\langle P \rangle$ . Clearly, the mere existence of this objective reason doesn't justify anyone in believing  $\langle P \rangle$ . To be justified in believing  $\langle P \rangle$ , one must also possess this objective reason to believe  $\langle P \rangle$ . So, what does this involve?

#### Access as Apt Seeming

Plausibly, part of it is having access to the apparent facts that Q and that R. What is access? We can give a direct account of access in terms of seemings and cognitive ability. Seemings and the attractions to assent that constitute them can display a familiar trio of properties. We can say that a subject's attraction to assent to  $\langle P \rangle$  is

*right*, if  $\langle P \rangle$  is true,

*competent*, if the attraction manifests a disposition of the subject's to be attracted to assent to propositions of some kind only if they are true,<sup>31</sup> and

---

<sup>31</sup>My earlier notes about fallibility and the competence/performance distinction extend here.

*apt*, if the attraction is right because competent.

With this pattern in mind, we can understand *accessing a fact* as hosting an apt seeming that the fact obtains—i.e., an apt attraction to assent to the proposition made true by the fact. Of course, to require a subject in possession of R to be *currently accessing* R is too demanding. We should only demand that the subject be *in a position* to access R. What does ‘in a position’ mean? This is vague, but so is the notion of possessing a reason and the notion of propositional justification. So this is no objection.

Some might object that the account is still too strong. Can we not possess merely apparent facts as good epistemic reasons for belief? My own inclination is to say that we just possess different facts in these cases. But if one finds this intolerable, it is easy to weaken the account. Just replace ‘apt’ with ‘competent’, and ‘fact’ with ‘apparent fact’.

Some might insist that we are conceding the failure of the Reasons First program. Isn’t the ideology from Sosa *normative*? If so, aren’t we conceding that there are more basic normative facts than facts about reasons? No. Competence and aptness can be stipulatively defined in non-normative terms, and so they are not really normative notions.

We can now understand part of what is required for one to possess the relevant objective reason in our inferential example. One must be poised to host an apt seeming that Q and R if one considers  $\langle Q \rangle$  and  $\langle R \rangle$ —or a competent seeming, if one prefers.

### **Sensitivity as Apt Treating**

This is not *all* a subject needs to possess  $\langle Q \rangle$  and  $\langle R \rangle$  as good reasons to believe  $\langle P \rangle$ . Suppose that the proof of  $\langle P \rangle$  from  $\langle Q \rangle$  and  $\langle R \rangle$  is arcane. Unless our subject has godlike acumen, she will not have propositional justification to believe  $\langle P \rangle$  simply in virtue of being aptly attracted to assent to  $\langle Q \rangle$  and  $\langle R \rangle$ . So, *mere access* to facts that constitute objective reasons for belief isn’t enough for possessing these reasons for belief.

What else does possession require? A subject also must be *sensitive to the favoring*

*relation* between the facts that Q and that R and the belief that P. What does such sensitivity involve? Some might require that the subject be in a position to see that  $\langle Q \rangle$  and  $\langle R \rangle$  constitute a deductive reason to believe  $\langle P \rangle$ . But this is an overintellectualization. To possess reasons, subjects don't need the concept of a reason. Even subjects who have this concept do not need to deploy it to possess reasons. Some subjects are competent enough to cleave reliably to good deductive patterns without representing these patterns to themselves.

This observation recommends a better approach to understanding sensitivity to favoring relations. It is possible *to treat* a consideration like a normative reason of some kind without having or deploying the concept of a normative reason. To treat a consideration like a normative reason of some kind is to be disposed to think in all or at least most of the ways that would be appropriate if this consideration were a normative reason of that kind.<sup>32</sup> Like attractions to assent, treatings can display a familiar triadic pattern of properties. A subject's treating a consideration R like a good reason of kind K to believe  $\langle P \rangle$  is

*right*, if R is a normative reason of kind K to believe  $\langle P \rangle$ ,

*competent*, if the treating manifests a disposition of the subject's to treat R-type considerations like normative reasons of kind K to believe P-type propositions only if they are normative reasons of kind K to believe P-type propositions, and

*apt*, if the treating is right *because* competent.

Another thing that our subject needs to possess the deductive reason constituted by  $\langle Q \rangle$  &  $\langle R \rangle$  to believe  $\langle P \rangle$  is to be disposed to *aptly treat*  $\langle Q \rangle$  &  $\langle R \rangle$  as a deductive reason to believe  $\langle P \rangle$ . This is an example of the kind of sensitivity that possession requires.

---

<sup>32</sup>Why 'all/most' rather than 'some'? To address pretense and supposition. If you suppose or pretend for the sake of argument that P is a conclusive objective reason to  $\phi$ , you are not irrational if you don't  $\phi$  given the appearance that P. But if you competently treat P like a conclusive objective reason to  $\phi$  in the sense relevant to rationality, you are irrational if you fail to  $\phi$  given the appearance that P. The difference between treating and supposition or pretense seems clear. If you merely pretend or suppose that P is a conclusive objective reason to  $\phi$ , there will be *many ways* that you are not disposed to respond to the appearance that P which would be correct P if P were a conclusive objective reason to  $\phi$ .



So, we can understand possessing a good reason  $R$  to believe  $\langle P \rangle$  as involving both access to  $R$  and sensitivity to the favoring relation between  $R$  and  $\langle P \rangle$ .

Why require an apt treating disposition rather than a merely competent one? My reason for preferring aptness is to explain cases like:

(MISFIRE) Bill is a competent but not infallible logician who can immediately grasp many sophisticated entailments. It strikes him tonight that there is a proof of  $\langle Q \rangle$  from  $\langle P \rangle$  involving certain steps, which he seems to see by exercising his competence. Alas, his fallible competence misfires and yields a mistaken impression: there is not *that* kind of proof of  $\langle Q \rangle$  from  $\langle P \rangle$ . There *is* a proof of  $\langle Q \rangle$  from  $\langle P \rangle$ , but it is far more difficult and beyond Bill's direct acumen.

$\langle P \rangle$  is a perfect deductive reason to believe  $\langle Q \rangle$ , since there is a proof of  $\langle Q \rangle$  from  $\langle P \rangle$ . Bill is competently attracted to treat  $\langle P \rangle$  like a perfect deductive reason to believe  $\langle Q \rangle$ , since his logical "vision" manifests his competence. This vision just misfires, as fallible vision can. So Bill does not possess the reason to believe  $\langle Q \rangle$  that  $\langle P \rangle$  *actually constitutes*, since he does not grasp the proof.<sup>33</sup> The apt treating requirement explains why.

### Non-Inferential Cases Again

It is clear enough now why a general theory of possessing  $R$  as a good reason to believe that  $P$  requires more than access to  $R$ . Access and sensitivity can come apart: one can have access to a fact that is an objective inferential reason to believe  $\langle P \rangle$  while having no grip on the inferential relation between this fact and  $\langle P \rangle$ .

One might doubt whether access and sensitivity can come apart in non-inferential cases. Consider arithmetic again. We saw that a person can possess the sheer fact that makes  $\langle 1 + 1 = 2 \rangle$  true as a non-inferential reason to believe  $\langle 1 + 1 = 2 \rangle$ . Here the ground of the favoring relation is the truth-making relation. It is hard to see how a subject could have access to the fact that  $1 + 1 = 2$  without being aware of this relation!

---

<sup>33</sup>This isn't to say that Bill has no reason to believe  $\langle Q \rangle$ . Just not the one actually constituted by  $\langle P \rangle$ .

Not all non-inferential cases are like this, however. In our nice world, perceptual experience reliably indicates reality. So, the fact that one has a perceptual experience can constitute an objective reason for beliefs about what that experience indicates. Yet this kind of perceptual case differs from the case of believing  $\langle 1 + 1 = 2 \rangle$  on the basis of the sheer fact that  $1 + 1 = 2$ . Access and sensitivity can come apart in this kind of case.

Speckled hens afford an illustration. Upon seeing a speckled hen, one is not automatically sensitive to the relation of support between one's having the experience and all the propositions represented by the experience. And so one doesn't *possess* all the objective reasons that exist in virtue of the experience. Only equivocation could make one think otherwise. One might say: "Look, you *have* the experience, and the fact that you're having the experience *is* an objective reason to believe those precise propositions. So, you *have* that reason." 'Have' means crucially different things in those sentences.

### 1.3.5 The Extended Account

We are now in a position to revise the account sketched in §3.3 to get a general theory of what it takes to possess a normative reason R to believe that P. The account is:

(L1+) S is propositionally justified in believing  $\langle P \rangle$  iff S possesses some fact F as a good epistemic reason to believe  $\langle P \rangle$  and possesses no defeating reasons, and

S is doxastically justified iff S also believes  $\langle P \rangle$  *because* S has that good reason.

(L2+) S possesses some fact F as a good reason to believe  $\langle P \rangle$  iff:

(*Access*) S is in a position to host an apt seeming that F obtains,

(*Sensitivity*) S is disposed to aptly treat F like an objective reason to believe  $\langle P \rangle$ .


One could weaken the account by replacing 'aptly' with 'competently', and 'fact' with 'apparent fact'. But remember: this is not necessary to avoid making justification factive. This is obvious in inferential cases where the inference rule is fallible: one can easily move from facts to non-facts. And in cases of non-veridical experience, one

could adopt the Williamsonian view that one does not possess the same reasons as in a case of veridical experience. Alternatively, one could distinguish between possessing objective reasons and possessing merely apparent reasons, and divorce justification and rationality.

### 1.3.6 Further Applications: Lyons on Perceptual Justification

Our account suggests an answer to another member of the Resistance: Jack Lyons (2009a). Lyons argues that reasons are unnecessary for justified perceptual beliefs. But like Sosa, what Lyons really shows is that we should reject a statist ontology of reasons and understand possession as an achievement grounded in cognitive ability.

Lyons begins by noting that there are two phenomena that talk of “experiences” can pick out: *sensations* and *percepts*. Sensations are “what a Lockean *tabula rasa* would experience”, while percepts are “the result of the mind’s unconscious and involuntary attempt to make sense of the world.”<sup>34</sup> Necker’s cube brings out the supposed distinction:



In one sense, the cube can appear as if viewed from above and to the right or appear as if viewed from below and to the left. Yet it is tempting to think that there is another level of appearance that remains constant. For Lyons, sensation is this constant level of appearance. One sensation is, he says, compatible with many distinct percepts.

Given this distinction, Lyons constructs a Sellarsian dilemma. Assuming that the defender of RP must identify reasons for perceptual beliefs with experiences (“experientialism”), she has two options: she can identify them with sensations or percepts. Sensations, Lyons insists, lack the right kind of content to be suitable for grounding all

---

<sup>34</sup>Lyons (2009a: 42).

justified perceptual beliefs. This content is too low-level, and it is not clearly propositional. Percepts, on the other hand, do have the right kind of content. But percepts cannot, Lyons believes, provide good epistemic reasons for perceptual beliefs. For percepts, he argues, are too belief-like.

Indeed, Lyons thinks that percepts occupy the same functional role as perceptual beliefs when perceptual beliefs are formed. And even when perceptual beliefs are not formed, perceptual seemings are plausibly viewed as attractions to assent. Here Sosa's thought applies: attractions cannot justify unless they independently qualify from the epistemic point of view. Given these observations, Lyons reasons as follows:

1. Experiences cannot be good epistemic reasons.
2. If (1), then only beliefs can be good epistemic reasons.
3. But not all beliefs are justified by other beliefs.
4. So, not all beliefs are justified by reasons.

But Lyons is wrong to assume that good reasons are either experiences or beliefs. That is compulsory only assuming that good reasons are mental states rather than facts or apparent facts *provided* as good reasons by mental states. This is not a trivial difference. To think you can justify a belief by appealing to its apparent truth-maker is not to think your belief can be self-justifying. Yet Lyons's argument for (1) rests on the thought that percepts in good cases are beliefs and the thought that beliefs cannot be self-justifying.

Lyons could reframe his argument. Even if we view normative reasons as facts or apparent facts, we still view their *providers* as states. So Lyons could instead argue:

- I. A percept with the content  $\langle P \rangle$  cannot provide the apparent fact that P as a reason to believe  $\langle P \rangle$  unless this percept qualifies epistemically.
- II. But the story about how the percept qualifies cannot be reason-based: after all, apparent facts cannot justify beliefs unless they are possessed as good epistemic reasons, but they are possessed as good epistemic reasons in virtue of percepts!
- III. If (II), RP is false.

IV. So, RP is false.

This argument fails for a more interesting reason: (III) is false.

We can adopt the following view. A percept cannot provide a reason unless it manifests perceptual ability. Still, perceptual beliefs are justified wholly by being based on reasons. What reasons? The facts or apparent facts *provided* as reasons by percepts that manifest perceptual ability. The percepts remain distinct from beliefs: percepts are attractions to assent that cause beliefs, when unresisted. This is all consistent with RP.

Lyons cannot respond without undermining his own view. He agrees that it is a truism that how things look helps to explain why our perceptual beliefs are justified.<sup>35</sup> He just gives a *non-experientialist* explanation of this truism: the epistemologically relevant looks are outputs of reliable perceptual systems, and a belief's being grounded in such a look makes it non-inferentially justified. Here Lyons effectively concedes that RP could be true as long as experientialism is dropped.

### 1.3.7 Taking Stock

Let's take stock. We considered one of the major objections to RP: namely, that reasons are unnecessary in certain cases of non-inferential justification. The objection rested essentially on two bad views: (i) a statist ontology of normative reasons, and (ii) a weak account of possession. By viewing normative reasons as facts or apparent facts and possession as an achievement that essentially manifests cognitive ability, we avoided the objection. En route, we saw that there are two achievements involved in possessing a normative reason: *access* and *sensitivity to support relations*. They boil down to different kinds of aptness: access amounts to *apt seeming*, and sensitivity to *apt treating*.

---

<sup>35</sup>See Ch. 4 of Lyons (2009a).

## 1.4 Is Belief for Sufficient Reasons Enough?

The objection in the last section is not the only objection to RP in epistemology. Many epistemologists worry that believing something because one has good reasons is *insufficient* for doxastic justification. If correct, this would undermine the following implication of RP:

- (\*) Believing that P in an epistemically justified way simply consists in believing that P *because* one has a sufficient epistemic reason to believe that P.

The lesson of this section will resemble the lesson of the last. The Resistance assumes a flawed picture of what it takes to believe something because one has good reasons to do so. In particular, the Resistance assumes that believing something because one has good reasons is something one can pull off by incompetent luck. If that assumption were true, (\*) would be false. But that assumption is not true. That is the real lesson of their cases.

### 1.4.1 The Insufficiency Argument

To see why the Resistance rejects (\*), consider the following case from Turri (2010):

(SPURS WIN) 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. Thus, they both possess sufficient epistemic reasons to believe that the Spurs will win. And they both infer that the Spurs will win from those two propositions. But Ponens uses *modus ponens*, while Lacy uses the wildly invalid *modus profusus* rule: for any P, Q, and R, infer R from P & Q.<sup>36</sup>

Ponens is doxastically justified. Lacy isn't. So, some members of the Resistance reason:

---

<sup>36</sup>Goldman (2012: 7) discusses a similar example and defends a similar claim on the basis of it. These sorts of examples predate both Turri and Goldman, and can be found in Armstrong (1973: 98), Swain (1988: 467), and Millar (1991: 57). While all of those examples are inferential, there are also non-inferential examples that are supposed to establish the same point—most notably, the case of the speckled hen, originally found in Chisholm (1942) but revived in an attack on reasons-based accounts in Sosa's contribution to Sosa and BonJour (2003).

1. Lacy believes the Spurs will win because he has a certain rationale—viz., <the Spurs will win if they play the Pistons> and <the Spurs will play the Pistons>.
2. This rationale happens to correspond to a sufficient epistemic reason to believe that the Spurs will win.
3. So, Lacy believes that the Spurs will win because he has this sufficient epistemic reason.
4. But Lacy’s belief that the Spurs will win is doxastically unjustified.
5. So, RP is false.

This can sound compelling, but the move from (1) and (2) to (3) is invalid.

#### 1.4.2 Believing Something Because You Have Good Reasons

The move works only if we accept the following view about what it takes to believe something because you have a good epistemic reason to believe it:

(*Coincident Rationale*) S believes that P because S has a good epistemic reason to believe that P iff (i) S has a rationale *r* for believing that P and (ii) *r* happens to correspond to a good epistemic reason that S has to believe that P.

This view is false. Believing something because you have a good epistemic reason to believe it is an *achievement*. If Coincident Rationale were true, it would be not be.

In order to believe something because one has a good reason to believe it, isn’t enough that one’s rationale merely corresponds to that good reason. The fact *that one has a good reason* must also play an explanatory role.<sup>37</sup> While the idea of a normative

---

<sup>37</sup>Notice how this differs from the suggestion that Turri (2010: 316) briefly considers. He considers a view that requires for doxastic justification that one believe something on the basis of good evidence “as good evidence”. But it becomes clear that what Turri means is merely that the subject *views* that evidence as good evidence. This is weaker than what I am requiring. I am requiring that the actual goodness of the reason play an explanatory role. This, I will maintain, is incompatible with incompetent basing when understood correctly. Merely requiring that someone *view* the reason as a good reason is compatible with incompetent basing, since people can *incompetently view* certain evidence as good. So there is a key difference between my proposal and the one Turri rightly dismisses as unhelpful. I require not that we merely *view* our reasons as good reasons, but that the *actual goodness* of these reasons play an explanatory role by triggering an exercise of competence.

fact's playing an explanatory role may sound wild, we can understand it unremarkably. Good reasons *qua* good reasons can play this role by triggering certain cognitive abilities.<sup>38</sup>

Here we again find a role for a familiar triadic pattern of properties. We can say that a subject's having some rationale for believing that P is

*right*, if the rationale corresponds to a sufficient epistemic reason to believe P,  
*competent*, if her having this rationale manifests a reliable disposition of hers to have a rationale for believing something only if it is a right rationale, and  
*apt*, if the subject's having the rationale is right *because* competent

Given these ideas, the following is a natural way to understand the achievement:

(*Apt Rationale*) S believes that P because S has a sufficient epistemic reason to believe that P iff S has an apt rationale for believing that P.

We can now explain what is amiss in SPURS WIN. Lacy does not believe that the Spurs will win *because* he has good reasons to believe that the Spurs will win. Why? Because his rationale for believing that the Spurs will win is not apt. It manifests wild incompetence!

The explanation is consistent with RP. RP implies that a subject is doxastically justified in believing that P if she believes that P because she has sufficient reasons to believe that P. SPURS WIN does not undermine this implication of RP. It *supports* this implication. Lacy is not doxastically justified in believing that the Spurs will win for a straightforward reason: he fails to believe that the Spurs will win *because he has sufficient reasons to believe that the Spurs will win*. One could disagree only if one

---

<sup>38</sup>Once again, reliabilists such as Goldman (2012) have recommended similar accounts of what they call *proper basing*. But the claim being made here is a stronger claim that, if true, would show that RP does not need to be modified at all. If one could only appeal to proper basing, RP could not be vindicated in its pure form. But I'm claiming that it can be vindicated in that form, so long as the defender of RP has the correct account of believing something because one has a good reason. Or, equivalently, I'm claiming that there is no difference between believing something because you have a good reason and proper basing—which I take to be an interesting claim that reveals a new avenue for friends of RP.



presupposed Coincident Rationale. But cases like SPURS WIN are counterexamples to Coincident Rationale, not to RP.

We can now precisify RP's story about doxastic justification:

(\*\*) Believing that P in an epistemically justified way consists in believing that P on the basis of an apt rationale.

(\*\*) explains why Lacy is doxastically unjustified. As far as I can see, it also explains all other cases resembling SPURS WIN that Turri and others have advanced.

### 1.4.3 Other Reasons to Reject RP's Account of Doxastic Justification?

Not all the cases that lead the Resistance to reject RP's account of doxastic justification clearly resemble SPURS WIN. Cases of *epistemically culpable ignorance* raise a further worry.<sup>39</sup> Consider:

(IGNORER) At  $t$ , Sebastian has only good epistemic reasons to believe P. Between  $t$  and  $t+$ , Sebastian receives a mass of counterevidence, but ignores it all. By  $t++$ , he has forgotten it, remembers only the original good reasons, and so believes P.

Some find it clear that Sebastian is doxastically unjustified in believing P at  $t++$ . They also find it clear that he believes that P because he has good epistemic reasons at  $t++$ .

How can the defender of RP respond? If we agree that Sebastian had sufficient epistemic reasons at  $t$ , it is hard to deny that he has sufficient epistemic reasons at  $t++$ . After all, quantitatively speaking, he has the same reasons. If so, then in the intermediate time, Sebastian could have insufficient reasons only because he has *more* reasons. If he lacks these further reasons at  $t++$ , and has at  $t++$  all the reasons he had at  $t$ , it is plausible that his reasons were sufficient at  $t$  iff they are sufficient at  $t++$ .

But note: Sebastian's behavior at  $t+$  suggests that he lacks the disposition that Apt Rationale requires. At  $t+$ , Sebastian is not disposed to have a rationale only if that

---

<sup>39</sup>See Greco (2005) and Goldman (2009).

rationale corresponds to a sufficient epistemic reason. He is radically negligent, after all. Yet this disposition is required for Sebastian's rationale to be apt. If so, a proponent of RP might agree that Sebastian has sufficient reasons at  $t++$ , but deny that he holds the belief *because he has sufficient reasons*. If so, (\*\*) can explain why he is unjustified.

Compare a more vivid case. Suppose Salvatore is honest about his aversion to counterevidence: "I don't care whether there is counterevidence. I'll believe anyway." Suppose he ignores countervailing reasons, forgets them, and now only possesses good epistemic reasons for his beliefs. He does not believe as he does *because he has sufficient epistemic reasons*. If so, (\*\*) explains why he is unjustified.

If Sebastian is like Salvatore, the friend of RP can use (\*\*) to deny that Sebastian believes for sufficient reasons at  $t++$ . We might deny that he does at  $t$  as well, if he lacks the dispositions necessary for having an apt rationale at  $t$ . If Sebastian is not relevantly like Salvatore, I lose my intuition that he is doxastically unjustified. Suppose he "ignored" the counterevidence just because he was too distracted. Suppose he would have responded if he had the attentional resources. If so, Sebastian strikes me as justified at  $t++$ .

I conclude that the defender of RP can either directly explain cases like IGNORER or dismiss the intuitions about them as not probative (because the cases are underspecified).

## 1.5 Two Levels of Explanation

I have addressed the main challenges to RP that I hear from externalists in epistemology. The challenges only reinforce (i) the virtues of the factualist ontology of normative reasons, and (ii) the plausibility of viewing the *possession* of good epistemic reasons and the business of believing something *because one has good epistemic reasons* as manifestations of ability.

Sometimes I hear this reply: "Your defense of RP just illustrates that reasons are epiphenomenal in the analysis of justified belief. Cognitive ability is what really matters." As I said earlier, I think that this reply conflates two levels of theorizing. In this

section, I deepen this theme by explaining why reasons-first epistemology and reliabilism contribute to the two different levels of explanation, and so cannot directly conflict.

### 1.5.1 Examples of the Two Levels: Buck-Passing and Hedonism

To appreciate the two levels of explanation, remember that there are two questions to ask about the nature of normativity. One question is:

(Q1) What are the fundamental units of normativity, and how are other normative facts grounded in facts about these units?

Q1 attracts great interest in meta-ethics. This owes partly to interest in:

(*Reasons First*) Facts about reasons ground all other normative facts.

Scanlon (1998) prompted interest in Reasons First. One of his contributions to Reasons First was his account of how facts about value are grounded in facts about reasons:

(*Buck-Passing*) For any valuable X, X's being valuable is grounded in X's having descriptive properties that give sufficient reasons for having a pro-attitude to X.

This view is compatible with views like:

(*Hedonism*) X's being valuable is grounded in X's being pleasure-conducive.

Views like Hedonism partly answer the other question I mentioned at the outset:

(Q2) How, if at all, are normative facts grounded in non-normative facts?

Q1 and Q2 both concern the metaphysics of normativity. But they can be independently answered without conflict. This is for a simple reason. One thing can be either *directly* or *indirectly* grounded in another. If X grounds Y and Y grounds Z, X can thereby ground Z—just indirectly. We can accept both Buck-Passing and Hedonism by taking the former to give us the first step in the grounding of value and the latter to summarize the final step. For any valuable X, we might say (where the 'because' is metaphysical):

X is valuable BECAUSE X is reasonable to desire (a Buck-Passer's claim)

and

X is reasonable to desire BECAUSE X is pleasure-conducive

And we can thereby say:

X is valuable BECAUSE X is pleasure-conducive (a Hedonist's claim)

So, Buck-Passing and Hedonism are compatible. They conflict if we replace 'is grounded in' with 'is directly grounded in' Hedonism. But we should not do this. This commits Hedonists to the claim that facts about value are normatively basic. For then no other normative facts would intervene between facts about value and facts about pleasure-conducivity in the grounding chain. Hedonism should have no such implications, since it only concerns the non-normative grounds of value.

### 1.5.2 Reliabilism : RP :: Hedonism : Buck-Passing

RP is similar in point to Buck-Passing, and reliabilism is similar in point to Hedonism. If so, it is a mistake to think that there is a debate between reliabilism and RP. Moreover, we can use broadly reliabilist ideas to *explain in descriptive terms* what it takes to possess good reasons for belief and to believe things because one possesses good reasons for belief. This is no blow to RP, just as the truth of Hedonism would be no blow to Buck-Passing. In parallel to the grounding chain for value, we can say:

S's belief that P is justified BECAUSE [insert RP's story]

and

RP's story is true BECAUSE [insert sophisticated reliabilist story]

And we can thereby say:

S's belief that P is justified BECAUSE [insert sophisticated reliabilist story]

We should not be surprised. RP is a view about the priority of certain normative units over others. Reliabilism is not a view of this kind. Reliabilism grounds epistemic normativity in the non-normative. Goldman (1979: 1) was clear from the start:

Any correct definition [of ‘justified’] would also feature evaluative terms, but I am not interested in them. I want a set of substantive conditions that specify when a belief is justified. Compare the moral term ‘right’. This might be defined in other ethical terms or phrases, a task appropriate to metaethics. The task of normative ethics, on the other hand, is to state substantive conditions for the rightness of actions. [...] I want a theory of justified belief to specify in non-epistemic terms when a belief is justified. This is not the only kind of theory of justifiedness one might seek, but it is one important kind and the kind sought here.

Reliabilism, then, resembles Hedonism. Hedonism and the buck-passing account are not rivals. They answer different questions about the nature of value. Reliabilism and RP are not rivals either. They simply answer different questions about the nature of justification.

### 1.5.3 Against the Deflationary Response

This undermines a deflationary response to the foregoing sections—viz., that reasons are epiphenomenal and reliability does the real work. That might be plausible if reliabilism and RP were rivals. But they are not rivals. RP contributes to the project of answering Q1, whereas reliabilism has always been a contribution to the project of answering Q2.

RP has upshots for Q2 *if conjoined* with particular substantive groundings of the reasons ideology. So, it is crucial to show that RP can jibe with the best views about how epistemic normativity is grounded in the non-normative. But once that is done, we must remember that Q2 takes us to a level of metaphysical explanation on which RP is silent. RP can guide our thinking about that level, by suggesting that we find the non-normative grounds of justification by finding the non-normative grounds of the reasons ideology. But that’s it.

## 1.6 Conclusion

*Recap.* We have covered a lot of ground, so let’s briefly review the highlights.

I answered two kinds of objections to RP. According to the first, reasons cannot explain why some foundational beliefs are justified. Members of the Resistance point to examples like the belief that  $1 + 1 = 2$ , and maintain that there are only two candidate justifying reasons: (a) the intuitive seeming that  $1 + 1 = 2$  and (b) the sheer fact that  $1 + 1 = 2$ . They dismiss (a) as unhelpful. Seemings reduce to attractions to assent. Attractions can manifest incompetence, and the people who host them may be incompetent at the second order, failing to possess any relevant defeaters. After dismissing (a), the Resistance insists that (b) not a serious option and leads to circularity problems.

There are two mistakes here. The first is to dismiss (b) as a serious option. The apparent truth-maker of one's belief is the perfect reason for that belief if one possesses that apparent truth-maker in the right way. This is illustrated by introspective examples, like believing that one is in pain on the basis of the clear fact that one is in pain. There is no circularity here. The basing is not inferential, and only inferences can be circular.

Of course, we cannot possess any old fact as a direct license to believe what it makes true. Indeed, we cannot possess the apparent fact that P as a direct license to believe  $\langle P \rangle$  just by finding  $\langle P \rangle$  intuitive, if our intuitive attraction is incompetent. But the intuitive case is not special! Not all introspective facts can be possessed in this way either. This is revealing, because reasons are clearly doing justifying work in good introspective cases: the trouble with people like Fred from §3.2 is that they don't possess the relevant reasons. What we need to possess intuitive facts as direct licenses is what Fred needed: ability.

The second mistake is to assume that defenders of RP must embrace low bar accounts of possession on pain of circularity. Defenders of RP can appeal directly to non-normative factors to explain when certain mental states can enable one to possess an apparent fact as a good epistemic reason for belief. I demonstrated this in §3.4, by giving an ability-based account of the two achievements involved in possession.

That account did not address the second major objection to RP. Some say cases like SPURS WIN show that a subject can believe P because she has good epistemic reasons to believe P and yet lack doxastic justification. We saw, however, that these

cases just refute an overly permissive account believing something because one has good epistemic reasons—viz., *Coincident Rationale*. I offered a better account: to believe that P because one has good epistemic reasons to believe that P, one needs an *apt rationale*.

These responses address the major objections I have heard from externalist crowds in epistemology. The Resistance might reply that this shows that reasons aren't doing the real explanatory work: aptness does that work. But this, I argued, is a mistake. It ignores the question that RP is designed to answer. The point of RP is to show how one normative property reduces to a more fundamental normative property. RP is unattractive if RP cannot mesh with the best non-normative grounding stories in epistemology. But once we have shown how RP can mesh, we must recognize that RP is not itself such a story. It is an account like the buck-passing account of value. It is in a different ballpark than reliabilism.

*Contrasts with Evidentialism/Reliabilism Syntheses.* That observation highlights one reason why I have not been reinventing the wheel forged by recent syntheses of evidentialism and reliabilism.<sup>40</sup> It is wrong to speak of a synthesis of reasons-first epistemology and reliabilism, just like it is wrong to speak of a synthesis of the buck-passing account of value and hedonism. These accounts operate at different levels: one contributes to a picture of the internal structure of the normative, whereas the other contributes to a picture of how the normative is grounded in the natural.

If I had the space, I would argue for a bolder conclusion: the idea of a debate between *evidentialism* and reliabilism involves the same mistake, since the concept of evidence is normative. There are real debates between *evidentialists* and *reliabilists*: Conee and Feldman's substantive views are clearly not Goldman's. But Conee and Feldman's substantive disagreements with Goldman go beyond the minimal evidentialist thesis, which is just a thesis like RP. Conee and Feldman would not disagree. From the beginning, they stressed that evidentialism was not a substantive theory so much as an indication of "the kind of justification we take to be characteristically epistemic".<sup>41</sup> And

---

<sup>40</sup>See Comesaña (2010) and Goldman (2011). Alston (1988) is a precursor.

<sup>41</sup>Conee and Feldman (1985: 15).

they agreed that a reliabilist analysis of well-foundedness is conceptually possible (just false).<sup>42</sup>

Here is not the place to defend this bolder conclusion. What matters is that I have not engaged in the same project that synthesizers of evidentialism and reliabilism take themselves to be engaged in. I did sketch substantive accounts of what it takes to *possess* good epistemic reasons and to believe something *because one possesses good epistemic reasons*. But I did this only to show that RP makes no bad predictions *per se* about cases. Only if a defender of RP embraces certain specific views will she get bad predictions about cases.

But my substantive views are preferable to existing evidentialism/reliabilism syntheses. For one thing, synthesizers frame their views by presupposing a statist rather than a factualist ontology of reasons, which invites all the problems that my view avoids. For another, existing syntheses neglect the second condition on possession that my account highlights (i.e., sensitivity to favoring relations). This is why they face problems with cases like:

(KOPLIK SPOTS) Albert lives in a period long before it was discovered that Koplik spots reliably indicate the presence of measles. Albert sees the spots for the first time ever in Bertrand's mouth and spontaneously believes on the basis of the appearance that Bertrand has measles without performing any inference.

The fact that there appear to be these spots in Bertrand's mouth is an objective epistemic reason to believe that Bertrand has measles. But Albert does not *possess* this reason.

Why not? Standard syntheses of evidentialism and reliabilism will have a hard time giving a satisfactory answer. For they agree that (i) experiences are possessed evidence for spontaneous, non-inferential beliefs *when* (ii) there is an objective probabilistic relation between the experience and the truth of the belief. But as we saw, having access

---

<sup>42</sup>“These certainly are *conceptually* different accounts of the grounds of well-founded beliefs. In spite of this conceptual difference, reliabilism and WF [the evidentialist account of well-foundedness] may be extensionally equivalent” (1985: 25). Also noteworthy: “[R]eliabilism may not even be a rival to WF” (27).



to something that is an objectively good reason for belief is not sufficient for possessing a good reason for belief: one must also be sensitive to the relevant favoring relation. By honoring this constraint, my view explains why Albert doesn't possess the reason.

*Other Objections?* One might be left wondering whether I have addressed all the main worries in epistemology about RP. One might, for example, wonder about stored beliefs for which the original evidence has been forgotten. Many of our justified beliefs are like this. Aren't these obvious examples of justified beliefs that are based on no reasons?

They are not. One could claim that the fact that one seems to remember that P is a good reason to believe that P *when* this seeming manifests competence. Since RP is not an essentially internalist view, this strategy is available to the defender of RP. If one wants to avoid overintellectualization, one could instead say that memory seemings can sometimes put one into contact with (apparent) facts just as perception and intuition can, enabling them to serve as licenses for believing what they (apparently) make true. "I just remember!" feels a lot like "I just see!"

There are other cases that worry people. Jack Lyons has pressed two other putative counterexamples. One case involves justified but false introspective belief, and the other involves beings that lack phenomenal consciousness. Lyons (2009b: 255)'s first example:

[T]here is no obvious way to make sense of justified but false introspective beliefs consistently with [RP]. Obviously our other beliefs fail to determine whether or not we are justified in some particular introspective belief.... So if we are to have any ground for the introspective belief, it would have to be a nondoxastic state. But which state? Suppose I am standing at the refrigerator because I have mistaken my boredom for hunger. What justifies my belief that I'm hungry? Not my hunger, since, by hypothesis, I'm not actually hungry. Not my boredom, surely?

I am unconvinced. Other beliefs alone can rarely if ever justify introspective beliefs. But some introspective beliefs could be justified by a combination of (i) introspectible facts, and (ii) facts about what other states we are typically in when these facts obtain. In Lyons's case, one might reason as follows, if rapidly or subconsciously: (1) I am in state X (=wanting to go to the fridge); (2) Whenever I am, I am typically in state Y (=being hungry); (3) So, I am probably in state Y. There must be something that

prompts one's belief. The prompt alone may not justify one's belief. But it could given certain background factors.

Lyons's other worry concerns zombies (i.e., beings that lack phenomenal consciousness). He thinks it is obvious that "zombies can have beliefs and... that some can have justified beliefs".<sup>43</sup> But he worries that a zombie's non-inferential beliefs about the world could not be justified according to RP, because he assumes that the only reasons to which a defender of RP could appeal would be experiences or beliefs.

This is wrong. Lyons is again assuming the statist ontology of reasons. If Lyons defined 'grounds' to mean 'beliefs or experiential states', he would have an argument against a stipulative version of RP. But not against mine. Someone who rejects experientialism can accommodate zombies. My view simply appealed to apt seemings as providers of reasons. There is no reason to think that zombies cannot have apt seemings.<sup>44</sup>

*Envoi.* I am doubtful about counterexample-based cases against RP. Since RP is antecedently attractive, I think we have sufficient reasons to remain attracted. At any rate, I have done what the epistemologist *qua* epistemologist can do. Work remains for Reasons Firsters. But I hope to have cleared the way by addressing the Resistance in epistemology.

---

<sup>43</sup>Lyons (2009b).

<sup>44</sup>There are other problems. What Lyons finds clear isn't clear. I agree that zombies might have justified external world beliefs. But many will lack Lyons's intuitions. One might deny that the zombies are justified but hold that they are excusable or structurally rational. And there is a big step from the claim that the zombies are *not unjustified* to the claim that they are *as justified* as us. Zombies could have weak reasons by having coherent doxastic systems. The isolation objection to coherentism does not show that coherence generates no reasons—just not good enough reasons to explain ordinary cases.

## Chapter 2

### What Apparent Reasons Appear to Be

#### Overview

Many meta-ethicists think that rationality only requires us to heed *apparent* normative reasons, not objective normative reasons. But what are apparent reasons? One hears two kinds of answer. On *de dicto* views,  $p$  is an apparent reason for S to  $\phi$  when it appears to S that  $p$  is an objective reason to  $\phi$ . On *de re* views,  $p$  is an apparent reason for S to  $\phi$  when (i)  $p$ 's truth would provide an objective reason to  $\phi$  and (ii) it appears to S that  $p$ . *De re* views are currently more popular because they avoid overintellectualizing rationality. But they face problems owing to the way in which they do so. Some assume we can escape the problems by requiring more descriptive facts to be apparent or by appealing to defeat. But these strategies fail. So I defend a new view that is closer in spirit to *de dicto* views but less demanding (at least granting the assumptions needed to support overintellectualization worries). On this view, apparent reasons are apparent facts that agents are competently attracted to treating like objective reasons, where competence is indirectly defined in terms of objective reasons and a competence/performance distinction is drawn.

#### 2.1 Reasons and Apparent Reasons

Meta-ethicists often view normative reasons as objective facts that count in favor of actions and attitudes. These facts include ones to which we do not necessarily have privileged access. For example, the fact that the lemonade is arsenic-laced is a conclusive reason for you not to drink it. This fact remains a conclusive reason for you not to drink the lemonade even if you are in no position to see that the lemonade is arsenic-laced.

Clearly, however, you need not be *irrational* in drinking the lemonade if you are in no position to appreciate this fact. So, meta-ethicists like to separate rational  $\phi$ -ing from  $\phi$ -ing that is supported by all the objective normative reasons. Rationality retains an indirect connection to objective normative reasons: it requires responding to *apparent* objective reasons. On many views, the ‘apparent’ is not a success term. There may appear to be sufficient objective reasons to  $\phi$  even if there are conclusive objective reasons not to  $\phi$ . The lemonade case illustrates this fact. Moreover, there may be no apparent reason to  $\phi$  even if there is an objective reason to  $\phi$ . The same case illustrates this fact: while there is no apparent reason to refuse to drink the lemonade, there is an objective reason to refuse (viz., that the lemonade is lethal).

Now, many meta-ethicists who draw this distinction believe that rationality is just coherence.<sup>1</sup> Accordingly, they assume that apparent reasons are apparent in a belief-relative sense. But we can understand ‘apparent’ more broadly without collapsing rationality into objective correctness. Epistemic rationality requires responding to the recommendations of perceptual appearance even if this appearance is radically non-veridical (e.g., in the Matrix or other skeptical scenarios). It does not require responding to all the objective evidence out in the world. Many epistemologists who embrace these thoughts would not cash them out in coherentist terms (e.g., by treating perceptual appearances as beliefs).

As an epistemologist, I am predisposed to see the requirement to respond to apparent normative reasons as stronger than any coherence requirement. One could try to reduce this requirement to a coherence requirement. I independently find this project suspect, but what matters for our purposes is that even on non-coherentist views rationality is weaker than the kind of correctness that goes with doing what the totality of objective facts would favor. For not all objective facts are apparent, and not all apparent facts are real.

The focus of this paper is on what it takes for a consideration to be an apparent reason. I think the main ways in which meta-ethicists answer this question are flawed,

---

<sup>1</sup>See, for example, Scanlon (1998), Broome (1999, 2005, 2007, 2013), and Kolodny (2005).

and I will offer a new answer. In §2, I introduce two main types of view that meta-ethicists have endorsed: *de dicto* views and *de re* views. *De re* views are currently more popular because many worry that *de dicto* views overintellectualize rationality. While I agree about this worry, I think moving to standard *de re* views is a mistake. In §3, I pose some problems for *de re* views and show how potential responses fail. I then turn to my own view in §4, explaining how it avoids the problems for *de re* and *de dicto* views. Finally, I answer an objection to my view—and to all other views discussed earlier—in §5.

I will stress in advance that my real foe in this paper is the *de re* theorist. While I don't accept the *de dicto* view, I also don't reject it. Non-acceptance is weaker than rejection: I merely don't want to commit to the views in the philosophy of mind that the *de dicto* view would force us to accept on pain of overintellectualization. While many important features distinguish my view from the *de dicto* views in the literature, my view would entail a (new) version of the *de dicto* view *if* certain extreme views about concept possession that I reject were true. My view simply doesn't require these extreme views to succeed. So, while it is less demanding (granting the assumptions about concept possession that drive the overintellectualization worry), my view is closer in spirit to the *de dicto* tradition. I will explain these nuances further in the next section.

## 2.2 Apparent Reasons *De Dicto* and *De Re*

### 2.2.1 *De Dicto* Views

When meta-ethicists first began to distinguish between rationality and correctness, apparent reasons were understood very narrowly. For example, Scanlon (1998: 25) held that irrationality in its “clearest sense” occurs when someone “judges that something is a reason but fails to be affected by it in one of the relevant ways”. Kolodny (2005) echoed him, holding that the core requirements of rationality are requirements to respond to our *de dicto* beliefs about normative reasons. These views are examples of a broader view:

DE DICTO: R is an apparent normative reason for S to  $\phi$  iff it appears to S that

R is an objective normative reason to  $\phi$ .

While Scanlon and Kolodny take the relevant sense of ‘appears’ to be belief-relative, it is not compulsory to do so. We have known at least since Chisholm (1957) that appearance words can be used in many ways. Besides the doxastic use, there are non-doxastic uses of ‘appears’. Even if you fail to take any doxastic stance on whether P, it can still perceptually appear to you that P. Non-doxastic appearances are not confined to perceptual cases. Propositions in *a priori* domains can appear true in this way too: consider inconsistent triads of *a priori* plausible claims.

Normative propositions can appear true in this way too. So, if one agrees that rationality demands less than conformity with all the objective normative reasons but dislikes coherentism, one could preserve DE DICTO but understand the ‘appears’ less narrowly. One could agree that there is a more restricted kind of rationality that is belief-relative—viz., structural rather than substantive rationality. But one could deny that it is the only kind of rationality to be distinguished from *fact-relative* statuses like objective correctness. Substantivity and fact-relativity are orthogonal: apparent reasons are a clear illustration!

### 2.2.2 Understanding the Overintellectualization Objection

Even understood in this wider way, many meta-ethicists would reject DE DICTO. A major reason concerns overintellectualization. Parfit (2011: 118) expresses the concern well:

We can have rational beliefs and desires, and act rationally, without having any beliefs about reasons. Young children respond rationally to certain reasons or apparent reasons, though they do not yet have the concept of a reason. Dogs, cats, and some other animals respond to some kinds of reason...though they will never have the concept of a reason. And some rational adults seem to lack this concept....

Parfit's thoughts are plausible, carefully understood.<sup>2</sup> But there are two unequally plausible ways to extract a challenge to DE DICTO. The most ambitious is a knock-down argument:

### **The Strong Overintellectualization Argument**

1. Rational subjects must possess the concept of a normative reason if DE DICTO is true, since in order to host appearances with contents of the form *R is an objective normative reason to  $\phi$*  one must have this concept.
2. But rational subjects need not possess the concept of a normative reason.
3. So, DE DICTO is false.

While there is a version of the overintellectualization objection worth taking seriously, the Strong Overintellectualization Argument is not it. Defenders of DE DICTO could easily reject (1) or (2). Doing so just requires some burdensome commitments on which a general theory of apparent reasons and the nature of rationality would ideally be neutral.

For example, one might insist that possessing the concept of a normative reason is easier than Parfit and others suppose. One need not be able to articulate one's thoughts involving a concept to possess it, or be able to use the words that express the concept. But perhaps all that is clear is that infants, toddlers, and animals cannot articulate any beliefs about normative reasons. Perhaps, one might insist, they still have these beliefs.

Obviously, a burden remains: friends of DE DICTO must offer a positive story about how it is that children and animals possess the concept of a normative reason. Friends of DE DICTO might maintain that possessing the concept simply consists in having

---

<sup>2</sup>Might it be insisted that while infants, toddlers and animals lack the concept of a normative reason, they have a related *proto-concept*? Isn't it plausible that even infants can believe that certain considerations *count in favor*? The concept of favoring may seem less sophisticated. But note that it is a *relational* concept: considerations count in favor of *attitudes and acts*. For us to explain the rationality of their beliefs, desires, etc., on the revised *de dicto* view, children must not only have the concept of favoring, but also the concepts of these attitudes. One cannot view the incomplete content <R counts in favor> as the relevant object of children's thoughts and seemings. The content would have the form <R counts in favor of  $\phi$ -ing>, where  $\phi$ -ing could be believing, desiring, etc. Yet it is even more doubtful that infants and some toddlers have these mental state concepts. Many animals certainly lack them. So a problem remains.

certain dispositions to treat considerations in ways that would be appropriate if they were real normative reasons. Proponents of the representational theory of mind (e.g., Fodor (1998)) will find this an implausible view about concept possession. Being in favor of the representational theory of mind myself, I would prefer a theory that does not force us to accept a purely dispositional account of concept possession.

One might take a different tack. One might deny that possessing the concept of a normative reason is necessary for hosting appearances with the content *R is a normative reason to  $\phi$* . For example, one might insist that one can host an appearance with the content *X is F* simply by being attracted to treat X like an F. I would find this an unattractive hypothesis in the theory of content. Does it really seem to a dog *that its owners are vending machines* simply because this dog treats its owners like vending machines? On many views, this idea would be implausible if the content attribution is intended seriously—as it must, if the view is a *de dicto* view. I would prefer a view that doesn't raise these worries.

Friends of DE DICTO can surely defend their view to their own satisfaction by adopting burdensome views in the philosophy of mind or theory of content. But this highlights a different way to understand the objection. The objection is not that one can *obviously* be rational without having normative concepts. Instead, we can put it as follows:

### **The Modest Overintellectualization Argument**

1. A view about rationality that is defensible without invoking burdensome views in the philosophy of mind or theory of content is *ceteris paribus* preferable to a view that is indefensible on pain of overintellectualization if such views fail.
2. DE DICTO avoids overintellectualizing only if some burdensome views in the philosophy of mind or theory of content are true.
3. So, unless there is no defensible view of rationality that can avoid invoking these views, we should not endorse DE DICTO.



The conclusion does not tell us to deny DE DICTO. It just tells us to withhold belief and search for a view that does not make it compulsory to endorse burdensome views in the philosophy of mind or theory of content. The view I will offer in §4 is such a view. The conjunction of my view and these burdensome views might entail DE DICTO. But the plausibility of my view does not depend upon the truth of these burdensome views.

A comparison vindicates the point. Consider the early literature on the Gettier problem. Clark (1963) thought that we could fix the justified true belief account of knowledge by suggesting that subjects who know that P do not essentially rely on inferences from falsehoods in arriving at their justified true beliefs that P. According to this picture, what is going awry in Gettier cases is that subjects are arriving at their justified true beliefs by essentially relying on inferences from falsehoods. The obvious worry about this proposal is that there are Gettier cases where subjects do not seem to perform inferences at all in arriving at their Gettiered beliefs (e.g., Fake Barn Country). One could try to attribute intricate subconscious inferences to these subjects, as Harman (1973) did. And if one did this, one could arguably explain all Gettier cases with a sophisticated Clarkeian theory.

But today, virtually no one accepts this as a solution to the Gettier problem. This is not because everyone finds it obvious that subjects perform no relevant subconscious inferences in cases like Fake Barn Country. It is because an ideal solution to the Gettier problem would not require views in the philosophy of mind that few epistemologists would embrace. Hence the greater popularity in the later 1970s of defeasibility theories. These theories subsumed the good predictions of Clarkeian theories without appealing to any controversial views in the philosophy of mind. The two kinds of theory would be extensionally equivalent if we accepted these controversial views. But defeasibility theories seemed preferable, since they did not *require* these controversial views to work.

The view that I will be developing in §4 stands to DE DICTO as defeasibility theories stand to theories in the Clarkeian tradition. Again, my view would be extensionally equivalent to some version of DE DICTO if certain highly controversial views about concepts and content were true. But the plausibility of my view does not depend upon

the truth of these views, and for this reason stands a better chance of survival.

This is my main reason for not affirming DE DICTO. Again, I do not reject DE DICTO. But I will only affirm a theory that can get the right predictions without relying on any strong views in the philosophy of mind or theory of content. If I believed these strong views, I might embrace DE DICTO in my heart of hearts. But it is not a view that I want to assert: I want to assert something more modest, at least in certain respects.

### 2.2.3 Other Concepts?

Before setting aside *de dicto* views, it is worth considering a final response. I have only discussed *de dicto* views that invoke the concept of a normative reason. But there are other normative concepts that might not raise such stark overintellectualization worries. Besides the concept of a normative reason, there are thin evaluative concepts like *goodness*. There are also thick concepts, like *gross*, *delightful*, *terrifying*, and so on. Isn't it more plausible that even animals can conceptualize things in these ways?

Perhaps. I would still prefer a theory that did not *require* us to attribute any such concepts to all beings that can act and believe rationally. But the analogy with Clarkean responses to the Gettier problem is less obvious with these possibilities on the table.

So, why can't we just relax? Because the worries return when we look carefully. My worries concern the implications for epistemology. Suppose that Alpha is competent with some valid pattern of first-order inference—say, *modus ponens*. And suppose that it appears to Alpha that P and that Alpha knows that if P then Q. This is a paradigm case of having an apparent reason to believe Q. How can a *de dicto* theorist explain this case? The theorist must suppose that it also appears to Alpha *that believing Q would be positively evaluable in some way*. But must any subject in Alpha's position be capable of hosting an appearance with this content, in order to gain an apparent reason to believe Q?

Not plausibly. This proposal involves a new overintellectualization. This time, the overintellectualization has nothing to do with requiring the possession of certain *normative* concepts. The problem lies in requiring the possession of certain *mental state*

*concepts*. The *de dicto* theorist must say that Alpha conceptualizes the doxastic response of believing Q in a positive light. But mental state concepts are sophisticated. Subjects do not need to be able to think about their own beliefs in order to perform rational first-order inferences as elementary as *modus ponens*.

One might again try to retreat to lax views about concept possession or otherwise argue that it is easier than it looks to think about one's mental states. But now the analogy with Clarkean theories is vivid. It is better to avoid making these strong claims in the philosophy of mind to save the *de dicto* theory. My view will have this advantage.

#### 2.2.4 The *De Re* View

But the view that I will offer is not the usual reaction to the overintellectualization objection. The usual reaction is to move from DE DICTO to a view of this form:

DE RE: R is an apparent reason for S to  $\phi$  when (i) R's truth would give S an objective normative reason to  $\phi$ , and (ii) it appears to S that R.<sup>3</sup>

This view is embraced by theorists like Parfit (2001, 2011), Schroeder (2007), and Way (2009). Like Scanlon and Kolodny, they often unpack 'appears' in a belief-relative manner. For example, Parfit (2011: 111) writes:

While reasons are given by facts, what we rationally want or do depends on our beliefs. If we have certain beliefs about the relevant, reason-giving facts, and what we believe would, if it were true, give us some reason, I am calling these *beliefs whose truth would give us this reason*. Such beliefs give us an *apparent reason*. When such beliefs are true, this apparent reason is also a real reason.... We ought rationally to respond to apparent reasons even if...these reasons are not real.

And Schroeder (2007: 14) says the following, where 'subjective reason' = 'apparent reason': "For R to be a subjective reason for X to do A is for X to believe R, and for it to be the case that R is the kind of thing, if true, to be an objective reason for X to do A."

---

<sup>3</sup>This is formulated in a way that makes it sound like I am assuming that reasons are propositions. But this assumption is not essential to anything in this paper. It just makes certain formulations verbally simpler. If one prefers, one can rewrite everything in terms of states of affairs *a la* Dancy (2000). Indeed, I prefer the Dancyean picture. It would just lead to clunkier formulations in this context.

Once again, the belief-relative formulations are not mandatory. The ‘appears’ in DE RE could be understood in a non-doxastic sense of the sort mentioned earlier. Indeed, when it comes to stating his epistemological views, Schroeder abandons the belief-relative formulation and requires only that R be the content of some *presentational mental state*.<sup>4</sup> Such states include things like perceptual experiences and intuitive seemings. If DE RE is to be taken seriously by epistemologists, it should be formulated in this less narrow way.

However formulated, DE RE avoids the overintellectualization of DE DICTO. It does not entail that rational subjects must possess any normative concepts, or that subjects must represent apparent reasons as reasons. So, given the Modest Overintellectualization Argument, it might seem that we should pursue DE RE.

## 2.3 The Failure of *De Re* Views

### 2.3.1 The Problem of Unapparent Reasons

But only if DE RE is plausible along other dimensions. It is not. Ironically, DE RE’s flaws stem from the way in which it reacts to the overintellectualization of DE DICTO.

Notice that DE RE does not merely suggest that rational subjects need not represent apparent reasons as objective reasons. If DE RE is true, R can be an apparent reason for a subject to  $\phi$  *even if it would not be rational for her to treat R like an objective reason to  $\phi$* . For even if it is not rational for a subject to treat R like an objective reason to  $\phi$ , it might remain apparent to her that R is the case, and R might still be a consideration that would be an objective reason to  $\phi$  if it were the case. And this is that all DE RE requires for R to be an apparent reason to  $\phi$ !

Worries about overintellectualization do not support this feature of DE RE. One can treat something like an F without representing it as an F or having the concept of an F. Some cat might treat its owners like vending machines without having the concept of a vending machine or representing its owners as vending machines. So the requirement that apparent reasons be considerations that it is rational to treat like good reasons

---

<sup>4</sup>See Schroeder (2011).

is not overly intellectual. This requirement is plausible. How could we be rationally required to respond to considerations that it is not rational for us to treat like good reasons?

Yet DE RE implies that we could. This is a serious problem. I call it the *Problem of Unapparent Reasons*, since the problem is that DE RE classifies as apparent reasons considerations that are not apparent reasons.

Many possible cases fit this bill. Even if we know all the relevant descriptive facts, there might be objective reasons given by these facts that it would not be rational for us to treat like objective reasons. Our sensitivity to normative reality is limited. For those of us with normative concepts, this limitation can manifest in our inability to rationally settle certain normative questions. In such cases, (i) and (ii) in DE RE can easily be satisfied with respect to some consideration while the consideration fails to be an apparent reason.

The fundamental limitation clearly does not extend only to those of us who possess normative concepts. It might be true that any rational being has the ability to correctly respond to some would-be objective reasons. But this ability has bounds. Even given perfect descriptive knowledge, an ordinary rational being's competence only goes so far. It may not reach so far as to enable this being to rationally treat like objective reasons all the apparent facts that would be objective reasons if these facts were real. When the bounds of competence are surpassed, conditions (i) and (ii) in DE RE can easily be satisfied with respect to some consideration while the consideration fails to be an apparent reason.

One can construct simple counterexamples to DE RE without appreciating these deeper points. But I prefer putting the deeper points first, because they allow us to see how to forestall the main responses to simple counterexamples.

Let's consider these simple counterexamples. Here is one drawn from epistemology. Suppose it appears to Jones that P—indeed, suppose Jones knows that P. And suppose that P logically entails Q. So P is an objective reason to believe Q. Both clauses of DE RE are satisfied: the relevant consideration is apparent, and this consideration would be an objective reason to believe Q if true (indeed, it *is* one). Compatibly with all this,

Jones's inferential abilities might be too limited for him to be sensitive to the logical relation between P and Q. Indeed, the entailment might be arcane, visible only to the greatest logicians. If so, Jones cannot rationally treat P like an objective reason to believe Q. And if so, it is wrong to view Jones as having P as an *apparent* normative reason to believe Q.

Notice that this differs from a standard objection to the view that believing that P is sufficient for *possessing* P as a normative reason. The standard objection to that view is that irrationally believing that P cannot put one in a better epistemic position with respect to P's consequences.<sup>5</sup> But we stipulated that Jones *knows* that P. So the problem has nothing to do with his epistemic position with respect to P. The problem lies in his inability to rationally treat P like an objective reason to believe Q.

### **Defeat?**

The problem is not plausibly addressed by an appeal to defeat. In the relevant sense of 'defeated', apparent reasons have to be defeated by other apparent reasons. In our case, the fact that explains why Jones cannot rationally treat P as a reason to believe Q is the fact that Jones's competence is limited. But this fact isn't one that *itself* must be apparent to Jones to prevent him from possessing P as an apparent reason to believe Q.

People can be deceived about their own abilities. They can regard themselves as geniuses when they are fools. If their incompetence is not apparent to them, it is not the sort of thing that could serve to defeat another *apparent reason*. Will they have other apparent defeaters owing to the mistakes their incompetence makes likely? Not necessarily. Some fools may arrive at mostly correct results by fluke. If so, they will not possess indirect evidence for believing themselves to be incompetent. So a defeat strategy will be inapplicable in some of these cases. Yet it is plausible that their incompetence prevents them from possessing apparent reasons of the sort that bear on substantive rationality.

---

<sup>5</sup>Schroeder (2011) addresses that standard objection, but not the one at issue here.

Obviously, there is defeat in an *objective* sense in these cases. But that fact is irrelevant. In the sense of ‘defeated’ relevant to rationality, apparent reasons have to be defeated by other apparent reasons. Might our intuitions be confused, conflating objective and apparent defeat? Not plausibly. There can be unpossessed objective defeaters for perfectly rational beliefs. Many Gettier cases are like this, and our intuitions do not lead us astray in these cases. Intuitively, our case is nothing like these cases. It is plausible in our case that it is not rational for Jones to believe that Q on the basis of P, but implausible that rationality is lacking in standard cases of unpossessed objective defeaters.

Let’s put a nail in this coffin. Imagine again that P entails Q, but that the entailment is an arcane one beyond Jones’s competence. But now imagine that Jones treats P as a reason to believe Q by means of an overgeneralizing incompetence that has, by sheer luck, yielded correct results so far. Imagine that he lacks the second-order competence to appreciate his first-order incompetence. According to DE RE, he has an apparent reason to believe Q. By stipulation, he also has no apparent defeaters. Now suppose Jones believes Q, and that Q is true. By the lights of DE RE, Jones has a rational true belief that Q. Clearly, Jones does not know Q. DE RE thus suggests that this case is a kind of Gettier case—a case of rational true belief without knowledge. This is absurd. This case is nothing like a Gettier case. The obvious reason is that Jones’s belief in Q is not rational.

Finally, notice that the basic intuition in the first case is that Jones cannot rationally treat P like an objective reason to believe Q, from which we infer that P is not an apparent reason for him to believe Q. So we are not committing the fallacy of assuming that there is no reason when there is a *weak* one.<sup>6</sup> We are not assuming this but *concluding* it.

These matters of degree can be made explicit. Given Jones’s limited abilities, it is not rational to any degree for him to treat P like an objective reason to believe Q.

---

<sup>6</sup>Schroeder (2007, Ms) calls this the ‘negative existential reasons fallacy’. I agree that it is a fallacy. I just do not see how this fallacy is driving the present argument. I also do not see how Schroeder’s usual strategy for explaining away the negative existential reasons intuition works here.

Might he have an apparent reason to believe Q of zero weight? This differs from our conclusion only verbally. A reason with no weight is not a normative reason.

### Why Not Demand the Enablers To Be Apparent?

The friend of DE RE could instead try to revise the view by requiring more descriptive information to be apparent to the subject. But while this strategy can help to address some cases, it cannot solve the fundamental problem.

To bring this out, let us consider the main version of this strategy. Notice that in the case involving Jones, P's truth would be an objective reason to believe Q *because* P logically entails Q. Facts about what is a reason for what are not brute facts. P will generally be a reason for some attitude toward Q *because* further descriptive facts obtain. This is not to say that these extra descriptive facts are *further reasons* to have the attitude toward Q. They are *enabling conditions*.<sup>7</sup> So one might revise DE RE thus:

ENABLED DE RE: R is an apparent reason for S to  $\phi$  iff (i) R's truth would be an objective normative reason to  $\phi$  given enabling conditions C, (ii) it appears to S that R, and (iii) it appears to S that C.

But this view also fails: it is both too strong and too weak.

The view is clearly too strong for perceptually apparent epistemic reasons. Either intrinsic or relational features of perceptual experiences enable them to provide reasons for belief. Intrinsic features might include the presentational character of perceptual experience; relational features might include the reliable links between experience and reality. We do not need to represent these enabling conditions to form rational perceptual beliefs. This would be a serious overintellectualization of the very sort that DE RE was designed to avoid. Children and animals do not form beliefs about these features of perceptual experience. Even adults with the relevant concepts need not do so.

ENABLED DE RE is too strong even for apparent deductive reasons. Having beliefs about entailments is *one* way to be sensitive to logical relations between propositions.

---

<sup>7</sup>Cf. Dancy (2004).



But it is not the *only* way. Another way is to have the ability to competently infer one proposition from another by following an *inference rule*. Instead of reasoning

- (i)  $\neg(P \vee Q)$
- (ii) If  $\neg(P \vee Q)$ , then  $\neg P \wedge \neg Q$
- (iii) So, by *modus ponens*,  $\neg P \wedge \neg Q$ ,

one could directly infer (iii) from (i) by relying on one of the DeMorgan rules.

We cannot replace all rules by extra premises. Even axiomatic systems need some rules. And most people simply do not reason axiomatically. People often use rules rather than conditionals that express logical truths plus *modus ponens*. Admittedly, most people are not so logically competent that they can use any potential rules as easily as they use *modus ponens*. But equally clearly, there is a spectrum of acumen. Some people do have the native logical abilities to cleave reliably to far more intricate patterns.

If we accept ENABLED DE RE, we could not describe all the apparent reasons these people can acquire. Competent inferrers need not represent the patterns to which they cleave to possess premises as apparent reasons for conclusions. Yet ENABLED DE RE entails that the premises fail to be apparent reasons for these subjects' conclusion beliefs.

So, it cannot be claimed that what Jones is missing *must* be a further conditional belief. Such a belief could help. But that is not the only way his situation could improve. Rather than gaining a new conditional belief, Jones could acquire a new ability. DE RE and ENABLED DE RE fail to accommodate this possibility. This is a problem. These theories all represent a subject's set of apparent reasons by a set of descriptive considerations that are apparent to the subject. ENABLED DE RE differs from DE RE only in requiring more descriptive considerations to be apparent. This strategy is clearly misguided.

While too strong in one respect, ENABLED DE RE is too weak in another. What fundamentally matters is not how much descriptive knowledge an agent possesses.

Agents might have equal descriptive knowledge, yet more considerations might be apparent reasons for one than for another. Among subjects with normative concepts, the difference can manifest in one's having greater normative knowledge. Among subjects without normative concepts, the difference can manifest in one's being in a position to rationally treat more considerations like objective reasons. And that may owe to a difference in competence.

As we trace beings down the spectrum of competence, fewer considerations will be apparent reasons for them. But what diminishes down the spectrum is not necessarily the number of apparent descriptive facts. We can imagine beings who grasp as many descriptive facts as us but for whom fewer reasons are apparent owing to their lesser abilities. ENABLED DE RE cannot capture these cases. It predicts that the range of apparent reasons for a subject is simply a function of the number of apparent descriptive facts for that subject. By neglecting the role of competence, it makes many bad predictions.

### **Morals**

I think the proper response to these observations is to abandon the approach represented by standard *de re* theories, and to find a theory that explicitly captures the role of competence. This should not involve a retreat to DE DICTO. Given some controversial views in the philosophy of mind, this better theory might entail DE DICTO. But there is no obligation to commit to these views. We can choose a strictly weaker theory.

Might defenders of DE RE have another way to capture the role of competence? Perhaps they could try to weaken the link between apparent reasons and rationality. They might claim that it is rational for S to  $\phi$  iff (i) there are sufficient apparent reasons for S to  $\phi$  and (ii) S is in a position to competently treat them like objective reasons to  $\phi$ .

But this would be unsatisfying. Does someone with little mathematical ability have *apparent reasons* to believe the most abstruse theorems of mathematics simply in virtue of knowing the axioms? Not plausibly: owing to this person's limited abilities, he cannot rationally treat these axioms like reasons to believe the theorems. It is better to place

the competence constraint directly on the having of apparent normative reasons to  $\phi$ .

### 2.3.2 The Problem of Wouldn't-Be Reasons

So far the objection has been that the conditions in *de re* theories are *insufficient* for a consideration to be an apparent reason. Are these conditions even *necessary*?

While its importance can be overstated, condition (ii) is necessary: R cannot be an apparent reason for S to  $\phi$  if it is not apparent to S that R. But we should doubt condition (i), which entails:

(i\*) R is an apparent normative reason for S to  $\phi$  only if R's truth would give S an objective normative reason to  $\phi$ .

(i\*) is stronger than it might seem. Notice, for example, that it entails:

(i\*\*) If R is true, R is an apparent normative reason for S to  $\phi$  only if R is an objective normative reason for S to  $\phi$ .

To see this, suppose that R is true. (i\*) entails that R is an apparent reason to  $\phi$  only if its truth would be an objective reason to  $\phi$ . Since R is true, it follows that R is an objective reason to  $\phi$ . So, (i\*) entails that true apparent reasons to  $\phi$  are objective reasons to  $\phi$ .

But we should reject this claim. Suppose that it appears to some scientists that it is a law of nature that Fs are Gs, but this appearance is misleading. Every time an F appeared to be a G involved an illusion; in reality, Fs are nomically guaranteed *not* to be Gs. Now suppose that the scientists have been correctly told that X is an F, but nothing more. The fact that X is an F is an apparent reason for the scientists to believe that X is a G. But the fact that X is an F is not an objectively good reason to believe that X is a G. Objectively speaking, the fact that X is an F is a conclusive reason to believe that X is *not* a G.

There is a more general problem. We can state it as an argument against DE RE:

#### The Argument From Wouldn't-Be Reasons

1. It is possible to rationally treat a consideration R like an objective reason to  $\phi$  even if R would not in fact be an objective reason to  $\phi$  if true.
2. If it is rational for one to treat R like an objective reason to  $\phi$ , R is an apparent reason for one to  $\phi$ .
3. So, it is possible for there to be apparent reasons for one to  $\phi$  that wouldn't in fact be objective reasons for one to  $\phi$  if true.
4. If (3), DE RE is false (because condition (i) is false).
5. So, DE RE is false (because condition (i) is false).

Cases like the two just mentioned support premise (1). Premise (2) seems like a truism. The rest of the argument is straightforward. This is a serious problem for DE RE, which we might call the *Problem of Wouldn't-Be Reasons*.

This problem undermines DE RE's central explanatory resource. It is obvious that not just any apparent fact is an apparent reason to  $\phi$ . So, we ask, what makes an apparent fact into an apparent reason to  $\phi$ ? If a theory of apparent reasons cannot answer this question, it is defective. But the only answer that *de re* theorists give is that this apparent fact must be one which would be an objective reason to  $\phi$  if it were real. The Argument from Wouldn't-Be Reasons suggests that this answer cannot answer our central question.

How can a *de re* theorist respond? The only options are to (A) question premise (2)'s status as a truism, or (B) deny that cases like the case of the misled scientists support (1).

**Against Strategy (A).** Consider strategy (A) first. One might complain that all that is clear in the case of the misled scientists is that they would be rational to believe that X is a G. One might claim that this belief would be rational *not* because the fact that X is F is an apparent reason for them to believe that X is a G. It would be rational simply because it is rational for them to treat X's being F like an objective reason to

believe X is a G. And that treatment would be rational on *de re* grounds: so far every F has appeared to be G.

I doubt that this saves DE RE. DE RE still implies that the fact *that X is F* is not an apparent reason for the scientists to believe that X is a G. This is itself an intolerable implication. So the starting point is flawed: it simply isn't true that all that is clear in the case of the misled scientists is that it would be rational for them to believe that X is a G.

Still, the second part of strategy (A) needs to be addressed—i.e., the suggestion that it is rational for the scientists to treat X's being F like an objective reason to believe that X is a G on purely *de re* grounds. For this suggestion contradicts a broader moral I drew.

It can sound plausible that there are *de re* reasons why it is rational for the scientists to treat X's being F like an objective reason to believe X is a G. Isn't the fact *that all Fs have appeared to be Gs* an apparent reason to treat X's being F like an objective reason to believe that X is a G? If so, the explanation of rational treating at least remains *de re*.

My worry about this response is that it trades on tacitly assuming the presence of something beyond mere *de re* appearance—viz., competence. The crucial question is whether any beings to whom (a) all Fs have appeared to be Gs would be (b) at least *prima facie* rational to treat the apparent fact that X is F like an objective reason to believe that X is a G, where (b) is explained wholly by (a). The answer to this question is 'no'. There are surely possible beings who lack the competence to project inductively from past experience. For such beings, all Fs could appear to be Gs, and yet it could fail to be even *prima facie* rational for these beings to treat the fact that a new X is F like an objective reason to believe that X is a G. The explanation is that they lack a kind of competence.

This echoes an earlier point. Any adequate theory of apparent reasons must honor the role of competence. Appealing to *de re* appearances of objective reasons cannot alone secure this desideratum. Hence, the deeper moral from the Problem of Wouldn't-Be Reasons stands.

**Against Strategy (B).** Strategy (A) fails. What about strategy (B), which involves denying that the case of the misled scientists supports premise (1)? This strategy is even less plausible. There are only two ways to pursue it. One way is to deny that it is rational for the scientists to treat the fact that X is F like an objective reason to believe that X is a G. This is clearly wrong. Another way would be to hold that the fact that X is F is, after all, an objective reason to believe that X is a G. But this is false by stipulation!

One might try to dispute the stipulation. But once we get clear on what it does not exclude, there is no clear reason to do so. We did not stipulate that there is no other objective reason for the scientists to believe that X is a G. The claim was that the fact that X is an F is not an objective reason to believe that X is G. This is the only claim that needs to be made. It entails that an apparent fact can be an apparent reason even if this apparent fact would not be an objective reason if it were real. We could still agree that the fact that all Fs have appeared to be Gs is a non-trivial objective reason to believe that X is a G. But this does not undermine my point, which concerned the fact that X is an F.<sup>8</sup>

There is no good solution, then, to the Problem of Wouldn't-Be Reasons. Given this problem and the Problem of Unapparent Reasons, I doubt any pure *de re* theory will do.

## 2.4 Competence, Appearance, and Treating

### 2.4.1 A Key Ingredient: Treating

A better theory is worth seeking. An ideal theory would avoid the problems for *de re* theories without regressing to the excesses of *de dicto* theories.

How can one walk this fine line? Well, remember that it is possible to treat a consideration like an objective reason without having any normative concepts or beliefs. A natural thought is to invoke this less demanding idea in a new theory. We could

---

<sup>8</sup>This forestalls application of the strategy from Lord (2010).

construct a theory that is structurally like DE DICTO, but replace *beliefs about objective reasons* with *competent attractions to treat considerations like objective reasons*.

The Problem of Unapparent Reasons recommends this approach. The problem arose because it is possible for a proposition whose truth would provide an objective reason to appear true without one's being in a position to rationally treat its apparent truth like an objective reason. It makes sense to construct a theory that honors this thought. It would, however, be hasty to construct a theory *just* on the basis of this thought. The following theory will not do:

RATIONALLY ATTRACTED TO TREAT (RAT): R is an apparent normative reason for S to  $\phi$  iff (i) it appears to S that R, and (ii) S is rationally attracted to treating R like an objective normative reason to  $\phi$ .<sup>9</sup>

As a necessary biconditional, RAT is plausible. But as a substantive theory, it is unsatisfying. Like other theorists, I want to explain rationality in terms of apparent reasons. RAT looks circular from this perspective.

Still, RAT is extensionally plausible. So it would be nice to find an extensionally equivalent theory that does not appeal to rationality.

### 2.4.2 Another Key Ingredient: Competence

To see what such a theory would look like, remember the role of competence. If Jones's logical competence is limited, then even if it is apparent to Jones that P and P logically entails Q, it will not follow that P is an apparent reason for him to believe Q. As I have stressed, what is at bottom needed to gain apparent reasons is greater competence, not just further descriptive beliefs or appearances.

How is competence here to be understood? If objective reasons are the basic normative items—as I've assumed with Parfit, Scanlon, et al.—competence here must be analyzed in terms of objective reasons and non-normative items. I think the competence to invoke is a competence

---

<sup>9</sup>Here and elsewhere, *being attracted to treat X like an F* = *feeling the pull to treat X like an F*.

to treat considerations like objective reasons of some relevant kind

only if

these considerations would, if true, be objective reasons of this kind

Call dispositions with this structure *objective reasons-sensitive competences*. Such competences need not be *infallible*. Something might have a disposition to  $\phi$  even if it does not always  $\phi$ . This does not mean that the disposition is a disposition to do something weaker than  $\phi$ -ing. It just means that it is a less than infallible disposition to  $\phi$ .

To see why such competences matter, consider again the beings that lack the competence to inductively generalize from experience. Even if P's truth would inductively support believing Q and it is apparent to one of these beings that P, this being does not *ipso facto* have an apparent inductive reason to believe Q. To gain P as an apparent inductive reason to believe Q, what else does this being need? What this being needs is a disposition

to treat considerations like objective inductive reasons for belief

only if

these considerations would, if true, be objective inductive reasons for belief

Together with the appearance that P, is this competence all that one needs to gain P as an apparent inductive reason to believe Q?

Not quite. To see this, consider Julia. Julia is highly inductively competent. Right now, the appearance that R is attracting Julia to inductively infer that S. But on this rare occasion, Julia's attraction does not *manifest* her inductive competence. While Julia is biased only with respect to a single topic, R and S concern this topic. We can even suppose that Julia's bias by sheer luck gets it right here: totally unbeknownst to Julia, R does greatly raise the objective probability that S. Still, the mere fact that Julia (a) is attracted to inductively infer S on the basis of the apparent fact that R and (b) possesses inductive competence doesn't (c) convert R into an apparent normative



reason for Julia to believe S. What is needed for (c) is for Julia's attraction to *manifest* her competence.

We can generalize from this point and earlier points to get a better theory:

COMPETENTLY ATTRACTED TO TREAT (CAT): R is an apparent reason for S to  $\phi$  iff (i) it appears to S that R, (ii) S is attracted to treating R like an objective reason to  $\phi$ , and (iii) this attraction manifests S's relevant reasons-sensitive competence, where a relevant reasons-sensitive competence is a competence to treat R-like considerations like objective reasons to do  $\phi$ -like things only if they are objective reasons to do  $\phi$ -like things.<sup>10</sup>

CAT is attractive. It avoids the Problem of Unapparent Reasons with its appeal to treating and competence, and hence captures the main advantage that DE DICTO has over DE RE. But it does so without any overintellectualizations.

### 2.4.3 CAT and the Problem of Wouldn't-Be Reasons

It is worth pausing, however, to explain why CAT also does not invite the Problem of Wouldn't-Be Reasons. There are two reasons why it does not.

**Reason 1: The Fallibility of Competence.** First, CAT does not assume that competences are infallible. One might be competently attracted to treat some consideration like an objective reason to  $\phi$  even if the truth of this consideration would not be an objective reason to  $\phi$ . This does not mean that the competence is a competence to do something weaker than  $\phi$ -ing. It is just a less than fallible competence to  $\phi$ .

---

<sup>10</sup>I use 'do' in a broad sense here to refer to attitudes as well as actions.

**Reason 2: The Competence/Performance Distinction.** Secondly and more importantly, there is a deeper *competence/performance distinction*.<sup>11</sup> Having the competence to succeed does not entail being in a position to reliably succeed if one tries. A competence is a *disposition* to succeed. Just like other dispositions, there are favorable and unfavorable conditions for its successful manifestation.

Unfavorable conditions do not destroy competence. An archer retains the competence to hit the bull's-eye even when surprise gusts prevent her from being in a position to do so. An agent with the same competence could be relocated to a systematically unfavorable environment and be hoodwinked about its favorableness by a trickster. Consider a color expert who is systematically tricked by unusual lighting induced by a trickster. Her expertise is not destroyed, though she will be unreliable. The expertise may even be *displayed* in a sense weaker than proper manifestation. Compare how a car can exhibit its virtues for the road even in a display room.

All these points apply to objective reasons-sensitive competences. We find an initial illustration in the case from Williams (1981) where Bernie gets duped in his favorite bar, receiving a glass of petrol and tonic that is visually indiscriminable from the gin and tonic that he typically orders. Bernie's prudence is not undermined by the fact that he is about to drink some petrol. Indeed, assuming that Bernie would benefit from the drink that he requested, his prudence is *displayed*, though not strictly manifested. Conditions are simply unfavorable: through no fault of his own, he does something bad for him.

It is easy to imagine an agent who systematically lands in unlucky circumstances through no fault of her own. Like other dispositions to succeed, reasons-sensitive competences may not even guarantee *actual reliability*. One could competently treat a vast range of apparent considerations like objective reasons even if they wouldn't be objective reasons if true. All we need is for the circumstances to be both *unfavorable* and *misleading*.

---

<sup>11</sup>In what follows, I am indebted to Ernest Sosa—in particular, to Sosa (2010)'s illuminating discussion of competence, and Sosa (1991, 1993)'s influential take on the new evil demon problem for externalist accounts of justification. What I will say simply generalizes what Sosa has already said in epistemology.

Appreciating these divergences between competence and reliability *simpliciter* lets us avoid the Problem of Wouldn't-Be Reasons. We see that it is possible for one to competently treat many considerations like objective reasons even if they would not be objective reasons if true. This is why many apparent reasons need not be objective reasons if true.

Although CAT nicely avoids the Problem of Wouldn't-Be Reasons, it does establish a stronger connection between rationality and objective reasons than many have assumed. While it correctly denies that all apparent reasons are would-be objective reasons, it does entail a weaker conditional—viz., that:

if conditions for the exercise of competence are favorable, then considerations that are competently treated like objective reasons are likely to be objective reasons

Since competences are not infallible even in favorable circumstances, it does not follow that any consideration that is competently treated like an objective reason *is* an objective reason. But it is likely to be, in an objective sense: in favorable conditions, the competence has a propensity to yield treatings that are correct.

There would be a problem if it were possible for (i) a consideration R to satisfy the conditions in CAT while (ii) it is not rational for the subject to treat R like an objective reason. But there is no compelling reason to believe this is possible. In every case in which it is not rational to treat R like an objective reason, it will be implausible that one's treating it like an objective reason can display the relevant competence. Consider an archer who is rationally misled into thinking there will be gusts that require shooting at a bizarre angle. If the archer did not shoot at this angle given this evidence, it would be implausible to think that she would display her competence. Still, there are conclusive objective reasons for her not to take the angle shot. After all, there will not be gusts!

I say the stronger connection with objective reasons that CAT establishes is a virtue. While rationality does not guarantee reliability, the *capacities* that rational subjects exercise in bad cases are the same capacities that actually reliable people exercise. The differences lie in the favorableness of the circumstances and our sensitivity to them. While we must recognize a notion of rationality that falls in between coherence and

correctness, rationality is not intelligible independently of correctness. It is a competence to achieve correctness, however fallible and unhelpful in unfavorable circumstances.

#### 2.4.4 Further Analysis of Key Ingredients

CAT is the best theory we've seen. It steals the virtues of DE RE and DE DICTO while avoiding their vices. But crucial concepts of the theory deserve further elucidation.

##### More on Treating and Attraction

Consider treating again. It is, I have stressed, intuitive that we can treat considerations like objective reasons without having any normative concepts. Still, what exactly is it to treat P like an objective normative reason to  $\phi$ ?

To treat P like an objective normative reason to  $\phi$  is to be disposed to respond to the appearance that P in all or at least most of the ways that would be favored if P were an objective normative reason to  $\phi$ . Why 'all/most' rather than 'some'? My reason for imposing this stronger requirement is to get the right results in cases of supposition and pretense. We do not want the relevant kind of treating to be compatible with pretense.

If you pretend for the sake of argument that P is a conclusive objective reason to  $\phi$ , you are not irrational if you fail to  $\phi$  given the appearance that P. But if you competently treat P like a conclusive objective reason to  $\phi$  in the sense relevant to rationality, you are irrational if you fail to  $\phi$  given the appearance that P. The fundamental difference is straightforward. If you merely pretend or suppose for the sake of argument that P is a conclusive objective reason to  $\phi$ , there will be *many ways* that you are not disposed to respond to the appearance that P which would be correct P if P were a conclusive objective reason to  $\phi$ . You will not, for example, be disposed to adopt the stable attitudes that would be appropriate if P were a conclusive objective reason to  $\phi$ .

If this is what treating involves, why does CAT say only that apparent reasons are considerations that one is *attracted* to treating like objective reasons? Why not require that apparent reasons be considerations that one actually treats like objective reasons?

Here is the main reason for requiring less. Plausibly, weak-willed agents can have conclusive apparent reasons to  $\phi$  and yet fail to be motivated to  $\phi$ . If CAT required

actual treating, it would imply that such agents are impossible. Since this kind of agent seems possible, we want to require less. Nevertheless, when one is weak-willed in the face of an apparent reason, one is not left entirely cold. Real akratic agents will at least feel some pull to treat the apparent reason like an objective reason, if it really is an *apparent* reason for them. The trouble with the akratic agent is that her feeling this pull is not sufficient to make her treat the relevant consideration like an objective reason to  $\phi$ .

So the weaker requirement is apt. Why does CAT require even this much? Couldn't there be a sociopath who has conclusive apparent reasons to  $\phi$  but who is left entirely cold, failing even to feel the pull to treat these considerations as normative reasons to  $\phi$ ?

I find it implausible that this particular kind of sociopath is possible. Obviously, there are conclusive objective reasons for these sociopaths not to act as they do. But we are talking about *apparent* normative reasons. I don't see how a consideration could be an apparent normative reason for someone to  $\phi$  if the apparent truth of this consideration doesn't even make her attracted to treat her situation like one in which there is an objective reason to  $\phi$ .<sup>12</sup> It seems better to describe sociopaths as lacking the ability to acquire certain objective reasons as apparent reasons—just as CAT implies.

"But can't it be apparent to sociopaths that what they are doing is morally wrong?" Of course. But CAT doesn't imply otherwise. What it implies is that the appearance that something is morally wrong fails to give the sociopath an apparent normative reason. What is apparently immoral needn't be equally apparently disfavored by objective reasons. Morality may seem to the sociopath just like etiquette seems to the enlightened. Indeed, even in the more substantive sense of 'rational' that falls in between coherence and correctness, it sounds clearly false to deem sociopaths irrational.<sup>13</sup> Yet this kind of rationality consists in correctly responding to apparent normative reasons. What sociopaths lack are, precisely, conclusive apparent normative reasons not to do

---

<sup>12</sup>Once again, 'is attracted to  $\phi$ ' just means 'feels the pull to  $\phi$ '. It can sound odd to use talk of attraction in the context of sociopaths and morality. But given what the language means, this is no objection.

<sup>13</sup>Cf. Williams (1981).

what they do. CAT's predictions here are plausible.

### More on Competence

Another crucial notion in CAT that deserves further comment is the notion of competence. Earlier I said that the competence needed is a competence to:

to treat considerations like objective reasons of some relevant kind

only if

these considerations would, if true, be objective reasons of this kind

I have already talked at length about why requiring objective reasons-sensitive competences in a theory of apparent reasons does not lead to a version of the Problem of Wouldn't-Be Reasons. Still, questions remain about these competences.

A central question concerns their individuation. It is clear that to have an apparent deductive reason to believe  $Q$ , it is not necessary that a subject be competent with every kind of deductive inference. But if 'reasons of the relevant kind' is understood in a coarse-grained way, CAT might seem to make this crazy demand: the relevant kind of reason might then just be *deductive reason for belief*.

Accordingly, we should want a more fine-grained picture. But we don't want too much fineness of grain. Consider the disposition to:

(a) treat the apparent conjunctive fact that Bob is 6 feet tall and Mary is 5 feet tall like an objective reason to believe that Bob is taller than Mary

This disposition is a competence: the reality of the apparent fact in question would be an objective reason to believe that Bob is taller than Mary. But we can imagine possible people who have this disposition but lack related dispositions, such as the disposition to:

(b) treat the apparent conjunctive fact that Sally is 5 feet tall and Billy is 4 feet tall like an objective reason to believe that Sally is taller than Billy

Consider, indeed, a subject who has no further dispositions of this form. Call him Weirdo. The only disposition Weirdo has is (a). Weirdo has the reverse of disposition (b): he will treat the apparent conjunctive fact mentioned in (b) like an objective reason to disbelieve that Sally is taller than Billy. Weirdo fails to manifest any grip on *why* the apparent conjunctive fact in (a) would be an objective reason to believe that Bob is taller than Mary.

We do not want to respond to this observation by requiring that a subject must have *beliefs* about why a consideration would be an objective reason in order to possess it as an apparent reason. This would be a gross overintellectualization. We can have competences to respond to certain sorts of reasons without having detailed knowledge of what makes them good reasons. So, what we want to require is that a subject have competence beyond the trivially narrow competence to treat one highly specific reason like an objective reason.

How should we require more? Note that there is a broader inferential pattern to which Weirdo fails to be sensitive:

(A) X is  $m$  units tall.

(B) Y is  $n > m$  units tall

(C) Y is taller than X.

Of course, not all objective epistemic reasons qualify as such in virtue of being subsumed under necessarily truth-preserving inference patterns. But even ampliative epistemic reasons and non-epistemic reasons sometimes fit into patterns. The fact that all observed ravens are black is an objective reason to believe that all ravens are black precisely because facts of the form *all observed Fs are Gs* are generally objective reasons to believe conclusions of the form *all Fs are Gs*. The fact that it is healthy to eat vegetables is an objective reason to intend to eat them precisely because facts of the form *it is healthy to eat X* are generally objective reasons to form intentions with contents of the form *I eat X*.

When reasons fall into patterns, we can better understand the competences invoked by CAT. When there is a good pattern of reasoning that subsumes the subject's reasoning, the competences will be competences to treat considerations like objective reasons in accordance with some relevant pattern.<sup>14</sup> Not all such reasoning must be inferential. Suppose that I see that there is a red sphere before me. If I form the belief that there is a red sphere before me in response to this visual experience, I believe for a reason. But this instance of believing for a reason is not inferential. Clearly, though, there is a general pattern of good reasoning that my transition instantiates: namely, treating the fact that I have an experience with a certain content like an objective reason to believe that content.

So CAT can be modified:

(CAT\*) R is an apparent reason for S to  $\phi$  iff

(i) it appears to S that R,

(ii) S is attracted to treating R like an objective reason to  $\phi$ ,

(iii) this attraction manifests S's competence to treat considerations like objective reasons to do  $\phi$ -like things only if they are objective reasons to do  $\phi$ -like things,

and

(iv) if there are patterns of good reasoning that subsume the favoring relation between R and  $\phi$ -ing, S is competent with some and manifests this competence.

CAT\* only requires that the subject be competent with *some* general pattern of reasoning *if* there is any such pattern. To see why this is right, note that there are often several good patterns of reasoning from R to  $\phi$ -ing. The fact that the glass contains orange juice is an objective reason to drink on two counts: it would be healthy, and it would be delicious. If I am attracted to treat this fact like an objective reason to drink only as a manifestation of my disposition to drink healthy beverages, I still have an apparent reason to drink if I see that this glass contains orange juice. Accordingly, if my reasoning instantiates only one pattern of competent reasoning, that is enough.

---

<sup>14</sup>Here and elsewhere, I use 'reasoning' in a broad sense to mean *responding to (possibly bad) reasons*.



Some might wonder whether there will always be broader patterns of good reasoning to which we can appeal. But I have not assumed that there always will be broader patterns. I have only assumed that when there are broader patterns, the subject must be competent with some of them. Sure, not every good piece of reasoning will be subsumed by some interesting broader pattern. But it would be implausible to deny that there are ever any interesting broader patterns. When there are, competence requires proficiency with some of them. How else can we explain what Weirdo was missing, for example?

#### 2.4.5 Taking Stock

With these clarifications in place, I think we should treat CAT\* like what it is: a superior theory to the others on offer. Let us briefly take stock of why.

CAT\* avoids overintellectualization without inviting the Problem of Unapparent Reasons. This is achieved by its appeal to the notion of treating something like an objective reason. CAT\* avoids overintellectualization because such treating does not require the possession of any normative concepts. But it still avoids the Problem of Unapparent Reasons, because it is natural to think that if one is attracted to treat R like an objective reason, R is an apparent *reason* and not merely a frictionless apparent fact.

We do not want just *any* considerations that subjects are attracted to treat like objective reasons to count as apparent normative reasons. Suppose I am attracted to treat P like an objective reason to believe Q *via* some wildly overgeneralizing incompetence, where I am unaware that I have this incompetence. I do not gain a serious apparent reason to believe Q. CAT\*'s appeal to competence addresses this problem for DE RE. Moreover, it does so without appealing to further primitive normative concepts beyond the concept of an objective reason. After all, we analyzed the competences to which CAT\* appeals in terms of dispositions to treat objective reasons like objective reasons.

Once again, this does not lead us back to the Problem of Wouldn't-Be Reasons. There is a gap between competence and performance. If one is excusably misled about the favorableness of one's circumstances, one might retain a competence to  $\phi$  but not be in a position to reliably  $\phi$ . This offers a way to diagnose what is happening in the cases

that support distinguishing between rationality and correctness. Consider Williams's petrol and tonic case. Bernie does possess an objective reasons-sensitive competence: he is just excusably misled about his circumstances, like the great archer who receives strong misleading evidence that there will not be gusts each time she tries to shoot.

Certainly, the connection between apparent reasons and objective reasons is stronger according to CAT\* than according to theories like DE DICTO. But this is a virtue, not a vice. It is true that rationality does not guarantee actual reliability. But it is not true that the capacities that rational subjects exercise are different from the capacities that many actually reliable subjects exercise. The capacities are the same: what differ are the favorableness of the circumstances and the subject's sensitivity to them. Even if we recognize a notion of rationality that falls between coherence and correctness, we do not need to think that rationality is divorced from correctness: it is a competence to achieve correctness, however fallible and unhelpful in unfortunate circumstances.

## 2.5 But Don't Envelopes Slice All the Theories?

CAT\* fares better than DE DICTO and DE RE. It does, however, share with both a feature that some find objectionable. Specifically, all three theories entail that apparent reasons are *apparently objective* reasons. These theories simply understand the 'apparently' in different ways. Call this the *Thesis of Apparent Objectivity* (TAO).

Some believe there is a decisive reason to reject TAO: namely, that we can rationally do what we know there is not most objective reason to do. This objection to TAO can be illustrated with a case from Jacob Ross:

*Three Envelopes.* Chester can choose one of three envelopes. He is correctly told by a reliable informant that there is \$800 in Envelope 1. He is also correctly told there is \$1000 in either Envelope 2 or 3, and that the envelope that lacks it is empty. But he cannot learn anything else about which might contain it (and knows this).<sup>15</sup>

---

<sup>15</sup>See Ross (2012).

Cases like Three Envelopes can seem to pose a problem for TAO. Chester ought rationally to pick Envelope 1. But Chester can know that there is more objective reason for him to do otherwise. For he can know that either (i) there is more objective reason to choose Envelope 2 or (ii) that there is more objective reason to choose Envelope 3. After all, there is \$200 more in one of them. So, Chester ought rationally to make a choice that there is apparently insufficient objective reason to make.

Is that inconsistent with TAO? Many automatically assume that it is. But it is not. We have only discussed theories about what it is for a consideration to be *an* apparent normative reason to  $\phi$ . They are not theories about the *comparative weight* of apparent reasons. Maybe, then, there are ways of weighing apparent reasons that (i) preserve the thought that all apparent reasons are apparently objective (i.e., TAO), but also (ii) affirm that there is sufficient apparent reason for Chester to choose Envelope 1.

Consider:

NO SPECIFIC DOMINANCE (NSD): R is a sufficient apparent reason to  $\phi$  iff

(i) R is an apparently objective reason to  $\phi$ ,

and

(ii) there is no consideration other than R that is apparently a stronger objective reason to pursue any *specific alternative* to  $\phi$ -ing,

\* where specific alternatives include things like *taking Envelope 2* but not things like *not taking Envelope 1*.

NSD is compatible with TAO. NSD entails that a sufficient apparent reason must be an apparently objective reason, and hence entails TAO.

NSD is plausible. Suppose that I am  $\phi$ -ing, and I know that I could never discover a better objective reason to pursue some specific alternative to  $\phi$ -ing. Suppose moreover that I know that there is a good objective reason to  $\phi$ . It seems natural for me to conclude that I am rationally permitted to continue to  $\phi$ . What else am I supposed

to do? I cannot rationalize pursuing any specific alternative to  $\phi$ -ing, because I lack any apparently stronger objective reason to pursue any specific alternative. Moreover, I know that there is at least an objective reason to  $\phi$ . This is all compatible with my knowing *generically* that there is something better I could do, if I cannot discover the *specific thing* that is better.

Similar points apply to Three Envelopes. The fact that there is \$800 in Envelope 1 is a strong objective reason for Chester to choose it. This fact is apparent to him. Moreover, the generic fact that there is an envelope that contains \$1000 fails to be an apparent objective reason for Chester to pursue any *specific alternative* to choosing Envelope 1. *Not choosing* Envelope 1 is not a specific alternative. There are two ways not to choose, and no apparent objective basis for discriminating. If Chester's apparent reason to choose Envelope 1 is insufficient, there must be an apparent objective basis for discrimination.

This is plausible, and it supports NSD. Isn't *withholding* from choosing any envelope a specific available alternative? Yes. But there are decisive objective reasons not to withhold: viz., the fact that there is \$800 within Chester's reach which he will lack if he withholds forever and the fact that he cannot know before acting which of the other two envelopes contains the \$1000. These are not just apparent facts. They are objective facts. The second fact is a fact *only because* Chester is ignorant. But it is still an objective fact!<sup>16</sup>

So NSD predicts that Chester has a sufficient apparent reason to choose Envelope 1. This is because Chester's apparent objective reason to choose Envelope 1 is not outweighed by any apparent objective reason for any specific alternative to Envelope 1.

---

<sup>16</sup>The fact that ignorance has objective effects is often ignored. But it matters: it destroys the tendency to assume that real rather than merely apparent reasons must be ones we would have *if we knew all the facts*. Dancy (2000: 69) anticipated this point: "Assessing rationality relative to incomplete information does not require us to think that the rationality we are assessing is subjective in any damaging sense. . . . It is worth remembering in this connection that the fact that I do not know something can itself be a reason. Suppose that my information is limited, and I know this, and that I have no time to make further enquiries. The action I have most reason to do might not be the one that I would have had most reason to do if I had had more time to find out more of what I needed to know. This should remind us of Prichard's example of slowing down at a junction. In that case, the objectivist seemed perfectly able to say that the main reason for slowing down is that one does not know whether any traffic is coming. If one had known that there was no traffic coming, perhaps one would have had no reason to slow down."

This prediction is plausible. And it is compatible with TAO.

Cases like Three Envelopes do not undermine TAO. Since TAO is plausible and CAT\* is the best theory on offer that entails it, we can retain our confidence in CAT\*.

## 2.6 Concluding Remarks

I turn to some loose ends and concluding remarks.

I have argued against two major types of views about apparent reasons and presented an alternative that absorbs their virtues while avoiding their vices. I have also showed that the larger family to which all these theories belong—namely, theories that entail TAO—are not undermined by cases like Three Envelopes.

One might be left wondering whether there are even better theories that fall outside this class. As far as I know, the only other obvious alternative is a *dualist* one that takes as normatively primitive both objective reasons and apparent reasons (or something in the ballpark).<sup>17</sup> According to such theories, the best that we can do is point to connections that apparent reasons bear to other “subjectively normative” phenomena like rationality. Since we ought to avoid multiplying normative primitives, my approach is preferable.

Another kind of theory that I have not explicitly considered is a theory on which rationality is to be analyzed in terms of possessed normative reasons. The status of this theory depends upon the status of what Schroeder (2008) calls the Factoring Account. On the Factoring Account, possessing a normative reason is a matter of (i) there objectively existing a normative reason and (ii) one’s bearing some possession relation to this pre-existing reason. Schroeder rejects this account and takes the idea of a possessed reason as a further normative primitive. This is, of course, just a form of dualism.

But if the Factoring Account is true, the possessed reasons theory does diverge from all the theories discussed so far. Originally I suggested that the ‘apparent’ in ‘apparent reasons’ is not a success term: something can be an apparently objective reason in spite of not being an objective reason. But if both the Factoring Account and the possessed

---

<sup>17</sup>See Ross (2006).

reasons theory are true, the ‘apparent’ is a success term.

This is a vice. Apparent reasons need not be objective reasons or would-be objective reasons. This was the lesson of the Problem of Wouldn’t-Be Reasons. Recall again the case of the misled scientists. It appears to the scientists that it is a law that Fs are Gs, but this appearance is misleading. Every time an F appeared to be a G involved an illusion; in reality, Fs are nomically guaranteed to be non-Gs. The scientists have been correctly told that X is an F, but nothing more. The fact that X is an F is an apparent reason for the scientists to believe that X is a G. But the fact that X is an F is not an objective reason to believe that X is a G. That fact is an objective reason to believe that X is *not* a G.

This is compatible with affirming that there is *an* objective reason for the scientists to believe that X is a G. The claim was only that the fact *that X is an F* is not the objective reason. This is the only claim that needs to be made: it alone shows that there can be apparent reasons that would not be objective reasons if true.<sup>18</sup> After all, it is true that X is an F in this case, but this truth is not an objective reason to believe that X is a G.

Compatibly with all this, one might say that the fact *that all Fs have appeared to be Gs* is itself an objective reason to believe that X is a G. I see no need to resist this claim in worlds where there is an otherwise generally reliable connection between appearance and reality. This doesn’t undermine the point, which concerned the status of the fact *that X is F*. Anyhow, we could always imagine skeptical scenarios where appearances are not reliably connected with the facts. In these worlds, there need not be decent objective reasons for beliefs or other attitudes that are perfectly rational.

We might also agree that possessed objective reasons have a role to play *other* than analyzing rationality. Indeed, I think this is true. I think we ought to distinguish between *justification* and *rationality*. Justification also falls short of objective rightness, but it is stronger than rationality: truly justifying reasons must be objective reasons.

---

<sup>18</sup>Lord (2010) does not address this worry.

This leads to a picture on which three reason-involving phenomena must be distinguished:

*Correct*  $\phi$ -ing, which consists in  $\phi$ -ing in conformity with all objective reasons.

*Justified*  $\phi$ -ing, which consists in  $\phi$ -ing in compliance with the objective reasons that one possesses.

*Rational*  $\phi$ -ing, which is  $\phi$ -ing in compliance with all one's apparent reasons.

Assuming that possession can itself be analyzed without appeal to any normative items other than objective reasons, this threefold distinction violates no canons of parsimony.

Possession can be so analyzed. Indeed, I would recommend an account of possession that exploits some of the same concepts used in CAT\*. Following Ernest Sosa, we should distinguish between *competence* and *aptness*. One  $\phi$ s competently when one's  $\phi$ -ing manifests a competence to succeed, relative to some standard of success. One can succeed and possess competence, but fail to succeed *due to* one's competence. Apt  $\phi$ -ing is then  $\phi$ -ing whose success manifests competence. We have considered attractions to treat considerations like objective reasons, and the competences at issue will be defined just with objective reasons and dispositions. The attractions will be:

*right*, when the considerations that attract are in fact objective reasons

*competent*, when they manifest dispositions to treat only objective reasons like objective reasons, and

*apt*, when their rightness manifests competence (in the earlier senses).

If being *competently attracted to treat* things like objective reasons is what we need to analyze *apparent reasons*, being *aptly attracted to treat* things like objective reasons is what we need to analyze *possessed objective reasons*. Aptness entails competence, so all possessed normative reasons are objective reasons. This sits well with the Factoring Account.

This is an extension of the approach that I lack space to further develop here. Apparent reasons were the focus, and there are powerful reasons to doubt that apparent

reasons are possessed objective reasons. But it is a virtue of my overall approach that it extends naturally to help us understand possessed objective reasons. Even if one does not accept my view, this is more evidence that my overall approach is worth taking seriously.



## Chapter 3

### Rationality and Justification: Reasons to Divorce?

#### Overview

Many epistemologists treat rationality and justification as the same thing. Those who don't lack detailed accounts of the difference, leading their opponents to suspect that the distinction is an *ad hoc* attempt to safeguard their theories of justification. In this paper, I offer a new and detailed account of the distinction. The account is inspired by no particular views in epistemology, but rather by insights from the literature on reasons and rationality outside of epistemology. Specifically, it turns on a version of the familiar distinction in meta-ethics between possessing *apparent* normative reasons (which may be merely apparent) and possessing *objective* normative reasons. The paper proceeds as follows. In §1, I discuss the history of indifference to the distinction between rationality and justification in epistemology and the striking contrast with meta-ethics. I introduce the distinction between apparent reasons and possessed objective reasons in §2 and provide a deeper basis for it in §3. I explain how the ideas extend to epistemology in §4 and explore the upshots for some central issues in §5.

#### 3.1 Indifference

##### 3.1.1 A History of Indifference

Examples abound of epistemologists treating rationality and justification as the same thing. In a classic attack on reliabilism, Stewart Cohen tells us:

‘[R]easonable’ and ‘rational’ are virtual synonyms for ‘justified’.<sup>1</sup>

---

<sup>1</sup>Cohen (1984: 283).

Michael Huemer writes:

Another word for what is justified...is 'rational'.<sup>2</sup>

Somewhat more cautiously, Richard Fumerton says:

The metaepistemological project I am interested in concerns the concept of justified or rational belief.... [T]he expression 'rational' might be somewhat less misleading than the expression 'justified', but I will continue to use the two terms interchangeably....<sup>3</sup>

Laurence Bonjour affords another example in discussing Samantha the unwitting clairvoyant: "[Her] irrationality...prevent[s] her belief from being epistemically justified".<sup>4</sup>

There are more recent examples. Sinan Dogramaci writes in a forthcoming paper:

Rationality, justification, reasonableness: same thing. Use whichever word you like.<sup>5</sup>

And in a recent defense of the accessibility of justification, Declan Smithies tells us:

To say that one has justification to believe a proposition is to say that it is rational or reasonable for one to believe it.<sup>6</sup>

The indifference also arises in less explicit ways. Often it surfaces in expressions like 'rationality or justification', as used in the following sentence from Ralph Wedgwood:

I propose that it is a necessary condition on the rationality or justification of your current enduring belief-states that they should meet certain conditions of coherence.<sup>7</sup>

---

<sup>2</sup>Huemer (2001: 22).

<sup>3</sup>Fumerton (1995: 19); a search for 'rational' in the index yields: 'Rational belief. *See* justified belief'.

<sup>4</sup>BonJour (1985: 39)

<sup>5</sup>Dogramaci (forthcoming: 3).

<sup>6</sup>Smithies (2012: 274).

<sup>7</sup>Wedgwood (2012: 280).

Just do a Google search for ‘rationality or justification’ or ‘justification or rationality’. You get many hits. Revealingly, most are from epistemologists in different generations who are treating the verbal difference as marking no real distinction.

### 3.1.2 Why Indifference Is Striking

The indifference would be unremarkable if it were paralleled in other normative sub-disciplines. But it isn’t. In ethics, many doubt whether the fact that it is rational for S to  $\phi$  entails that there are truly good reasons for S to  $\phi$ . Many defend views on which the entailment fails. Derek Parfit, for example, holds that rationality consists in correctly responding to *apparent* normative reasons, where ‘apparent’ is not a success term: “We ought rationally to respond to apparent reasons even if...these reasons are not real.”<sup>8</sup> And Parfit’s view is just one member of a family of views that analyze rationality in terms of apparent or subjective normative reasons, in contrast to objective normative reasons.<sup>9</sup> On many views, an apparent normative reason is not a special kind of objective normative reason.<sup>10</sup>

Other meta-ethicists doubt that there are always genuine reasons to comply with requirements of rationality. John Broome writes: “I doubt that, necessarily, we ought to satisfy each of the individual requirements of rationality. Indeed, I doubt that, necessarily, we have any reason to satisfy each of these requirements.”<sup>11</sup> And Niko Kolodny argues that there is no genuine reason to be rational *as such*, debunking intuitions to the contrary with an error theory. This error theory rests essentially on an account of rationality on which it consists in heeding what *appear to be* good reasons.<sup>12</sup>

Admittedly, no discontinuity between epistemology and ethics follows *immediately* from these observations. This is because talk of justification is somewhat uncommon in contemporary ethics. Normative reasons have become the most popular currency for

---

<sup>8</sup>Parfit (2011: 111).

<sup>9</sup>See also Parfit (2001), Schroeder (2008, 2009, 2011) and Way (2009).

<sup>10</sup>One exception is Lord (2010), whose view I’ll discuss in §2. Reflections on Lord’s view provide the basis for my distinction between apparent reasons and possessed objective reasons.

<sup>11</sup>Broome (2005: 321).

<sup>12</sup>See Kolodny (2005), whose view was inspired by Scanlon (1998).

cashing out normative claims. When one asks what normative reasons are supposed to be, one usually hears that they are considerations that *count in favor* of acts or attitudes.

Still, many ethicists would regard justification as going hand in hand with the possession of real rather than merely apparent normative reasons. Normative reasons are often explicitly agreed to be the kinds of things which can help to justify our attitudes and acts if we possess them. Indeed, many people use the motivating/normative distinction interchangeably with the motivating/justifying distinction.<sup>13</sup>

Moreover, the main context in which talk of justification arises in contemporary ethics is when people are distinguishing between justifications and excuses. In recent work, distinctions involving reasons draw the line. John Gardner, for example, treats justifications for  $\phi$ -ing as consisting in normative reasons to  $\phi$  and excuses for  $\phi$ -ing as consisting not in normative reasons to  $\phi$  but rather in the appearance that there were reasons to  $\phi$ .<sup>14</sup>

Now, there was a time when excuses were understood merely as *denials of responsibility*.<sup>15</sup> But this was a mistake. Excuses can express our responsibility, by showing that we rationally reacted to the sadly misleading appearance that there were justifying reasons. As Gardner writes, “in this respect the action remains rational, and the agent who offers it claims rational competence. And this, in turn, is where an excuse differs fundamentally from a denial of responsibility.”<sup>16</sup> But a difference with justification remains, because

it is one thing to have a reason to defend oneself and quite another to have every reason to believe that one has a reason to defend oneself that in reality one does not have (e.g. because one strayed accidentally and without warning onto the set of an action movie)... In that case the most we can hope for is an excuse.<sup>17</sup>

---

<sup>13</sup>Lenman (2009) provides an overview of the literature on the motivating/normative contrast that uses the term ‘justifying reason’ instead of ‘normative reason’. Before the 2000s, this was more common. Dancy (2000: 6–7) takes exception to the identification of normative and justifying reasons, but he seems to be relying on the dialectical sense of ‘justify’ in this passage—a sense that epistemologists often avoid.

<sup>14</sup>Gardner (2007).

<sup>15</sup>See, e.g., Hart (1968).

<sup>16</sup>Gardner (2007: 86).

<sup>17</sup>Gardner (2007: 87).

Gardner's points address internalists who insist that the distinction between justification and blamelessness cannot explain away their intuitions.<sup>18</sup> Crucially, excusability is *more than blamelessness*. It is a positive status—indeed, a product of the same competences that enable us to correctly respond to objective normative reasons when conditions are favorable. But it can fall short of a real justification, since the fact that one's reason appears to be objectively good doesn't always guarantee that it is objectively good.

### 3.1.3 Why Resisting Indifference Matters

The mismatch between ethics and epistemology might be worth ignoring if nothing in epistemology turned on it. But this is not the case. Among other things, the distinction proves important for debates between internalists and externalists about justification in epistemology just as it proved important for debates between internalists and externalists about normative reasons in meta-ethics.

Consider a watershed moment in meta-ethics to see the resemblance. T. M. Scanlon was one of the first to distinguish between rationality and correctly responding to genuine normative reasons. The distinction played a significant role in his response to Bernard Williams's rejection of externalism about normative reasons. Williams famously said:

There are...many things that a speaker may say to one who is not disposed to  $\phi$  when the speaker thinks that he should be, as that he is inconsiderate, or cruel, or selfish, or imprudent; or that things, and he, would be a lot nicer if he were so motivated.... But one who makes a great deal out of putting the criticism in the form of an external reason statement seems concerned to say that what is particularly wrong with the agent is that he is *irrational*.

Scanlon replied: "Williams is quite right that this claim would be implausible, but wrong...to hold that his opponent is committed to it."<sup>19</sup> Wrong, Scanlon insists, because it needn't be *irrational* to fail to respond to good reasons if the quality of these reasons is not apparent. Williams's internalist strictures might apply to factors that it would be irrational not to heed. But the externalist's theory need not concern these factors.

---

<sup>18</sup>See especially Pryor (2001).

<sup>19</sup>Scanlon (1998: 27).

This thought has become common among externalists in meta-ethics. It is no surprise that the resurgence of externalism has been matched by a wave of skepticism about the significance of rationality. One way to put a nail in the internalist coffin is to agree that internalists are right about rationality but to deny that rationality matters as such. The absence of similar thoughts in epistemology is striking—especially striking if we recall the passage that follows my opening quote from Cohen:

If the Reliabilist wants to distinguish ‘justified’ from ‘reasonable’ or ‘rational’ he may do so. But clearly the important epistemic concept, the one epistemologists have been concerned with, is what the Reliabilist would call ‘reasonability’ or ‘rationality’.

Reliabilists would do well to mimic Scanlon and dispute both parts of the last sentence. Grant to Cohen that rationality can come apart from reliability, and that it can thwart the fundamental epistemic goal of believing truly. One could then ask: “Why not take this to show that rationality lacks importance as such from the epistemic point of view?” One could push further: “Why think this concept is ‘the one epistemologists have been concerned with’ in giving theories of justification?” ‘Justification’ is often stipulatively defined by its role in JTB+ theories of knowledge. It is hardly clear that rationality will be fit to play this role if we understand it on its own terms.

So, the distinction between rationality and justification is not merely a pedantic one. It matters if it resembles the distinction drawn in meta-ethics. The parallel distinction in meta-ethics marked a triumph for externalists. If externalists are right about normative reasons, the internalists’ *idée fixe* arguably has derivative significance at best. Apparent reasons should attract as much underivative interest as apparent wealth. Externalists in epistemology could reasonably claim that internalists are looking for fool’s gold unless they somehow merge clarity and distinctness of appearance with truth-conducivity.

### **3.1.4 Antecedents and the Need for Detail**

Some epistemologists have been more careful than the ones I quoted at the outset. The only trouble is that these epistemologists have not given detailed accounts of what the rationality/justification distinction is supposed to be. This invites Cohen’s worry that

the distinction is merely a way to save their theories of justification from counterexamples.

Goldman (1986: 27), for example, writes:

I will not attempt to analyze...all terms of epistemic evaluation. The salient omission here is rationality, which has figured prominently in epistemology.

Since Goldman spends several chapters giving a process reliabilist account of justification, he clearly acknowledges a distinction between justification and rationality. More recently, Alex Jackson (2011) and Clayton Littlejohn (2012) have relied on the distinction. Jackson relies on it in critiquing *seemings internalism* about justification. Littlejohn's project requires such a distinction, since he denies that there are justified false beliefs! Jackson and Littlejohn are aware of the meta-ethics literature on which I'm drawing. But neither offers a detailed account of how rationality and justification differ.

Greater detail is needed. One reason is that the most common distinction in meta-ethics is between correctly responding to apparent reasons and correctly responding to *all* the objective reasons in the world. But justification is not a function of *all* the objective reasons, if we follow meta-ethicists in taking these to include virtually any facts. It is at most a function of the objective reasons that one *possesses*. We need an argument that possessed objective reasons and apparent reasons are distinct.

Another reason why more care is needed is that many meta-ethicists understand rationality very narrowly. While they hold that rationality requires correctly responding to apparent reasons, many will understand 'apparent' in a *belief-relative* way and view this requirement as a coherence requirement. Epistemologists are unlikely to find this interesting. Coherentism about anything other than coherence is now widely rejected.

This does not show that the meta-ethics literature is irrelevant or that the distinction can't be drawn. It just shows the need for greater detail and care—which I'll provide.

### 3.1.5 The Plan

Here is the plan. In §2, I explain which concepts from meta-ethics should be imported into epistemology. I use the distinction between apparent reasons and possessed objective reasons to explain the distinction between *substantive* rationality and justification, and I separate both of these statuses from *structural* rationality. In §3, I provide a deeper account of the distinction between apparent and possessed objective reasons that falls directly out of an account I've elsewhere given of the nature of apparent reasons. In §4, I explain how my distinctions extend to epistemology and improve on prior attempts to bifurcate epistemic evaluation. In §5, I explore the implications for major issues in epistemology.

## 3.2 Objective Reasons, Possession, and Apparent Reasons

### 3.2.1 Objective Reasons and the Early Distinction

Much of the conceptual progress in recent meta-ethics owes to the way meta-ethicists have come to understand *normative reasons*. Since the late 1990s, it has been popular to view normative reasons as objective facts that count in favor of acts and attitudes. Such facts are not ones to which we necessarily have access. For example, the fact that the lemonade is arsenic-laced is a conclusive normative reason for one not to drink it even if one is unaware of this fact. Call reasons of this sort *objective reasons*.

What is the connection between these reasons and rationality? There is clearly no *direct* connection. It is not irrational to drink the lemonade if one cannot see that it is arsenic-laced. And it may be irrational for one to drink some lemonade even if it is not arsenic-laced but merely appears to be. So, a consideration P's being an objective reason to  $\phi$  seems neither necessary nor sufficient for P to exert rational pressure to  $\phi$ .

This fact initially led people in meta-ethics to distinguish between rational  $\phi$ -ing and  $\phi$ -ing that is supported by *all* the objective normative reasons. And when this distinction was first drawn, many meta-ethicists understood rationality very narrowly. The requirements of rationality got identified with coherence requirements such as:



(*Enkrasia*) Rationality requires that if you believe you ought to  $\phi$ , you  $\phi$ .<sup>20</sup>

(*Means-End*) Rationality requires that if you intend to E and believe that M-ing is a necessary means for E-ing, you intend to M.

In the early 2000s, it was popular to view ‘rationality requires’ as taking wide scope over the conditionals, so that these requirements could be complied with in several ways.<sup>21</sup> On this wide scope view, one could comply with Enkrasia by  $\phi$ -ing or by dropping one’s belief that one ought to  $\phi$ . Even those who rejected the wide scope account agreed that the pressures of coherence are essentially *hypothetical* and *rationally escapable*, and differ from substantive pressures in this respect.<sup>22</sup> This picture led meta-ethicists to regard the key distinction as a distinction between *structural* and *substantive* evaluations.

The distinction between structural and substantive evaluations is important. But the contrast between the pressures of rationality and the pressures of objective reasons isn’t fully captured by this distinction. For we cannot ground the pressures of all apparent reasons by appeal to coherence requirements. Consider:

(\*) If it appears to you that there is arsenic in the glass, the apparent fact that there is arsenic in the glass is an apparent reason for you not to plan to drink from it.

Even if the perceptual appearance is misleading, it is *prima facie* irrational to ignore the apparent reason it provides. That is true regardless of whether one takes a doxastic stance on the presence of arsenic in the glass. So, (\*) is not merely a descendant of Enkrasia.

Might one claim that appearances can always give *objective* reasons strong enough to explain the relevant rational pressure? Not plausibly. You might reside in an empty world being fed pure illusion by some demon. The appearances here bear no objective probabilistic relation to extra-mental facts. If so, they cannot provide serious objective reasons for beliefs about these facts in these worlds.

---

<sup>20</sup>‘Enkrasia’ is John Broome’s term; see, e.g., Broome (2013).

<sup>21</sup>Thanks to Broome (1999)’s influence. Schroeder (2004) and Kolodny (2005) marked a departure.

<sup>22</sup>See Lord (2011) for a lucid explanation of how a narrow scoper can capture these features.

Call the rationality that consists in heeding these apparent reasons ‘substantive’ if you like. If you like, deny that it has anything in common with the structural rationality exhibited by compliance with Enkrasia and Means-End. Still, the rational pressure exerted by these apparent reasons is just as divorceable from the presence of objective reasons as the rational pressure exerted by coherence requirements.

### 3.2.2 Apparent Reasons vs. Possessed Objective Reasons

Some recent writers have appreciated this fact. Mark Schroeder has drawn attention to *subjective normative reasons*, which are the same things I am calling apparent reasons.<sup>23</sup> And Schroeder allows presentational mental states other than beliefs to ground possession of these reasons.<sup>24</sup> But he has also argued that these reasons are not just objective normative reasons to which we bear some privileged relation.<sup>25</sup>

Not everyone agrees with Schroeder. Some argue that apparent reasons are a proper subset of the objective reasons: they are the objective reasons *possessed* by the agent.<sup>26</sup> We have already seen one reason to worry about this view. It is hard to see how it can explain distant demon worlds where the reliable links between appearance and reality are wholly broken. We’ll see presently how the view falters in practical cases. But as we will also see, it does not follow that possessed objective reasons have no important role to play.

To see the problem in the practical case, consider cases of objective undercutting defeat. Suppose Benedict promised to pick Margaret up from the airport tomorrow. Unbeknownst to him, some maniacs ensure that it will be impossible for him to pick up Margaret tomorrow by creating an impenetrable forcefield around his house tonight. Plausibly, the fact that one promised to  $\phi$  at  $t$  can constitute an objective reason for one to plan to  $\phi$  at  $t$  only if it is *not impossible* for one to  $\phi$  at  $t$ . A fact of impossibility undercuts whatever objective reasons might have been provided by the good properties

---

<sup>23</sup>See Schroeder (2007: Ch. 1) and (2008).

<sup>24</sup>See Schroeder (2011).

<sup>25</sup>See Schroeder (2008).

<sup>26</sup>See Lord (2010).

of some imagined choice. This is why there is no objective reason to plan to travel back in time to prevent the crusades, slavery, and World War II.<sup>27</sup> But before Benedict discovers the forcefield, the fact that he promised to pick Margaret up tomorrow remains an *apparent* reason for him to plan to go tomorrow. So, not all apparent reasons are objective reasons.

Some might insist that even if the consideration that explicitly motivates an agent to  $\phi$  is not an objective reason to  $\phi$ , the agent will possess other objective reasons to  $\phi$ . But this strategy is inapplicable here. Any other reasons Benedict might have had to plan to go to the airport tomorrow are also objectively undercut by the fact of impossibility. Might it be claimed that there is a weak objective reason for Benedict to plan to go to the airport tomorrow? Not plausibly. The fact of impossibility is a total undercutting defeater. There is not even a weak objective reason to plan to travel back in time to prevent the crusades. Might Benedict have an objective reason to plan to go to the airport with zero weight? It is hard to see how this differs from our conclusion. A reason with no objective weight is not an objective reason. Accordingly, we should reject:

*(The Factoring Account of Apparent Reasons)* R is an apparent reason for S to  $\phi$  iff R is an objective reason for S to  $\phi$ , and it is apparent to S that R.<sup>28</sup>

Crucially, however, this does not show that there is nothing important that defenders of the Factoring Account are tracking. It only shows that possessing an apparent reason is not the same thing as possessing an objective reason.

We care about possessing objective reasons. It is not enough if there merely exists something that could objectively support us. We want to *receive* this support and *be poised to act on the basis of it*. And even if our acts and attitudes fall short of ideal correctness, we can have more than excuses. Imperfect conduct is not always unjustified. Possessing objective reasons matters because justification matters.

---

<sup>27</sup>This is Bart Streumer's great example. See Streumer (forthcoming) for a defense of this plausible idea.

<sup>28</sup>"Factoring Account" is Schroeder (2008)'s term for this type of view.

Now, possessing *some* objective reason to  $\phi$  is not enough to have justification to  $\phi$ . One might possess stronger objective reasons on the other side. The fact that these objective reasons outweigh the objective reason to  $\phi$  does not destroy that reason. Undercutting defeaters can destroy objective support relations, but to outweigh is not to destroy. So what does justification require? I suggest that to be justified in  $\phi$ -ing is to possess an objective reason to  $\phi$  and to possess no stronger objective reason for any specific alternative to  $\phi$ -ing. This is less than ideal correctness but more than rationality.

The distinctness of justification and ideal correctness is illustrated by this case:

*Three Envelopes.* You are correctly told by a reliable authority that you can take one of three envelopes. You are also correctly told that there is \$800 in Envelope 1. Finally, you are correctly told there is \$1000 in either Envelope 2 or 3, and that the envelope that doesn't have the \$1000 is empty. But you are told nothing to indicate which envelope might contain it, nor can you find out before choosing.<sup>29</sup>

The fact that there is \$800 in Envelope 1 is a strong objective reason to take it. Since you know that there is \$800 in this envelope, you possess this objective reason. You do know that you are not doing the *ideal* thing, since you know that there is more objective reason to choose either Envelope 2 or 3. But note that the disjunctive fact that either Envelope 2 or Envelope 3 contains \$1000 is not an objective reason to choose one rather than the other. Yet this is the only objective reason *not* to take Envelope 1 that you possess. So, you possess an objective reason to choose Envelope 1 and possess no objective reason that favors a specific alternative to Envelope 1. This is why choosing Envelope 1 is justified.

Having justification remains distinct from having merely apparent reasons to  $\phi$ . To see this, make some inaccessible changes to *Three Envelopes*. Suppose the reliable authority got things wrong: a sneaky trickster replaced all the money with Monopoly money. Moreover, unbeknownst to you, all the envelopes are surrounded by impenetrable force fields. Suppose you reach out and burn your hands in the force field. It

---

<sup>29</sup>Ross (2012) introduced this case. I reject Ross's own diagnosis of this case for reasons that will soon become clear.

becomes clear that the envelopes are inaccessible: the forcefield boundaries light up as you get burned. If you encounter the trickster, what can you say? He may say: “There could be no objective reason to try to accomplish something you couldn’t possibly have accomplished!”

While the trickster’s exclamation is apt, you can show that you were not *crazy* by telling the story about the authority who was wrong about the accessibility of the envelopes. The story shows that you had apparent normative reasons to try to pick Envelope 1. These apparent reasons excuse you in the best way, by showing that your rationality was in working order. But you could reasonably wish that you had been in the original Three Envelopes case. If you do, what you wish for is a real justification.

So, there are two phenomena, significant in different ways. Possessed objective reasons matter for justification. Apparent reasons matter for rationality. The objective reason given by the \$800 in Envelope 1 in the original case partly explains why you were justified in choosing Envelope 1. But while you did not do the *ideal* thing here, this case is different from the revised case. You had less by way of a defense there.

### 3.2.3 Taking Stock: A Spectrum of Reason-Based Evaluations

The last two subsections bring out a spectrum of useful reason-based evaluations:

THE EVALUATION	INVOLVES COMPLYING WITH
ideal correctness	<i>balance of all existing objective reasons</i>
justification	<i>balance of possessed objective reasons</i>
substantive rationality	<i>balance of apparent reasons</i>
structural rationality	<i>balance of believed reasons</i>

The early meta-ethics literature focused on the distinction between the first and the last entries. The need for an intermediate evaluation is clear. But as we have seen, there is actually need for more than one intermediate evaluation.

I will discuss how these distinctions extend to epistemology in §4 and §5. Given

my goals, the distinction between justification and substantive rationality is the most important. It has the greatest implications for epistemology. Still, the distinction between structural rationality and justification gets blurred in epistemology too, with some striking results. So, I will also push for greater recognition of that distinction.

### 3.3 A Deeper Rationale for the Distinction

First, I want to provide a deeper basis for the distinction between apparent and possessed objective reasons. To reveal this basis, I will rehearse a framework for understanding apparent reasons that I've developed in work in meta-ethics. Part of the reason for rehearsing it is that the distinction between apparent reasons and possessed objective reasons falls directly out of the framework. Since my framework is independently defensible, this further illustrates that my distinction is not *ad hoc*.

My distinction is a new instance of a pattern familiar from the work of Ernest Sosa. Apparent reasons, I will suggest, are apparent facts that we *competently treat like objective reasons*, where competence is indirectly defined in terms of objective reasons and a competence/performance distinction is honored. Possessed objective reasons are apparent facts that we *aptly treat like objective reasons*. Since one can aptly treat a consideration like an objective reason only if it is one, these apparent facts are also real.

This basis for the distinction clarifies the relationship between rationality and justification. While not all apparent reasons are possessed objective reasons, all possessed objective reasons are apparent reasons, since aptness entails competence. So, while substantive rationality does not entail justification, justification entails substantive rationality. This is important: I want to vindicate the thought that subjects in demon worlds may not be justified, but I deny that unenlightened clairvoyants are justified.

#### 3.3.1 Apparent Reasons: Two Views

Let's turn to consider what apparent reasons might be. Some meta-ethicists have suggested that for R to be an apparent reason for S to  $\phi$  is for it to appear to S that R is

an objective reason to  $\phi$ .<sup>30</sup> Call this view the *de dicto* view.

One worry about this view is that it overintellectualizes rationality.<sup>31</sup> Children and animals can be evaluated for rationality. But it is doubtful that they have the concept of a normative reason. If they lack that concept, it cannot strike them *de dicto* that anything is a normative reason. This is a reason to avoid *de dicto* views. Of course, *de dicto* theorists could retreat to permissive accounts of concept possession. Or they might say that it can appear to someone that X is an F even if this person lacks the concept of an F. But these burdensome commitments are worth avoiding if possible.

So we should not accept the *de dicto* view unless forced. This is not to say that we should reject the *de dicto* view. Non-acceptance is weaker than rejection. My own view would be a *de dicto* view if permissive accounts of concept possession were true. But my view does not require these accounts to succeed. This is why it is preferable.

But my view is not the standard alternative. The standard alternative regards apparent reasons as apparent facts that would be objective reasons if they were real facts.<sup>32</sup> Call this the *de re* view. We can state it more officially as follows:

(*De Re*) A consideration P is an apparent reason for S to  $\phi$  iff (i) it appears to S that P and (ii) P would be an objective reason for S to  $\phi$  if P were the case.

Alas, this view is unacceptable. It is easy to imagine cases where (a) someone knows that P, (b) P is an objective deductive reason to believe Q because P entails Q, but (c) the entailment is so arcane that the person gains no apparent reason to believe Q. The *de re* view entails that as long as (a) and (b) are true, the person gains an apparent reason to believe Q. This is wrong: subjects with weak mathematical abilities lack apparent reasons to believe the most arcane theorems even if they know the relevant axioms.

There are some replies to this objection. But they are ultimately unconvincing. To see the first reply, consider the common distinction between *reasons* and *enabling conditions*. The fact that P entails Q is not itself a reason to believe Q. It is a fact that

---

<sup>30</sup>This view is defended by Scanlon (1998) and Kolodny (2005).

<sup>31</sup>See Parfit (2011).

<sup>32</sup>This view is defended by Parfit (2001, 2011), Schroeder (2007), and Way (2009).

enables P to be a reason to believe Q. The defender of the *de re* view might revise her view by demanding the relevant enabling conditions to be apparent, and say:

P is an apparent reason for S to  $\phi$  iff (i) P's truth would be an objective reason for S to  $\phi$  given conditions C, (ii) it appears to S that R and (iii) it appears to S that C.

But this view is too strong. Consider perceptual beliefs. Certain properties of perceptual experiences enable them to provide apparent reasons for belief. Intrinsic features might include the presentational character of experience. Relational features might include the reliable connection between experience and reality. We do not need to represent these enabling conditions to form rational perceptual beliefs.

The view is too strong even in some deductive cases. Having beliefs about entailments is *one* way to be sensitive to logical relations between propositions. But it is not the *only* way. Another way is to have the ability to competently infer one proposition from another by using an *inference rule*. Instead of reasoning:

- (i)  $\neg(P \vee Q)$
- (ii) If  $\neg(P \vee Q)$ , then  $\neg P \wedge \neg Q$
- (iii) So, by *modus ponens*,  $\neg P \wedge \neg Q$

one could *directly* infer (iii) from (i) by relying on one of the DeMorgan rules.

We cannot replace all rules by extra premises. Even axiomatic systems need rules, and most people do not reason axiomatically. Of course, most people are not so competent that they can use any rules like they use *modus ponens*. But equally clearly, there is a spectrum of acumen. Some people do have the native logical abilities to cleave reliably to more intricate inferential patterns without forming beliefs about the entailments that legitimate them. The revised *de re* theory cannot explain all the apparent reasons these people can acquire.

There is a different reply that *de re* theorists might pursue. They might embrace the conclusion that knowing that P gives one *some* apparent reason to believe all of



P's implications. But they might add that this apparent reason is defeated when the implications are arcane, so that it is never rational to believe these implications.

But this reply does not withstand scrutiny. Apparent reasons must be defeated by other apparent reasons if the defeat is to affect degrees of rationality. In the earlier case, it was the fact *that the subject's logical abilities were limited* that explained why the subject could not rationally believe the most arcane consequences of certain axioms. *This* fact is not one that itself must be apparent to prevent the subject from being able to rationally believe these arcane consequences. People can be deceived about their abilities, taking themselves to be geniuses when they are fools. By a fluke, their incompetence might have happened to land them on correct results, so that they lack apparent defeaters. But they are not rational. Incompetence alone can preclude rationality.

### 3.3.2 Treating, Competence and Appearance: A Better View

A better theory is worth seeking. But how can we avoid the problems with *de re* views without overintellectualizing rationality? We can see how by considering the most general description of the problem for *de re* views. The overarching problem was that these views imply that a person can have P as an apparent reason to  $\phi$  *even when that person cannot competently treat P like an objective reason to  $\phi$* . We should profit from this observation.

Notice that it is possible to treat something like an F without having the concept of an F. My cat can treat me like a vending machine without having the concept of a vending machine. Similarly, one can treat a consideration like an objective reason without having the concept of an objective reason. My cat can treat the sound of food going into its dish like an objective reason to walk into the room containing the dish.

What is it to treat R like an objective reason to  $\phi$ ? To treat R like an objective reason to  $\phi$  of some kind is to be disposed to form attitudes or act in all or at least most of the ways that would be correct if R were an objective reason to  $\phi$  of that kind. Barring a permissive theory of concept possession, this disposition is weaker than a *de dicto* belief about objective reasons.

The obvious idea is to invoke this weaker notion in a new theory. Apparent reasons,

I suggest, are apparent facts that we competently treat like objective reasons. More officially:

*Competent Treating* (CT): A consideration P is an apparent reason for S to  $\phi$  iff

- (i) it appears to S that P,
- (ii) S treats P like an objective reason to  $\phi$ , and
- (iii) this treating manifests a relevant objective reasons-sensitive competence of S's,

where the competence is a disposition of S's to treat P-like considerations like objective reasons to do  $\phi$ -like things only if they are objective reasons to do  $\phi$ -like things.

The core ideas are simple. For P to be an objective reason for one to  $\phi$ , it is not enough that P appears to be the case and would be an objective reason to  $\phi$  if it were the case. P must also be something that one is disposed to treat like an objective reason to  $\phi$ . The view captures the intuition that apparent reasons must “look like” objective reasons from someone’s perspective. But it does so without overintellectualization.

Of course, one might incompetently treat a consideration like an objective reason. One might randomly treat Q and R like objective reasons to believe S. By luck, there might be a proof of S from Q and R. But if this proof is arcane, it is not rational for one to believe S on the basis of Q and R. This is why CT invokes competence, which is indirectly defined in terms of objective reasons and non-normative factors.

### 3.3.3 Grounding the Distinction: Competence and Performance

If we accept CT, we can better understand the difference between apparent reasons and possessed objective reasons. CT predicts and explains this difference in two ways.

First, competences can be fallible. A competence to  $\phi$  is not necessarily an infallible ability to  $\phi$ . So, one can competently treat P like an objective reason to  $\phi$  even if P wouldn't be an objective reason to  $\phi$  if P were true.

Secondly and more crucially, there is a *competence/performance distinction*. Having a competence to succeed does not entail being in a position to succeed if one tries. Competences are dispositions. Like other dispositions, there is a range of favorable conditions for manifestation. Unfavorable conditions don't destroy competence. An archer, for example, retains the competence to hit the bull's-eye even when some trickster creates unexpected gusts that blow the arrow off target. An agent with the same competence could be relocated to a systematically unfavorable environment and be hoodwinked about its favorableness by a trickster. If so, this agent's skill is not destroyed. But it won't yield reliability.

Like other competences, reasons-sensitive competences do not guarantee actual reliability. So, in principle, one could competently treat a vast range of apparent considerations like objective reasons even if they wouldn't be objective reasons if true. All we need is for the circumstances to be both *unfavorable* and *misleading*.

So, CT predicts that apparent reasons can fail to be objective reasons. But an indirect connection to objective reasons remains. While rationality does not give us the ability to reliably  $\phi$  for objective reasons, the capacities rational subjects display are the same capacities that yield success in favorable cases. Rationality is a competence to get connected to objective reasons in favorable environments, however unhelpful in unfavorable ones.

### 3.3.4 Possessing an Objective Reason to $\phi$

A consideration can be competently treated like an objective reason without being one. Indeed, a consideration can be competently treated like an objective reason without being objectively likely to be an objective reason. In these cases, the consideration is only an *apparent* reason. If so, what more does it take to possess an objective reason R to  $\phi$ ?

My answer involves a spin on CT that will sound familiar to epistemologists:

*Apt Treating* (AT): S possesses an objective reason R to  $\phi$  iff

(i) S correctly treats R like an objective reason to  $\phi$ —i.e., R is an objective reason

to  $\phi$  and S treats R like an objective reason to  $\phi$ ,

(ii) S's treating manifests S's relevant objective reasons-sensitive competence, and

(iii) the correctness of S's treating manifests that reasons-sensitive competence.

To see why AT is plausible, consider an example:

(*Misfire*) Bill is a competent but not infallible logician who can directly see many entailments. It strikes him tonight that there is a proof of Q from P involving certain steps, which he seems to see by exercising his competence. Alas, his fallible competence misfires and yields a mistaken impression: there is not *that* kind of proof of Q from P. As it happens, there *is* a proof of Q from P. But it is arcane and far beyond Bill's direct grasp.

Here P is an objective reason to believe Q because there is a proof of Q from P. Bill is competently attracted to treating P like an objective reason to believe Q: his logical 'vision' is a product of his mathematical competence. But Bill's vision misfires, as fallible vision can. For this reason, Bill does not possess the objective reason to believe Q that P *actually constitutes*, since he does not grasp the real proof of Q from P.

Here is why. While Bill's attraction to treat P like an objective reason to believe Q is *correct* in the minimal sense that P is an objective reason to believe Q, the correctness of Bill's attraction does not manifest his competence. The attraction is correct for reasons that are beyond the reach of his competence. The attraction manifests competence, but its correctness has nothing to do with that competence.

None of this is to deny that Bill might possess other objective reasons to believe Q. We are only claiming that Bill does not possess the objective reason to believe Q that P actually constitutes, since he cannot see the real proof. P is for Bill only an apparent reason to believe Q: the connection Bill sees between P and Q is not the connection that actually exists. There is an analogy with veridical hallucination: there appears to Bill to be a connection and there is one, but the connection does not explain the appearance.

### 3.3.5 The R/J Distinction: A New Instance of an Old Pattern

We can now see that the distinction between rationality and justification is a new instance of an old pattern. Ernest Sosa has made us familiar with a triadic pattern of evaluation for attitudes and activities that aim at some target. An attitude or act A is:

*correct*, when A hits the relevant target,

*competent*, when A manifests a disposition to hit the relevant target, and

*apt*, when A's correctness manifests the relevant competence.

What have interested us are *treatings of considerations like objective reasons*. The same triadic pattern arises for these treatings. These treatings are:

*correct*, when the treated considerations are in fact objective reasons,

*competent*, when the treatings manifest a competence to treat only objective reasons like objective reasons, and

*apt*, when the correctness of these treatings manifests this competence.

Apparent reasons are considerations that one competently treats like objective reasons. Possessed objective reasons are considerations that one aptly treats like objective reasons. Since apt treating entails correct treating, possessed normative reasons are objective reasons. It also falls out that possessed normative reasons are apparent reasons, since apt treatings are also competent. But it doesn't fall out that all apparent reasons are objective reasons. For competent treating is not necessarily correct treating.

My distinction between rationality and justification is an upshot of these ideas:

(J) S has *ex ante* justification to  $\phi$  iff  $\phi$ -ing is supported by the balance of considerations that S aptly treats like objective reasons to  $\phi$ .

(R) It is *ex ante* rational for S to  $\phi$  iff  $\phi$ -ing is supported by the considerations that S competently treats like objective reasons to  $\phi$ .<sup>33</sup>

---

<sup>33</sup>*Ex ante* justification is justification *to believe*, as opposed to justified *believing*; similarly for *ex*

Since apt treating entails competent treating, *ex ante* justification entails *ex ante* rationality. But since competent treating doesn't entail apt treating, *ex ante* rationality does not entail *ex ante* justification. So while rationality is separable, the two are not orthogonal. This is why unenlightened clairvoyants are unjustified. They fail to possess objective reasons *precisely because* they lack apparent reasons.

The connection with Sosa's triadic pattern reinforces the objectivity of my distinction. The triadic pattern arises for any attitude or activity with an aim. Even if one doesn't want to organize all of one's normative theorizing around this pattern, one cannot deny that it is a real pattern. The connection also allows us to see how my view improves on Sosa's virtue epistemology. My appeal to reasons-based competences yields two improvements: (i) greater fineness of grain, and (ii) a more obvious explanation of why unenlightened clairvoyants, Truetemp, and the like are not only not rational but unjustified. But the views are similar in spirit. One might say that my view is a *reasons-based* virtue epistemology.

### 3.4 Extending the Distinctions to Epistemology

Extending my central distinctions to epistemology is now easy. We need only draw our attention to the existence of objective epistemic reasons and apply the framework.

#### 3.4.1 Objective Epistemic Reasons: Unpossessed and Possessed

What are objective epistemic reasons? Clear examples include pieces of evidence that are not necessarily possessed by anyone. Call such pieces of evidence *objective evidence*. Evidence in this sense is, as Tom Kelly puts it, a "reliable sign, symptom, or mark of that which it is evidence of."<sup>34</sup> While many epistemologists ignore this kind of evidence, the concept is a clear one that often features in our ordinary thinking.

Talk of objective evidence applies to items in many different ontological categories.

---

*ante* rationality. *Ex post* evaluations also matter. I assume the orthodox view about the relationship between the *ex ante* and the *ex post*. Elsewhere I defend this view against objections of the sort from Turri (2010).

<sup>34</sup>Kelly (2006).

But much of this talk is elliptical. Evidence must play certain roles: it must be able to probabilify, to be explained, and to figure in our reasoning, for example.<sup>35</sup> Plausibly, only things with a proposition-like structure can play these roles. Hence, when we say:

(a) these fingerprints on the gun are evidence that he is the murderer,

what we really mean is something like:

(b) the fact that these fingerprints are on the gun is evidence that he is the murderer.

Why call objective evidence ‘objective’? Because of its connection to objective probability: a piece of objective evidence has an objective probabilistic connection to truth, by being a reliable indicator of truth. ‘Objective’ does not mean ‘public’. Many examples of objective evidence are public. But a fact needn’t be public to be objective evidence.

Here is an illustration. Facts about your current feelings are objective evidence for certain conclusions about the broader mood you are in. Not all of our mental life is transparent. Sometimes we must reason inferentially about it. There are better and worse ways. In the case of feelings and moods, reliable correlations help to explain why some ways are better than others. But facts about your feelings are private.

Epistemologists often ignore objective epistemic reasons, regarding them as irrelevant to justification. But this is a mistake. Yes, we must *possess* epistemic reasons for them to justify us. But it does not follow that what we possess cannot be an objective epistemic reason. If the name on the driver’s license is not yours, it cannot give you legal permission to drive. That does not show that you need a different kind of thing to drive. You just need to bear a different *relation* to the same kind of thing.

Here we find another place where epistemologists get worried. Many assume that we will have to analyze the relation of possession in terms of justification or knowledge. So, many assume that giving an account of justification in terms of objective epistemic reasons would be circular. But this is hasty. We can appeal to non-doxastic mental

---

<sup>35</sup>See Williamson (2000) for a now classic discussion of some of these roles.

states, and to normative concepts more primitive than the concept of justification. My theory does this. My theory does appeal to one normative concept. But it is the concept of an objective reason. Alternatively, one could directly ground the relation of possession in non-doxastic mental states that meet certain reliabilist constraints, thereby avoiding circularity.

Many epistemologists have thought that possession is a matter of access. Access need not be understood in internalist terms. Does that mean that access must be understood in terms of knowledge or justification? No. Note that non-doxastic seemings can be evaluated under the triadic scheme noted in the last section. A visual seeming, for instance, can be

*accurate*, when what visually seems so is so,

*adroit*, when the seeming manifests the subject's visual ability, and

*apt*, when the seeming's accuracy manifests the subject's visual ability.

Apt non-doxastic seemings are not knowledge, since knowledge entails belief. Moreover, these seemings cannot be justified or unjustified, though they can manifest a more primitive sort of ability. Nevertheless, they afford access to reality.

Access to objective epistemic reasons is not sufficient for possessing these reasons. Having access to a fact E that constitutes objective evidence for P is not sufficient for possessing E as good evidence to believe P. Here we can consider variations on the counterexamples to *de re* theories. The fact that someone has Koplik spots is objective evidence that this person has measles. But seeing that this person has Koplik spots is insufficient for one to have reason to believe that this person has measles. After all, the connection between measles and Koplik spots has not always been apparent.<sup>36</sup>

Does this mean that one must also be justified in believing that E is objective evidence for P to possess E as evidence for P? No. We found a path between *de dicto* and *de re* views by attending to the notion of treating something like an objective reason. We can invoke this notion again. If one is aptly attracted to treat the presence

---

<sup>36</sup>Adam Marushak alerted me to this example and its usefulness in this context.



of Koplik spots as evidence that the person has measles, one possesses that piece of objective evidence to believe that she has measles. This is what we need, beyond access. More officially, then:

*Objective Evidence Possession:* S possesses objective evidence E to believe P iff

- (i) it aptly seems to S that E, and
- (ii) S aptly treats E like objective evidence to believe P.

If one thinks that objective evidential relations can hold between false propositions, one could relax the account by replacing ‘aptly’ in (i) with ‘competently’. But my own inclination is to say that E is not objective evidence to believe P when E is false.

That is compatible with allowing that one has *other* objective evidence for P when E is false. Since ‘objective’ doesn’t mean ‘public’, this could include the sheer fact that it non-doxastically appears to one that E when there is a reliable link between appearance and reality. So, it is not as if one is always unjustified in believing P on the basis of E if E is false.

### 3.4.2 Apparent Epistemic Reasons

But if all reliable indication relations between appearance and reality are broken, we should deny that the fact that one is appeared to in some way is objective evidence that things are that way. One can, however, have apparent reasons in these cases. This falls out of the account I offered in the last section. The following is an upshot:

*Apparent Evidence Possession:* S possesses apparent evidence E to believe P iff

- (i) it appears to S that E, and
- (ii) S competently treats E like objective evidence to believe P.

Here E could be some false proposition about the external world that appears to be true in virtue of S’s visual experience. Accordingly, one can possess an external world proposition E as an apparent epistemic reason to believe P, despite E’s being false.

Can we explain why one has apparent epistemic reasons even when appearance-reality correlations are annihilated? Yes. Recall the competence/performance distinction. My envatted brain retains my epistemic competences. Like other dispositions, competences can be retained in inhospitable situations.

### 3.4.3 Rationality vs. Justification in Epistemology

Given the distinction between apparent and possessed objective epistemic reasons, it is easy to distinguish between substantive epistemic rationality and justification. For *ex ante* rationality and justification, the distinction would go as follows:

( $R_e$ ) It is *ex ante* epistemically rational for S to have doxastic attitude D(P) iff D(P)-ing is favored by the balance of S's apparent epistemic reasons.

( $J_e$ ) S has *ex ante* epistemic justification to have doxastic attitude D(P) iff D(P)-ing is favored by the balance of the objective epistemic reasons that S possesses.

*Possessing* an objective epistemic reason and being an *apparent* epistemic reason would then be analyzed in the way that I suggested earlier, yielding the more precise:

( $R_e^*$ ) It is *ex ante* rational for S to have D(P) iff D(P)-ing is favored by the considerations S competently treats like objective epistemic reasons.

( $J_e^*$ ) S has *ex ante* justification to have D(P) iff D(P)-ing is favored by the balance of considerations S aptly treats like objective epistemic reasons.

Because I am drawing a distinction that epistemologists in many different camps should acknowledge, I have not given necessary and sufficient conditions in descriptive terms for being an objective epistemic reason or for being favored by an objective epistemic reason. The distinction should be neutral on these issues.

I did suggest examples of objective epistemic reasons—viz., objective evidence. I also suggested that favoring could be understood in terms of reliable indication. But ultimately I want to be neutral on whether all objective epistemic reasons are evidence and whether favoring must be understood purely in terms of reliable indication.

### 3.4.4 Further Divergences: Structural Rationality

The difference between  $(R_e)$  and  $(J_e)$  doesn't exhaust the distinction between rationality and justification. Structural rationality was what meta-ethicists first contrasted with support by objective reasons. While many epistemologists reject coherentism about justification, confusion remains about the relation between structural pressures and justification.

Here is one illustration. Several epistemologists claim that merely believing that one's belief was unreliably formed can defeat one's justification.<sup>37</sup> This is a mistake. It is incoherent for one to hold some belief while believing it to be reliably formed. But the problem lies in the *set of attitudes*. If one's first-order belief was reliably formed and one has no reason to believe that it was unreliably formed, one should resolve this conflict of attitudes by dropping that meta-belief and retaining the first-order belief.

We can deny that one's first-order justification can be defeated by the meta-belief. We can instead use requirements like the following to explain what is going awry:

*(Epistemic Enkrasia)* Structural rationality requires that if one believes that some doxastic attitude that one holds is epistemically unjustified, one abandons that belief.

We could follow Broome (1999) in taking 'requires' to take wide scope over the conditional, so that Epistemic Enkrasia is equivalent to a prohibition against a certain conjunction of states. But even if we understand Epistemic Enkrasia in a narrow scope way, we could agree that this requirement is rationally escapable in the sense that one can permissibly exit this narrow scope requirement by abandoning the meta-belief.<sup>38</sup>

Exactly what are the requirements of structural epistemic rationality? This is a huge question. What matters here is that these requirements can be rationally escaped in a way that the requirements of justification cannot. Of course, one could escape certain requirements of justification by losing certain evidence. But this escape will be a non-rational process or a process that makes one epistemically blameworthy.

---

<sup>37</sup>See Goldman (1986), Plantinga (1993), and Bergmann (2006).

<sup>38</sup>See Lord (2011) on narrow-scoping and the distinction between exiting and violating a requirement.

### 3.4.5 Comparisons with Other Bifurcating Proposals

I have now illustrated how the distinction applies in epistemology. This is not the first time someone has tried to bifurcate epistemic evaluation. Why prefer my bifurcation?

A large reason is that it is an instance of a more general distinction that is defensible across the board, drawing on insights outside of epistemology. Indeed, it follows from simple reflections on objective reasons and our relations to them. Similar points do not hold for bifurcation proposals that other epistemologists have advanced.

Consider Goldman (1988)'s distinction between strong and weak justification. While the examples he used to motivate this distinction do illustrate two different properties, there is a worry that weak justification is not really a kind of *justification*. Many ethicists would balk, seeing weak justification as mere excusability. By contrast, the distinction between rationality and justification can be gleaned by inspecting the nature of rationality and justification. It is no surprise that we already have separate words for them.

Similar worries arise for attempts to distinguish between subjective and objective justification. Why is "subjective justification" a kind of justification? When our undergraduates talk about "subjective truth", they are not talking about a kind of truth. Why think that talk of subjective justification should be taken more seriously?

Other bifurcations are stipulations that fall out of certain frameworks of evaluation. Ernest Sosa has separated two kinds of justification on the basis of a distinction between exercises of competence in hospitable environments and displays of competence in inhospitable environments.<sup>39</sup> I like the underlying idea. But I dislike the labels, which invite the objection that one of these properties is not a kind of justification at all.

Similar remarks apply to the attempt to distinguish between personal and doxastic justification.<sup>40</sup> There is a distinction between evaluations of persons and evaluations of attitudes or acts. But it is not a distinction between two kinds of justification. Rather, it is marked by the distinction between evaluations like *blameworthy*, *praiseworthy*,

---

<sup>39</sup>See Sosa (1993).

<sup>40</sup>See Bach (1986), Engel (1993), and Littlejohn (2012).

*excusable*, etc. and evaluations like *wrong*, *justified*, *permissible*, etc.

Some have recognized this fact and instead appealed to the distinction between justification and blamelessness. But blamelessness is too weak to be interesting. Denials of responsibility establish blamelessness. But there is no important positive epistemic property that is compatible with total insanity, though total insanity can make one blameless.

Excusability is a stronger contender than blamelessness if we accept Gardner's picture and view excuses as expressions of our capacity for responsibility rather than denials of responsibility. So understood, my inclination is to see excusability and rationality as cognate notions (like Gardner). So if you want to appeal to excusability, you are my ally.

But there are reasons to appeal to rationality rather than excusability. Many worry that talk of epistemic responsibility presupposes an implausible doxastic voluntarism. I don't share these worries, but they are pervasive enough to be worth sidestepping.

### 3.5 Implications for Epistemology

Having defended my distinction, I turn to consider some implications for epistemology. Drawing the distinction in epistemology makes a difference—indeed, several differences.

#### 3.5.1 How the Internalism/Externalism Debate Shouldn't Proceed

First of all, my distinction clarifies the terms of debates between internalists and externalists in epistemology and shows how these debates should not proceed.

Disagreements between internalists and externalists in epistemology are usually framed as disagreements about a single property—justification, most centrally. People assume we can apply the methodology of systematizing intuitions about cases to make progress. Many internalists support their views by appealing to cases. Consider BonJour's clairvoyance cases and Cohen's demon world.<sup>41</sup> Similarly, many externalists have tried to undermine internalism just by appealing to cases. Consider Goldman's

---

<sup>41</sup>See Cohen (1984) and BonJour (1985).

forgotten evidence cases.<sup>42</sup>

Once the distinction between rationality and justification is appreciated in its full generality, this familiar way of conducting the disputes becomes unsatisfying. Internalist intuitions are liable to conflate these phenomena. It is no surprise that almost every epistemologist I quoted at the outset is an internalist. We can grant the internalist's intuitions as right for rationality but argue that they undershoot the intended target.

Notice that this leads to no irenic dissolution of the debate. Instead, it leads to an advantage for the externalist. Because justification requires possessing objectively good reasons and rationality does not, internalists about epistemic justification are in trouble. The property they care about may have a role. But it is irrelevant as such to justification.

Because some internalists uphold the value of epistemic blamelessness, they might accept this verdict but see it in a different light. They might insist that blamelessness is what we really care about. But not all internalists will agree. Consider Jim Pryor:

Many philosophers share Cohen's intuition that it's possible for a brain in a vat, if he conducts his affairs properly, to have many justified...beliefs about his environment... [T]his intuition survives the recognition that being epistemically blameless does not suffice for being justified. It doesn't seem merely to be the case that the brain in a vat can form beliefs in a way that is epistemically blameless. It also seems to be the case that he can form beliefs in a way that is epistemically proper, and that the beliefs...would be fully justified....<sup>43</sup>

Deeper arguments are needed. If our paradigms of blamelessness include "lack of muscular control..., subjection to gross forms of coercion by threats, and types of mental abnormality",<sup>44</sup> it is implausible that the envatted are just epistemically blameless. But that just illustrates the need to separate blamelessness and excusability. The envatted are better off than people bereft of the capacity for responsible thought (e.g., the mad). But the claim that the envatted achieve more than epistemic excusability is a stronger claim.

---

<sup>42</sup>See Goldman (1999a) and (2009) for two different cases of this sort.

<sup>43</sup>Pryor (2001: 117).

<sup>44</sup>From Hart (1968)—a target in Gardner's attack on the view that excuses are denials of responsibility.

### 3.5.2 A Challenge for Internalists

To defend that claim, internalists must establish that the epistemic domain differs from others in admitting of no distinction between apparent reasons and possessed objective reasons. It is unclear that internalists can do this unless they deny that truth is the basic epistemic goal. If a belief essentially aims at any properly epistemic target, objective epistemic reasons are going to have something objective rather than merely apparent to do with helping it hit that target. An epistemic reason R to believe P will be objectively good to the extent that believing P for R makes it objectively more likely that one will achieve the aim of belief with respect to P, and bad if it does not.<sup>45</sup> If the aim is truth, unreliable indicators of truth will not be objectively good epistemic reasons for belief.

Here we find a reversal of the standard dialectic on the new evil demon problem. Given a richer set of concepts, it is more natural to take the demon world to support a reliabilist externalism about justification, not to undermine this view. The demon-worlder's reasons just don't look objectively good if the fundamental epistemic goal is true belief or knowledge. This undermines Cohen's reply to reliabilists. Cohen insisted that rationality is what obviously matters from the epistemic point of view. This is false: it is *unobvious* how rationality matters from the epistemic point of view. This conclusion echoes recent doubts about the significance of rationality in meta-ethics.<sup>46</sup>

It is unclear why we necessarily have objective epistemic reasons to be epistemically rational if we embrace the fundamental epistemic goals that many embrace. There are contingent links. But these links are irrelevant to the issue at hand. The internalist claims that the demon-worlder is epistemically justified, which entails that she possesses objectively good epistemic reasons. That claim is not bolstered by contingent links. Internalists owe us a deeper explanation.

---

<sup>45</sup>I formulate this in a proposition-relative way not only because this is the most natural way to formulate it, but because formulating it in a global way leads to problems well-documented by Berker (2013).

<sup>46</sup>See especially Kolodny (2005).

### 3.5.3 Challenging Concessive Externalist Views

There are challenges for epistemologists other than internalists. While externalists deny that mere beliefs and seemings are sufficient for positive justification, they often concede that these factors can defeat justification regardless of their externalist features.

Our distinction casts doubt on this concession. There is no reason to assume that the reasons *to*  $\phi$  that are relevant to justification are more objective than the reasons *against*  $\phi$ -ing that are relevant to justification. Consider prudence: if there are powerful objective prudential reasons not to perform some act, the mere fact that one's laziness makes them seem weak doesn't defeat these reasons. The same goes for morality, law, chess, etc.

It is more natural to deny that beliefs and non-doxastic seemings can as such defeat justification. They may affect some belief's degree of rationality or affect the overall rationality of certain sets of doxastic attitudes. But they do not necessarily reduce the justifying power of any objectively good reasons for particular doxastic attitudes.

Yet many externalists have not agreed. Bergmann (2006) and Plantinga (1993) hold that believed defeaters are real defeaters. While Goldman (1979) had a uniform reliabilist account of positive and negative justification-relevant factors, he later holds something closer to this concessive view. Discussing a case where someone believes her visual powers to be impaired, he said: "What she believes, then, is such that if it were true, the beliefs in question... would not be permitted by a right rule system. Satisfaction of this condition, I now propose, is sufficient to undermine permittedness."<sup>47</sup> Something similarly troubling holds for other concessions to internalists. Externalists will allow experiences *per se* to be defeaters. Consider Goldman (2011: 272) discussing an example involving Sidney, who continues to believe

that it is sunny right now...despite the fact that he is walking in the middle of a rainstorm. Surely his current perceptual experience is a defeater for this belief [...].

---

<sup>47</sup>Goldman (1986: 111).



If we focus on fortunate subjects, this is plausible. Our visual experiences are reliable indicators of the facts. But it is easy to imagine subjects whose experiences are inaccessibly unreliable. The intuition remains that Sidney's twin would be irrational if he kept the belief while finding himself faced with experiences like the ones had by Sidney. Could this render the twin's belief unjustified if it was previously justified? If the experiences are unreliable indicators of the facts, it is hard to see why we should agree.

It is better to explain the defeat intuitions by appeal to rationality and apparent reasons. Appearances of objective reasons *against* attitudes and acts are no less capable of being misleading than appearances of objective reasons *for* attitudes and acts.

### 3.5.4 Light Shed on Puzzling Cases

My distinction also helps to resolve conflicting intuitions about some puzzling cases.

**Checkered Experience.** Consider cases of experiences with checkered etiologies:

(*Fearful Sight*) Before seeing Jack, Jill fears that Jack is angry at her. When she sees him, her fear causes her to have a visual experience in which he looks angry.<sup>48</sup>

It is plausible that this experience's etiology undercuts its justifying force. Not only is it plausible that Jill would not be *ex post* justified in believing that Jack is angry: it is also plausible that the reason that experience provides for her *to believe* that Jack is angry is defeated because of the defective etiology. The intuition has nothing to do with Jill's access to the fact that her experience was merely caused by baseless fear. We can stipulate that she is in no position to appreciate this fact. The intuitions remain.

We can explain the intuitions in a framework where justification is understood in terms of possessed objective reasons. While unpossessed *rebutting* defeaters have no influence on whether one has an objective reason to believe something, unpossessed *undercutting* defeaters do. In *Fearful Sight*, the fact that Jill's experience is grounded in fear explains why that experience does not add to the stock of objective reasons she

---

<sup>48</sup>This is from Siegel (forthcoming). See also Siegel (2012).

has. It is clear why: experiences that are merely products of baseless fear are unreliable indicators of reality.

But it is unsatisfying to stop here. It is rational for Jill to increase her confidence that Jack is angry in *Fearful Sight*. My distinctions let us honor this intuition.

**Rational Self-Doubt and Supposed Examples of Higher-Order Defeat.** Our distinction casts doubt on the ubiquity of higher-order defeat. Some alleged examples of higher-order defeat rest on intuitions about the irrationality of not responding to certain higher-order appearances. It is unclear that agents in these cases automatically acquire objective epistemic reasons to modify their first-order attitudes. Moreover, even if the higher-order appearances *are* objectively good evidence for certain higher-order beliefs, worries about conflation remain. Let's consider these points in reverse.

Higher-order evidence against one's first-order competence generates no familiar kind of undercutting defeat. Consider a case from David Christensen (2010: 187):

*(Drugs)* "I am asked to be a subject in an experiment. Subjects are given a drug, and then asked to draw conclusions about simple logical puzzles. The drug has been shown to degrade people's performance in just this kind of task quite sharply. [...] I accept the offer, and, after sipping a coffee while reading the consent form, I tell them I'm ready to begin. Before giving me any pills, they give me a practice question:

*Suppose all bulls are fierce and Ferdinand is not a fierce bull. Which of the following must be true? (a) Ferdinand is fierce; (b) Ferdinand is not fierce; (c) Ferdinand is a bull; (d) Ferdinand is not a bull."*

I become extremely confident that the answer is that only (d) must be true. But then I'm told that the coffee they gave me actually was laced with the drug. My confidence that the answer is "only (d)" drops dramatically."

Clearly, one is irrational if one maintains the same degree of confidence in (d). This is Christensen's main intuition. Yet the objective probability that my belief in (d) is true conditional on the first-order evidence and the apparent fact that I took the drug is no less than the objective probability that my belief is true conditional on just the first-order evidence. So, it is puzzling how the higher-order evidence has any bearing

on what first-order beliefs I have objective epistemic reasons to have. It is less puzzling to think that there is a different issue here, separate from defeat at the first order.

We can reinforce this verdict by noting that intuitions of irrationality remain even if the subjects in these cases are not defective in the ways the appearances suggest, and even if the higher-order appearances are unreliable. Suppose that in a variation on *Drugs*, the experimenters lied and never drugged the coffee. This has no effect on the intuition that it is irrational for me to have the same degree of confidence in (d).

So there is a worry about conflation. Reasoning as the subject reasons in *Drugs* is correct! How can the irrationality of his total mental state have any bearing on whether he is permitted to reason in this way? The fact of irrationality has implications for whether he could be epistemically praised for doing the epistemically right thing. But just as there can be excusable wrongdoing, there can be praiseless and even blameworthy rightdoing.<sup>49</sup>

Those who believe in the ubiquity of higher-order defeat need deeper arguments. This is pressing, since they appeal to the irrationality of not responding to higher-order appearances to argue against plausible claims about justification. Schechter (2013), for example, rejects single-premise closure for justification on the basis of higher-order considerations. We should consider alternatives before rejecting such principles!

**Forgotten Evidence.** Our distinction also resolves conflicting intuitions about cases of forgotten evidence. Much of what we claim to know and justifiedly believe rests on reasons that have faded from view. One would invite skepticism if one insisted that we don't know or justifiedly believe in these cases. Externalists are right to insist on this.<sup>50</sup>

Nevertheless, there is a real tension in these cases that externalists neglect. They fail to explain, for example, why it would be irrational to keep believing P with a similarly high degree of confidence when one is asked why one thinks that P and realizes that one cannot bring to mind anything except the seeming that P.

---

<sup>49</sup>Cf. Zimmerman (1997).

<sup>50</sup>Again, Goldman (1999a) provides a classic illustration of how forgotten evidence cases raise troubles for internalists; Goldman has extended this critique in his (2009) with a new kind of forgotten evidence case (involving forgotten negative evidence rather than forgotten positive evidence).

Our distinctions ease the tension. When memory seemings are reliable, the fact that one seems to remember that P is a good objective reason for believing P. Still, when one is pressed, this seeming can easily be destroyed, and the objective force of the reason can cease to be apparent. When I consider whether I am justified in believing that Elizabeth I was born in 1533, doubts may arise. Ignoring these doubts is irrational.

Still, it is incredible to think that justification is so easily destroyed. So, it is better to apply distinctions. When doubts arise, the force of the objective reasons one has will cease to be apparent. So, it will be irrational to host the belief while the objective force ceases to be apparent. But one could revert to trust—whereupon the force will become apparent again. Hence, one faces no irreversible obligation to drop one's beliefs.

### 3.6 Concluding Remarks

Let's take stock. I've defended a distinction between justification and rationality. My distinction fell out of a general distinction between possessed objective reasons and apparent reasons. The distinction is grounded in a more fundamental distinction between competently treating something like an objective reason and aptly treating something like an objective reason. It is a consequence of the need to draw a competence/performance distinction with respect to our sensitivity to objective reasons.

One can see the distinction without seeing this more basic explanation. Examples of objective undercutting defeaters afford illustrations. These defeaters prevent apparent reasons from being objective reasons by destroying the objective favoring relations they apparently bear to relevant attitudes or actions. But one is still *rationally required* to take these apparent reasons into account, since the objective undercutter is not apparent.

It is worth remembering that unpossessed undercutting defeaters differ from unpossessed rebutting defeaters. The fact that there is a mountain of objective evidence against P that one has not yet discovered does not render one's belief that P unjustified. This is because the objective reasons that one possesses continue to be objective reasons, and are not outweighed by any other objective reasons that one possesses. To outweigh a reason is not to destroy it. Only total undercutting destroys.

This distinction explains why some rational beliefs are not justified beliefs without making the requirements for justification too strong. If an apparent reason that one has is objectively undercut, one's relevant attitude cannot be justified by that reason. But it is false that one is unjustified just because there is an unpossessed objective reason that outweighs the objective reasons that one possesses. This is why justified wrongdoing is possible, and why it is possible to have justified false beliefs.

It is a virtue of my version of the distinction that it does not make justification so demanding that we cannot have it unless our beliefs are true or constitute knowledge. This is what we should expect on general grounds. Justifiedly doing the right thing is best. But justified wrongdoing is still better than excusable wrongdoing.

The implications of these points for epistemology are significant, as we have seen. This is unsurprising. Parallel distinctions have caused important shifts in the dialectical terrain in recent ethics. I have only scratched the most obvious surfaces in this paper.

Before drawing things to a close, it is worth considering how one might try to resist my distinction. To do so, one must argue that there is something special about the norms of epistemology that reduces the distance between objective and apparent epistemic reasons. How could one defend this conclusion? How could there fail to be an appearance/reality distinction with respect to complying with a given norm?

There are many systems of norms for which the distinction applies undeniably—e.g., legal norms, norms of etiquette and prescriptive grammar, and many moral norms. But there are *some* cases where the distinction is hard to draw. Consider the norms of loyalty that are constitutive of some relationships. The following example illustrates why the distinction is hard to draw with respect to these norms:

(*Disguise*) A and B agreed to have a monogamous relationship. But A worries that B would cheat if B got the chance. A decides to test this hypothesis. With the help of some extraordinary costuming, A manages to dress up like a totally different person on whom B would have an instant crush. Disguised, A has been showing up around B's workplace to make advances. B believes on the basis of this misleading evidence that this is a fascinating person distinct from A. B now

seems to be having a date with this person when A had planned to be out of town....

Suppose now that A suddenly reveals the truth to B and demands an explanation. The following would not be a convincing response on B's behalf: "But look, I was not *unloyal* to you. After all, it was *you* I just showed to a fine evening!"

Why is this unconvincing? Because whether we manifest loyalty is determined by how we respond to the appearances. One really fails to manifest loyalty if it appears to one that some option would involve disloyalty but one pursues the option anyway.

There is a further point. Suppose B is a loyal partner. Would it cast doubt on B's loyalty if B had a romantic evening with someone who managed to look and act just like A, and whom B rationally believed on this basis to be A? No. If A were unaware that the person B was having the romantic evening with was dressed up this way, A could reasonably demand an explanation. But B *could* show that there was no failure of loyalty.

Are there other norms like this norm of loyalty? What is the basis for such norms? Remember that all norms are value-based: the point of complying with a norm is to respond to the value of something. Not all values are to be valued in the same way. Some values call for an internal kind of valuing. They call for us to hold certain attitudes in response to the appearances. Consider beauty. It calls primarily for admiration, and whether we admire beauty is determined entirely by how we respond to the appearances. Admiration thus differs from external forms of valuing, such as instrumental promotion.

Are there values that fundamentally call only for internal recognition? Some have thought so. Kant took the value of humanity to call most fundamentally for a kind of recognition by the good will. One could imagine a structurally similar view about epistemic value. One could imagine a theory of epistemic value that would affirm that accuracy is the fundamental epistemic value but hold that the fundamentally proper response to this epistemic value is respect, not promotion. What is respect for accuracy? Precisely what we manifest when we comply with requirements of epistemic rationality, one might say.

If accuracy is a fundamental epistemic value, I suspect one must be an epistemic Kantian if one wants to collapse the distinction between apparent and objective epistemic reasons. The only other option is to insist that rationality has fundamental epistemic value. But I think it is clear that we care about epistemic rationality because we care about accuracy. It is better to agree that accuracy is the fundamental epistemic value but to adopt a Kantian view about its value. This is an internalist way to make epistemology truth-oriented.

This view might be defensible. But clearly, it will take something radical to collapse the distinction between apparent and objective epistemic reasons. Such radical views have yet to be defended. Accordingly, work remains for those who want to collapse the distinction between epistemic rationality and justification.

## Chapter 4

### Veritism without Instrumentalism

#### Overview

It is attractive to view true belief as the fundamental epistemic value (“Veritism”). Epistemologists often take this view to entail that all other epistemic items can only have worth by standing in certain instrumental relations—e.g., by tending to produce a high ratio of true to false beliefs or by being products of sources with this tendency. Yet many value theorists deny that all *non-fundamental* value is grounded in *instrumental* relations to fundamental value (“Instrumentalism”). I argue that Veritists can and should reject Instrumentalism. By doing so, they can solve a generalized version of the swamping problem. Indeed, the moral of that problem is that Instrumentalism is false: if we disagree, similar problems will confront *any* economical epistemic axiology. To make things concrete, I sketch a non-Instrumentalist version of Veritism inspired by Thomas Hurka’s axiology and show that it avoids the swamping problem. While this is not the only way to be a Veritist without being an Instrumentalist, it is a promising way that needn’t rest on psychologically exacting views about rationality or knowledge.

#### 4.1 Veritism, Instrumentalism, and Fundamental Epistemic Value

“[T]he proper appreciation of a beautiful object is a good thing,” wrote G. E. Moore.<sup>1</sup> Yet while beauty and its proper appreciation are both good, it is implausible that they are equally fundamental goods. Plausibly, appreciating beauty is good *because* beauty is good. After all, it is only good to appreciate what merits appreciation. Appreciating trash is no good. So, although it is valuable, the appreciation of beauty seems to have

---

<sup>1</sup>Moore (1903: Ch. VI, §114).



a *parasitic* or *non-fundamental* kind of value.<sup>2</sup>

This is an example of a more general fact. In all evaluative domains, some values are more fundamental than others, in the sense that their value explains the value of the others. So, for any domain, it becomes natural to ask the *Fundamentality Question*:

(FQ) Which value or set of values is the most fundamental in the domain?

Ethicists have long addressed FQ as it arises in the practical domain. Epistemologists have recently taken interest in FQ as it arises in the epistemic domain. There are, after all, many epistemic values: accuracy, rationality, justification, coherence, knowledge, etc. But it is doubtful that they are equally fundamental. We admire some of them from the epistemic point of view *because* we admire others from the epistemic point of view.

So, which of these values is fundamental? Many have found it attractive to think that truth is at the bottom of it all. For plenty of items we value from the epistemic point of view, it is plausible that we value them *because* we value accuracy in belief. Inspired by this idea, one might favor:

(*Veritism*) True belief is the sole fundamental epistemic value.<sup>3</sup>

Many epistemologists have accepted Veritism, though it finds considerable opposition in recent literature. I also accept the view. But I think opponents and proponents alike understand it in an unjustifiably narrow way. Veritism is defensible if and only if it is understood less narrowly. The aim of this paper is to explain and defend this claim.

**The Narrow Assumption: Instrumentalism about Derivative Value.** Many epistemologists assume that there is only one kind of way in which we can explain one epistemic value in terms of a more fundamental epistemic value. The explanation, they assume, must proceed by invoking *instrumental relations*, so that for any X, X is

---

<sup>2</sup>I will be using ‘non-fundamental’, ‘parasitic’, and ‘derivative’ synonymously throughout.

<sup>3</sup>While this view is held by many epistemologists, its most prominent defender in recent years has been Alvin Goldman; see especially Goldman (1999b) and “The Unity of the Epistemic Virtues” in Goldman (2002).

derivatively epistemically good only by (i) tending to produce fundamental epistemic goods or (ii) being the product of something with feature (i).<sup>4</sup> Call the kind of value that something has just in virtue of (i) *production value*, and the kind that something has just in virtue of (ii) *mere product value*. More officially, then, many epistemologists have assumed:

*(Instrumentalism about Derivative Value)* All derivative epistemic value is just instrumental epistemic value of either the production or mere product kind.<sup>5</sup>

As a result, many have presupposed that Veritism is trivially equivalent to:

*(Instrumentalist Veritism)* True belief is the sole non-instrumental epistemic value. Everything else at best has some species of instrumental epistemic value.

None of this is trivial, however. Many value theorists have rejected Instrumentalism, including some who are consequentialists about rightness. Accordingly, it is only natural to wonder whether Veritism can take a less narrow form.

But what could derivative value be, if not instrumental value? The example with which I opened provides one illustration. Appreciating beauty is good *because* beauty is good. But this ‘because’ signals no instrumental relation. Appreciating beauty does

---

<sup>4</sup>One might reserve ‘instrumental epistemic value’ for (i). But I use it to capture the kind of derivative value that something has *just by standing in some instrumental relation*. The *product of* relation is just as much an instrumental relation as the *caused by* relation is a causal relation. Since many epistemologists allow that (ii) can ground derivative epistemic value, it is only fair to define Instrumentalism this way.

<sup>5</sup>While this thesis is typically presupposed rather than explicitly stated (and it is presupposed by virtually everyone writing about epistemic value), Goldman and Olsson (2009) and Goldman (2012) do explicitly favor a view that they call *type instrumentalism*. I intend the arguments in what follows to cover this doctrine. See Appendix B for a detailed discussion of whether the type/token distinction will help with the problems I’ll be discussing; while I do think it makes a difference for some problems, I don’t, as I argue, think it affects the points made in this paper.

It is, however, unclear to me whether Goldman (and Olsson) really do accept the sweeping version of Instrumentalism I’m attacking here, for reasons I explain in the Appendices. They recognize the category of extrinsic final value, and some examples of this type of value are plausibly understood as examples of derivative value that is not instrumental. To the extent that Goldman (and Olsson) would accept that characterization, we may disagree less than it seems.

Even if that is so, the importance of separating Instrumentalism from Veritism is far from being appreciated in the literature. As we’ll see, Pritchard (2010, 2011) crucially assumes that Veritists must be Instrumentalists, and many other discussions of the swamping problem also make this assumption (e.g., Zagzebski (2004)). There is a widespread belief that Veritism is the source of the swamping problem, but this is simply false, as I’ll be arguing: if Veritists reject Instrumentalism about Derivative Value, they can avoid the problem.

not reliably cause more beauty to exist, and is rarely the product of anything beauty-conducive. Luckily, we are not forced to appeal to instrumental relations. On a more natural model, the instance of appreciation derives value from its object because (a) its object is good and (b) it is the way to value that good object. The result is appealing: appreciating beauty seems parasitically but *also* non-instrumentally good.

Here is another way to understand what is going on. For X's value to be explained by Y's, X must stand in some relation to Y. Call such relations *value derivation relations*. Instrumental relations comprise one species of value derivation relations. But not the only species. In the case of beauty and its appreciation, the relation that mediates the derivation is not any instrumental relation but rather the relation of *being a way to value*.

**Hurka's Principle.** Some ethicists have constructed axiologies that honor this idea. Thomas Hurka, for example, defends a "recursive" account according to which there are *basic* non-instrumental values such as beauty and pleasure, and a principle for grounding *derivative* non-instrumental values in the basic ones:

(*Hurka's Principle*) When V is a non-instrumental value (in some domain), proper ways of valuing V (in that domain) and their manifestations have derivative non-instrumental value relative to V (in that domain).<sup>6</sup>

Given this principle, Hurka suggests that consequentialists about rightness can embrace the non-instrumental value of virtue while also respecting the thought that virtue has only a parasitic kind of value. They can do so by identifying virtues with proper ways of valuing more fundamental values and by making use of Hurka's Principle.

Note that there are two applications of Hurka's Principle. Consider an example to see both in play. Suppose Alice performs a charitable act because she cares about charity for its own sake, while Beatrice performs the same act as a PR stunt. Alice's

---

<sup>6</sup>A worry: can't everyone agree that proper ways of valuing are valuable, since 'proper' just means 'valuable'? Well, propriety is a kind of rightness. Rightness and goodness represent different species of normativity. There are long-standing debates about how they are related. The principle cannot be trivial for this reason. It is *non-trivial* for this reason! And the principle does not only say that proper ways of valuing are valuable. It says that they are derivatively valuable in a specific way—which is quite non-trivial.

act has greater worth than Beatrice's, and this extra worth is not merely instrumental. But the fact that Alice's act has greater worth is not just a brute fact. There is a two-stage explanation of this fact: (i) Alice's act has greater worth by manifesting real valuing of charity, and (ii) this valuing, in turn, is good because it is properly directed at something good (viz., charity). In stage (i), the relation of *manifestation* is the value derivation relation: the act inherits some of its glow from the valuing of charity that it manifests. In stage (ii), the relation of *valuing* is the value derivation relation. Charity is the more fundamental value.

Hurka's model is not the only model in value theory that rejects Instrumentalism. Other value theorists have rejected it.<sup>7</sup> I highlight Hurka's model only because it inspired the specific alternative to the instrumental model that I will defend. There are surely other ways for Veritists to avoid Instrumentalism than the way that I will pursue. I pursue this way mainly for purposes of detailed illustration. My fundamental point is more abstract, though the hard work lies in finding detailed implementations.

**A Caption of the View and the Plan.** My view will appeal to the two-part pattern illustrated above to explain why rational belief, justified belief, and knowledge have worth from an accuracy-oriented perspective. On my view, relations like rationally believing and knowing are epistemically valuable because they manifest certain *ways to place value on accuracy in thought*. I will understand the ways of placing value on accuracy in thought in a psychologically unexacting way: they need not be explicitly voiced, conscious, emotional, or global in the way that character traits are global. Indeed, they are things that children and animals can display, just by being disposed to think in certain ways.

The view is compatible with many substantive pictures of the nature of rational belief, knowledge, and other epistemic desiderata. I see the best pictures as simply disagreeing about what it takes to place value on accuracy in thought. My own view is that any way of placing value on accuracy in believing P involves a disposition to hold

---

<sup>7</sup>See Anderson (1993), Herman (1993), Kagan (1999), Korsgaard (1983), Moore (1903), Parfit (2011), Rabinowicz and Ronnow-Rasmussen (1999), Rashdall (1907), Scanlon (1998), and Zimmerman (2001).

the belief that P only if it is likely to be accurate relative to the epistemic reasons. There are *several* ways to place value on accuracy in thought because there are importantly different ways of understanding *likelihood* and *epistemic reasons*, which correspond to different epistemic desiderata. These ways and the corresponding desiderata qualify as derivatively epistemically good from an accuracy-oriented point of view thanks to Hurka's Principle.

With this capsule statement in mind, here is the plan. In §2, I explain why Veritism should not be identified with Instrumentalist Veritism. Instrumentalist Veritism faces a generalized version of the swamping problem. But this problem undermines Instrumentalism, not Veritism. For if we grant Instrumentalism, similar problems arise for *any* economical epistemic axiology. I show in §3 how Veritism could take a less narrow form and avoid the swamping problem. After answering objections in §4, I show in §5 why the apparent alternatives either fail or collapse into my view.

**Disclaimers.** Before proceeding, let me make a few remarks about what I am *not* doing.

Denying Instrumentalism is compatible with affirming that some derivative epistemic values are purely instrumental. Reliable types of processes have significant instrumental epistemic value, and reliabilists are right to place value on them.

Denying Instrumentalism is also compatible with the idea that some items might have both instrumental epistemic value and a different kind of derivative epistemic value. Being instrumental to accuracy, I agree, is a *necessary* condition for some epistemic virtues. I only deny that Instrumentalism can explain all the facts about derivative epistemic value. I will sketch a different model of epistemic value derivation. But I intend this model to *supplement* the instrumental model, not replace it.

Furthermore, I am not arguing in this paper against reliabilism. Reliabilism is a thesis about the nature of justified belief or—in some permutations—knowledge. Instrumentalism is a thesis about how derivative epistemic value is grounded in fundamental epistemic value. Reliabilism does not entail this thesis. Indeed, reliabilism does not by itself entail any claims about the worth that any beliefs possess from the epistemic point

of view. One would need to appeal to auxiliary assumptions about epistemic worth and the relationship between derivative and fundamental epistemic worth to derive any such conclusions.

Many have worried that reliabilists can *only* appeal to an Instrumentalist model of derivative epistemic value. But I think this is wrong: there are sophisticated forms of reliabilism that are not forced to accept a purely Instrumentalist model. Admittedly, I will discuss a simple form of reliabilism when I introduce the generalized swamping problem for Instrumentalism in the next section. But this is only as a nod to the literature.

Finally, my focus is on epistemic *value*. Here I will not defend a non-consequentialist view about epistemic *rightness*. Indeed, Hurka's view in ethics was intended to provide a way for consequentialists about ethical rightness to consistently uphold the non-instrumental value of virtue. By accepting a richer axiology, consequentialists can uphold a less revisionary account of our duties. This idea is old, found in ideal consequentialists of the early 20th century like Hastings Rashdall and G. E. Moore. While I do reject consequentialism about epistemic rightness, this paper is not an argument against it.

## 4.2 Instrumentalism and the Swamping Problem

Why should Veritists reject Instrumentalism? A large reason is that this is the best way to avert a generalized version of the swamping problem. To bring this out, I will rehearse the original swamping problem and explain why we should still take it seriously. I will then explain why it undermines Instrumentalist Veritism. But I will argue that the Instrumentalist half is the culprit. This is because *any* modest epistemic axiology that embraces Instrumentalism will face a relative of the swamping problem, and Veritists can easily solve the problem by rejecting Instrumentalism.

### 4.2.1 The Old Problem (and Why It Remains Important)

Originally, the swamping problem was presented as a problem for a simple kind of reliabilism. The following thoughts prompted the problem. The epistemic value of a reliable type of belief-forming process *per se* is just instrumental epistemic value relative to the goal of producing a high ratio of true to false beliefs. While reliable belief-forming processes have great epistemic value, one also wants to evaluate their products. And unfortunately, the following claim is plausible:

(A) A belief's having been produced by a reliable type of belief-forming process does not *as such* make that belief epistemically better if that belief is already true.

(A) is made plausible by an analogy from Zagzebski (1999). The mere fact that some good coffee came from a reliable type of coffeemaker does not make that *coffee* better. But considered merely as such, a reliably produced true belief seems analogous to a reliably produced cup of good coffee. Thus:

(B) If knowledge = true belief produced by a reliable type of belief-forming process, then knowledge is not as such epistemically better than true belief.

But knowledge is as such epistemically better than true belief. Hence the problem.

Some might try to resist this argument by observing that we do sometimes place greater value on products of reliable sources. For example, the market value of a good watch produced by Rolex would be far higher than that of an unusual, qualitatively identical watch produced by Casio. Doesn't this undermine the swamping argument?<sup>8</sup>

No. Intuition pumps like this do support the conclusion that some products of reliable sources are better than intrinsically similar products of unreliable sources. But this conclusion does not undermine the reasoning behind the swamping problem. Here is why. The key thought behind the swamping argument is this:

---

<sup>8</sup>For this type of reply, see Goldman and Olsson (2009) and Goldman (2012: Introduction). Goldman's ultimate reply contains some further details that I discuss in the Appendices.

(\*) A good F's having come from a reliable source of good Fs cannot **as such** make it better.

The 'as such' matters. It is consistent with (\*) that other properties can *contingently accompany* a good F's reliable ancestry and add value to that F. The Rolex/Casio example is an illustration: Rolexes are status symbols. That is why the Rolex is preferable. If it were not a status symbol (say, because we all had Rolexes), it would be irrational to prefer the Rolex to a qualitatively identical but cheaper Casio.

Proponents of the swamping argument will agree that reliably produced good Fs are sometimes better than unreliably produced good Fs. Zagzebski (1996), for example, agrees that reliability is a necessary condition for complete epistemic virtue.<sup>9</sup> Thus, she agrees that some reliably produced true beliefs are better than mere true beliefs: some reliably produced true beliefs will satisfy the *other* necessary conditions for manifesting complete epistemic virtue, after all! What Zagzebski and others deny is that reliable ancestry *as such* is sufficient to explain why knowledge *as such* is epistemically better than true belief.

That was the problem all along. So, the conclusion (i.e., (B)) from the original literature stands. It is insufficient to observe that a reliably produced F is sometimes better than an unreliably produced but intrinsically similar F. Zagzebski and others will agree but deny that the explanation of why one is better than the other proceeds *via* the bare fact of reliable ancestry. That was the point of the coffee analogy. The analogy forces us to ask what else could make the difference, since reliable ancestry *as such* seems insufficient.

This isn't to deny that reliable processes have great epistemic value. The problem has always concerned their products. As Jones (1997: 425) put it:

When we ask about the value of knowledge over mere true belief, we are asking not about belief-forming methods but about their products. We have clear reason to care about reliable methods of belief-formation, and the reliabilist is right to emphasize them. But it is unclear, on reliabilism, why we should divide the desired products of belief-formation,

---

<sup>9</sup>See Zagzebski (1996: 165-194), and especially the sections entitled "The Success Component of the Intellectual Virtues" and "Montmarquet on the Virtues and Truth-Conduciveness."



true beliefs, into those which were brought about by reliable methods and those that were not. The epistemic instrumentalist gives us no way of evaluatively dividing true beliefs into those which have been brought about by justifying methods versus those which have not.

This challenge remains a serious one.

#### 4.2.2 The Deeper Problem

The swamping problem runs deeper than this, however. It is easy to see that the problem is not just a problem about knowledge or just a problem for reliabilists.

To see the first point, note that reliabilists identify a belief's being justified with its being formed by a reliable type of belief-forming process. Assume they are right for the sake of argument. We can use the same coffee analogy to argue that *justification* cannot as such add epistemic value to true belief. That is bad: justification as such does add epistemic value to true belief! This is not a restatement of the problem about knowledge: due to the Gettier problem, no reliabilist will equate knowledge with justified true belief.

There is an even more general structure that makes the problem not just of limited interest to reliabilists. On any view on which

being justified : true belief :: being made by a good coffeemaker : good coffee

there is a worry that justification cannot as such add epistemic value to true belief. So, there is a worry for any view on which the epistemic value of justified belief consists in its being the mere product of some type of source that is only instrumentally good relative to true belief. Reliabilists are hardly the only epistemologists who accept this assumption.

These observations lead to a general argument against Instrumentalist Veritism, as Duncan Pritchard (2010, 2011) in effect observed. Instrumentalist Veritism entails that:

(A) Items other than true belief are epistemically good iff they (i) instrumentally promote true (and not false) belief or (ii) are products of a type of source with property (i).

And the following is a natural generalization of the points about the coffeemaker case:

*(B: Swamping Premise)* If X has its source in something that is only instrumentally good relative to property F and X already exemplifies F, the mere fact that X has that type of source cannot as such make X better.

But according to Instrumentalist Veritism, justifying and rationalizing sources only have instrumental epistemic value relative to true belief. So (A) and (B) will entail:

- y. A true belief's being justified cannot as such make that belief epistemically better.
- z. A true belief's being rational cannot as such make that belief epistemically better.

Now, Pritchard took this to undermine Veritism *simpliciter*. But this is only because he presupposes Instrumentalism. Indeed, in setting up his argument, Pritchard (2011: 245) says he will assume that “[a] good is fundamental...if its value is not instrumental relative to further goods of the same type.” This is equivalent to Instrumentalism.

If Pritchard were right to make that assumption, Veritism would suffer. It would collapse into Instrumentalist Veritism, and the generalized swamping problem would undermine it. But Veritists can reject Instrumentalism and view the generalized swamping problem as an argument against it. To support this tactic, I will show that the Swamping Premise does not extend to other species of derivative value. This will show that the problem is a local one for Instrumentalists. To drive the point home, I will show that there are similar problems for all modest alternatives to Veritism that cling to Instrumentalism.

### 4.2.3 Why Other Forms of Derivative Value Help

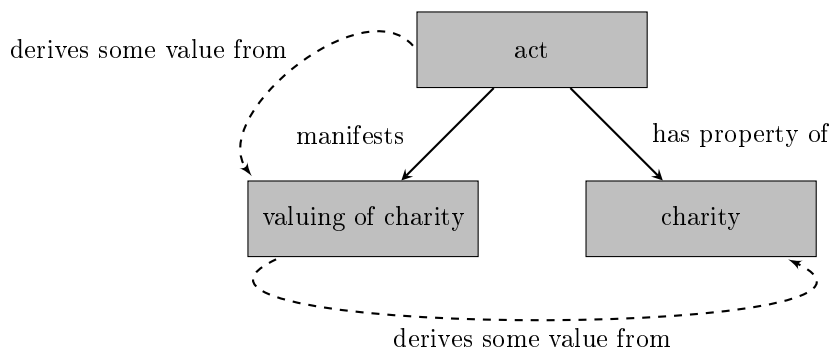
Let us first understand why rejecting Instrumentalism can help Veritists. Doing so can help them, I suggest, because other forms of derivative value are not subject to swamping by the values from which they derive value. To see this, consider:

(*Stronger Swamping Premise*) For no type of derivative value is it true that:

if X has its source in something that only has derivative value relative to property F but X already has F, X's having that source *as such* makes X better.

This is false. Recall the phenomena that motivate Hurka's Principle. Alice performs an act of charity because she values charity, while Beatrice performs the same kind of action as a PR stunt. Plausibly, Alice's action is better than Beatrice's because it has its source in real concern for charity. Yet just as appreciating beauty is good *because* beauty is good, so valuing charity is good *because* charity is good. It is just that the 'because' here signals no instrumental explanation. Instead, the explanation goes by Hurka's Principle.

Here is a case where an act derives value from a source that is only derivatively good relative to another property that the act exemplifies. The source is (1) Alice's valuing of charity, and the property exemplified is (2) charity. Yet it is plausible that the sheer fact that Alice's act has its source in (1) makes that act worthier *per se*. In a picture:

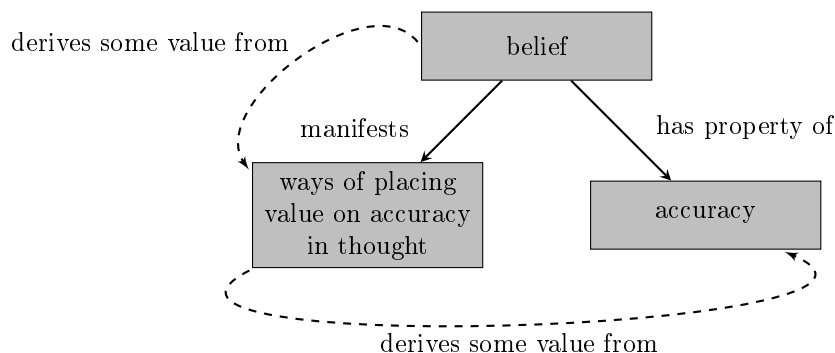


This would be impossible if the Stronger Swamping Premise were true. So it is false.

If the foregoing points are right, Veritism *per se* generates no problem: only its conjunction with Instrumentalism does. If Veritists adopt a different model of value derivation, they can avoid swamping just as it was avoided in the case of charitable action.

On the view I develop in §3, rational belief and knowledge gain epistemic worth by manifesting ways of placing value on accuracy in thought. On this view, the case

of knowing parallels the case of Alice’s charitable action. Knowledge requires a belief whose accuracy manifests a disposition to hold beliefs only if there is sufficient objective and subjective evidence that they are true. This disposition *just is* a way to place value on accuracy in thought. Ways of placing value on accuracy are epistemically good because accuracy is epistemically good. But this ‘because’ is not purely instrumental. Paralleling Alice’s case:



I’ll explain the ways of placing value on accuracy in thought in §3. The tactic is what matters now. Veritists can claim that true beliefs can gain epistemic worth by manifesting ways of placing value on accuracy, just as charitable actions can gain moral worth by manifesting valuing of charity. Both cases undermine the Stronger Swamping Premise.

So Veritists can avoid the swamping problem if they appeal to forms of derivative value beyond the instrumental. But they must avoid the problem in this way. For the restricted Swamping Premise *is* plausible. Instrumentalist Veritism *is* false. But since it is plausible that epistemic evaluation is truth-oriented, we should only reject the Instrumentalist half.

#### 4.2.4 A General Problem for Instrumentalists

We can strengthen this advice by seeing that if Instrumentalism were true, other economical epistemic axiologies would face relatives of the swamping problem.

Observe that there is another side to Zagzebski’s analogy. We do not only think that coming from a reliable coffeemaker *per se* cannot improve good cups of coffee. We also think that coming from a reliable coffeemaker *per se* cannot improve bad cups of

coffee. If you drink some vile coffee and cringe, it is not comforting to be told: “Hey, at least it was produced by a reliable coffeemaker.” These facts support a more striking sibling of the restricted Swamping Premise:

*(Dud Principle)* If X was produced by a source that is only good because it produces good Fs, that fact as such can’t make X better if X is otherwise a bad F.

Carter and Jarvis (2012) took this to undermine the intuitions behind the swamping problem. For they thought that the Dud Principle would imply, crazily, that “non-factive epistemic properties—most saliently justification—are never epistemically valuable”. But they were hasty: that crazy conclusion follows only if we grant Instrumentalism.

Like the original Swamping Premise, the Dud Principle concerns a kind of instrumental value. The crazy conclusion would not follow from the Dud Principle if non-factive epistemic properties had a different kind of derivative epistemic value. Only if the Dud Principle extends to other kinds of derivative value is Carter and Jarvis’s conclusion fair.

But the Dud Principle does not generalize, for the same reasons why the Swamping Premise did not generalize. Consider someone trying to perform a charitable act just because this person values charity and failing only due to bad luck. Her efforts remain better than the failed efforts of someone merely looking for a PR boost.

What the Dud Principle really suggests is that Instrumentalism is an incomplete model. After all, it is not as if Carter and Jarvis can convince us that vile coffee is better if it comes from an otherwise reliable coffee machine. Like the original Swamping Premise, the Dud Principle captures a fact about products of merely instrumentally good sources.

But it is then easy to see that any modest axiology that embraces Instrumentalism will face a relative of the swamping problem. Suppose, for example, that knowledge is one’s fundamental epistemic good. Given Instrumentalism, how can one explain the epistemic value of justified beliefs? One must claim that such beliefs are good

by being products of knowledge-conducive types of processes. But the Dud Principle will make it mysterious why justified false beliefs are epistemically good. From an Instrumentalist point of view, they are duds just like bad cups of coffee from otherwise reliable coffeemakers.

Expanding the stock of fundamental epistemic values doesn't really help. Even if one adds justified belief, knowledge, understanding, and true belief to the list, there remain epistemic values that (a) no one can reasonably take to be fundamental, but (b) admit of no Instrumentalist explanation. Consider the epistemic value of *trying one's best to form beliefs accurately*. This is a *paradigmatically derivative* value: trying to do something good is generally admirable *because* it is intentionally directed at something good. But if we accept Instrumentalism, it is hard to explain why it is derivatively epistemically good. Merely trying to form one's beliefs accurately is not reliably instrumental to accuracy, knowledge, justification, etc. One might be an amateur whose best efforts can't reliably yield these results. But there remains something admirable in one's best efforts. This is not evidence of a new fundamental value. This value is paradigmatically derivative.

Any modest epistemic axiology that endorses Instrumentalism—even a pluralist one that places justified belief, knowledge, understanding, and true belief in the bedrock of value—will face a relative of the swamping problem. There are paradigmatically derivative epistemic values that admit of no general Instrumentalist explanation. So, everyone should reject Instrumentalism. Accordingly, it is not *ad hoc* for Veritists to do so.

### 4.3 Veritism without Instrumentalism

Of course, work remains for Veritists. They must offer us a more specific view that takes advantage of our observations about forms of value derivation beyond the instrumental, and explain how this view can capture central intuitions about epistemic value.

My view will secure these desiderata. I'll start by explaining two key ideas. The first idea is that there are more ways to value something than by instrumentally promoting

it. This holds for accuracy as much as it holds for other values. The second idea is the part of Hurka's Principle which says that manifestations of proper ways of valuing X are derivatively valuable relative to X. With these ideas in hand, I'll suggest that central epistemic values like coherent belief, rational belief, and knowledge can be viewed as manifesting different ways of placing value on accuracy, and as deriving accuracy-oriented epistemic value thanks to Hurka's Principle. I'll then turn to answer some objections in §4.

### 4.3.1 Other Ways of Valuing in General

Let's start with a more general fact. It is a truism that values are items that it is proper to value. Given the truism, one should ask for any fundamental value V in any domain:

(FVQ) Which ways of valuing V are the proper ways in the domain?

Instrumentalists can be cornered into accepting a surprisingly narrow answer to FVQ:

*(The Teleological Answer)* For any fundamental value V, the only basically proper way to value V is to instrumentally promote V.<sup>10</sup>

For suppose the Teleological Answer is false. Then there are basically proper ways to value something good other than producing it. Since proper ways of valuing are derivatively good, it follows that there are derivative values that are not instrumental values. This contradicts Instrumentalism. So, if the Teleological Answer is false, Instrumentalism is false. Thus, if Instrumentalism is true, the Teleological Answer is true.

But the Teleological Answer is implausible. There are many ways to value: being dedicated to, protecting, respecting, being loyal to, taking delight in, etc. This answer regards only one as basically proper to fundamental value. This is hardly a default view. For some values, instrumental promotion is not the basically proper response. If friendship were fundamentally "to be promoted", we could properly spend less time

---

<sup>10</sup>I say 'basically proper' because the instrumentalist can obviously allow that other ways of valuing V are non-basically proper in virtue of helping to instrumentally promote V.

caring about the friends we have and more time amassing friends or causing others to have more friends. This misunderstands the value of friendship, as others have stressed. We can properly value friendship without taking ourselves to have reasons to produce more instances of friendship.<sup>11</sup> Fans of the Teleological Answer could multiply fundamental values in reply, and claim that loyalty, commitment, respect, dedication, etc., are all fundamental values. But this pluralism is implausible. Besides violating canons of parsimony, it fails to explain asymmetries. Dedication to one's friends matters *because* friendship matters.

These points highlight an insight needed for a version of Veritism that rejects Instrumentalism. The Veritist should capitalize on this insight and suggest that there are more ways to place value on accuracy in thought than by producing a high ratio of true to false beliefs by any means, including means that disrespect accuracy.

### 4.3.2 Ways to Place Value on Accuracy in Thought

How can we implement the insight? What are the ways to place value on accuracy in thought? I understand the ways that are suited to play a role in traditional epistemological projects as different ways to honor the following ideal of accuracy:<sup>12</sup>

(AI) It is correct to believe P iff P is true.

AI is not a directly belief-guiding norm. But we can honor this norm indirectly. By doing so, we place value on accuracy in our thinking. Coherent beliefs, rational beliefs, and knowledge can be viewed as epistemically good from a truth-oriented point of view by manifesting different ways to place value on accuracy in thought.

---

<sup>11</sup>Cf. Scanlon (1998: 89): “[T]he primary reason to be loyal to one's friends is not that this is necessary for the friendship to continue to exist.... We would not say it showed how much a person valued friendship if he betrayed one friend in order to make several new ones, or in order to bring it about that other people had more friends.”

<sup>12</sup>I make no claims about ways of valuing accuracy that are not suited to play a role in these projects. Here I distance myself from responsibilists less conservative than Zagzebski. I also distance myself from any standard kind of responsibilism, since responsibilists typically stress intellectual character traits. My ways of valuing are not global like character traits. They are quite local. Here again I am close to Hurka, who also stressed in his (2006) that virtues need not amount to character traits.



I will discuss three ways that we can use to ground the epistemic value of coherent belief, rational belief, and knowledge from an accuracy-oriented perspective: commitment, respect, and compliance. I will draw comparisons with other ideals to illustrate the more general pattern of evaluation. It is worth noting that the three ways of honoring ideals that I will discuss are the ones intuitively missing in cases of *recklessness*, *negligence*, and *non-compliance*, respectively. While I emphasize the connection with these pre-theoretically important concepts in discussing the examples, the connection is worth stressing separately to forestall the worry that the discussion is purely stipulative.

### **Commitment**

To see the first way to honor AI, consider other ideals, like this ideal of politeness:

(PI) It is correct to say Q to some casual interlocutor only if saying Q would not make this interlocutor pointlessly uncomfortable.

Imagine that Edward is trying to make you pointlessly uncomfortable by saying Q. He may fail: unbeknownst to him, Q may put you at ease. If so, he conforms to PI. Still, he is clearly reckless with respect to PI, and manifests a lack of commitment to PI.

Being committed to an ideal is part of what it takes to honor it. Commitment to an ideal in  $\phi$ -ing is grounded in a disposition to  $\phi$  only if one takes there to be sufficient evidence that  $\phi$ -ing is in conformity with that ideal. When we criticize people for recklessness or hypocrisy with respect to some ideal, it is this we find lacking.

Like any ideal, the ideal of accuracy also calls for commitment. To see what commitment to AI requires, consider someone who takes himself to have conclusive evidence that he believes P inaccurately but believes P anyway. He does not place sufficient value on accuracy in his thinking. He might believe accurately anyway. But he fails to manifest commitment to AI in believing P, like how Edward failed to manifest commitment to PI in saying Q. This is why he doesn't fully honor AI.

Commitment to AI lines up with one interesting epistemic value—viz., structural epistemic rationality. Complying with all the norms of structural epistemic rationality

is the way to be committed to the ideal of accuracy. Coherent agents may fail to place value on accuracy in some ways, but not by lack of commitment to the ideal of accuracy.

## Respect

Coherence is not the whole of epistemic rationality. Epistemic rationality also has a substantive side. Substantive epistemic rationality involves a stronger way to honor the ideal of accuracy than commitment, though it does not require conformity.

To understand this, it will again prove helpful to compare other ideals. Consider the ideal of politeness again. This time imagine Edna, who seems to succeed by her own lights with respect to PI. She is confident that her interlocutor likes talking about X and so intends to bring up X. Unfortunately, Edna neglects certain apparent evidence. She has a seeming memory that X makes her interlocutor uncomfortable but disregards it as misleading. This apparent memory might really be misleading. If so, Edna conforms to PI. But she is negligent with respect to PI, and fails to *respect* PI in one natural sense.<sup>13</sup>

Respect is another way to honor an ideal. To respect an ideal in  $\phi$ -ing is to manifest a disposition to  $\phi$  iff the apparent evidence indicates that  $\phi$ -ing would be in conformity with the ideal. When we criticize people for being negligent with respect to an ideal, it is this we find lacking. Like any ideal, the ideal of accuracy calls for respect. Consider someone who mistakenly takes himself to have conclusive evidence for P by randomly disregarding obvious evidence against P as misleading. While he might still believe something true, he partly fails to place sufficient value on accuracy in believing P, just as Edna partly failed to place sufficient value on politeness in bringing up X.

Respect for AI lines up with another interesting epistemic value—viz., substantive epistemic rationality. Being substantively epistemically rational is the way to respect the ideal of accuracy. Substantively rational thinkers may fail to place value on accuracy in their thinking in other ways, but not by a lack of respect, in one intuitive sense of ‘respect’.

---

<sup>13</sup>There are affinities here with Darwall (1977)’s general notion of recognition respect.

## Compliance

Even if we are fully committed to some ideal and respect it, we may not honor it fully. To honor an ideal fully, one must actually conform to it. Not all non-conformity manifests disrespect or lack of commitment to the ideal. Imagine someone falsely telling you that a topic does not make her uncomfortable. Perhaps she recognizes that you enjoy this topic and wants to oblige you. If you bring it up, you fall short with respect to PI. But not by disrespect or lack of commitment.

Of course, mere conformity is not a way to honor an ideal. Conformity can be the lucky product of negligence or recklessness. If you fall short because you do not fully honor PI, it is not just because of non-conformity. Rather, what you fail to do is something you are sadly in no position to do: namely, to *comply* with the ideal.

What is it to comply with an ideal? It is to conform with the ideal *by* respecting it. In cases like the one imagined, compliance is not open to you: only if it became apparent that the topic makes your interlocutor uncomfortable would compliance be open to you. This is why you are excusable. But compliance sometimes is open to us. When it is, it is the most proper way to honor the ideal. Like other ideals, AI also calls for compliance. Sometimes we can conform to AI *by* respecting AI, thereby complying with AI. When we do, we believe accurately *by* believing rationally. This way of placing value on accuracy lines up with another interesting epistemic value—viz., knowing. Knowing that P just is a way to comply with the ideal of accuracy with respect to P.

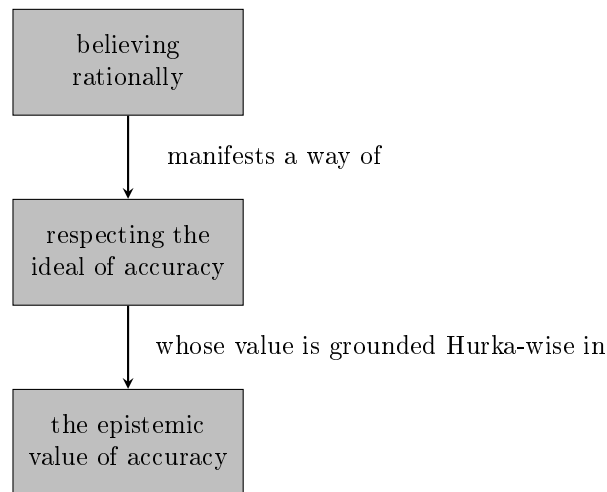
### 4.3.3 Deriving Epistemic Value via Hurka's Principle

Let's summarize the connections from the last few subsections in a table:

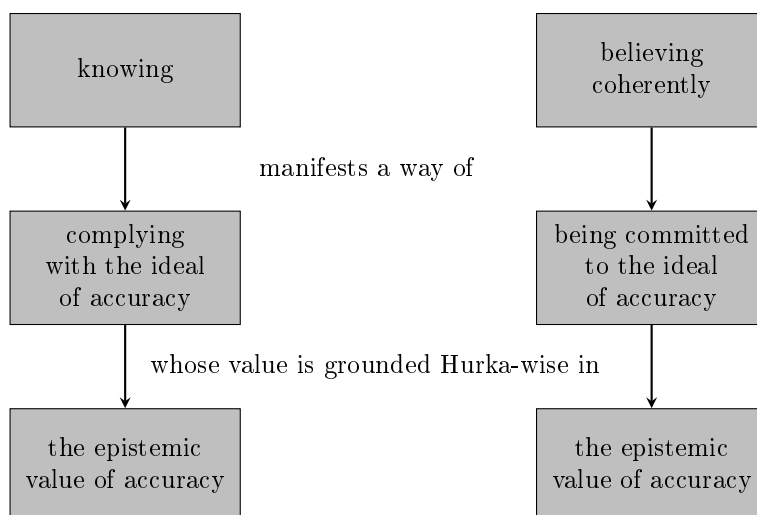
Way of Valuing	Manifested By	Epistemic Property
<i>Commitment</i>	Belief-Relative Likely Accuracy	Coherence
<i>Respect</i>	Appearance-Relative Likely Accuracy	Substantive Rationality
<i>Compliance</i>	Real Accuracy <i>via</i> Respect	Knowledge

Once we see that these central epistemic values reflect different ways of placing value on accuracy in thought, we are only a step away from an accuracy-oriented account of their epistemic value. The other ingredient we need is Hurka's Principle.

Given Hurka's Principle, we can ground the epistemic value of other properties from the perspective of accuracy. Consider substantive rationality. It requires sensitivity to the apparent evidence bearing on one's accuracy. As I suggested in the last section, that is a way to place value on accuracy in thought. If so, we can use Hurka's Principle to explain the epistemic value of substantively rational belief in accuracy-oriented terms. In a picture:



The same story explains why coherence and knowledge have derivative epistemic value. Being coherent is the way to be committed to the ideal of accuracy. Knowing is the way to comply with the ideal of accuracy. Being committed to and complying with the ideal of accuracy are epistemically proper ways to place value on accuracy in thought. Accordingly, we can explain the epistemic value of coherence and knowledge from the perspective of accuracy just like we explained the epistemic value of substantive rationality. In pictures:



There are epistemic values other than coherence, rationality, and knowledge. But I hypothesize that every properly epistemic value that is not just instrumentally valuable relative to the others that I discussed will be explicable in something like this framework. In §4, I will address objections to this hypothesis. For now, I am happy to have derived the epistemic value of three central properties from the perspective of accuracy.

#### 4.3.4 Comparative Propriety and Rankings of Value

The story so far leaves open the ranking of epistemic values. Commitment, respect, and compliance are all epistemically proper ways to honor AI. If so, coherent belief, rational belief, and knowledge are all derivatively good from the perspective of accuracy.

Yet one would like to rank these epistemic values. Coherent belief, rational belief and knowledge are not equally epistemically good. One wants to say that knowledge is better than substantively rational belief, and substantively rational belief is better than coherent belief. How should we explain these comparisons?

Although commitment, respect and compliance are all epistemically proper ways to honor the ideal of accuracy, they are not equally proper. For any ideal in any domain, it is most proper to comply with an ideal. If one must fail to comply, it is more proper to do so with respect than without. If one must fail to respect an ideal, it is more proper to do so while remaining committed to the ideal. These general differences explain why

knowledge > substantive rationality > coherence

After all, for any ideal, compliance is a more proper way to honor an ideal than respect, and respect is more proper than commitment. Since what is best from the perspective of an ideal is explained by what is most proper to that ideal, we can say that knowledge is best from the perspective of AI. So, compatibly with accuracy's being the *fundamental* epistemic value, we can agree that knowledge is epistemically *best*.

This is not paradoxical. Compliance is the most proper response for any system of ideals. This does not mean that there is a further ideal of compliance. In addition to a given ideal X, one does not need another ideal that says: "Comply with X!". There are just ideals, and it is a conceptual truth about ideals that compliance is more proper than mere respect and respect is more proper than mere commitment. In the epistemic domain, the fundamental thing is AI. Just as compliance is not a further fundamental ideal, knowledge is not a further fundamental value.

Is this illuminating? Yes. We can see the ranking of epistemic value as a special case of a general propriety ordering for ways of recognizing ideals. So, the account is far more illuminating than some view on which it is a bedrock truth that knowledge is better than substantive rationality and substantively rationality is better than coherence.

Have I presupposed what I was trying to explain? No. I appealed to general truths about the comparative propriety of certain ways of honoring ideals, not particular truths about the comparative goodness of the states that manifest these ways of honoring the ideal of accuracy. Was the appeal to truths about comparative propriety circular? Not as far as I can see. Propriety is a deontic concept, not an evaluative one. It is not trivial to draw a conclusion about comparative goodness from a comparative propriety judgment.

### 4.3.5 Advantages

The version of Veritism I've sketched has many advantages. Let's consider them.

Among versions of Veritism, mine is the first to explicitly honor the fact that value derivations need not be grounded in instrumental relations. This yields advantages. The Swamping Premise does not generalize to other forms of derivative value. Just as acts that manifest ways of valuing ethical ideals are *ipso facto* ethically better, so instances

of believing that manifest ways of valuing accuracy are *ipso facto* epistemically better.

My view inherits the broader virtues of Veritism. Veritism is a maximally simple axiology. Parsimony is as much a virtue in axiology as it is elsewhere. Supplemented with more sophisticated views about value derivation, Veritism also has great explanatory power. And its explanatory power is intuitive. It is intuitive that we care about things like rationality because we care about accuracy.

Since I deny that the ‘because’ must be understood instrumentally, I can vindicate this intuitive claim more directly than Instrumentalists. Forming beliefs in a way that appears likely to achieve accuracy does not entail reliable achievement of accuracy. Just consider demon worlds, where our experiences are the same but the demon ensures that they radically fail to match external reality. These worlds do suggest that some epistemic values cannot be explained in terms of objective truth-conducivity. But this does not show that not all epistemic values are truth-oriented. It just shows that truth orientation should not be understood merely instrumentally.

Isn’t rationality a way of “subjectively promoting” accuracy? Sure. But if we are Instrumentalists, it is unclear why this is more than a merely apparently good thing. Expected value is perhaps connected to obligation. But it is not a kind of value, just as expected wealth is not a kind of wealth. Yet rationality does not just appear to have epistemic value. It has real epistemic value, even in demon worlds. The only way to explain this short of an error theory is to view the “subjective promotion” of accuracy as *constituting something further*: a way to place value on accuracy.

#### 4.4 Five Objections Answered

A theory can have many virtues while facing conclusive objections. Having sketched my theory and explained its virtues, I will now answer five major objections.

#### 4.4.1 Too Demanding?

I claimed that coherence, rationality and knowledge are epistemically valuable in virtue of manifesting certain ways to place value on accuracy in thought. But must people really place value on accuracy in thought to believe coherently, rationally, or knowingly? Isn't this overly demanding?

*Reply.* This objection is itself nourished by overly demanding assumptions. In particular, it assumes that placing value on accuracy in thought is more demanding than it is.

It is plausible that when we think carefully by heeding all the apparent evidence bearing on the accuracy of our beliefs, we are placing value on accuracy in our thinking. This fits under a more general pattern. When we reason carefully and heed the apparent evidence bearing on whether we would conform to any standard by  $\phi$ -ing, it is plausible that our  $\phi$ -ing exhibits respect for that standard. It is plausible, then, that respecting accuracy is not psychologically demanding.

This idea is not strained. If Alfonso thinks carelessly and forms beliefs iff they are comforting, it is entirely natural to fault him for placing insufficient value on accuracy in his thinking. Still, when we do so, what we are calling for is something more modest than love or passion for accuracy. Love, passion and the like are not the only forms of valuing. They did not figure in my account. We need not view respect, commitment or compliance as constituted by emotions or states with any fancy qualitative feel.

They are just ways to be attuned to different factors that bear on the accuracy of particular beliefs. The ways to place value on accuracy that I invoked consist in forms of sensitivity to truth-oriented reasons: the reasons are *belief-relative* in the case of commitment, *appearance-relative* in the case of respect, and *fact-relative* in the case of compliance. In this unexacting sense, even children may place value on accuracy in thought—e.g., by being disposed to abandon their beliefs when and because they are probably inaccurate and to have beliefs when and because they are probably accurate. This is not implausible.



Indeed, ways of valuing accuracy could simply be grounded in competences. This is true, at any rate, if competence is understood as epistemically good *because* its exercises constitute ways to place value on accuracy. Consider Greco (1999: 289):

[S]ubjective justification can be understood in terms of the dispositions a person manifests when she is thinking conscientiously.... [B]y ‘thinking conscientiously’, I do not mean thinking with an explicitly voiced purpose of finding out the truth. Neither do I mean thinking with this as one’s sole purpose. Rather, I intend the state that most people are in as a kind of default mode—trying to form one’s beliefs accurately.

Properly framed, this could be a version of my picture.

#### 4.4.2 Doesn’t the View Then Collapse into Familiar Views?

My reply dovetails with another objection. Doesn’t this reply show that my view collapses into familiar views? Aren’t the properties that subvene my ways of placing value on accuracy just familiar ones from other theories of rationality and knowledge?

*Reply.* This objection misses the point of my view. I can agree that my ways of valuing accuracy are grounded in familiar properties from other accounts of the nature of rationality and knowledge. My view is not about the nature of rationality or knowledge, but about why these properties are epistemically good from the perspective of truth. What matters for this purpose is that what constitutes these properties constitutes something further—viz., ways to place value on accuracy.

Here we see what is new. My view captures unappreciated truth-oriented unity in epistemic value by its appeal to Hurka’s Principle and the ways of placing value on accuracy in thought. These ideas are not familiar in epistemology, since most epistemologists assume without argument that all derivative value must be grounded *via* instrumental relations to fundamental value. So, the key part of my view collapses into no familiar tack.

One might try to restate the objection: “But if ways to place value on accuracy are grounded in familiar properties, then ways to place value on accuracy will only be epistemically valuable if these familiar properties are epistemically valuable. But it is your job to explain why these familiar properties are epistemically valuable!”

I reply that they are epistemically valuable because they constitute ways of valuing accuracy, which in turn derive epistemic value from the perspective of accuracy *via* Hurka's Principle. This sounds question-begging only if we are hoodwinked by a fallacy.

Consider a beautiful painting. It is grounded in a bunch of atoms. These atoms are not beautiful. If they have value, it is because they ground something else. We should not insist that the atoms must have antecedent value and worry that the painting is trash because we cannot find this antecedent value. Their value can be explained in a trickle-down fashion: they ground the painting and it is valuable, so they are valuable. Downwards causation may be incredible. But downwards value inheritance is the norm.

#### 4.4.3 What About Other Epistemic Values?

One might concede my account of the epistemic value of coherence, rationality, and knowledge but worry that the approach cannot stretch farther. How can it ground all epistemic values, including rich ones like open-mindedness and intellectual courage?

*Reply.* I have a two-pronged response. First, I think these rich properties have a special glow partly by having extra value that is not properly epistemic. Beliefs that exhibit these properties might be somewhat epistemically better than mere true beliefs. But—second prong—my framework can explain their properly epistemic value.

A warm-up to the first prong. Epistemic value theorists often try to explain too much by not distinguishing two things we could mean by 'epistemic value'. I would separate

(a) being good from the epistemic point of view

from

(b) being good *simpliciter* and also epistemic.

There are parallels. Consider talk of a good chess strategy. This may pick out

(c) something that is good *from the point of view of chess strategy*

or

(d) something that is good *simpliciter* and also a chess strategy.

If your chess strategy will humiliate your opponent, it exemplifies (c) but not (d). Similarly, something might be epistemically superb but of little worth absolutely.

Must the chess expert explain the goodness of chess strategies from the Point of View of the Universe? No. It is similarly hard to see why the epistemologist must explain which epistemic items matter from that point of view. A theory of epistemic value should not make it impossible to understand why certain epistemic properties are desirable from the Point of View of the Universe. But this is a weak constraint. If we reject Instrumentalism, there is no reason to believe that Veritism violates this constraint.

Intuitions voiced in the literature on epistemic value often conflate (a) and (b). Once we remove obstacles to seeing how knowledge and understanding can possibly have value *simpliciter*, we must recognize that intuitions about (b) are best left for non-epistemic axiologists to explain. A theory of what is valuable from the epistemic point of view need not explain them, just as the chess strategist need not vindicate chess fanaticism.

With this in mind, I approach some items with suspicion. Consider understanding. While special in absolute terms, its significance from the epistemic point of view is unclear. Note that while (A-C) are defensible epistemic requirements, (\*-\*\*) aren't:

(A) Don't believe P while also believing that your evidence for P is insufficient.

(B) Don't believe P while also believing that you lack justification to believe P.

(C) Don't judge that P while also believing that you don't know that P.

(\*) Don't believe P while also believing that you do not understand why P.

(\*\*) Don't have beliefs about X while also believing that you don't understand X.

Of many facts I know, I am sure I do not understand why they are so. Of many topics about which I have beliefs, I am sure I do not understand them. Sometimes this is

not my fault. We may hope to understand the facts. But intelligibility is up to them. They may *be* unintelligible. We exhibit no epistemic flaw if we take some facts to obtain while finding them incomprehensible. By contrast, we exhibit a flaw if we judge while admitting that we do so irrationally, unjustifiedly, etc.

So it is unclear that a theory of epistemic value must explain our thirst for understanding. Still, understanding has some accuracy-oriented value. Like Pritchard (2010: 75-6), I think one cannot understand a topic without having many rational true beliefs about it. I just suspect that whatever value understanding has beyond this is not properly epistemic.

I am less moved yet by other items that are sometimes claimed to have some truth-disconnected epistemic value. Some point to features like open-mindedness and intellectual courage.<sup>14</sup> While the goodness of these features is not wholly truth-disconnected, I am unconvinced that their extra value is properly epistemic. This is not *ad hoc*: being willing to take the ideas of others seriously is good, but it can be epistemically counterproductive. The same goes for the intellectual courage associated with finding alternatives to common views even when they face no serious objections.

Again, it is important not to misunderstand the claims here. I'm not claiming that epistemic value theorists can wholly ignore the non-epistemic value of broadly epistemic properties. Epistemologists should not make it impossible to understand how these properties could be non-epistemically valuable. But the main challenge to the possibility of understanding this is the swamping problem, and I explained how to avoid it.

I am also not claiming that epistemologists shouldn't be interested in the nature of understanding, intellectual courage, open-mindedness, and so on. These properties are broadly epistemic properties. It is just that some of their glow is not properly epistemic.

#### 4.4.4 Trivial Truths and the Valuing of Accuracy

There is a different objection worth discussing in connection with the last. Many have objected that accuracy cannot be the fundamental epistemic value because the best way

---

<sup>14</sup>Cf. Montmarquet (1987: 484-5) and Baehr (2012).

to promote that value would be to do absurd things like memorizing telephone books.

But accuracy-first views need not encourage such nonsense. My view, for example, only implies that *when* one is taking a stance on some proposition, it is epistemically desirable to hold a doxastic attitude that manifests certain ways of honoring the ideal of accuracy:

(AI) It is correct to believe P iff P is true.

AI is logically weaker than:

(AI+) One ought to believe P iff P is true.

Granting AI+, we would be epistemically obliged to memorize telephone books, count blades of grass, and the like. But AI calls for no such thing: AI only implies that if P is true, it is correct to believe that P. The claim that it is correct to believe truths does not imply the claim that it is incorrect to not believe truths—i.e., obligatory to believe truths. But only that second claim yields the absurd conclusions.

This reply is preferable to the common one. Many suggest that we trade AI+ for:

(AI-) For any important proposition P, a person ought to believe P iff P is true.

While (AI-) avoids bad implications about trivial truths, it creates new problems. It fails to explain why it can be as epistemically criticizable to have wildly irrational beliefs about phone numbers as about philosophy. (AI), by contrast, renders all inaccurate beliefs equally incorrect. So, beliefs that equally disrespect (AI) are equally epistemically flawed.

#### 4.4.5 Does Hurka's Principle Really Generalize?

While Hurka's Principle is intuitive for value *simpliciter*, one might worry that it doesn't generalize to insulated domains like epistemology, aesthetics, grammar, and so on. Sometimes I hear the following version of this worry: "Suppose I love the rules of prescriptive grammar. My love has no *grammatical* value. Yet doesn't your generalized

version of Hurka's Principle predict otherwise?"

*Reply.* But this is not a counterexample to the generalized version of Hurka's Principle.

To see this, let me underline a crucial part of the principle:

When V is a non-instrumental value in some domain, proper ways of placing value on V in that domain have derivative non-instrumental value.

Is it grammatically proper to love the rules of grammar? No. But if so, Hurka's Principle does not predict that love for the rules of grammar is grammatically good. For the principle cannot predict that love for X is valuable in any domain unless it is also true that love is a proper response to X in that domain. This is what the underlined part says.

What about the epistemic domain? Respect, commitment, and compliance in my sense are epistemically proper ways to honor AI. If so, I can use Hurka's Principle compatibly with denying that loving the rules of grammar is grammatically good.

The fact that our principle appeals to domain-relative propriety does not trivialize it. Instrumentalism is inconsistent with our principle. So, our principle is clearly non-trivial: it conflicts with orthodoxy! Might one insist that 'proper' means 'valuable' and insist that the principle is circular for this reason? No: this insistence rides roughshod over non-trivial issues. Propriety is a kind of rightness. Rightness and value represent different species of normativity. There are long-standing debates about how rightness and value are related. So, Hurka's Principle is not circular in linking propriety and goodness.

#### **4.5 On the Competition**

I have displayed the advantages of my view and answered a host of objections. But even if a view has only advantages and avoids objections, it is only worth endorsing if it outperforms the competition. Does my view do so?

As we saw earlier, no modest form of pluralism about fundamental epistemic value that clings to Instrumentalism can avoid every relative of the swamping problem. With

Instrumentalism dropped, there is no reason related to the swamping problem to prefer these forms of pluralism. Since parsimony is a virtue, we should prefer my view.

What about other forms of monism? Granting Instrumentalism, they too will face relatives of the swamping problem. If the desire to avoid the swamping problem attracts one to a knowledge-first epistemic axiology, one is short-sighted: this view will face related problems, as we saw earlier. One could embrace a non-Veritist version of monism while also rejecting Instrumentalism. But because I am optimistic about the analyzability of knowledge, I doubt that the main monist alternative—i.e., the knowledge-first view—is well-motivated on other grounds.

#### 4.5.1 What About Reliabilist Virtue Epistemology?

Still, this does not exhaust the alternatives.

One alternative is virtue epistemology of the sort defended by Ernest Sosa. On this view, beliefs gain epistemic value by manifesting the epistemic agent's competence. While competence here is reliabilist, it is also agent-level in a way that reliable processes need not be. Owing to this difference, Sosa's model makes a narrower prediction than the process reliabilist's: only beliefs that manifest the agent's epistemic competence gain extra epistemic worth. This can seem to give Sosa a more principled response to the swamping problem. Apt belief is an achievement, while products of reliable processes are not necessarily achievements and may be mere products.

Compare this with my view. On my view, beliefs gain epistemic worth by manifesting certain ways of placing value on accuracy in thought. Structurally, the views are similar: in both cases, the belief derives epistemic worth by manifesting a truth-connected person-level feature. Verbally, the feature differs: it is a reliabilist competence in one case, and a way of placing value on accuracy in thought in the other. But one might think Sosa has the advantage, since his feature sounds less lofty. Why prefer my view?

I have two responses. First, this question rests partly on a false dilemma. I would view some kinds of competent belief as constituting ways to place value on accuracy in thought. Since I do not understand ways of placing value on accuracy in thought in a demanding way, there is no divide between Sosa and me here.

The difference rests not in the demands we make, but in our views of epistemic value derivation. On my view, it is only because some competences constitute ways to place value on accuracy in thought that their manifestations have derivative epistemic value. Sosa explains why manifestations of competence have epistemic value without regard to whether they constitute ways to place value on accuracy in thought.

This brings me to the second response. It is crucial for the success of virtue epistemology that competences be able to ground ways of placing value on accuracy in thought. We can see this by seeing why Sosa's picture leaves us with no convincing way to understand the difference in *epistemic worth* between the first-order beliefs of the clairvoyant and those of the sighted. While there are subtleties in Sosa's view that give him advantages over other reliabilists, the advantages are insufficient for the purposes of epistemic axiology. Perhaps clairvoyants have animal knowledge and achieve something epistemically. The question is whether Sosa can explain the difference in epistemic worth between the clairvoyant's first-order beliefs and those of sighted children. We will see that Sosa's animal/reflective distinction does not help: there is a first-order difference.

Sosa could instead revise his account of first-order competence in a way that John Greco has recommended. But the revision is attractive because it turns competences into the sorts of things that plausibly ground ways of valuing accuracy. The revision is not motivated unless viewed as conceding that the reliabilist part of the view does not do the axiological work. So, there is a dilemma: this form of virtue epistemology either (i) is lacking in explanatory power or (ii) collapses into a version of my view.

#### 4.5.2 First Horn

Competences for Sosa are reliable belief-forming dispositions of the agent relative to favorable conditions in the actual world. This account predicts symmetry between the first-order beliefs of clairvoyants and the first-order beliefs of sighted subjects. Their beliefs can equally manifest reliable belief-forming dispositions. If the facts of epistemic value are explained by such competences, the first-order doxastic attitudes of clairvoyants must be equal in epistemic value to the parallel attitudes of sighted perceivers.

But even if we agree that the doxastic attitudes of clairvoyants have some epistemic



worth, it is implausible that the degree is the same. We might agree that the clairvoyant knows. We are doing epistemic axiology now. Our objection is just as well expressed as one about the difference in epistemic worth between two instances of knowledge.

Some reliabilists try to deem the clairvoyant incompetent by adding a ‘no defeaters’ clause to their theory. But the best versions of this strategy count as defeaters factors that should look irrelevant from a purely reliabilist perspective—e.g., mere beliefs or appearances about the quality of one’s belief-forming processes.<sup>15</sup> I am with John Greco in finding this unprincipled: “Reliabilism insists on a reliabilist account of evidence in favor of belief. . . . [H]ow can the same theory plausibly understand evidence against belief differently? Such a strategy seems at best *ad hoc*.”<sup>16</sup>

Sosa’s tack differs.<sup>17</sup> He thinks clairvoyants might have animal knowledge and holds that what they lack is reflective knowledge. Reflective knowledge is understood as apt belief that one aptly believes, where aptness is understood in the same externalist way at both orders. What clairvoyants lack is an externalistically parsed grasp of the status of their first-order beliefs. This gives Sosa’s explanation more unity than the defeat strategy.

But the explanation falls short axiologically. Consider a child seeing a red apple on the table in good light and then judging that there is a red apple on the table. This child might lack the conceptual resources to form a belief about the propriety of her first-order belief. Indeed, she might lack the mental state concept BELIEF and the ability to engage in any second-order reflection. Nevertheless, her perceptual beliefs about the apple are epistemically worthier than those of the clairvoyant 400 miles away

---

<sup>15</sup>Some reliabilist accounts of defeat are supposed to address clairvoyance. But they encounter troubles in modified cases. Consider the alternative reliable process account of defeat. On this account, having a defeater for one’s belief that *p* consists in having available to one an alternative reliable process that, if used in addition to or instead of the one actually used, would have led to one’s not believing *p*. Goldman (1986: 112) suggests that this addresses Norman the clairvoyant, claiming that Norman could reliably reason as follows: ‘If I had a clairvoyant power, I should surely find some evidence for this. Since I lack any such evidence, I apparently have no such power.’ As Lyons (2009: 124) observes, this doesn’t go far enough: “[I]t is not difficult to modify the counterexample so that it precludes such a reply. We can just stipulate that Norman is so bad at reasoning about such matters that there is no [alternative] process that is both reliable and available to him. This doesn’t seem to change the intuitive verdict.”

<sup>16</sup>Greco (2009: 158–9).

<sup>17</sup>I rely on Sosa (1991)’s discussion of clairvoyance. Since then he hasn’t explicitly discussed it.

who equally reliably forms beliefs about the child's surroundings. This difference is not at the reflective level.

Some might bite the bullet, claiming that the child's beliefs are on a par with the clairvoyant's. But the reply is costly, and it underestimates the scope of the problem. Many adults lack much of a reflective stance on their perceptual beliefs and so have little reflective perceptual knowledge. The fact that some are in a position to gain reflective knowledge is irrelevant to the quality of their actual beliefs. By analogy: suppose I believe something for bad reasons, though I am in a position to believe for good reasons. That I can do better is irrelevant to the quality of my actual beliefs.

The real difference is not at the second order. If one wants to explain the difference between clairvoyants and the sighted, the story should come at the first order.

### 4.5.3 Second Horn

John Greco has offered a story by enriching the account of complete first-order competence. But his proposal faces his own objections to liberal defeat strategies unless his view is understood as a version of my view. To see this, consider the details. Greco (2010: 43) first suggests that we analyze epistemic responsibility in terms of the intellectual dispositions that one manifests when one is motivated to believe the truth:

S's belief that  $p$  is epistemically responsible if and only if S's believing that  $p$  is properly motivated; if and only if S's believing that  $p$  results from intellectual dispositions that S manifests when S is motivated to believe the truth.<sup>18</sup>

He then analyzes epistemic virtue in terms of reliability *and* epistemic responsibility:

S's belief that  $p$  is epistemically virtuous if and only if both (a) S's belief that  $p$  is epistemically responsible; and (b) S is objectively reliable in believing that  $p$ .<sup>19</sup>

The unjustified clairvoyant, Greco suggests, fails condition (a).

---

<sup>18</sup>Greco (2010: 43).

<sup>19</sup>Greco (2010: 43).

But Greco's explanation confronts his own objection to the process reliabilist's appeal to an unreliabilist kind of defeat. Greco insists that reliabilists cannot treat positive and negative justification-relevant properties differently without a unified explanation of why both matter. He faces a similar question. He *says* that epistemic responsibility is a virtue. But we must ask him *why*. What unified story can deem both objective reliability and motivation to believe the truth as epistemic virtues?

Greco has offered no answer. By contrast, my view provides a simple explanation of why being motivated to believe the truth is epistemically virtuous. It is a way to place value on accuracy in thought, and it matters from the epistemic point of view because accuracy matters (*via* Hurka's Principle). This confirms my prediction of convergence. Indeed, there is not just convergence but collapse. For the best explanation of why the dispositions constitutive of being motivated to believe the truth are epistemically virtuous is that they ground epistemically proper ways of valuing accuracy. There is no general reliabilist explanation. If there were, there would have been no need to revise the account!

Thus the second horn: revising the account of first-order epistemic virtue is principled only if underpinned by a view of epistemic virtue like mine. So, virtue epistemology either cannot explain all the facts of epistemic value or collapses into my view.

#### 4.5.4 A Middle Way Between Reliabilism and Responsibilism

Reliabilist virtue epistemology is not the only kind of virtue epistemology. The main alternative is responsibilism. Since responsibilism is typically defined as a view that emphasizes the role of intellectual character traits, I doubt that it can succeed. I doubt that any enduring traits of intellectual character are necessary for justification, rationality, knowledge, or other items of traditional epistemological interest. Since I agree with existing critiques, I will not dilate on these doubts.<sup>20</sup>

Here is what matters. The common distinction between reliabilist and responsibilist virtue epistemology is not exhaustive. There is much unoccupied space between, since

---

<sup>20</sup>In particular, I agree with many of Dougherty (2011)'s criticisms of responsibilism.

there are pictures of virtue that virtue epistemologists ignore. Hurka (2001) understands virtues as ways of valuing more basic values. But ways of valuing are neither faculties nor character traits. While I have not understood them (as Hurka (2006) does) as occurrent states, I have understood them as less than character traits or faculties.

So, my view could be understood as a new version of virtue epistemology. For my purposes, the label is inessential. But it highlights a direction for future research and positions my view in an illuminating way with respect to the competition. Examined from this perspective, one can also see that it strikes a golden mean between two ends of the spectrum. It inherits some advantages from existing traditions in virtue epistemology while avoiding their flaws. This is another reason to suspect it is on the right track.

#### 4.6 Conclusion

Let's take stock. I argued that Veritists can and should reject Instrumentalism. This is the moral of the swamping problem. We can see that it is the moral by seeing that if Instrumentalism were true, the points that led to the swamping problem would lead to related problems for other economical epistemic axiologies. This should not convince us that a radical form of pluralism is true. It should convince us that Instrumentalism is false.

I have shown that Veritism is defensible once we abandon Instrumentalism. Veritists can avoid the swamping problem if they reject Instrumentalism. As I argued in §2, the kind of derivative value generated by Hurka's Principle is immune from swamping. A charitable act that manifests the agent's valuing of charity is *as such* better than a similar act performed just as a PR stunt. This is plausible despite the fact that the goodness of valuing charity is parasitic on the goodness of charity.

Of course, this fact would be useless if we could not apply Hurka's Principle to the relevant cases in the epistemic domain. But we can. I argued that central epistemic values like structural rationality, substantive rationality, and knowledge essentially manifest

certain ways to place value on accuracy in thought. None of this is an overintellectualization. To think otherwise is to implicitly overintellectualize the modest achievement of placing value on accuracy in thought. It is natural to fault subjects for failing to place sufficient value on accuracy in thought, and plausible that people who disregard the evidence are open to criticism *precisely because* they fail to place sufficient value on accuracy in thought.

So, there is a straightforward way for Veritists to avoid the swamping problem and provide explanations of the epistemic value of central epistemic properties like rationality and knowledge. Since it is antecedently attractive to think that epistemic evaluation is fundamentally truth-oriented, I think we should continue to be attracted to Veritism.

#### 4.7 Appendix A: Must Reliabilists Be Instrumentalists?

I argued that Veritists should reject Instrumentalism. I do not take this to be an argument against reliabilism. Indeed, I have defended a sophisticated version of reliabilism elsewhere. My views are not inconsistent. Let me explain why. Instrumentalism is an axiological thesis about how non-fundamental value is grounded in fundamental value. Reliabilism is a first-order account of the nature of justified belief or—in some permutations—knowledge. Reliabilism does not by itself entail any claims about the worth that any beliefs have from the epistemic point of view. One must appeal to auxiliary assumptions to get such conclusions.

Obviously, many have worried that reliabilists are forced to appeal to Instrumentalist explanations of the value of knowledge. This worry was behind much of the early literature on the swamping problem. But this worry is mistaken, for several reasons. For one thing, a reliabilist might consistently hold that knowledge is a further fundamental epistemic value. Reliabilism and Veritism are independent. No first-order account of knowledge entails that knowledge is not fundamentally epistemically good. For another, even if reliabilists agree that knowledge is not a fundamental epistemic value, it is unclear why they are forced to appeal to Instrumentalist explanations.

There are two reasons why it is unobvious. To see the first, let's distinguish two

different types of theory that might be called “reliabilist”. *Strong* reliabilist theories identify a belief’s all-things-considered justifiedness with (i) its having come from a reliable process and (ii) there being no alternative reliable process available to the believer that would lead to her abandoning the belief. It might be true that strong reliabilist theories are forced to appeal to instrumentalist explanations. But reliabilists do not always hold theories of this strong kind. Although Goldman (1979) upheld a strong account, he later adopted a more concessive picture of defeat and all-things-considered justification. Goldman (1986) allowed, for example, that beliefs about the unreliability of one’s belief-forming processes can defeat justification regardless of their historical properties; Goldman (2011) has allowed that experiences can defeat regardless of their historical properties.

Call a theory a *weak* reliabilist theory if it only entails that a necessary condition for a belief’s being justified is its having its origin in a reliable process. Some of Goldman’s later views exemplify weak rather than strong reliabilism. Moreover, many virtue epistemologists agree that reliability is a necessary condition for justified belief and knowledge. They simply add that there is something else important that helps to explain why justified belief and knowledge have the value that they have. Indeed, even the virtue epistemologists who have persistently pressed the swamping problem are reliabilists in the weak sense. Zagzebski (1996) explicitly agreed that reliability is a necessary condition for complete epistemic virtue and distanced herself from virtue epistemologists such as Montmarquet who reject this claim. Even so, Zagzebski rejects Instrumentalism.

So it is simply false that reliabilists must appeal to Instrumentalist explanations, since weak reliabilists can avoid these explanations. But we can go even deeper. I think some strong reliabilists can consistently reject Instrumentalism. My own view about justified belief is a strong reliabilist view. I hold that justified beliefs are beliefs that one holds because one possesses sufficient epistemic reasons. But I add that the possession and sufficiency of objective epistemic reasons can be grounded purely reliabilistically. This view meshes with my non-Instrumentalist story about the epistemic worth of justified belief. I can give a top-down, constitutive explanation: the reliabilist properties

that ground believing for objective epistemic reasons have a non-instrumentally truth-oriented kind of value by constituting something else that has such value. Believing something for sufficient epistemic reasons plausibly just is a way to place value on accuracy in thought. Placing value on accuracy in thought also has some non-instrumentally truth-oriented epistemic value, thanks to Hurka's Principle. Since whatever constitutes something valuable inherits its value, forming doxastic attitudes because one has sufficient epistemic reasons inherits this value. But forming doxastic attitudes for sufficient epistemic reasons is constituted by reliabilist properties. So, these properties also have a non-instrumental epistemic value because of what they constitute, beyond their instrumental epistemic value.

Is this circular? No. It is an example of a common direction of value inheritance. The atoms composing a painting are good because the painting is good. If we look at the atoms singly, it is hard to see why they matter. But that only shows that we must "look up" at what they constitute to grasp their value. The same holds for the epistemic value of the reliabilist features that constitute believing for objective epistemic reasons. The only reason to believe that a strong reliabilist picture is inconsistent with the rejection of Instrumentalism involves a fallacy. It is the same fallacy that would deem a painting valueless because the atoms that constitute it are, considered singly, valueless. If we reject this fallacy, we can embrace strong reliabilism but reject Instrumentalism.

#### **4.8 Appendix B: Does Type Instrumentalism Make a Difference?**

Goldman and Olson (2009) are not just reliabilists but self-styled instrumentalists. They claim to avoid the swamping problem. I mentioned some doubts about this claim earlier. They take the problem to rest on the assumption that true beliefs come from reliable sources cannot be better than true beliefs that are products of unreliable sources.

But as we saw, the key assumption is more fine-grained. The key assumption is that the mere fact that a good item comes from a reliable source cannot as such make it better. This assumption is compatible with the claim that some reliably produced true beliefs are better than true beliefs. Even Zagzebski agrees. Zagzebski (1996) agrees

that reliability is a necessary condition for complete epistemic virtue. It follows that she agrees that some reliably produced true beliefs are better than mere true beliefs, since some reliably produced true beliefs satisfy the other necessary conditions for manifesting complete epistemic virtue. Goldman and Olson have not, accordingly, addressed the fundamental problem.

Goldman (2012) has recently made some further suggestions, however. He claims that type instrumentalism is the key to avoiding the swamping problem. Type instrumentalism is the view that a token instance of an instrumentally valuable type of source can have a kind of instrumental value merely by being a token of that type, irrespective of whether that token produces anything good. I agree that type instrumentalism may help to explain how the epistemic value of token processes. But it is unclear how this view helps with the fundamental problem. This is because the fundamental problem was never about the value of the processes but rather about the value of the beliefs that are their products. Type instrumentalism by itself makes no headway on this issue.

Goldman discusses this further problem after emphasizing the importance of type instrumentalism. But the key moves in this discussion are not consequences of type instrumentalism. Indeed, some of Goldman's examples are suggestive of something other than an Instrumentalist model. Given the points made in Appendix A, this is not really a problem for him qua reliabilist: reliabilists are not forced to be Instrumentalists.

The driving example is a work by Rembrandt. It seems that the sheer fact that the painting is Rembrandt's work makes the painting better. But this does not vindicate Instrumentalism, as virtue epistemologists will be keen to insist. While we should agree that a painting by Rembrandt is better simply in virtue of being the work of Rembrandt, we should not agree that a cup of coffee is better simply in virtue of coming from a well-functioning coffeemaker. What is the difference? The difference is simple, according to virtue epistemologists: Rembrandt's work manifests the person-level abilities of an agent. It is a doing of Rembrandt's that manifests his painterly ability. The explanation is not an instrumental explanation. If virtue epistemology has anything in common with virtue ethics, it is its ability to appeal to this type of explanation. As I have repeatedly suggested, Goldman does not need to confine himself to Instrumentalist



explanations. Given his recent efforts to synthesize evidentialism and reliabilism, he could easily appeal to the model suggested at the end of Appendix A. But even if he dislikes that model, there are probably other forms of value derivation to which he can appeal, especially given the many sophisticated turns that reliabilism has taken over the years. So, Goldman and I need not disagree.

## Chapter 5

### Rationality and Fundamental Epistemic Value

#### 5.1 A Puzzle about the Value of Rational Belief

Consider Alpha, who inhabits a world like ours. Alpha knows just about everything there is to know about this world, and Alpha's beliefs and reasoning are flawless in other epistemologically interesting respects. Now compare Beta. Beta's mental life is internally just like Alpha's. Any experience, apparent memory, or intuitive seeming that Alpha has is equally had by Beta. Alas, Beta is a victim of skeptical disaster. Beta inhabits a world whose only other occupant is a demon. This demon ensures that Beta's beliefs about extra-mental reality are radically false.

While Beta is less accurate, reliable, and knowledgeable than Alpha about the contingencies of extra-mental reality, most epistemologists will agree that there is some property of epistemic worth that Beta's beliefs share with Alpha's. Different epistemologists just have different labels for it. Some externalists will deny that Beta's beliefs about extra-mental reality are as justified or warranted as Alpha's. But they will have *some* term of appraisal for these beliefs of Beta's, like "rational" or "blameless".<sup>1</sup> My term is "rational", and I will call Beta's rational beliefs about extra-mental reality "Beta's relevant beliefs" to save words.

Beta's relevant beliefs appear to retain genuine worth from the epistemic point of view. This appearance is not a mere byproduct of our imaginatively placing ourselves in Beta's shoes. Even if we occupy a third-person perspective and attend to Beta's

---

<sup>1</sup>See Littlejohn (2012) for the first option. See Goldman (1988) and Littlejohn (2009) for something best understood as the second option. Others have proposed that while Beta's *beliefs* are unjustified, Beta *herself* has some good epistemic property. This leads some to distinguish between doxastic and personal justification; see Bach (1986) and Engel (1992). Littlejohn (2012, 2009) also favors this distinction, though he would call Beta's beliefs "rational".

radical unreliability with respect to questions about extra-mental reality, it remains plausible that Beta's relevant beliefs have epistemic worth. So the epistemic worth of Beta's relevant beliefs seems to go beyond merely expected value relative to Beta's perspective. Let us give this appealing thought a more official statement:

**Retention of Worth:** Even in bad worlds like Beta's, rational beliefs about extra-mental reality retain some genuine worth from the epistemic point of view.

While many epistemologists can share this intuition, few have tried to explain how it could be true. If it is true, it calls for explanation: it is not a brute fact. But Retention of Worth is harder to explain than one might think. Two common assumptions make it hard to explain.

To see the first common assumption, note that many epistemologists think there is a *fundamental* epistemic value—i.e., some epistemic value in terms of which we can explain the epistemic value of everything else. The typical candidates are true belief and knowledge. Most epistemologists think that properties like rational belief and justified belief are only derivatively significant relative to the fundamental epistemic value. This is plausible: we seem to care about rational beliefs because we care about something else, like truth or knowledge. Hence the first assumption:

**Derivativeness:** Rational beliefs are only derivatively epistemically good.

There is a second common assumption that is sometimes conflated with Derivativeness. To see it, note that most epistemologists view the fundamental epistemic value as an epistemic goal. It is supposed to be something that we ought epistemically to produce as an end. Most epistemologists take it to follow that derivative epistemic values can only have epistemic value in virtue of *standing in instrumental relations* of two sorts: (i) in virtue of being types of things that tend to produce the end (“production values”), or (ii) in virtue of being products of types of sources that are likely to produce the end (“product values”).<sup>2</sup> Hence the second common assumption:

---

<sup>2</sup>One might reserve ‘instrumental epistemic value’ for (i). But I use it to capture the kind of derivative value that something has just by standing in some instrumental relation. The *product of*

**Instrumentalism:** Derivative value in any domain is just instrumental value in that domain, of either the production or mere product kind.<sup>3</sup>

Good belief-forming processes are supposed to have the first kind of instrumental value, and their products are supposed to have the second kind of instrumental value.

Given these assumptions, it is hard to see how Retention of Worth could be true. The problem is simple, I will argue in §2: if Derivativeness and Instrumentalism are true, Retention of Worth is false. There is no satisfactory Instrumentalist explanation of Retention of Worth. All the obvious explanations fail. So, if Retention of Worth is true, one of the other claims must go.

I think we can and should explain Retention of Worth by rejecting Instrumentalism. Reflection on the worth of rational belief reveals that not all derivative epistemic value is just instrumental epistemic value of either the production or mere product kind. We can reject that claim without difficulty by looking beyond epistemology, as I showed in the last chapter. Non-epistemic value theorists have rejected Instrumentalism, and there are epistemic analogues of their ideas. Given these views, we can explain Retention of Worth consistently with Derivativeness.

Is it overkill to reject Instrumentalism to explain Retention of Worth? No. As we saw in the last chapter, there is a much more general problem for Instrumentalism—the swamping problem. While the problem was initially presented as a problem for reliabilists, it is easy to show that any economical epistemic axiology that embraces Instrumentalism will face a relative of the problem. Some theorists—e.g., Carter and Jarvis (2012)—conclude from this fact that the intuitions that drive the swamping problem are flawed. But they arrive at this conclusion by presupposing Instrumentalism. There is nothing amiss with the intuitions behind the swamping problem. The moral is that Instrumentalism is false.

---

relation is just as much an instrumental relation as the *caused by* relation is a causal relation. Since many epistemologists assume that (ii) can generate derivative epistemic value, it is only fair to define Instrumentalism this way.

<sup>3</sup>I borrow the term from Jones (1997), one of the first to give the swamping problem a detailed presentation. Like me, Jones regards Instrumentalism as the root cause of the swamping problem. It is a shame that his points are neglected in the new literature that blames the problem on reliabilism or Veritism.

You might think that we could live with the conclusion that rationality is a fundamental epistemic value. You might think that we could learn to live with a modest form of pluralism about fundamental epistemic value: besides true belief or knowledge, rational belief is also on the list of fundamental epistemic values. But as we saw in the last chapter, the reasoning that supports placing rational belief on the list of fundamental epistemic values also supports radically inflating that list in ways that no one should tolerate. If we reject Instrumentalism, we can keep our list of fundamental epistemic values sparse. The view I defended in the last chapter provides an illustration. This view has a unified explanation of the derivative epistemic value of rational belief, justified belief, coherence, knowledge, and more. The crucial upshot for the purposes of this chapter is that rational beliefs matter in exactly the same way that these other epistemic properties matter. This depuzzles the epistemic worth of rational beliefs.

This upshot will bring out another lesson that I discuss near the end of the paper. Many epistemologists think that reliabilists about justified belief have a unique ability to explain why justified belief matters owing to the instrumental link that their view forges with true belief. But this thought depends essentially on Instrumentalism. As it turns out, Instrumentalism actually makes it harder to explain the value of justified belief. Without Instrumentalism, there is no reason to think that reliabilism has this advantage, and good reason to think that non-reliabilists can secure a truth-oriented explanation of the epistemic value of both justified and rational belief. This undermines one major argument for reliabilism—though I would independently embrace a sophisticated version of reliabilism for other reasons.

With these themes in mind, here is how I will proceed. I will begin by explaining more carefully why it is hard to give a plausible instrumentalist explanation of the value of rational belief, considering and rejecting two more sophisticated forms of Instrumentalism that I will call *Counterfactual Instrumentalism* and *Ingredient Instrumentalism*. I will then proceed to remind the reader of the main result of the last chapter: namely, that Instrumentalism is generally-ill equipped to explain derivative epistemic value. Even if there *were* a necessary instrumental connection between rational belief and true belief,

we could not explain the derivative value of rational belief! This fact provides independent support for thinking that the lesson of our puzzle about the value of rationality is that we should reject Instrumentalism. Having reminded the reader of this fact, I will turn to show how the alternative model of epistemic value derivation that I sketched in the last chapter can provide *the same explanation* of the value of rational belief that can and should be given of the value of justified belief and knowledge. I will go farther than I did in the last chapter and also show why non-Veritists (e.g., knowledge-first epistemic axiologists) can provide the same style of unified story. The story is robust. I will then conclude by discussing the broader upshots for epistemology, and especially for long-standing debates between reliabilists and internalists.

### 5.1.1 Disclaimer

Before I proceed, a disclaimer is in order. In this paper, I am not discussing the familiar new evil demon problem for externalist accounts of justification. The focus is rather on how to explain the intuition that rational beliefs about non-mental reality have genuine epistemic *worth* even in the demon world. This intuition presupposes that these beliefs can be rational in such worlds. We can separate the question of worth from the question about the mere presence of the property of rationality. Indeed, we could in principle agree that the demon-worlder is rational but conclude that this undermines the worth of rationality *as such*, perhaps by analogy with Kolodny (2005)'s skepticism about the normativity of rationality.

I agree that the correct response to the new evil demon problem is to distinguish between justification and rationality. Cohen (1985)'s dismissal of this response is hasty—though I will not dispute it here. But even if we are externalists and agree that the demon worlder's beliefs are rational, we are left with unanswered questions about the epistemic worth of these rational beliefs. Everyone needs to answer these questions, internalists included. This is especially salient because some internalists accept both Instrumentalism and Derivativeness (e.g., BonJour (1985)).

## 5.2 Instrumentalism $\wedge$ Derivativeness $\rightarrow \neg$ Retention of Worth

The first claim I want to defend is that Instrumentalism and Derivativeness entail the negation of Retention of Worth. There is a simple argument for this claim, but the work lies in knocking down more sophisticated Instrumentalist replies. As we will see, it is not difficult to show that the main possible Instrumentalist replies are unsuccessful.

To get us started, the *Simple Argument* can be stated as follows:

1. If Instrumentalism and Derivativeness are true, rational beliefs have real epistemic worth in a world  $w$  *only if* they tend to instrumentally promote fundamental epistemic value in  $w$  or are products of types of sources that tend to do so in  $w$ .
2. Beta's relevant beliefs do not tend to instrumentally promote fundamental epistemic value in Beta's world and are not products of types of sources that instrumentally promote fundamental epistemic value in Beta's world (assuming Derivativeness).
3. So, if Instrumentalism and Derivativeness are true, Beta's relevant beliefs do not retain real epistemic worth in Beta's world. (From (1 and 2))

Since the consequent of (3) is in effect the negation of Retention of Worth, (3) is equivalent to:

- 3 restated. If Instrumentalism and Derivativeness are true, Retention of Worth is false.

This argument is valid and (3) is just a consequence of (1) and (2). So if the argument fails, either (1) or (2) must be false. Accordingly, if one is both an Instrumentalist and a fan of Derivativeness, what can one say against (1) and (2)?

Here are the main candidate replies:

- A. Subjective Instrumentalism.** While rationality is not objectively instrumental to true belief or knowledge in Beta's world, it is subjectively instrumental to true belief in Beta's world: rationality maximizes expected accuracy relative

to Beta's perspective. But this is all we need to explain Retention of Worth. So, (2) is false: Beta's relevant beliefs do in a relevant sense instrumentally promote fundamental epistemic value in Beta's world.

**B. Counterfactual Instrumentalism.** What makes a belief-forming disposition instrumentally epistemically good is that it would be instrumental to fundamental epistemic value in favorable conditions. But Beta's belief-forming processes are epistemically good in that way and Beta's relevant beliefs are also epistemically good by manifesting those dispositions. Hence, (1) is false: although rationality does not instrumentally promote fundamental epistemic value in Beta's world, Beta's rational beliefs about extra-mental reality have real instrumental epistemic worth in an extended sense even in Beta's world.

**C. Ingredient Instrumentalism.** Even if X is not sufficient for promoting fundamental value, X can remain instrumentally good in an extended sense. Necessary ingredients of sufficient conditions for promoting fundamental values are instrumentally good in a natural sense. But rationality is necessary for knowledge, and knowledge is a plausible candidate for a fundamental epistemic value. So (1) is false: although Beta's relevant beliefs are not sufficient for promoting fundamental epistemic value in Beta's world, they still have a kind of instrumental value in Beta's world.

Unfortunately, none of these replies is ultimately plausible.

### 5.2.1 The Irrelevance of Subjective Instrumentalism

Let's consider reply (A) first. It is true that rationality subjectively promotes the goal of believing accurately even in bad worlds. So, it is plausible that rationality promotes expected epistemic value in Beta's world. It is also plausible, then, that in the subjective sense of 'ought', Beta ought to have the beliefs that Beta has. The trouble is that these claims are irrelevant to Retention of Worth. Retention of Worth concerns the *real* epistemic worth of Beta's relevant beliefs. Expected value is not a kind of real value, just as expected wealth is not a kind of wealth.



The Subjective Instrumentalist might reply that it is just obvious that Beta's relevant beliefs lack any real epistemic worth: they obviously only have expected epistemic value relative to Beta's perspective. But this reply is implausible.

We can look squarely at the fact that Beta's relevant belief-forming processes are unreliable and still find it plausible that there is genuine worth in the output beliefs, not just apparent-worth-relative-to-Beta's-perspective. If the Subjective Instrumentalist were right, the only explanation of why we find this plausible would be confusion: we imagine ourselves into Beta's shoes and get deceived, temporarily thinking that Beta's relevant beliefs are objectively promoting fundamental epistemic value. We are not guilty of that. We think that in spite of Beta's unreliability, Beta's relevant beliefs retain genuine worth from the epistemic point of view.

Some of the Subjective Instrumentalist's claims are true but unhelpfully so. The subjective 'ought' *just is* the 'ought' of rationality. It is trivially true that Beta ought rationally to be rational. But the question that we are asking—"How can Beta's relevant beliefs retain real epistemic worth?"—is not answered by this trivial claim. We are not asking about whether there is a sense in which Beta *ought* to hold the relevant beliefs.

So, Subjective Instrumentalism does not address our question, and it does not yield a plausible error theory. This is not to deny that there is a sense of 'ought' in which Beta ought to be thinking as Beta is thinking. Beta ought rationally to be thinking that way. Beta ought subjectively to be thinking that way too: the subjective 'ought' coincides with the 'ought' of rationality. But these facts do not help. So, reply (A) is ultimately unhelpful.

### 5.2.2 The Failure of Counterfactual Instrumentalism

Framed properly, reply (B) can appear more helpful than reply (A). To bring out the appeal, we can consider Ernest Sosa's view about subjects like Beta.<sup>4</sup> For Sosa, the belief-forming dispositions that constitute epistemic rationality are *competences*. Sosa's notion of competence is broadly reliabilist: what makes a belief-forming disposition a

---

<sup>4</sup>See Sosa (1991, 1992, 2003) for this, and Sosa (2007, 2010) for a great discussion of competence.

competence is some instrumental relation to true belief. But Sosa's view differs importantly from other versions of reliabilism: by using competence as the core notion, Sosa can appeal to the *competence/performance distinction*.

Competences are dispositions. Like other dispositions, there is an understood set of favorable conditions for their manifestation. If you take a box of matches into a room without oxygen and strike them, none will light. That does not show that none is flammable. If you take an archer into a hurricane, she will not hit any targets. That does not show that she is an incompetent archer. The reliability of competence is relative to favorable conditions. Sosa exploits this idea, suggesting that worlds like Beta's only illustrate that reliability of *this* sort is necessary for rationality.

One might agree with these claims about the necessary conditions for rationality but wonder how they explain why Beta's relevant beliefs have worth from the epistemic point of view in Beta's world. To address this question, Sosa appeals to an analogy with a temperature control device:

Suppose it is taken off the shelf in the display room for a demonstration, and a situation is simulated wherein it should activate the cooling trigger, and then a second situation is simulated wherein it should activate the warming trigger. In such a test the device might either perform well or not. But the quality of its performance is not to be assessed through how well it actually brings about the goods that it is meant to bring about in its normal operation. For in the display room it brings about neither the cooling nor the heating of any space.... What we are doing is quite obvious: we are assessing whether it performs in ways that would enable it to bring about the expected goods once it was properly installed....<sup>5</sup>

Sosa's analogy suggests that something can be derivatively good in a way worth calling "instrumental" in some situation even if the situation is not one where it is conducive to a more fundamental good. One might conclude that this yields a satisfactory Instrumentalist explanation of why Beta's relevant beliefs retain epistemic worth.

But this conclusion is incautious. The analogy with the temperature control device only works in *mild* deception examples.<sup>6</sup> Mild examples feature people in worlds that are

---

<sup>5</sup>Sosa (2003: 175–6).

<sup>6</sup>Foley (2004) makes a point related to the one to follow. He tries to argue that Sosa's account of the *nature* of rationality is defective. I think Sosa's account of the nature of rationality is right at some level. I am only objecting to the ability of Sosa's point to explain the epistemic *worth* of rational belief.

otherwise favorable who become victims of deception—e.g., people like us whose brains get envatted by evil neuroscientists. In the worlds of these examples, unreliability is not the norm. Here the temperature control device yields an analogy:

- (1) hosting rational beliefs about the world in favorable conditions in  $w$  :
- (2) hosting rational beliefs about the world in unfavorable conditions in  $w$  ::
- (3) activating cooling trigger in display room in  $w$  :
- (4) activating that trigger when properly installed in  $w$

But the analogy dissolves when we consider *extreme* deception examples, where deception is the norm in the world and all nearby possible worlds. Beta's world was supposed to be such an example. Indeed, we could even imagine that there is a law of nature which ensures that Beta's relevant beliefs will be never be true. Sosa's strategy suggests that when we regard Beta's relevant beliefs as good from the epistemic point of view, what we find good is that these beliefs are formed in ways that would be reliable if Beta were in a completely different kind of world.

But the original analogy does not support this picture. Imagine a qualitative replica of the temperature control device in a distant world with such different laws of nature that it could never reliably control the temperature of anything there. Imagine that the same holds for all nearby worlds. Should we say that the device still exhibits a good feature in the world it inhabits because if it *were* suitably installed in some very distant world it *would* reliably control the temperature?

This is implausible. So the analogy does not extend. The first half of the new analogy would have to be:

- (1\*) hosting rational worldly beliefs in extremely bad world  $w$  :
- (2\*) hosting rational worldly beliefs in a normal case in distant good world  $w^*$

This does not align with the second half of the original analogy. It must change:

- (3\*) activating cooling trigger in unfavorable world  $w$  with its different laws :

(4\*) activating that trigger when properly installed in distant favorable world  $w^*$

But there is nothing good *in*  $w$  corresponding to (3\*). If there are no conditions in  $w$  or any nearby worlds in which a mechanism would reliably bring about success, it is no good feature of that mechanism in  $w$  that it would bring about such success in a distant world  $w^*$ . It would be a good feature in  $w^*$ . But we are asking what is good in  $w$ . If anything, this analogy makes it harder to understand how Retention of Worth could be true, not easier.

This is not to deny that a belief-forming disposition is rational in a world like Beta's only if there is another favorable world where it is reliable. It is not to deny that being formed by dispositions that would be reliable in the actual world is necessary for a belief to be rational. It is just to deny that these necessary conditions explain the epistemic worth of rational beliefs about extra-mental reality in bad worlds. The analogy suggests that they don't explain this worth.

Counterfactual Instrumentalism provides no satisfying explanation of Retention of Worth. There might be a good analogy between the temperature control device and people from an otherwise good world who become deception victims. But there is no good analogy between the temperature control device and Beta, who inhabits a distant, wholly bad world surrounded by similar worlds. The comparison makes it harder to understand why Retention of Worth is true.

### 5.2.3 The Failure of Ingredient Instrumentalism

The earlier strategies are the most obvious ones for Instrumentalists to pursue if true belief is the candidate for fundamental epistemic value. But even if Instrumentalists accept Derivativeness for rationality, they might deny that true belief is the only candidate. Knowledge is a contender after Williamson (2000). A new strategy is possible with knowledge on the table: *Ingredient Instrumentalism*, which tries to profit from rationality's status as a necessary condition for knowledge.<sup>7</sup>

---

<sup>7</sup>'Ingredient' does not mean 'independent factor', so this idea is consistent with Williamson's attack on composite analyses of knowledge. Various ingredients for a dish can be blended together by cooking so that none can be factored out of the cooked dish. In any case, I just use the term for vivid

The plausibility of Ingredient Instrumentalism depends on the status of this argument:

### The Ingredients Argument

- a. Necessary conditions for more basic goods retain instrumental value *qua* ingredients for these goods even when the other necessary conditions are missing.
- b. If (a) is true, then rational beliefs can retain instrumental epistemic value as ingredients for knowledge even when other necessary conditions (e.g., truth) are missing.
- c. If the consequent of (b) is true, then rational beliefs about extra-mental reality retain real instrumental epistemic value even in bad worlds like Beta's.
- d. So, Beta's relevant beliefs retain real epistemic value. (From (a–c))

If successful, the Ingredients Argument undermines (1) in the Simple Argument: while Beta's relevant beliefs do not yield knowledge and are not products of dispositions that would be knowledge-conducive in Beta's world, these beliefs retain a kind of instrumental value in Beta's world.

But the Ingredients Argument does not explain why Beta's relevant beliefs retain epistemic value. The flaw in the argument resembles the flaw in the attempt to support Counterfactual Instrumentalism by analogy with the temperature control device.

To see the flaw, consider when (a) looks plausible. (a) looks plausible when we imagine a kind of world where the other necessary ingredients stand a chance of existing. Imagine one has half the ingredients for some great dish. It is plausible that even if one is missing the other ingredients, one's ingredients retain instrumental value. This is plausible because the other necessary ingredients stand a chance of existing in worlds like one's own—and presumably do exist.

Suppose, on the other hand, that your friend says of some seemingly useless things: "If we were just to have some further ingredients, we could use these objects to make

---

suggestiveness.

a neat thing X. So let's buy them." You reply: "But the other ingredients for X can't exist in a world like this one!" Your friend replies: "That doesn't matter. They are still instrumentally good in our world as ingredients for X, even if we can't use them in a world like ours to make X." This reply is absurd.

There is a similar problem with using the Ingredients Argument to explain the value of Beta's relevant beliefs. By stipulation, Beta's relevant beliefs cannot be converted into knowledge: these beliefs are false products of processes that the demon guarantees to be unreliable. By stipulation, this is true of all nearby possible worlds. From an Instrumentalist point of view, Beta's relevant beliefs are like the objects that your friend wants to buy. So, the Ingredients Argument fails.

Necessary ingredients for goods do retain instrumental value in worlds where the other ingredients stand a chance of existing. But this does not support the claim that Beta's rational beliefs retain instrumental value as ingredients for knowledge worlds like hers. In such worlds, the other necessary ingredients stand no chance of existing.

### 5.3 What Should We Conclude?

Three major strategies for securing an Instrumentalist explanation of Retention of Worth fail. The core problem is simple. There are only two general approaches appealing to instrumental relations that could explain why rational beliefs have real epistemic value:

**The Production Strategy.** Rational beliefs have derivative epistemic value by being things that produce fundamental epistemic values like true belief or knowledge.

**The Product Strategy.** Rational beliefs have derivative epistemic values as products of types of sources that are conducive to more fundamental epistemic values.

Beta's case makes both strategies look implausible. In Beta's world and nearby worlds, rational beliefs are guaranteed to be reliably false and rational belief-forming dispositions are guaranteed not to be conducive to more fundamental epistemic values.

We found it unhelpful to reply that Beta's rational belief-forming dispositions and beliefs are subjectively conducive to fundamental epistemic value. We also found it unhelpful to reply that Beta's rational belief-forming dispositions would be conducive to fundamental epistemic value if Beta were in another world. This helps only if the analogy with the temperature control device extends. It does not, as I demonstrated above.

Since the Ingredients Argument also failed, it is hard to see how one could resist the Simple Argument. I think the simple appearance is not misleading: Instrumentalism and Derivativeness are inconsistent with Retention of Worth. Since Retention of Worth is what we want to explain, we should reject either Instrumentalism or Derivativeness.

How should we choose? Some pluralists about fundamental epistemic value might add rationality to the stock of fundamental epistemic values.<sup>8</sup> But I think Derivativeness is not what we should reject. There is, as we saw in the last chapter, a clear and unified story about how epistemic properties have derivative epistemic value once we take note of how non-epistemic value theorists reject Instrumentalism. Since this story is simpler and more explanatory than pluralism about fundamental epistemic value and avoids other problems for Instrumentalism, it is *ceteris paribus* preferable.

Moreover, there are deeper reasons to reject Instrumentalism, as we saw in the last chapter. Ironically, Instrumentalism makes it impossible to understand the epistemic value of properties and states that *are* instrumentally connected to true belief. I will now proceed to briefly remind the reader of the case for this point.

#### 5.4 Reminder: Independent Reasons to Reject Instrumentalism

Why does Instrumentalism leave us generally ill-equipped to explain derivative epistemic value? Because, as I argued in the last chapter, Instrumentalism is responsible for a generalized version the swamping problem.

The problem, remember, was first raised as a problem for a naïve reliabilism. It begins with the observation that

---

<sup>8</sup>See, e.g., Brogaard (2009), DePaul (2001), Kvanvig (2005), and Zagzebski (2004).

(\*) a belief's having been produced by a reliable belief-forming process does not as such make that belief epistemically better if that belief is already true.

(\*) is made plausible by an analogy. The mere fact that some good coffee came from a reliable coffeemaker does not make that coffee better. Given (\*), it follows that:

If knowledge = true belief produced by a reliable type of belief-forming process, then knowledge is not as such epistemically better than true belief.

The problem does not end here. It arises for states other than knowledge and for theorists other than reliabilists. Reliabilists will identify a belief's being justified with its being formed by a reliable type of belief-forming process. We can use the same coffee analogy to argue that justification cannot as such add epistemic value to true belief. That is also a bad conclusion. But this is not a restatement of the problem about knowledge. Knowledge is not justified true belief.

The problem is more general than this. On any view on which being justified : true belief :: being made by a good coffeemaker : good coffee, there is a worry that justification cannot as such add epistemic value to true belief. So there seems to be a worry for any view on which the epistemic value of justified belief consists in its being the mere product of something that is only instrumentally epistemically good. It is natural to suspect that Instrumentalism is the cause of the problem.

Of course, one might blame this problem on Veritism rather than Instrumentalism. But this would be a mistake, as we saw in the last chapter. For there is another side to the coffeemaker analogy. We do not only think that coming from a reliable coffeemaker cannot improve good cups of coffee. We also think that coming from a reliable coffeemaker cannot improve bad cups of coffee. Now, suppose for the sake of argument that knowledge is the fundamental epistemic value. How could one explain the epistemic value of justified beliefs? One must claim that such beliefs are good in virtue of some instrumental connection to knowledge. But then it will be puzzling why justified false beliefs are epistemically good.

Even if one regards justified belief, knowledge, and understanding as fundamental epistemic values, there will remain epistemic values that are clearly derivative but admit



of no Instrumentalist explanation. Consider, again, the epistemic value of trying one's best to form beliefs accurately. This is a paradigmatically derivative value. But it is not reliably instrumental to any plausible FEVs. There is something admirable in one's best efforts, even if they always fail. This suggests that any reasonably pluralist view that endorses Instrumentalism will face a relative of the swamping problem. So, as we saw, Veritism isn't the cause of the problem. Instrumentalism is the cause of the problem.

This provides independent reason to think that rejecting Instrumentalism is a reasonable response to our puzzle about the value of rational belief. This is important for two reasons. First of all, without this independent argument, it would be unclear why we should not simply embrace a pluralist view that adds rationality to the list of fundamental epistemic values and deny Derivativeness. But as we have seen, the very same argument that recommends adding rational belief to the list of fundamental epistemic values also supports inflating that list in ways that no one should tolerate.

Another reason why the independent argument is important is that it undermines a different kind of response to our puzzle. Someone who is convinced of Derivativeness might simply deny on Instrumentalist grounds that rational beliefs really do retain worth in worlds like the demon world. And they might proceed to try to tell an error theory, perhaps along the lines of Kolodny (2005). This move comes with a cost, of course. Retention of Worth is extremely intuitively plausible. Indeed, it seems very clear that there is something very bad from the epistemic point of view about irrational beliefs even in a world like the demon world. By seeing that there are quite general reasons to reject Instrumentalism, we see that there is no reason to accept this cost.

## 5.5 The Value of Rational Belief Depuzzled

Of course, without a model of how epistemic value derivation works other than the Instrumentalist model, we are left with a puzzle. But I showed in the last chapter how we can provide a unified story about derivative epistemic value that avoids swamping problems. As we will now see more explicitly, this framework also reveals how we can give the same account of the value of rational belief that we can and should give of

other epistemically significant kinds of belief. There is no need for a new story: we already needed a different story than the Instrumentalist story, and that alternative story applies with no difficulty to the case of rational belief. This solves our puzzle.

Recall our alternative principle for grounding derivative value:

**Hurka's Principle:** If V is a fundamental value, ways of placing value on V and their manifestations have derivative value (relative to V).

Value derived via Hurka's Principle is, as we saw in the last chapter, immune to swamp-ing. As a reminder: suppose Gamma performs an act of charity because she values charity, while Epsilon performs the same kind of action as a PR stunt. Gamma's action seems better than Epsilon's. Yet just as appreciating beauty is good because beauty is good, so valuing charity and manifesting this valuing is good because charity is good. Here an act derives value from a source that is only derivatively good relative to another feature the act has. The source is (i) Gamma's valuing of charity. The property exemplified is (ii) charity. The fact that Gamma's action has its source in (i) makes that action worthier, though it has property (ii).

How do we apply Hurka's Principle in the epistemic domain? The view defended in the last chapter provides one illustration: the beliefs we admire from the epistemic point of view as manifestations of *ways to place value on accuracy in thought*, and are derivatively valuable relative to accuracy *via* Hurka's principle. What are these ways to place value on accuracy? They are, I suggested earlier, different ways of honoring the following norm of accuracy:

(AI) It is correct to believe P iff P is true.

Now, there are four increasingly demanding ways to honor any norm N:

**commitment** to N is a disposition to  $\phi$  only if  $\phi$ -ing seems relative to one's beliefs to confirm to N;

**weak respect** for N is a disposition to  $\phi$  only if  $\phi$ -ing seems, relative to how things appear in a non-doxastic sense, to conform to N;

**strong respect** for N is a disposition to  $\phi$  only if  $\phi$ -ing is objectively likely, conditional on how things appear in a non-doxastic sense, to conform to N;

to **comply** with a norm N is to conform with N by respecting it in the strong sense.

The following connections are, I argued in the last chapter, plausible:

one is coherent iff one manifests full commitment to AI;

one's belief is substantively rational iff this belief manifests weak respect for AI;

one's belief is justified iff this belief manifests strong respect for AI;

one's belief is knowledge iff one complies with AI in holding this belief.

There is unity here: rational belief, justified belief, and knowledge all manifest different ways of placing value on accuracy in thought. Given this unity in the nature of these kinds of belief, we can provide a unified explanation of their epistemic value. We can offer the following kind of story, which I illustrate with rational belief:

1. Rational beliefs manifest one way to place value on accuracy in thought.
2. Accuracy has fundamental value in the epistemic domain.
3. Manifestations of ways of placing value on fundamental values have derivative value. (From Hurka's Principle)
4. So, rational beliefs have derivative value in the epistemic domain.

The same kind of story can (and should) be told about the epistemic value of justified belief and knowledge. We just appeal to more demanding ways of placing value on accuracy. This is an important result: *it shows that no new story is needed to explain the epistemic worth of rational beliefs.*

Although I am a Veritist, my explanatory strategy could be exploited by proponents of other fundamental epistemic values. If one prefers a knowledge-first approach, one could appeal to the following fundamental ideal rather than (AI):

(KI) It is correct for one to believe P iff one knows that P.

The structure of the explanations remains the same. We can distinguish between commitment to, strong and weak respect for, and compliance with KI. The first three ways of honoring KI would correspond, as before, to structural rationality (or coherence), substantive rationality, and justification.

Since knowledge is the fundamental epistemic value on this picture and it thus does not have derivative epistemic value, we would not claim that knowledge is valuable in virtue of the fact that it constitutes compliance with KI. Indeed, that would not be plausible anyway unless one embraced the KK principle. But there is a distinct, derivative epistemic value that we could claim corresponds to compliance rather than mere conformity with KI with this version of the framework: *reflective knowledge*. I take it to be a virtue of my explanatory strategy that, when combined with a knowledge-first epistemic axiology, we get an explanation of why reflective knowledge is better than mere first-order knowledge. This is not something well-explained by existing frameworks. Indeed, there is a familiar worry that Ernest Sosa's account of the distinction between first-order knowledge and reflective knowledge provides no satisfying story about why reflective knowledge is epistemically better.<sup>9</sup>

Anyhow, what matters is that my conclusion is robust. Take either of the main candidates for the fundamental epistemic value, and we can provide a unified story about derivative epistemic value that solve the puzzle about the value of rational belief, revealing that no new story is needed to explain its value.

## 5.6 Implications for Epistemology

If sound, the overarching argument in this chapter has some important implications for epistemology. The most important ones affect some common thoughts about the “truth connection” in epistemology and the advantages of reliabilist views over non-reliabilist views. Reliabilist forms of externalism in epistemology are often thought to have a unique ability to give truth-oriented grounds for epistemic evaluation. Indeed, the

---

<sup>9</sup>See Fumerton (2004) for this complaint.

ability to secure a truth-oriented picture of epistemic evaluation is sometimes understood as the defining feature of externalist views. Consider Ernest Sosa:

Central concepts of epistemology—knowledge, for example, or epistemic justification or intellectual virtue—derive in one way or another from that of truth. *Assigning such a central role to truth makes one an externalist in epistemology.*<sup>10</sup>

It is commonly thought that this advantage yields a further unique ability: externalists can explain why we admire the beliefs we admire from the epistemic point of view by citing the instrumental connection between epistemic normativity and truth. Richard Fumerton—himself no externalist—expresses the thought nicely:

The reason epistemologists want justified beliefs has something to do with having true beliefs. At the same time, we must understand justification in such a way that we allow the possibility of justified false belief. How do we establish a connection between justification and truth without making it impossible to have a justified false belief? The [reliabilist's] answer is to focus on the processes that produce beliefs. [According to the reliabilist], [t]he beliefs that are “good” from the epistemic perspective are those that are produced by reliable belief-forming mechanisms.<sup>11</sup>

The argument behind this thought is simple:

### **The Truth Promotion Argument (TPA)**

1. Justified beliefs are good but not fundamentally good from the epistemic point of view.
2. (1) is true only if having justified beliefs is objectively conducive to having true beliefs.
3. So, having justified beliefs is objectively conducive to having true beliefs.
4. But only externalists can accept (3).
5. So, only externalists can explain how (1) is true.

---

<sup>10</sup>Sosa (2001: 641).

<sup>11</sup>Fumerton (1995: 97).

TPA obviously depends on Instrumentalism, since (2) is questionable without it.

Curiously, the externalists who favor TPA will agree that there is some epistemically admirable kind of belief that is not connected to truth in this way. They will just give it a different name than “justified belief”, like “rational belief”. The question for them is simple: why do beliefs of this kind retain genuine epistemic worth in worlds like Beta’s?

This question admits of no optimistic answer if Instrumentalism is true. The externalists who advance TPA might conclude that rational beliefs are not epistemically good *as such*. They might claim that they are contingently epistemically good at best and epistemically worthless in worlds like Beta’s. Yet once we realize that Instrumentalism cannot even explain why beliefs that *are* instrumentally related to truth are epistemically good, this becomes an unattractive route to take. It would be better to find an alternative model that provides a unified explanation of the value of rational belief, justified belief in the externalist’s sense, and knowledge.

The upshot is that TPA collapses. (2) is false. There are truth-oriented explanations of the value of justified belief and rational belief that do not appeal to the idea of instrumental conducivity to true belief. Indeed, my truth-oriented account did not appeal to instrumental relations at all in explaining the epistemic value of any epistemically admirable kinds of belief.

Once we appreciate this, the terrain shifts. There is no good argument from the epistemic value of justified belief to the truth of reliabilism or any kind of externalism about justified belief. The only argument that could privilege reliabilism here would assume Instrumentalism. But Instrumentalism makes it harder to explain the value of justified belief even if we agree that justified belief is necessarily a product of truth-conducive sources. It doesn’t make it easier.

Indeed, reliabilists are not the only ones who can vindicate the idea that we value justified belief from the epistemic point of view *because* we value accuracy from the epistemic point of view. Internalists who equate justified belief with rational belief can vindicate that idea too. Obviously, none of this is to say that we should become internalists about justification, or that there are no good arguments for externalism. There are many good arguments for externalism about justification. But they have

nothing to do with the idea that epistemic value is about instrumentally promoting true belief and demoting false belief. This cannot be what divides internalists and externalists. Internalists and externalists alike should reject the idea.

## Bibliography

- Alston, W. 1988. "An Internalist Externalism." *Synthese* 74: 265–83.
- Anderson, E. 1995. *Value in Ethics and Economics*. Harvard UP.
- Armstrong, D. M. 1973. *Belief, Truth, and Knowledge*. Cambridge: Cambridge University Press.
- Bach, K. 1985. "A Rationale for Reliabilism." *The Monist* 68: 246–63.
- Baehr, J. 2012. *The Inquiring Mind*. Oxford: OUP.
- Bergmann, M. 2006. *Justification without Awareness*. Oxford: OUP.
- Berker, S. 2013. "Epistemic Teleology and the Separateness of Propositions." *Phil. Review* 122: 337–93.
- Bernecker, S. 2007. "Remembering without Knowing". *Australasian Journal of Philosophy* 85: 137–156.
- BonJour, L. 1985. *The Structure of Empirical Knowledge*. Cambridge MA: Harvard UP.
- Brogaard, B. 2009. "The Trivial Argument for Epistemic Value Pluralism" in Haddock, A., Millar, A. and Pritchard, D. (eds.) *Epistemic Value*. Oxford: OUP.
- Broome, J. 1999. "Normative Requirements." *Ratio* 12: 398–419.
- Broome, J. 2005. "Does Rationality Give Us Reasons?" *Philosophical Issues* 15: 321–337.
- Broome, J. 2007. "Does Rationality Consist in Correctly Responding to Reasons?" *Journal of Moral Philosophy* 4: 349–374.
- Broome, J. 2013. *Rationality through Reasoning*. Oxford: Blackwell.
- Chisholm, R. M. "The Problem of the Speckled Hen." *Mind* 51: 368–373.
- Chisholm, R. M. 1957. *Perceiving: A Philosophical Study*. Ithaca: Cornell University Press.
- Christensen, D. 2010. "Higher-Order Evidence." *PPR* 81: 185–215.
- Clark, M. 1963. "Knowledge and Grounds." *Analysis* 24: 46–48.
- Comesaña, J. 2010. "Evidentialist Reliabilism." *Nous* 44.4: 571–600.
- Cohen, S. 1984. "Justification and Truth." *Phil. Studies* 46.3: 279–95.
- Conee, E. and Feldman, R. 2004. *Evidentialism: Essays in Epistemology*. Oxford: OUP.
- Dancy, J. 1995. "Supervenience, Virtues, and Consequences." *Phil. Studies* 78: 189–205.
- Dancy, J. 2000. *Practical Reality*. Oxford: Oxford University Press.
- Dancy, J. 2004. *Ethics without Principles*. Oxford: Oxford University Press.
- Darwall, S. 1977. "Two Kinds of Respect." *Ethics* 88: 36–49.
- DePaul, M. 2001. "Value Monism in Epistemology" in Steup, M. (ed.) *Knowledge, Truth, and Duty*. Oxford: OUP.
- Dogramaci, S. Forthcoming. "Communist Conventions for Deductive Reasoning." *Nous*.
- Dougherty, T. 2011. "Reducing Responsibility." *European J. Phil.* 18: 1–14.



- Engel, M. 1992. "Personal and Doxastic Justification in Epistemology." *Philosophical Studies* 67: 133–50.
- Feldman, F. 2000. "Basic Intrinsic Value." *Phil. Studies* 99: 319–46.
- Feldman, R. 2005. "Justification Is Internal" in *Contemporary Debates in Epistemology*.
- Fodor, J. 1998. *Concepts: Where Cognitive Science Went Wrong*. Oxford: OUP.
- Foley, R. 2004. "A Trial Separation between the Theory of Knowledge and the Theory of Justified Belief" in Greco, J. (ed.) *Ernest Sosa and His Critics*. Oxford: Blackwell.
- Fumerton, R. 1995. *Metaepistemology and Skepticism*. Lanham: Rowman & Littlefield.
- Fumerton, R. 2004. "Achieving Epistemic Ascent" in Greco, J. (ed.) *Ernest Sosa and His Critics*. Oxford: Blackwell.
- Gibbons, J. 2006. "Access Externalism." *Mind* 115: 19–39.
- Goldman, A. I. 1976. "Discrimination and Perceptual Knowledge." *Journal of Philosophy* 73: 771–91.
- Goldman, A. I. 1979. "What Is Justified Belief?" in Pappas, G. and Swain, M. (eds.) *Justification and Knowledge*. Dordrecht: D. Reidel.
- Goldman, A. I. 1986. *Epistemology and Cognition*. Cambridge MA: Harvard UP.
- Goldman, A. I. 1988. "Strong and Weak Justification." *Philosophical Perspectives* 2: 51–69.
- Goldman, A. I. 1999a. "Internalism Exposed." *Journal of Philosophy* 96: 271–93.
- Goldman, A. I. 1999b. *Knowledge in a Social World*. Oxford: Oxford University Press.
- Goldman, A. I. 2002. *Pathways to Knowledge*. Oxford: Oxford University Press.
- Goldman, A. I. 2009. "Internalism, Externalism, and the Architecture of Justification." *Journal of Philosophy* 106: 309–38.
- Goldman, A. and Olsson, E. 2009. "Reliabilism and the Value of Knowledge" in Haddock, A., Millar, A., and Pritchard, D. (eds.) *Epistemic Value*. Oxford: OUP.
- Goldman, A. I. 2011. "Toward a Synthesis of Reliabilism and Evidentialism?" in Dougherty, T. (ed.) *Evidentialism and its Discontents*. Oxford: Oxford University Press.
- Goldman, A. I. 2012. *Reliabilism and Contemporary Epistemology*. Oxford: OUP.
- Greco, J. 1999. "Agent Reliabilism." *Phil. Perspectives* 13: 276–96.
- Greco, J. 2004. "How to Preserve Your Virtue While Losing Your Perspective" in *Ernest Sosa and His Critics*. Oxford: Blackwell.
- Greco, J. 2005. "Justification Is Not Internal" in *Contemporary Debates in Epistemology*.
- Greco, J. 2010. *Achieving Knowledge*. Cambridge: CUP.
- Harman, G. 1973. *Thought*. Princeton: Princeton University Press.
- Hart, H. L. A. 1968. *Punishment and Responsibility*. Oxford: Clarendon Press.
- Herman, B. 1993. *The Practice of Moral Judgment*. Cambridge MA: Harvard UP.
- Huemer, M. 2001. *Skepticism and the Veil of Perception*. Lanham: Rowman & Littlefield.
- Hurka, T. 2001. *Virtue, Vice and Value*. Oxford: OUP.
- Hurka, T. 2006. "Virtuous Act, Virtuous Dispositions." *Analysis* 66: 69–76.
- Jackson, A. 2011. "Appearances, Rationality and Justified Belief." *PPR* 82: 564–93.
- Jones, W. 1997. "Why Do We Value Knowledge?" *American Phil. Quarterly* 34: 423–439.
- Kelly, T. 2006. "Evidence." *SEP* (Fall 2008 Edition).

- Kolodny, N. 2005. "Why Be Rational?" *Mind* 114: 509–560.
- Korsgaard, C. 1983. "Two Distinctions in Goodness." *Phil. Review* 92: 169–195.
- Lenman, J. 2009. "Reasons for Action: Justification and Explanation." *SEP* (Winter 2011 Edition).
- Littlejohn, C. 2009. "The Externalist's Demon." *Canadian Journal of Philosophy* 39: 399–434.
- Littlejohn, C. 2012. Justification and the Truth Connection. Cambridge: CUP.
- Lord, E. 2008. "Dancy on Acting for the Right Reason." *Journal of Ethics and Social Philosophy*. <http://www.jesp.org/articles/download/ErrolLordDiscussion.pdf>
- Lord, E. 2010. "Having Reasons and the Factoring Account." *Phil. Studies* 149: 283–296.
- Lyons, J. C. 2009a. Perception and Basic Beliefs. Oxford: Oxford University Press.
- Lyons, J. C. 2009b. "Perception and Virtue Reliabilism." *Acta Analytica* 24: 249–261.
- McDowell, J. 1995. "Knowledge and the Internal." *Philosophy and Phenomenological Research* 55: 877–93.
- McDowell, J. 2002. "Replies" in Smith, N. (ed.) Reading McDowell. London: Routledge.
- Millar, A. 1993. Reasons and Experience. Oxford: Clarendon Press.
- Montmarquet, J. 1987. "Epistemic Virtue." *Mind* 96: 482–497.
- Moore, G. E. 1903. Principia Ethica. Cambridge: Cambridge University Press.
- Parfit, D. 2001. "Rationality and Reasons" in Dan Egonsson, Bjorn Petterson and Toni Ronnow-Rasmussen (eds.) Exploring Practical Philosophy. Aldershot: Ashgate.
- Parfit, D. 2011. On What Matters. Oxford: Oxford University Press.
- Plantinga, A. 1993. Warrant and Proper Function. Oxford: OUP.
- Pritchard, D. 2010. The Nature and Value of Knowledge. Oxford: OUP.
- Pritchard, D. 2011. "What Is the Swamping Problem?" in Reisner, A. and Steglich-Petersen, A. (eds.) Reasons for Belief. Cambridge: Cambridge University Press.
- Pritchard, D. 2012. Epistemological Disjunctivism. Oxford: Oxford University Press.
- Pryor, J. 2000. "The Skeptic and the Dogmatist." *Nous* 34: 517–49.
- Rabinowicz, W. and Ronnow-Rasmussen, T. 2000. "A Distinction in Value: Intrinsic and For Its Own Sake." *Proceedings of the Aristotelian Society* 100: 33–51.
- Rashdall, H. 1907. The Theory of Good and Evil. Oxford: Clarendon Press.
- Raz, J. 1999. "Explaining Normativity." *Ratio* 354–79.
- Ross, J. 2006. Acceptance and Practical Reason. Ph.D. Thesis, Rutgers University.
- Ross, J. 2012. "Rationality, Normativity, and Commitment" in Oxford Studies in Metaethics 7. Oxford: Oxford University Press.
- Ross, W. D. The Right and the Good. Oxford: Oxford University Press.
- Scanlon, T. M. 1998. What We Owe to Each Other. Cambridge MA: Harvard University Press.
- Schechter, J. 2013. "Rational Self-Doubt and the Failure of Closure." *Phil. Studies* 163: 429–452.
- Schroeder, M. 2007. Slaves of the Passions. Oxford: Oxford University Press.
- Schroeder, M. 2008. "Having Reasons." *Philosophical Studies* 139: 57–71.
- Schroeder, M. 2009. "Means-End Coherence, Stringency, and Subjective Reasons." *Philosophical Studies* 143: 223–248.

- Schroeder, M. 2011. "What Is It To 'Have' A Reason?" in Reisner, A. and Steglich-Petersen, A. (eds.) *Reasons for Belief*. Cambridge: Cambridge University Press.
- Schroeder, M. 2012. "Stakes, Withholding, and Pragmatic Encroachment on Knowledge." *Philosophical Studies* 160: 265-285.
- Schroeder, M. Ms. "The Negative Existential Reason Fallacy." Unpublished MS, USC.
- Siegel, S. 2012. "Cognitive Penetration and Perceptual Justification." *Nous* 46: 201-222.
- Smithies, D. 2012. "Moore's Paradox and the Accessibility of Justification." *PPR* 85: 273-300.
- Sosa, E. 1991. *Knowledge in Perspective*. Cambridge: Cambridge University Press.
- Sosa, E. 1993. "Proper Functionalism and Virtue Epistemology." *Nous* 27: 51-65.
- Sosa, E. 2001. "Epistemology and Primitive Truth" in Lynch, M. (ed.) *The Nature of Truth*. Cambridge, MA: The MIT Press.
- Sosa, E. 2003. "The Place of Truth in Epistemology." in DePaul, M. and Zagzebski, L. (eds.) *Intellectual Virtue*. Oxford: Oxford University Press.
- Sosa, E. and Bonjour, L. 2003. *Epistemic Justification*. Oxford: Blackwell.
- Sosa, E. 2007. *Apt Belief and Reflective Knowledge, Volume I*. Oxford: OUP.
- Sosa, E. 2010. *Knowing Full Well*. Princeton: Princeton University Press.
- Sosa, E. 2011. "How Competence Matters in Epistemology." *Philosophical Perspectives* 24: 465-475.
- Swain, M. 1988. "Alston's Internalistic Externalism." *Philosophical Perspectives* 2: 461-73.
- Turri, J. 2009. "The Ontology of Epistemic Reasons." *Nous* 43: 490-512.
- Turri, J. 2010. "On the Relationship between Propositional and Doxastic Justification." *Philosophy and Phenomenological Research* 80.2: 312-26.
- Way, Jonathan. 2009. "Two Accounts of the Normativity of Rationality." *JESP* 4.2: 1-8.
- Wedgwood, Ralph. 2003. "Choosing Rationally and Choosing Correctly" in Stroud, S. and Tappolet, C. *Weakness of Will and Practical Irrationality*. Oxford: Oxford University Press.
- Williams, B. 1981. *Moral Luck*. Cambridge: Cambridge University Press.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Wright, C. 2004. "On Epistemic Entitlement." *Aristotelian Society Supplementary Volume* 78.1: 167-212.
- Zagzebski, L. 1999. "From Reliabilism to Virtue Epistemology." *Proceedings of the Twentieth World Congress of Philosophy* (Boston, 1998) v.: Epistemology, ed. R Coff-Steven. Bowling Green, OH: Philosophy Documentation Center.
- Zagzebski, L. 2004. "Epistemic Value Monism" in Greco, J. (ed.) *Ernest Sosa and His Critics*. Oxford: Blackwell.
- Zimmerman, M. 1997. "A Plea for Accuses." *American Phil. Quarterly* 34: 229-243.
- Zimmerman, M. 2001. *The Nature of Intrinsic Value*. Rowman & Littlefield.